

The Mind, The Brain, and Systemic Functional Linguistics

James Williams, B.A (hons), M.Phil

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy

November 2015

Abstract

Modern cognitive science is characterized by a number of different proposals that aim to provide a means of conceptualizing cognitive systems, and stating the kinds of information that are needed to study the phenomenon from a number of different perspectives. Recently, there has been a growing trend to think about cognitive systems in terms of *embodiment* and *embeddedness*. These two factors emphasize, respectively, the role the body and the environment in shaping the nature of cognitive processes. However, what these theses currently lack is a well-articulated *conceptual framework* for guiding their work. Specifically, it lacks an explicit account of the types of *levels of analysis* that are needed to study cognition from these perspectives. This explicitness is a feature of alternative conceptions of the mind, the classical approaches and connectionist approaches that despite their differences, both draw upon the notions of *computation* and *representation* in explaining cognitive life. Where these proposals fall short is their lack of focus of the kinds of factors that motivate the *embedded* and *embodied* view.

This thesis aims to fill this gap in the literature by beginning to develop an explicit *conceptual framework* for embodied and embedded cognition. It is proposed that such a framework can be based upon the principles of Systemic Functional Linguistics (SFL), a theory of language that considers communication in its eco-social context. Specifically, it proposes a number of *levels of analysis* that are useful for such a framework. Motivation for such a proposal is that: (1) both language and general cognition are **semiotic systems**, (2) they can both be described as **dynamic, open systems**, and (3) all the types of *levels of analysis* previously identified in cognitive science are also found in the SFL literature. This thesis, therefore, investigates whether such a cross-disciplinary application has the potential to be successful.

In order to achieve the above goals, this thesis examines the following possible applications of notions from SFL to conceptual frameworks in cognitive science.

Firstly, this thesis examines the theoretical notions of *stratification*, *instantiation*, *delicacy* and *rank* in SFL and then continues to outline their application to this neighbouring field. These are the *levels of analysis* indicated above. *Stratification* allows cognitive systems to be considered simultaneously at a number of different *levels of interpretation*. Specifically, it considers cognitive systems to consist of a **neurophysiological system** and a **psychological system**, i.e. a system of possible neurological states, and a system of possible cognitive acts. *Delicacy* describes these systems in more or less specific terms. *Rank* identifies the structural unit at which meaningful behaviour is taking place, and *instantiation* describes the interplay between observed instances and the underlying system. This last factor enables a **dynamic** view of cognition to be proposed.

Furthermore, this thesis examines whether ideas from SFL concerning text-context interactions can prove useful for considering organism-environment interactions. The main proposal here is that the environment need not be considered as a vast, heterogeneous entity, but instead is itself a system involved in the creation of meaning. In this way, it can also be described along the dimensions discussed above, although at this stage only the notions of *stratification* and *instantiation* can be discussed in any depth. Also, it proposes that the environment can be described according to three *environmental parameters*: (1) the *objective*, that describes the objects and relations between them, (2) the *relational,* which describes relations between the individual and said objects, and (3) the *goal*, which provides the situation the individual is attempting to bring about. In this way structure can be brought to studies of cognitive *embeddedness*.

Finally, this work considers the notion of *abstraction* in SFL as it relates to the dimensions of *stratification*, *delicacy* and *rank*. All of these have been described as operating over this relation, but this thesis argues that its meaning is different in each case. It argues this by considering the links between *abstraction* and the related notions of *omission*, *generalization* and *decontextualization*. In line with the overall goal of this work, it is proposed that a more explicit understanding of these terms can help the process of beginning to build the framework under discussion, by linking data to the overall framework.

In sum, this thesis begins to develop a novel framework for cognitive science that is based upon the principles of Systemic Functional Linguistics. It is hoped that this work can bring some extra clarity to the notions of *embeddedness* and *embodiment* in cognitive science and show how cognitive systems may be studied from this perspective.

Acknowledgements

Undertaking the task of writing a PhD thesis is a daunting one, and one that cannot be completed alone. This is especially true for work that attempts to cross discipline boundaries, as this one does. For this reason, there are a large number of people who, in one way or another, contributed to this work. They all deserve my heartfelt thanks.

Firstly, I have been extremely fortunate to have the support of two extremely patient and thoughtful supervisors. It was them who thought up the initial idea for the doctoral thesis, and I hope the work I have attempted at least resembles their initial hopes for the project.

Thanks to Dr Derek Irwin for constantly fitting me into an extremely busy schedule (preferably over beer!), for helping me through my first foray into SFL theory, for putting up with my niggling questions, and patiently listening to my half-thought through ideas.

Thanks to Dr Noah Russell for guiding me through the baffling world of neuroscience, and engaging with me on some of the philosophical issues in this work. We may not have always agreed, but I benefited immensely from his knowledge and guidance in coming up with the ideas put forward here.

I have a near inexpressible debt to the members of the Neurophotonics Lab (both past and present) at the UK campus: Alex Johnstone, Carmel Howe, Ben Paul, Solomon Idinyang, Chris Towlson, Sunil Rajput, Sebastiaan van Nuffel, Nitzan Herzog, Houda Sahaf, Hala Dhowre, Katharina Reusch, Rob Clay and Sid Visser. From the beginning they have been extremely welcoming to a student from a very different background to themselves and their patience in tolerating my often basic, and inevitably frustrating, questions about their subject matter has been a great help. I've learned so much from all of them, and I hope they learned at least a little from me.

At the Malaysian campus a big thanks goes to my fellow first cohort of English PhD students, especially Zhang Jing and Liu Ning for listening to my ramblings about the superiority of formal linguistics, and confusing tangents about neuroscience. Thanks as well to Dr Shivani Sivagurunathan, Dr Thomas Barker, Dr Ahmad Izadi and Dr Mohammad Safaei for all their support: both professionally and personally during my stays in Malaysia. I am also incredibly indebted to Dr Irwin's family, his wife Susan and children Georgia and Jamie, for welcoming me into their home when I first arrived. It made the move so much easier.

No-one gets into the position to try and complete a PhD without some inspiring teachers in the past. For this reason I would like to thank Dr Andrew Koontz-Garboden, Dr Martina Faller and Dr Wiebke Brockhaus-Grand at the University of Manchester for first installing in me a love of all things linguistic, and Professor Ian Roberts, Dr Napoleon Katsos and Dr Theresa Bierbauer at the University of Cambridge for imparting their knowledge and enthusiasm during my masters course. This may not be quite the thesis I thought I'd be doing, but hopefully it does your teaching justice.

To friends both in Nottingham and elsewhere for taking me away from it all (if only for a bit): Carmel and Alex for many nights forgetting our worries at the Johnson Arms, and David for silly pub quizzes and our never-ending quest to finally win one. To Alice for emotional support and Sean for oscillating between acting as a semi-consenting sounding board and a fellow gin enthusiast.

I owe more than I probably know to my parents, John and Anne for their emotional (and financial!) support during what has been, at times, a testing period, as well as my siblings Ben and Becky for reminding me there's more to life than writing a thesis.

And finally to Emma, who sacrificed the most.

Table of Contents

Abstract	i.
Acknowledgements	iii.
Table of Contents	iv.
List of Tables and Figures	ix.
Introduction	**1**
1.1. Conceptual Frameworks and cognitive science	4
1.1.1. Cognitive Science as a discipline.	4
1.1.2. Levels in Cognitive Science	5
1.2. Issues with current proposals	15
1.3. Motivation for using SFL theory as a basis for a cognitive framework	17
1.3.1. Cognitive Systems, meaning and SFL	17
1.3.2. Cognitive systems as open, dynamic systems	18
1.3.3. Richness of Levels of Analysis	20
1.4. Research Questions	22
1.5. Outline of this thesis	22
Chapter 2 – Previous conceptual frameworks in cognitive science	**26**
2.1. Conceptual frameworks and cognitive models	26
2.2. Philosophical Views and the 'Mind-Body Problem'	28
2.2.1. Supervenience	29
2.2.1.1. Reductionism and Mental-Physical identity	30
2.2.1.2. Supervenience: definitions	33
2.2.1.3. Supervenience and Neuroscience	36
2.2.2. Levels and Supervenience	37
2.2.3. Problems with Supervenience as an informative model	39
2.3. Classical Approaches	42

2.3.1. Representation 45

2.3.2. Computation 49

2.3.3. Multiple Realizability 54

2.3.4. What do you study on a classical view? 57

2.4. Connectionist Framework 58

2.4.1. Applications of the Connectionist Framework 63

2.4.2. Connectionism and the Brain 65

2.4.3. Studying Cognition in a Connectionist Framework 70

2.5. 4E approaches 71

2.5.1. Embedded Cognition 72

2.5.2. Embodied Cognition 76

2.5.3. Extended Cognition 82

2.5.4. Enacted Cognition 84

2.5.5. 4E approaches and the absenceof conceptual 87

2.6. Summary 89

Chapter 3 – Thoughts on a Stratificational model of cognitive and behavioural processes 97

3.1. Foundations of Systemic Functional Lingistics 98

3.1.1. Language as systemic 98

3.1.2. Language as function 101

3.2. Stratification 105

3.3. Delicacy 118

3.4. Rank 122

3.5. Instantiation 126

3.6. Relations between these dimensions 130

3.6.1. Rank and Delicacy 131

3.6.2. Stratification and Instantiation 133

3.7. The different types of levels in Systemic Functional Linguistics 135

3.8. Towards a stratificational framework of cognitive processes 139

3.8.1. Previous accounts of relations between SFL and cognitive science 139

3.8.2. Stratification 141

3.8.2.1. Towards a more general view of stratification 142

3.8.2.2. Towards identifying the strata 151

3.8.2.2.1. Neurophysiological stratum 152

3.8.2.2.2. Psychological stratum 154

3.8.2.3. Realization: comparison with previous uses 155

3.8.3. Delicacy 157

3.8.3.1. Delicacy at lower strata 158

3.8.3.2. Delicacy at the higher stratum 163

3.8.3.2.1. An example: memory 164

3.8.4. Rank 165

3.8.4.1. Rank at the neurophysiological stratum 166

3.8.4.2. Rank at the psychological stratum 168

3.8.5. Instantiation 169

3.9. Summary 173

Chapter 4 – Taking context seriously in cognitive science **176**

4.1. Introduction 176

4.2. Context in Systemic Functional Linguistics 177

4.2.1. The material situational setting versus the semiotic context 177

4.2.2. Contextual parameters: Field, Tenor and Mode 181

4.2.3. The status of context as a theoretical stratum 184

4.2.4. Stratification of the contextual stratum: Genre and Register 188

4.3. Organism-Environment relations through the lens of SFL 194

4.3.1. Material vs. Semiotic Cognitive Context 195

4.3.1.1. Evidence 1: Sensory Illusions 196

4.3.1.2. Evidence 2: Cross-species differences 197

4.3.2. Contextual Parameters in Cognitive Science 200

4.3.2.1. What to include? 200

4.3.2.2. Case Study 1: Attention Studies 202

4.3.2.3. Case Study 2: Goal-oriented studies 207

4.3.2.4. Environment Parameters: Objective, Relational and Goal 210

4.3.3. The contextual stratum in relation to the others 211

4.3.3.1. Theoretical dimensions at this level of interpretation 211

4.4. The story so far 215

Chapter 5 – Putting the ideas to use: Abstraction in Systemic Functional Linguistics 219

5.1. Introduction 219

5.2. Philosophy of Science and Systemic Functional Linguistics 223

5.3. Abstraction in the philosophy of science 226

5.3.1. Abstraction and “Omission” 227

5.3.2. Abstraction and Generalization 233

5.3.3. Abstraction and Decontextualization 238

5.4. Abstraction in Systemic Functional Linguistics 242

5.4.1. Delicacy 242

5.4.1.1. Omission? 243

5.4.1.2. Generalization? 248

5.4.1.2. Decontextualization? 250

5.4.2. Instantiation 252

5.4.2.1. Generalization? 254

5.4.2.2. Omission? 255

5.4.2.3. Decontextualization? 256

5.4.3. Stratification 258

5.4.3.1. Omission? 258

5.4.3.2. Generalization? 261

5.4.3.3. Decontextualization? 263

5.4.4. A proposed typology of abstraction types in SFL 265

5.4.5. A return to stratification 266

5.5. Conclusion 268

Chapter 6 – Conclusions and Looking Forward 270

6.1. Proposals in this thesis 270

6.1.1. Proposals in Theoretical SFL 271

6.1.2. Proposals in Cognitive Science 272

6.2. Putting it all together – Morris Water Maze Experiments 273

6.3. Future Work 277

6.3.1. Mathematical Formalization 279

6.3.2. Further steps towards the 'dynamic' nature of cognition 280

6.3.3. Case-Studies of particular phenomena 281

References **284**

List of Tables, Figures and Networks

Chapter 1

Figure 1.1: Marr's three-level framework 7

Figure 1.2: NETtalk connectionist architecture 9

Figure 1.3: Levels of interpretation in the connectionist framework 10

Figure 1.4: Supervenience 11

Figure 1.5: Levels of Analysis in Morton and Frith (1995) 14

Figure 1.6: Conceptual framework in an embedded, embodied approach 14

Chapter 2

Figure 2.1: Supervenience Outline 37

Figure 2.2: Marr's three-level framework 43

Figure 2.3: A simple feed-forward network 60

Figure 2.4: NETtalk model (Sejnowski and Rosenberg, 1987) 64

Figure 2.5: Deep learning neural architecture for natural language processing 65

Figure 2.6: Metalanguage and levels in interpretation in connectionist frameworks 66

Table 2.1: Relations between the neural, mathematical and conceptual levels in connectionism 67

Figure 2.7: Causal modelling framework 75

Figure 2.8: Three levels in an embodied, dynamic framework 80

Table 2.2.: Overview of conceptual frameworks in cognitive science 91

Chapter 3

Network 3.1: Simplified network showing paradigmatic relations for the MOOD system 100

Figure 3.1: The stratificational framework 106

Figure 3.2: The linguistic sign in structural linguistics 107

Figure 3.3: Content and Expression in the stratificational framework 108

Table 3.1: Realization as a top-down relationship 111

Figure 3.4: Stratification as levels of interpretation 112

Network 3.2: SPEECH FUNCTION system network 113

Network 3.3: Lexicogrammatical MOOD system 114

Figure 3.5: Realization: a blueprint 117

Network 3.4: Basic system network to one degree of delicacy 119

Network 3.5: System of MOOD at a low level of delicacy 119

Network 3.6: Basic system network to two degrees of delicacy 120

Network 3.7: System of MOOD with increased delicacy 120

Network 3.8: A simultaneous system network 121

Network 3.9: Consonantal potential 122

Figure 3.6: Rank based analysis for sentence (8) 123

Figure 3.7: The cline of instantiation 126

Figure 3.8: Variant of the system of FIGURE with equal probabilities 129

Figure 3.9: Illustrative example of the system of FIGURE in the sub-potential of cooking recipes 129

Network 3.10: System of grammatical MOOD 131

Network 3.11: System of THEME 132

Figure 3.10: The stratification-instantiation matrix 134

Table 3.2: The variety of levels in Systemic Functional Linguistics 138

Network 3.12: Intonation system in English 144

Network 3.13: Consonantal potential 146

Network 3.14: Lexicogrammatical MOOD system 147

Network 3.15: System of THEME 148

Network 3.16: System network of SPEECH FUNCTION 150

Figure 3.11: Spatial-Temporal Hypothesis of Stratification 151

Figure 3.12: The Action Potential 153

Figure 3.13: Oscillations in neural systems 159

Figure 3.14: Frequency classes for neural oscillations 160

Network 3.17: Oscillation network 162

Network 3.18: System of MEMORY types 165

Figure 3.15: Structural levels of organization in neural systems 166

Figure 3.16: Levels of Organization in spatial memory, based on Craver (2007: 166) 167

Figure 3.17: Cline of instantiation at the neurophysiological stratum 171

Figure 3.18: Cline of instantiation at the psychological stratum 173

Chapter 4

Figure 4.1: Matthiessen's systems of nature 179

Table 4.1: Differences between the material situational setting and the semiotic context 180

Figure 4.2: Instantiation/ Stratification matrix 185

Network 4.1: Field network (adapted from Bowcher, 2014) 186

Figure 4.3: Trajectory from use, through context, to language 187

Figure 4.4: Language, Register and Genre 190

Figure 4.5: Realization relations between genre, register and language 191

Figure 4.6: Mach band effect 196

Figure 4.7: Experimental set up in Demeter et al. (2010) 206

Figure 4.8: Cline of instantiation at the contextual stratum in SFL 213

Figure 4.9: Preliminary cline of instantiation at the stratum of semiotic environment 214

Figure 4.10: Cline of instantiation at the stratum of semiotic environment 214

Chapter 5

Figure 5.1: Model to predict the trajectory of a cannonball 229

Figure 5.2: The abstract concept of QUADRILATERAL formed by generalizing over experienced objects 235

Figure 5.3: Abstraction as generalization: a blueprint 236

Figure 5.4: Abstraction and the scale of delicacy 242

Network 5.1: System of MODALITY TYPE 244

Network 5.2: System network of MOOD 249

Figure 5.5: Abstraction and the cline of instantiation 252

Figure 5.6: Abstraction and the stratificational network 258

Figure 5.7: From text to stratification 260

Table 5.1: Abstraction types in SFL 265

Chapter 1

Introduction

> *"Owl," said Rabbit shortly, "you and I have brains. The others have fluff. If there is any thinking done in this Forest – and when I say thinking I mean thinking – you and I must do it."*
>
> *A.A.Milne. The House on Pooh Corner*

The aim of this thesis is to develop and defend a conceptual framework for cognitive science that utilizes notions previously developed in the context of Systemic Functional Linguistics (Halliday and Matthiessen, 1999; 2014). Specifically, it concerns itself with answering the question of **what and how many kinds of levels of analysis**[1] are necessary for cognitive science. Although modern cognitive science only emerged as a discipline in the mid-1950s (Miller, 2003), several different *conceptual frameworks* for studying cognitive systems have been developed, defended and ultimately criticized by others. As discussed below, a *conceptual framework* provides a means of conceptualizing the phenomenon under study, and states what kinds of facts are relevant to its understanding. Examples of such frameworks in cognitive science include the **symbolic** approach (Newell and Simon, 1976; Fodor, 1975; Marr, 1982; Pylyshyn, 1984), the **connectionist** approach (McCelland et al., 1986; Rumelhart et al., 1986), and a group of more recent proposals that have been grouped together as the **4e approaches** (Rowlands, 2010), which include the notions of **embedded**, **embodied**, **extended**, and **enacted** cognition. Ideally, such frameworks should achieve two aims: firstly, they should provide an account of how to conceptualise the phenomenon under study. The different approaches all present a different way of thinking about cognitive systems, often by means of a metaphorical analogy (Eliasmith, 2003). Secondly, they should

[1] Typographical note. This thesis will include terms and ideas from a number of different disciplines, and in some cases, the same term will be used differently in different subjects. To aid with understanding, I will employ *italics* when referencing a theoretical notion, and will give a definition, or point to where a definition will be given. In order to accommodate this choice, emphasis in the main text will be noted using **boldface**.

make explicit the type of information needed in order to understand cognitive systems. As will be seen, adopting a particular conceptual framework influences greatly the kinds of facts you include in your explanations. Crucial to both of these goals is the notion of *levels of analysis*. The exact nature of this notion will be discussed in the next section, but in general they specify the means of structuring these sources of information into one coherent picture.

The issue taken up by this thesis is that, although many different kinds of conceptual framework have been proposed in the literature, none of them have been universally accepted, and each of them has certain disadvantages. Some of these problems are **conceptual**, i.e. they pertain to the first aim of conceptual frameworks stated above. In particular, there has been increasing criticisms of the conceptual foundations of both the symbolic and connectionist frameworks from some quarters (Lakoff and Johnson, 1999; van Gelder, 1995; Thibault, 2004; Rowlands, 2009; Hutto and Myin, 2013; c.f. Lynott et al., 2014 for a recent view). The main criticism that has been raised against these approaches is their reliance upon the *computational metaphor*, the idea that we can properly understand cognitive systems only by treating them as computational systems (Marr, 1982). They argue that this favours a view of cognition that is detached from potentially important factors of cognitive life, including the role that the body, environment, and social systems all might have on the way that cognitive systems operate (c.f. Thagard, 2014: §5.2 for a list and discussion of these criticisms). However, the advantage of these approaches is that they are explicit about the number and type of *levels of analysis* that are needed for studying cognitive systems. Many of the criticisms raised against the frameworks just mentioned come from the developers of the various **4e approaches**, who attempt to incorporate these additional facts into their ideas. As will be seen in §2.5.5., however, the major issue with these approaches is that they are less concerned with giving precise articulations of their ideas in terms of *levels of analysis*. In fact, some researchers within these approaches seem to think that this omission is beneficial (Wilson and Clark, 2008: 55). In this respect, the failure of this approach is that they fail to satisfactorily achieve the second aim of conceptual frameworks.

This thesis aims to address this gap by beginning to develop and defend a novel conceptual framework for cognitive science. The strategy taken in this work is to investigate whether

the conceptual framework of language as described by the theory of *Systemic Functional Linguistics* (SFL) can provide the basis for a framework in cognitive science that can maintain many of the conceptual ideas associated with the 4E approaches, while providing the means of outlining the kinds of *levels of analysis* that are required to study cognition from a similar viewpoint. It might initially seem suspect that a framework for linguistics may be usefully applied to a framework for cognition more generally, but as will be discussed in §1.3, there are motivating reasons for thinking such a cross-disciplinary application plausible. These are: (1) the idea of **meaning** has played an important role in the discussion of cognitive systems. Some researchers think that cognitive systems are themselves **meaningful systems** (Fodor, 1975; Dretske, 1981; Millikan, 1984; Eliasmith, 2003; 2013). Given this fact, cognitive systems can be considered to be *semiotic systems*, in other words systems of meanings. The framework developed in SFL is not taken just to be a framework for language, but of semiotic systems more generally, and has already been applied to non-linguistic examples of semiotics (c.f. work in Mahoob and Knight, 2010). (2) Both cognitive systems and language can be described as *open, dynamic semiotic systems*. (3) The SFL framework includes a number of different *levels of analysis*, some of which have already been claimed to be important for considering cognitive systems.

In sum then, the proposal put forward in this thesis is that ideas and theoretical notions from SFL can be applied more generally to cognitive science. It is claimed that the conceptual framework that arises out of this application has advantages over some of the previous proposals in the literature. Specifically, this thesis argues that this application can identify a number of *levels of analysis* that are required to study cognition from a *semiotic*, *embedded* and *embodied* viewpoint.

The rest of this introductory chapter will further detail the major points of this thesis just discussed, followed by the research questions this thesis hopes to answer. The chapter will end with an outline of the rest of the thesis.

Much of the information in the following sections will be covered in more depth later in this thesis. However, it is worth overviewing some of the issues at this introductory stage, as they are directly relevant for the motivation of this thesis and the development of the research questions.

1.1. Conceptual frameworks and cognitive science

1.1.1. Cognitive Science as a discipline

This thesis concerns *conceptual frameworks* and *levels of analysis* in cognitive science, so it is worth initially briefly describing this discipline. Most concise definitions highlight two fundamental aspects: firstly, the explanatory aim of cognitive science is the *mind*, to be defined shortly, and secondly, the approach taken by researchers in this field is an interdisciplinary one (Frankish and Ramsey, 2012: 1; Friedenberg and Silverman, 2012: 2). With regards to its interdisciplinary stance, as reported by Miller (2003), the following disciplines can be identified as being able to contribute to the enterprise of cognitive science:

- Psychology
- Neuroscience
- Linguistics
- Artificial Intelligence
- Anthropology
- Philosophy

These six disciplines remain today the partners that aim to collaborate within the field, although, as we will see in §2, different frameworks in cognitive science foreground different members of this group in their work.

The subject matter of the enterprise is usually characterized as the study of the *mind*. However, giving a precise definition of what the *mind* actually **is** proves difficult. A look at the topics found in the collected edition by Frankish and Ramsey (2012), though, illuminates the range of phenomena that cognitive scientists seek to explain. These include: **perception**, **action**, **learning and memory**, **reasoning**, **decision making**, **language**, **emotion** and **consciousness**. In broad terms, therefore, we can describe the mind as the collection of cognitive abilities that are evident in humans, other organisms, and potentially artificially intelligent systems: the ability to perceive the world around us, the ability to act to bring about changes on that environment, the ability to learn and remember, the ability to

communicate (in humans the ability to possess and use language), and so on. This definition will not be exact enough to satisfy everyone, but it will suit our purposes here.

1.1.2. Levels in Cognitive Science

An additional concern of this thesis is the notion of *levels of analysis*. This theoretical term has attracted much attention in discussions about the foundations of cognitive science and the relationships between the disciplines mentioned above (Fodor, 1975; Marr, 1982; Pylyshyn, 1984; Smolensky, 1988; Horgan and Tienson, 1993; Bechtel, 1994; Craver, 2007; Marraffa and Paternoster, 2012). However, as has been noted by cognitive scientists, a precise characterization of the term is difficult to give. For example, in discussing the debate between symbolic and connectionist approaches to cognitive science (which will be discussed in detail in §2), Bechtel (1994: 2, emphasis added) says the following:

> *The use of the term **levels** in this discussion has not always been clear. If we are to appraise the contributions of connectionism to the development of scientific enquiry into cognition, we need greater clarity on how the term **levels** is used and what is entailed by locating connectionism **at a given level**.*

Craver (2007: 164), in his discussion on explanations in the brain, appreciates the important role that this idea plays in modern neuroscience, but notes that: "the term **levels** is multiply ambiguous […] Not surprisingly, then, the term *level* has several common uses in contemporary neuroscience.", while Opie (2010: 270) notes that in many cases "the relevant notion of *level* is notoriously hard to pin down."

Although this ambiguity exists, examples of the notion being used in modern academic practice are not difficult to find, since the notion of *level* is not unique to cognitive science. In fact, the whole of modern scientific and humanistic enquiry is often thought to operate at different *levels* with different disciplines standing at a *higher* or *lower level* than another. By this I mean the modern disciplines of *physics*, *chemistry*, *biology*, *psychology*, *sociology* and so on. A common view is that there is an ordering of theories, in which they are given a

precise order. For instance, consider the views of Sobert (1999: 543) concerning the ordering of scientific disciplines:

> *The distinction between higher and lower of course requires clarification, but it is meant to evoke a familiar hierarchical picture; it runs (top to bottom) as follows – the social sciences, individual psychology, biology, chemistry and physics*.

On this view then, we can talk about particular disciplines operating at a *higher* or *lower* ***level*** to another. For instance, *physics* is generally considered to be the lowest, or most fundamental of the sciences, with the social sciences and the humanities at the top (Oppenheim and Putnam, 1958).

However, this is just one use of the term, and as will be seen, there are a large number of different **kinds of levels** in operation in cognitive science (Wimsatt, 1974; Craver, 2007: 174-176). To begin our discussion of the topic, and to hopefully make the notion clearer, this introductory chapter will now briefly overview the different *conceptual frameworks* that are found in modern cognitive science, and the different sorts of *levels of analysis* that have been proposed in them. Hopefully by overviewing the work at this early stage, these two crucial notions can be elucidated further.

One of the most influential conceptual frameworks has been the *classical* or *symbolic* framework of cognitive systems. Specifically, with respect to the idea of *levels of analysis*, one of the main proposals comes from Marr (1982), given below:

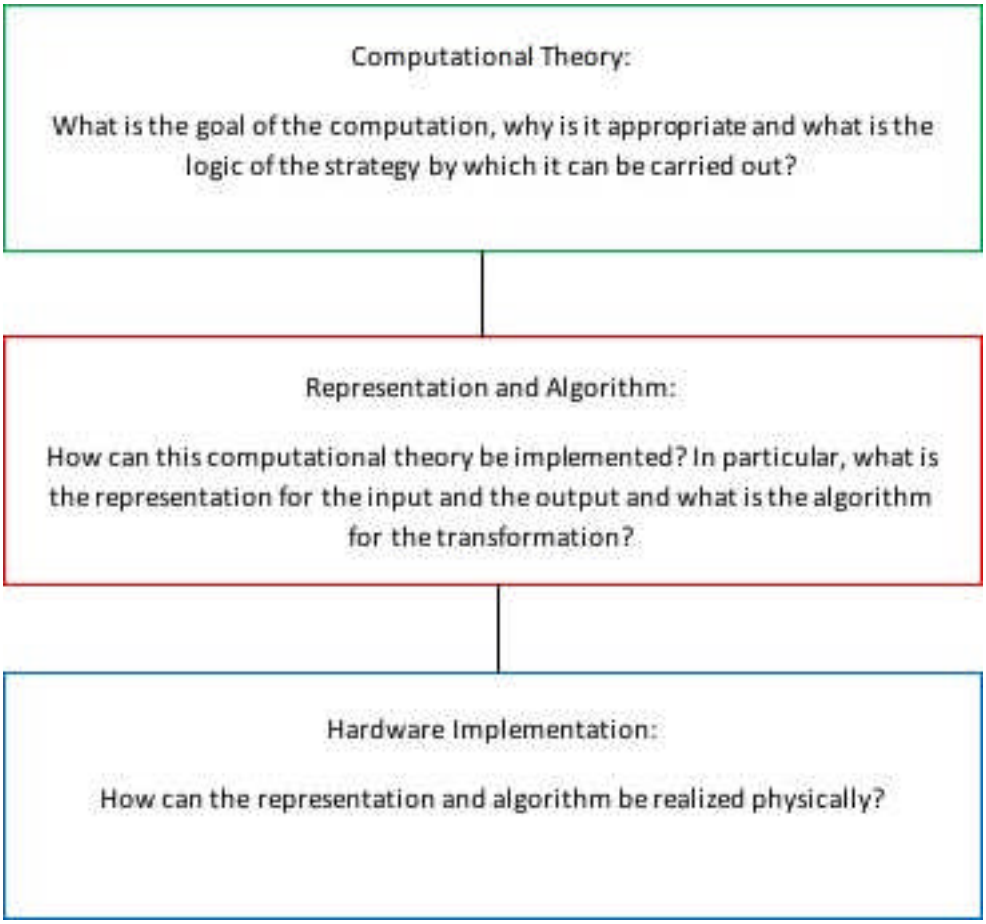


Figure 1.1: Marr's three-level framework

From the above picture, we can see the two aspects of conceptual frameworks mentioned above in operation. Firstly, the above figure provides us with a means of conceiving of cognitive systems. In this case, the guiding principle is that cognitive systems should be properly treated as a type of *computational* system, similar (but not identical) to other computational systems, such as mechanical computers. This can be seen by the kinds of questions the conceptual framework asks us to consider. The highest level poses the question of **what kind of computation the system is trying to achieve**? The secondlevel asks us **what representations and algorithms does the system need to possess in order to achieve that goal**? Finally, the lowest level enquires about **how those computations are realized in a physical system**.

In terms of the second facet of conceptual frameworks, it explicitly provides three *levels of analysis* at which cognitive systems can be studied. As discussed in §2.3, in this case we have an example of a series of ordered *levels of interpretation*. This involves each level in the framework representing a different perspective taken by the observer; there is only one

phenomenon under study, but it can be viewed in a number of different ways. On one view it is a computational system performing a certain task; on another it is a system manipulating a set of symbolic representations according to a specific set of rules, and on the last view it is a physical system in which signals are passed between certain units. Although this framework was first proposed in the 1970s – 1980s (Marr and Poggio, 1977; Marr 1982), it still remains an influential and widely debated framework today (see, Poggio, 2012; Neeleman, 2013; Eliasmith, 2015).

Other accounts within this framework point towards another type of *levels of analysis* being important. As can be seen in the above figure, one of the *levels of interpretation* requires that a certain set of *representations and algorithms* are specified. These representations can take the form of *symbolic structures*. One of the most prominent, and controversial, of these proposals is the *Language of Thought*, proposed by Fodor (1975; 2008). On this view, mental representations take the form of a language-like propositional logic. Important for this initial account is that these symbols can enter into larger structural units; symbolic propositions can themselves be built up from other propositions. It is therefore possible to consider a series of *levels of* organization, that is, a mereological notion that holds between parts and wholes, linking complex symbols to their constituents.

This idea of *levels of organization* is also important when one comes to consider work carried out under the **connectionist** framework. This approach took issue with some of the work carried out under the classical, symbolic framework of cognitive systems. Particularly, it was argued that these proposals were unduly detached from neurobiological reality, especially with regards to the serial, sequential processing strategies that they believed were entailed by such a view. Furthermore, they felt that the previous proposals were too concerned with 'passive' cognitive abilities such as the nature of thought and belief (McCelland, Rumelhart and Hinton, 1986). In order to counter-act what they saw as weaknesses in previous proposals, these researchers proposed the **parallel distributed processing** framework, now known as the **connectionist approach**. The main hallmark of this approach is the heavy use of *artificial neural networks* (ANNs); that is computer simulations using neurally inspired computational architectures, organized into a number of layers (c.f. Gurney, 2003). With training programs, these ANNs can adapt and learn to complete a

number of cognitive-related tasks. An example of such a computational architecture is given below. Below is the NETtalk model: an early connectionist simulation that translated written to spoken language (Sejnowski and Rosenberg, 1987):

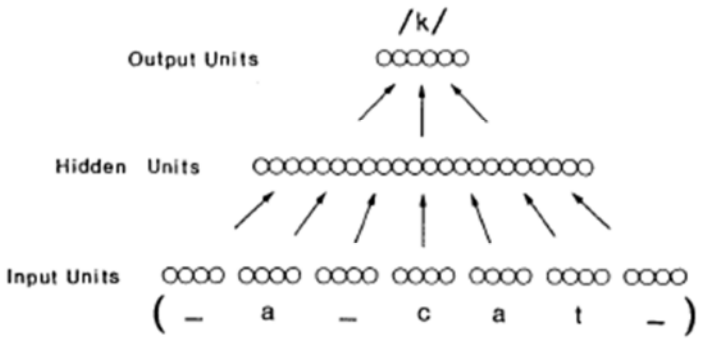


Figure 1.2: NETtalk connectionist architecture

In this way, connectionists attempt to understand cognitive systems by designing and training ANNs that can emulate certain cognitive abilities, for example face recognition or language processing (c.f. §2.4.1.). By doing so, they hope to hypothesise how such tasks are carried out in biological neural systems.

In terms of *levels of analysis*, the same types that were just seen in the symbolic framework are also in evidence in connectionist proposals. In terms of *levels of interpretation*, Smolensky (1986) states that computational simulations created can be interpreted in two different ways, as follows:

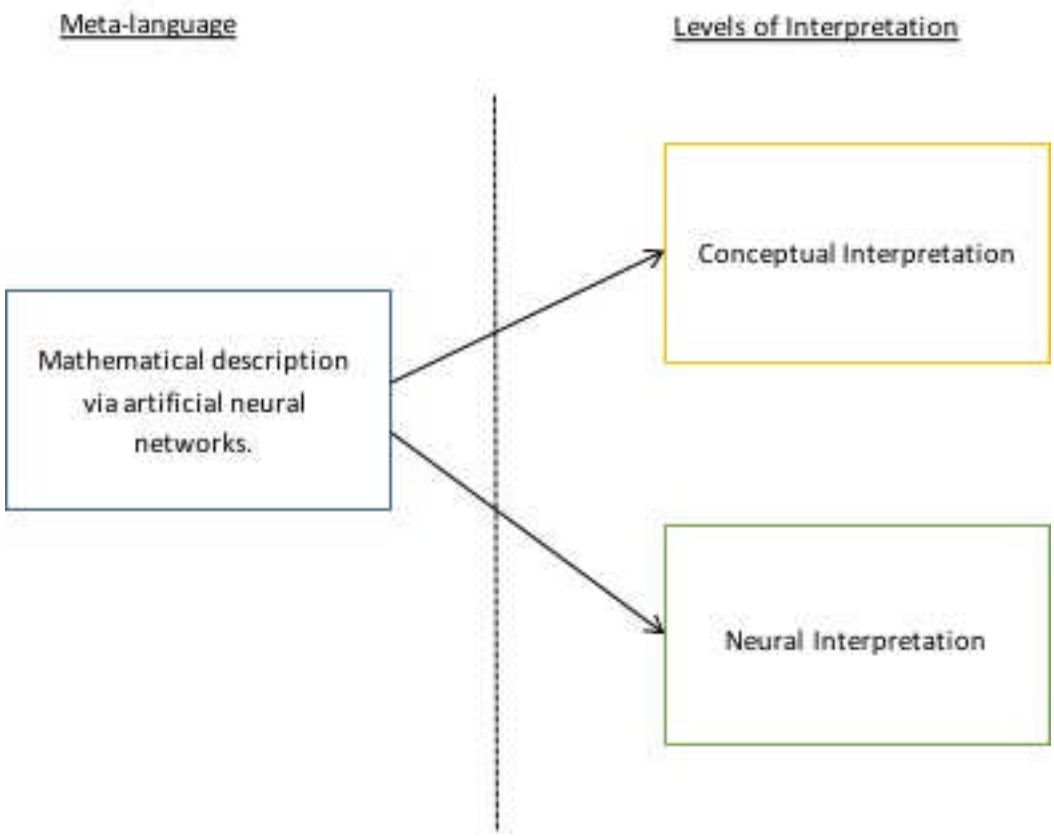


Figure 1.3: Levels of interpretation in the connectionist framework

Simulations, therefore, form a meta-language that has two different interpretations: on the *higher-level* view a certain simulation acts as a meta-language for how a certain conceptual or behavioural task may be achieved. Alternatively, it can be interpreted as an account for how biological neural systems operate in a similar situation. On the ideal view, a certain ANN will simultaneously act as both, and as a result link these two *levels of interpretation*. However, as will be seen in §2.4.2, this ideal situation is yet to be achieved.

As can be seen from the network presented in figure 1.2., however, these ANNs are not given in an unstructured form, but instead one crucial part of their design is that they are arranged according to a series of *levels of organization*. In many ANNs the relevant *levels of organization* are the *input layer*, which receives stimulation from an outside source, a *hidden layer*, which is important in the nature of parallel processing, and the *output layer*, on which the success of the task is assessed. In §2.4.2 we will see examples from the recent field of *deep learning* that employ rather more complex architectures. In sum then, the connectionist approach attempts to understand cognition by creating computational

simulations that attempt to achieve a range of cognitive tasks. By doing so, they hope to generate hypotheses on how such tasks may take place in biological organisms.

Questions concerning the nature of the *mind,* are not the sole concern of psychologists and neuroscientists, but have also long attracted the attention of philosophers. Within the philosophy of mind, a major issue is the *Mind-Body Problem*, that is, the question of how to specify the relation between mental and physical facts. Concerning the notion of *levels of analysis*, one prominent idea is that of **supervenience**, which states a dependency relation between *mental properties*, such as having a headache, and *physical properties*, such as being in a certain neurophysiological state (Kim, 1984; 1987; McLaughlin and Bennett, 2014). On this view, mental properties are not identical to physical properties, but certain constraints exist between the changes in the mental and physical properties of organisms. Although there are certain subtleties associated with *supervenience*, which will be discussed in §2.2, in broad strokes the relation specifies that any change in mental properties must be linked to an associated change in physical properties. In diagrammatic form, it can be conceptualized broadly as follows:

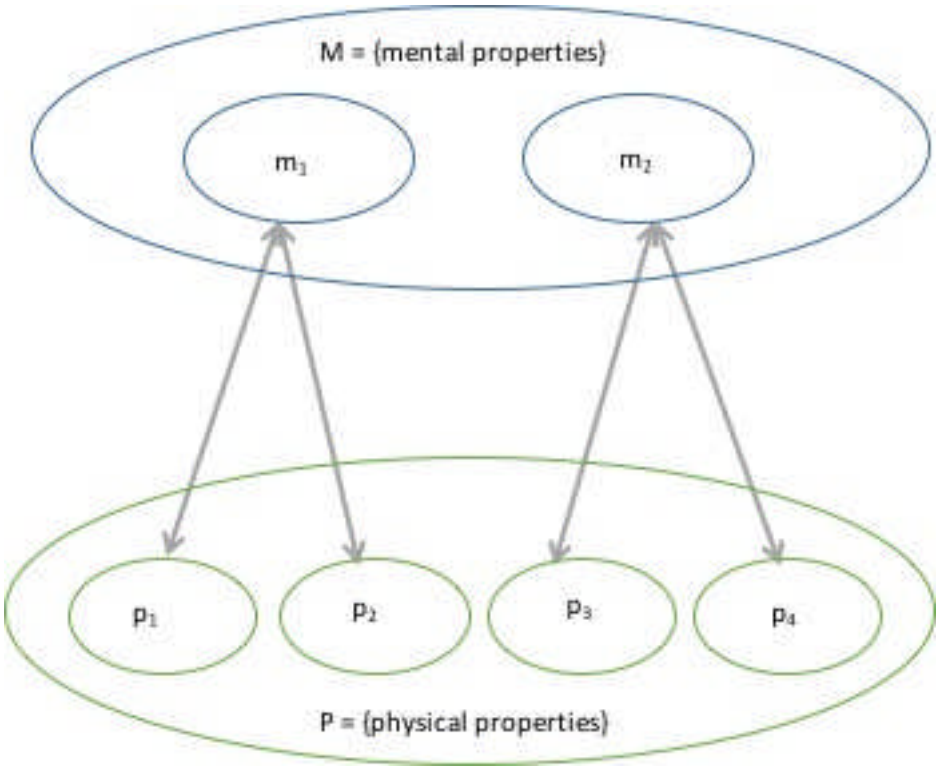


Figure 1.4: Supervenience

In brief, the above picture allows one mental property to be associated with more than one physical property, thereby allowing for variation at the physical level. But, crucially, two different mental properties cannot be associated with the same physical one.

Determining the type of *level* in this example is a difficult issue, but the majority consensus seems to be that the higher level of mental properties and the lower level of physical properties are two *levels of ontology*. This is due to the fact that supervenience is compatible with the notion of *property dualism*, the philosophical view that although there is only one kind of *substance* in existence (usually considered to be physical substance), there are two types of *properties*: mental and physical properties and the former cannot be reduced to the latter (Kim, 1996). *Levels of ontology*, then, deal not with different interpretations of a phenomenon, but instead state that there are two fundamental different *types* of properties in nature.

The final set of proposals to be discussed in this thesis can be grouped together as the **4E approaches** (Rowlands, 2010: 3). Although the various members of this approach differ in many ways, and in some cases not all of the proposals are necessarily compatible with one another, they share a similar set of grievances with previous approaches. These criticisms focus upon the traditional view of cognition involving the manipulation of *representations*, the appeal to the computational metaphor, and that the majority of important facts for understanding cognition solely concern the brain (Menary, 2010: 459). It should be noted, however, not all proposals within this broad group completely reject the idea of *mental representations*, but they do believe that the notion should undergo a certain amount of reformulation (Clark and Chalmers, 1998; Clark, 2008).

The various accounts that combine to make the 4E approach are as follows:

1. **Embedded Cognition/ Dynamicism** (c.f. van Gelder, 1995; van Gelder and Port, 1995; Beer, 2000; Eliasmith, 2013).
2. **Embodied Cognition** (c.f. Sharipo, 2004; Lakoff and Johnson, 1999; Beer, 2000)
3. **Extended Cognition** (c.f. Clark and Chalmers, 1998; Clark, 2008; 2010a).
4. **Enacted Cognition** (c.f. O'Regan and Noë, 2001; 2002; Noë, 2004; 2009; Hutto and Myin, 2013).

The conceptual ideas behind each of these accounts are as follows: **embedded cognition** and **dynamic** ideas on cognitive processes highlight, firstly, the important role of the external environment, and secondly, that they unfold over time. Taken together, these views emphasise how organisms are coupled to an external environment, and are interested in how the bidirectional relations between these two unfold and change over time. **Embodied cognitive** accounts highlight instead the way in which the organism's body can influence the nature of cognitive processes. This is more than saying that being in possession of some kind of body is necessary to perform actions, but that bodily features shape the nature of cognitive life, as well as providing the means for external implementation. **Extended cognition** argues that the barrier between the 'inside' and the 'outside' is artificial, and processes that they view as fully cognitive can take place wholly outside of an individual's brain/ body. Finally, **enacted cognition** purports that cognition is heavily dependent on action; they argue against the 'passive' view of cognition whereby neural systems simply represent some pre-given world, but that the nature of an individual's perceptual experience arises from the ways in which she interacts with it.

In terms of motivating the content of this thesis, it is important to note that the issues of *levels of analysis* and *conceptual frameworks* have not received as much attention in this latter group of proposals. Much of the focus, so far, has been to elucidate their position and motivate it from specific philosophical points, or from individual psychological cases (Wilson and Clark, 2008). For this reason, Northoff (2014: 240) states that the idea of "extended brain […] characterized by an intrinsic relationship between the brain and both the body and the environment" as one major area of future research.

That is not to say that there are no proposals in terms of *levels of analysis*. For instance, in a framework designed to better understand the aetiology of dyslexia, Morton and Frith (1995) present the following *levels of analysis*, in a broadly embedded viewpoint:

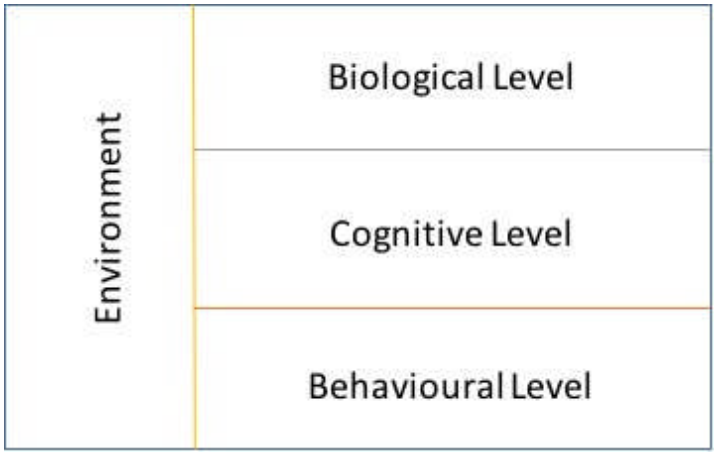


Figure 1.5: Levels of analysis in Morton and Frith (1995)

The above framework attempts to provide a means of mapping the various different causes for dyslexic behaviour. They say that these causes can take place at any of the three levels of *biology, cognition* and *behaviour*, and may or may not have the source of the cause partially in the *environment*. As will be discussed later, it is difficult to fully understand the nature of the *levels* in this case. At first glance, it looks like a series of *levels of interpretation*, similar to above, but as will be discussed the notion of causality in this work makes such a statement problematic.

Another example that has arisen out of the 4E approach comes from Beer (2000), who presents an embedded, embodied cognitive framework as follows:

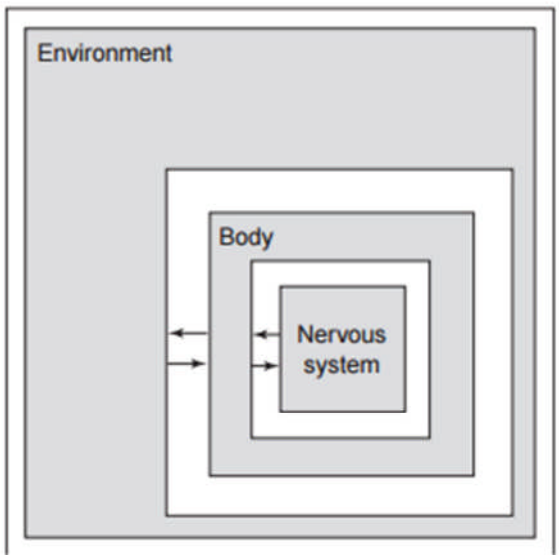


Figure 1.6: Conceptual framework in an embedded, embodied approach

The above framework is within a sub-group of **dynamic** approaches. On this view, each of the *levels* above is considered to be a space of possible states, and the goal of the framework is to track state changes over time. In terms of specific types of *levels of analysis*, it seems that this is a case of *levels of organization*: nervous systems are one part of the body, and the organism's body is one part of the wider environment.

The purpose of this section was to give a brief overview of previous approaches, so as to help explain the motivating factors that make up this thesis. Further details will be given in the following chapter. In order to build up towards the application of SFL in cognitive science, and in order to state the research questions of this thesis, we need to outline some of the issues with the above approaches, and therefore identify the gap in the current literature this work hopes to address.

1.2. Issues with current proposals

This thesis maintains that there are issues with all of the above cases. As before, the details of these criticisms will be discussed in detail in §2, but we can sketch out the general view as follows: both the **classical** and **connectionist approaches** have as an advantage explicit proposals in terms of *levels of analysis*. However, this thesis agrees with the criticisms levied against these views from the **4E** perspective. In the case of the **4E approaches**, the reverse is true: although it is conceptually on firmer ground, it lacks a **satisfactorily explicit** account of the types of *levels of analysis* needed to study cognition from such a viewpoint.

As will be argued in the next chapter, studying cognition along the lines of a broadly computational view unnecessarily biases one to underestimate the fact that organisms are constantly performing actions and cognitive processes within a rich external environment. That is, these views are predominately **decontextualized**. Usually, these computational views on cognition place the majority of their explanations on the manipulation of amodal, wholly internal *representations*, with external factors regulated to the background. The conceptualization is predominately is of a passive processor of information, with factors such as action not playing a prominent role on such views.

On the other hand, the lack of an explicit, or an adequate conceptual framework is a problem with this second set of proposals. Regardless of whether you agree with the conceptual foundations behind Marr's proposal, it does set out the strategy behind studying cognition from this perspective: firstly, you need to understand the nature of the cognitive task, secondly, you need to specify the required set of mental representations, as well as the set of operations that need to be performed upon them. Finally, the relevant physical system needs to be of a type that can implement such facts. Furthermore, these facts are ordered relative to one another, and certain thoughts about their interactions are given.

A similar account is missing from the accounts that are more 'outward looking', in the sense given above (c.f. §2.5.5.). For example, consider again the set of levels given by Beer (2000) in figure 1.6. Although, this thesis agrees that it is necessary to consider facts about nervous system, the body and the environment, there is little detail in this account about how exactly this relation should be specified. It is noted that the relationship between two adjacent levels is bidirectional, but there is less details about how exactly this relation should be conceived. Furthermore, the details within each level are sparse, and as will be seen in §2.5.1., there are issues with the details as they are given. Overall, the situation is summed up by Wilson and Clark (2008: 55) as follows:

> *Importantly, researchers who self-identify with the "situated movement" [read 4E approach – JW] in the cognitive sciences have not been paralysed by a relative lack of attention to the conceptual articulation of their paradigm. They have instead gotten on with the task of showing how the situated perspective leads to interesting and novel approaches to understanding particular cognitive abilities.*

There is nothing wrong, up to a point, with continuing with such an approach. One of the ways to motivate an alternative viewpoint is to show how it can improve our understanding of particular cases, without worrying initially with how the fully fledged details of the proposal may look. But the lack of "conceptual articulation" is a potential issue for the movement advancing. An explicit framework provides a common reference point for different examples to be brought together in one unified picture, lays out the kinds of facts

that are important for developing ideas, and sets out the underlying assumptions behind the approach.

In sum, then, in modern cognitive science there is a potential gap in previous frameworks of cognition. On one hand, there is a set of proposals that contain articulated and explicit frameworks, but suffer from particular issues in their conception of the phenomenon. On the other hand, another set of proposals improves upon the conception of cognitive systems, but lacks a detailed conceptual framework to guide their work. This thesis attempts to start to bridge this gap, by beginning to develop a conceptual framework for cognitive science that takes its influence from the theory of *Systemic Functional Linguistics* (Halliday, 1961; 1979; Halliday and Matthiessen, 1999; 2014). In the next section, we will consider why this viewpoint is initially plausible.

1.3. Motivation for using SFL theory as a basis for a cognitive framework

As will be stated in §1.4 the main research question of this thesis will be to assess whether the application just mentioned is appropriate, and what picture of cognitive systems it presents us with. Therefore, at this introductory stage, it is necessary to show that such an approach is at least plausible, and that there are motivations for thinking that it might be successful. This section outlines these initial considerations.

1.3.1. Cognitive systems, meaning, and SFL

Important to many accounts and explanations of cognitive systems is the idea of **meaning**. For instance, Adams (2010: 325, emphasis in original) considers the issue as follows: "suppose you want to build a computer that can **think for itself**. To do so the symbols in the machine **have to be meaningful** to it". Further appeals to the notion of meaning, or semantic content, being important for understanding the nature of cognitive systems are found in Fodor (1987); Dretske (1988; 1996), Searle (2000), Eliasmith (2005; 2013), amongst others. Many of the appeals to meaning in cognitive systems come from advocates of the representational view of cognition (i.e. manly classical and connectionist accounts), as the representations by definition need to possess meaningful content. This aspect of meaning is thought to be one means of differentiating cognitive from non-cognitive systems (Adams, *op cit.*).

If, as is considered by many, meaning is fundamental for cognition, then cognitive systems are by definition **semiotic**, since this is the discipline that concerns itself with the study of meaning broadly concerned. Language is but one example of a large number of semiotic systems that exist. SFL defines itself as a *social semiotic* theory of language (Halliday, 1979; Halliday and Matthiessen, 1999; 2014), that is, a theory of language that draws upon certain ideas in semiotics more generally. Specifically, it is a theory of *social semiotics*: a study of language as it is used within social and cultural settings (Halliday, 1979).

Importantly for the motivation of the content and research questions of this thesis, ideas and heuristics from SFL have already been applied to other semiotic systems. Mahboob and Knight (2010) provide a collection of such applications, which include pedagogy, multimodal studies, visual semiotics, and social structures more generally. Further examples come from the ideas on *visual grammar* (Kress and van Leeuwan, 1996; Royce, 1999), and SFL principles applied to the field of architecture and art (O'Toole, 1990). Furthermore, it has been applied to non-human forms of communication, especially in the base of primates and primate-human interactions (Benson and Greaves, 2009). That is, although ideas from SFL were developed for the semiotics of language, they have found uses in other semiotic explanations.

The first motivation for the approach taken in this thesis is as follows: if, as is assumed by many researchers, meaning is important for understanding cognition, then cognitive systems are semiotic systems. SFL has been shown not just to be useful for studying language, but various other semiotic systems. This link between the cognitive systems and semiotics has already been noted by Halliday and Matthiessen (1999: 607), who say that "neural events […] are themselves in the broadest sense semiotic", and by Matthiessen (2012: 451) who in terms of his idea of *orders of nature* (c.f. §4.2.1) states that "this always leaves open the possibility of doing analysis within all four systemic orders (for example in the contexts of neuroscience and medicine)." This possibility of applying semiotic considerations found in SFL to cognitive science more generally is the first motivating factor behind this work.

1.3.2. Cognitive systems as dynamic, open systems

Although both cognitive and linguistic systems can be described as **semiotic**, it would be beneficial for the perspective taken in this thesis if further similarities can be found between

these two systems. This thesis takes the position that both language and cognition can be thought of as *dynamic, open systems*. Lemke (1995: 162) describes such systems as those "which survive only through continuous interaction with their environments through processes that exchange matter, energy and information with those environments". He goes on to give an example of a flame (*ibid*) that "must be open to its environment to survive; if we close off the flame from its source of fuel or oxygen, it dies".

Organisms in the most general sense can be described in such a manner, since their existence is dependent upon a constant interaction with an external environment. Taking away from such settings, without the ability to take on food or any other source of sustenance, they will soon die. They are also constantly exchanging energy and information with such an environment; they are constantly having an effect on the world, and are in turn affected by what is happening around them. Other examples include communities, cities and other social structures (Melrose, 2003: 432).

There is also evidence that pertains directly to perceptual systems that a constant need for interaction in order to survive. This is important for the stated goals of this thesis, because, as was seen above, perception is one of processes that are the focus of cognitive science. As reviewed by Merabet and Pascual-Leone (2010), there have been studies that have shown cases of sensory deprivation that have resulted in loss of a certain perceptual function, with the resources previously used for such a process being taken up by others. This is known as *cross-modal neuroplasticity*. One infamous study was undertaken by Hubel and Wiesel (1964), whereby they sutured shut one eyelid of cats at various ages. They found that if the operation was performed on new-born kittens, when the eyelid was reopened after a period of around 6 months, optical stimulus to the previously shut eye resulted in reduced neural activity, whereas similar stimulus to the always opened eye resulted in neural activity that would normally be associated with both visual pathways. That is, the lack of informational or energy exchange had fundamentally altered neurophysiology. Similar findings were found by Riesen and Zilbert (1975), who raised primates in darkness between the ages of 3-6 months. After re-emergence into a non-sensory deprived environment they had significant difficulty with simple identification tasks. It took them weeks or months to achieve such skills, instead of a few days, as was typical. It is possible to see in these cases the importance of environmental interaction in cognitive processes.

This thesis maintains that cognitive systems are not just **semiotic**, but also **dynamic, open systems** that depend on constant interaction with an external environment for their operation and maintenance. It therefore concurs with the work of Thibault (2004), who also maintains such a view. This property is also a property of language as conceived of in SFL, and the theoretical notions in the theory are supposed to handle such systems.

1.3.3. Richness of *levels of analysis*

Given the above two points: firstly, that cognitive systems are semiotic systems, and secondly, that they can be considered as open, dynamic systems, SFL potentially can provide a good comparative theory in order to hypothesize about the nature of cognition. Specifically concerning the issues of *levels of analysis*, SFL theory proposes a number of different *levels of analysis* for studying systems that have both of these qualities. They are listed below:

> 1. *Stratification*: Language can be studied at multiple *levels of interpretation*. That is, the system of language itself is made up of a series of sub-systems. These systems are *phonology* (the system of sounding), *lexicogrammar* (the system of wordings) and *semantics* (the system of meanings). A particular text is an instance of all these systems simultaneously. Furthermore, there exists a system of *context of culture* which language operates in, and is important for understanding the operation and nature of language. These different levels are ordered in a particular way, running from *phonology*, to *lexicogrammar*, with *semantics* and *context* being the two highest levels on this dimension. These strata are linked by a bi-directional relationship of *realization*.
>
> 2. *Delicacy*: SFL considers each *level of interpretation* to be a system, that is, a set of possible values that a text may take. However, this system of potential values can be described in more or less specific ways. The notion of *delicacy*, therefore, describes the potential of language in terms of *levels of specificity*, ranging from the most general to the most specific. The benefit of this is that it allows researchers to

describe particular linguistic facts at a level where significant observations can be made, without including too much detail to obscure such patterns.

3. *Rank*: The units that make up language have the ability to be combined into larger units. Within the stratum of *lexicogrammar*, this is the ordering of *morpheme*, *word*, *group/phrase*, *clause*, *clause complex*. Within SFL theory, the hypothesis is that each linguistic stratum comes with its own constituency ordering, although the precise details are not agreed in all cases. The scale of *rank*, then, describes a series of *levels of organization* that comes with each *level of interpretation*.

4. *Instantiation*: Unlike some other theories of linguistics, especially Chomskyian grammar, SFL does not posit a dichotomy between knowledge of language and use of language. These two aspects are linked by the *cline of instantiation*, which has as its two extremes the language *system* and an *instance* of language through text: any instance of spoken, written or signed language. There is interplay between these two extremes: the nature of the linguistic system is influenced by the texts that are produced, and the system governs the features of texts that are produced in certain situations. Between these two ends it is possible to conceive of a number of intermediate points. These capture, firstly, the regularities that are seen in certain texts, and also represent a variant of the overall system that is applicable to a certain situation. In this thesis, this will be described as operating over a number of *levels of regularity*.

All of the above will be considered in more detail in §3. As far as motivating factors are concerned, it can be seen from the above sketches that the theoretical dimensions in SFL already include some of the kinds of *levels of analysis* that have previously been thought to be important for understanding cognitive systems. Specifically, the ideas of *levels of interpretation*, and *levels of organization* have previously been identified as being important

in cognitive frameworks. Although, of course, the details between these previous proposals and the dimensions above may well differ, the fact that similar types of theoretical notions can be found in both motivates the work undertaken here.

With the motivations for this thesis outlined, it is now possible to state the research questions.

1.4. Research Questions

In sum, then, this thesis aims to investigate and provide answers for the following research questions:

- Can theoretical notions from SFL, especially as they pertain to the kinds of *levels of analysis* important for studying semiotic systems, be of use to cognitive science more generally?
- More specifically, is it (a.) possible, and (b.) useful to describe cognitive systems along the dimensions of *stratification*, *delicacy*, *instantiation* and *rank*?
- What are the features of such a framework? What view of cognitive systems does it provide us with? How does this picture differ from current proposals in cognitive science?
- Can such a framework overcome some of the potential problems with current approaches?

To conclude this introductory chapter, the structure of this thesis will be outlined. The details of what applications of SFL to cognitive science are discussed in each section will be outlined, as well as how they are claimed to help answer the research questions of this work.

1.5. Outline of this thesis

The structure of this thesis differs slightly from the way in which doctoral theses in Linguistics are usually laid out. Instead of the traditional progression of literature review, methodology, data and discussion, the nature of this PhD necessitates a different approach. The aim of this introductory chapter has been to outline, step-by-step, the questions that this thesis will investigate and the rationale into why these questions should be considered plausible.

The second chapter contains a critical review of current conceptual frameworks in cognitive science. It will outline current approaches and critically evaluate them in terms of their advantages and shortcomings. This will include further exploration of the various points raised above in §1.1 and §1.2. It will provide the details of the **classical**, **connectionist, supervenience** and **4E approaches** that were briefly overviewed above. The aim of this chapter is to explain how each approach conceptualizes cognitive systems, and the kinds of conceptual frameworks that have been proposed in each case. It will focus upon the notion of *levels of analysis*, in line with the stated research aims above. This will form a basis of comparison for the later work in this thesis, where an alternative conceptual framework will be explored.

Chapters 3-5 each contain essays that investigate different aspects that pertain to the overall research goals. However, these essays should not be necessarily considered to be independent of each other; each one will build upon arguments presented in the ones that precede it.

The third chapter contains a presentation of a conceptual model of cognition and behaviour that is based upon theoretical notions from Systemic Functional Linguistics. It will be based upon a discussion of the conceptual dimensions of *stratification*, *realization, rank*, *instantiation* and *delicacy*, in the context of the wider field of cognitive science. The chapter will begin with a review of these notions as they have been used in SFL. From this point, all of these dimensions will be discussed in the context of cognitive science more generally. The aim of this chapter is to consider whether there are any motivating reasons for including such notions in a conceptual framework for cognitive science, and, where applicable, how it differs from previous accounts. It will be argued that all of the dimensions mentioned above can usefully be applied in this way, as they all capture different aspects of cognitive systems that are thought to be important, and furthermore, unify them into one coherent framework. Initial proposals into the details of some of these notions will be given, and avenues for further research will be identified. Overall, however, the aim of this chapter is to motivate their inclusion as potential *levels of analysis* in cognitive science. After this discussion, it will be possible to provide a general statement concerning how applying such notions can affect out conceptualization of cognitive systems. This will be in direct answer to one of the research questions above.

The forth chapter explores whether discussions of the interaction between text and context in SFL can be used to develop a more systematic treatment of the interaction between neuronal systems and the external environment. This is an important step towards the *embedded* picture of mental processes this thesis aims to argue for. In doing so, it will consider the following factors of context-text interactions and argue for possible applications in cognitive science more generally:

1. The distinction between the *material situational setting* and the *semiotic context*
2. The contextual parameters of *field, tenor* and *mode*.
3. The contextual stratum as a system and its theoretical description.

It will argue that there are possible applications of many of these features. Specifically, it will show how the distinction in 1 maps on to well-known facts about perceptual systems, where organisms' construal of 'reality' may differ from what we know to be the case. The notion of *contextual parameters* (point 2) will be used to propose a means of systematizing descriptions of the environment in cognitive science. The parameters argued for are (i.) the *objective*, which describes the various objects the organism can sense, and relations between them, (ii.) the *relational*, which describes relations between the individual and said objects, and (iii.) the *goal*, which describes the situation the organism is attempting to bring about. It is argued that by adopting these groupings, it is possible to bring some structure to what can seem a heterogeneous concept. Finally, the theoretical notions described in §3 will be shown to apply to the level of the environment as well, with particular focus upon the notions of *stratification* and *instantiation*.

Chapter 5, "On the Notion of Abstraction in Systemic Functional Linguistics", concerns the issue of how it might be possible to start building the framework from future findings. It does so by examining the notions of *abstraction* as applied to all the dimensions of *stratification*, *instantiation* and *delicacy*. In previous elucidations of the foundations of SFL, all of these notions have been described as operating over a dimension of *abstraction*. The main argument in this chapter will be that, although the same term has been applied in each case, on closer examination there is a different set of conceptual notions at play in each case. The upside to this is that, by achieving a higher degree of explicitness about this notion, it can provide us with a means of considering how the dimensions can be practically

applied in cognitive science. This will contrast with the more conceptual discussions that are the hallmark of §3 and §4.

The final chapter concludes the discussion by reiterating the proposals made in this thesis. In doing so, it will provide a concrete example of the experimental technique of the *Morris Water Maze* to frame the discussion. As much as is possible at this stage, the proposals put forward in the previous chapters will be applied to this case. This will also allow a further comparison to previous ideas identified in §2, in order to emphasise the novel aspects of this work. The proposals that this thesis makes are also explicitly given, as they apply both to theoretical work in SFL and cognitive science more generally. It will end with an identification of avenues of further work that will hopefully build upon the ideas presented here, and hopefully address any gaps in the proposals as they are found here.

Chapter 2

Previous Conceptual Frameworks in Cognitive Science

> *"When […] you Think of Things, you find sometimes that a Thing which seemed very Thingish inside you is very different when it gets out into the open and has other people looking at it."*
>
> A.A.Milne. The House at Pooh Corner.

Before this thesis turns to the investigation of whether theoretical notions from SFL can prove useful for the wider field of cognitive science, a critical discussion of previous frameworks in the field is necessary. The purpose of this chapter, therefore, is to lay out the various proposals in detail, building upon the brief overview given in §1.2. This will form the background to the following chapters in this work.

The approaches discussed in this section will be: (1) Philosophical accounts, especially those based upon the notion of **supervenience**, (2) **Classical approaches** to cognition, and Marr's **three-level framework**, (3) **Connectionist** frameworks, and (4) the **4E** theories of cognition: **Embodied**, **Embedded**, **Extended**, and **Enactive** cognition. This chapter aims to not just give an overview of these proposals, but also to be critical in its outlook. Therefore, advantages and criticisms that have been levied against each of the above will be discussed.

Before we turn to the literature, an issue of nomenclature needs to be raised, so as to avoid certain confusions that may arise. This concerns a distinction between *conceptual frameworks* and *(cognitive) models*.

2.1 Conceptual Frameworks and Cognitive Models

In many disciplines, researchers often use the terms *model* and *framework* more-or-less interchangeably. However, this overlooks a distinction that is important for this work in this thesis. As a result, I will use these two terms as follows:

1. A *model* will refer to a mathematically defined, and usually computationally implemented, description of a certain phenomenon. Within cognitive science, the term *cognitive model* usually refers to simulated artificial neural architectures that aim to show how certain psychological or behavioural tasks might be implemented in neural systems. This would include cognitive architectures such as LISA (Hummel and Holyoak; 2003) and Soar (Newell; 1990). Also included in the umbrella term *model* will be mathematical descriptions of neurological phenomena, such as mathematical models of neural encoding and decoding (see Rieke et al., 1997 for examples). Within SFL, this term would also refer to computational generation and parsing systems, for instance the COMMUNAL system (Fawcett and Tucker, 1990) and SysAm (Wu, 2000).

2. A *conceptual framework* will refer to an abstract, theoretical framework that provides the means of conceptualizing a particular phenomenon. Ideally, these frameworks include the underlying assumptions behind a particular theoretical approach, and specify the relevant *levels of analysis.* Within Linguistics, examples of conceptual frameworks would be the stratificational framework in SFL (c.f. Halliday and Matthiessen, 2014), and the Minimalist framework within generative linguistics (Chomsky, 1995). Examples of conceptual frameworks from the wider field of cognitive science are discussed in this chapter.

Occasionally, a particular proposal may be considered to be both a *mode*l and a *framework*. For instance, connectionist models of cognition, as will be discussed in §2.4, contains aspects of both. However, this distinction is meant to highlight the intentions of this chapter and this thesis as a whole: in the discussions which follow, the aim is to enquire into *conceptual frameworks* in cognitive science, and not necessarily the mathematical or computational implementations of such approaches. In other words, the focus of this review is the underlying assumptions and conceptualizations of cognition, behaviour and the 'mind'.

With this distinction in mind, this chapter will now review the previously proposed conceptual frameworks in cognitive science, focussing on the notion of *levels of analysis*. We will begin with a discussion of ideas put forward in the discipline traditionally known as the philosophy of mind, and specifically the relation between *levels* known as *supervenience*, to be defined below.

2.2 Philosophical Views and the 'Mind-Body Problem'

Attempts to provide an understanding of the mind have long occupied the thoughts of philosophers, from the period of philosophical antiquity through to the present day. For example, in the *Pheado*, Plato writes about the division of the body and the soul, and that the soul aims to discover the nature of *Ideas*[2] through rational means. Later, Aristotle stated that the soul was the seat of the mind in *De anima*. In more modern times, this issue is usually discussed in terms of the *Mind-Body Problem*: the issue of how to link the mental to the physical (Lowe, 2000: 8). Discussions concerning this topic have their roots in Descartes, especially his writings in *Meditations on the First Philosophy* (Descartes, 1641/1999).

In this work, Descartes argues that the Mind and the Body are essentially of different kinds, or made of different *substances*. This means that reality consists of two different types of objects with fundamentally different properties (Robinson, 2014). He reached this conclusion by arguing that there was only one fact that he could know for certain: that he himself existed. Everything else, from the existence of an objective outside world to even the knowledge that he possessed a body, can be doubted, since these facts all come via fallible sensory systems. Therefore, this 'I' that existed was something different from the body, since the former cannot be doubted whereas the latter can. Further arguments he proposes include that the body, like other material things is divisible and has extension, whereas the mind is neither divisible nor extended: it has no parts but only exists as a unified simple substance. He motivates this point by noting that parts can be

[2] The term *Ideas* (also known as *Forms*) refers to Plato's doctrine that there exists an ideal realm of existence that is populated by timeless, non-changing *Ideas.* In most cases, humans only sense an imperfect form of these Ideas, but through philosophical and mathematical investigation we are able to overcome this deficiency in our understanding.

lost from the body while the mind remains intact. This viewpoint is known as *Cartesian Dualism*; or the ascription of the Mind and Body as separate substances.

Regardless of the strength of Descartes' arguments from a modern perspective, the influence he has had upon the philosophy of mind in Western philosophy cannot be understated. Much later work in the field has either implicitly or explicitly attempted to either modify the Dualist picture presented by Descartes, or alternatively to deny the existence of two distinct substances and show how the mental and the physical can be unified into one metaphysical picture. For instance, the philosophical position of *physicalism* or *materialism* argues that there exists only one kind of substance, that of the physical. Although a broad doctrine, with regards to the philosophy of mind, it is the denial of the existence of mental substances. Instead, it argues that everything about the mind and cognition can be brought under the remit of the physical sciences. Modern proponents of the physicalist thesis within the philosophy of mind include Block and Stalnacker (1999), Hill and McLaughlin (1999) Searle (1999) and Kim (2005), amongst others.

This thesis, however, is not specifically about the metaphysical picture of the mind and body, although it will be touched upon. The rest of this section, therefore, will outline the notion of *supervenience*, which directly relates to a certain type of *level of analysis*. It is worth keeping in mind the ideas of *dualism* and *physicalism*, however, as it will make some of the following discussion clearer and place it in its proper context.

2.2.1. Supervenience

Within the philosophy of mind, *supervenience* is important to many modern physicalist ideas. Although its proponents consider there to be only one kind of substance, they consider there to be two different kinds of *properties*: mental properties and physical properties. *Supervenience* describes a particular dependency relation that holds between these properties. It was first brought to the attention of the philosophical community by Davidson (1970), although it has a richer history, and was already used in the field of ethics before this date (McLaughlin and Bennett, 2014). *Supervenience* is an example of a particular type of *physicalist* thesis: a *non-reductive physicalist* treatise (Davidson, 1970; Beckermann, 1992; Heil, 2004: 100; Hoffmann and Newen, 2007; Schneider, 2013, but

see Mantik, 2011 for an alternative view). Initially, then, it is necessary to discuss this notion of *non-reductive physicalism*, to fully appreciate how some philosophers of mind conceptualize the relationship between the *Mind* and the *Body*.

2.2.1.1. Reductionism and Mental-Physical Identity

The issue of *reductionism* is an extremely important one in the contexts of the philosophy of mind and the philosophy of science. It concerns the question of whether it is possible to link terms, facts and other abstract objects from one theory to those of another. If possible, one can say one theory has been *reduced* to a more fundamental one. The dominant modern conception was proposed by Nagel (1961), and Beckermann (1992: 107) gives its definition as follows:

> *[R]eduction is a relation which holds between two theories T_1 and T_2 if the laws of theory T_1 can directly or with the aid of so-called "bridge laws" be deduced from the laws of theory T_2.*

Beckermann (*ibid*) provides an example: the reduction of thermodynamics to statistical mechanics. His example makes use of the Boyle-Charles law in thermodynamics, which concerns the properties of gaseous substances, and how it can be deduced from, or reduced to, the laws of statistical mechanics. Boyle-Charles' law is given as follows[3]:

(1) $p * V = N * k * T$

The corresponding law in statistical mechanics is as follows[4]:

(2) $M * V = \frac{2}{3} * N * E$

To *reduce* the theory of thermodynamics to that of statistical mechanics some way of linking the equations given in (1) and (2) is required. However, the complication with reduction is that certain terms only have meaning within their respective discipline (Lewis, 1970). For instance, the terms *temperature* and *pressure* have no meaning within the context of statistical mechanics, whereas the reverse is true for the notions for

[3] Where p = pressure, V = volume, N = number of molecules, k = the Boltzmann constant and T = absolute temperature.

[4] Where M = average amount of momentum transferred from the molecules in the gas to the walls of the container, and E = the mean kinetic energy. Other symbols defined as in fn2.

molecular momentum and *kinetic energy* in thermodynamics. The issue is, essentially, one of linking concepts and terms in theories that use different 'vocabularies'. In order to achieve reduction, it is necessary to define certain *bridging laws*, which specify those links. In the above case the appropriate laws can be stated as follows:

(3) $p \leftrightarrow M$

(4) $\frac{2}{3} * E \leftrightarrow k * T$

In (3) and (4) the terms in both theories are linked by the logical biconditional relation. With such laws in place, it is now possible to reduce equation (1) to (2) by logical deduction. As a result, it is possible to say that we have *reduced* (1) to (2).

Mumford (2003: 173) gives a similar definition of reductionism, while highlighting the fact that reduction should be considered to be a *one-to-one* mapping between theoretical terms. Within the context of cognitive science, it is discussed by Fodor (1974; 1975). Fodor takes reductionism to be "the view that all special sciences reduce to physics" (1975*:* 10), where within the special sciences he is mainly focussing on psychology. His characterisation of reductionism (one which he goes on to deny is applicable in the case of cognitive science) also utilizes the notion of bridge laws linking laws in psychological theory with those in physical theory. Crucially, these mappings between theoretical laws are also taken to be *one-to-one*.

With this brief outline given, we can now outline *supervenience* as a **non-reductive *physicalist thesis***. This refers to the position that denies that one can link psychological theory with more fundamental theories in a similar way to above. That is, there is not a one-to-one mapping between mental and physical properties. Reasons for such a view stem from a discussion of *mind-brain identity,* or whether it is possible to say that mental properties **do not just correlate** with certain physiological properties, but are *numerically identical* with them. *Numerical identity* conveys a statement that a certain mental property **just is** a certain physiological property.

The debate about mind-brain identity is usually raised in terms of the philosophical distinction between *types* and *tokens*. Although the distinction between these two terms is a complicated issue, an approximate understanding of the distinction is sufficient for

our purposes here. The clearest way to understand the difference is by considering linguistic examples such as the following:

(5) The man hit the nail with the hammer

Considering (5), we can ask: how many words are in this sentence? There are two possible answers to this question: in one sense there are eight, determined by simply counting the individual words. In another sense there are only 6, since the word <the> is repeated three times. We can say that there are 8 word *tokens,* but only 6 word *types*.

Relating to *mind-brain identity*, it is possible to re-phrase the question about whether mental states or properties are identical to brain states or properties by asking whether the identity relation should be thought in terms of *token-token* identity, or *type-type* identity. These different perspectives have large consequences on whether researchers consider mind-brain identity to be a coherent thesis or not. *Token-token identity* is taken by many to be relatively uncontroversial. For this variety of identity to hold, all that needs to be argued is that a certain mental property at a certain time is identical with a certain physical property at the same time. Considered in this sense, not many cognitive scientists would object to its validity, but instead comment that it does not provide a particularly rich picture of the relationship between mind and body. To illustrate this, consider the following situation. Say two people, Sarah and Gareth, wake up one morning with a headache. To put it in other words, both Sarah and Gareth can be said to be in two mental property *tokens* of "having a headache". Further, suppose that it is possible to measure accurately the neurophysiological state of Sarah and Gareth. In order for token identity to hold, it must be claimed that Gareth's mental property *token* is identical to his physiological property *token*, and similarly for Sarah. Crucially, however, there need not to be any coherent link between Gareth's and Sarah's physiological properties. Intuitively, there is something missing here. Any cognitive theory would like to say that there would be some kind of similarity between the physical *tokens* of both Gareth and Sarah. This, however, is not assured through *token-token identity*.

Type-Type mind-brain identity, on the other hand, is a stronger thesis. It states that mental state *types* are identical to physical state *types*. To continue our example, we can say that both Gareth's and Sarah's headaches are *tokens* of the same *type* of 'having a

headache'. The *type-type mind-brain identity thesis* states that this mental *type* is identical to a certain physiological *type*, say activation of certain areas of the cingulate gyrus (c.f. Vogt, 2005, although the precise physiological details are not crucial). In other words, 'having a headache' **simply is** 'activation of certain regions of the cingulate gyrus'. So in order to 'have a headache', both Gareth and Sarah must exhibit activation in this brain region; any other response cannot be present in subjects having headaches. It is clear that this is a stronger idea than token identity, and, as such, Fodor (1975:12-13) notes that type identity entails token identity, but not vice-versa. It is thought that type-type identity is untenable, for the reasons that it does not allow for any variance between subjects, and it necessarily entails that only entities with neuronal systems could be in possession of mental properties. For these reasons, *type-type* identity is not widely accepted (Putnam, 1975, but see papers in Gozzano and Hill, 2012 for recent defences of type-identity). Such issues with *type-type identity* have led to the development of the dependency relation of *supervenience* between mental and physical properties.

2.2.1.2 Supervenience: definitions

Supervenience goes against the idea that we can reduce the psychological sciences to neurophysiological ones. Its basic formulation can be stated as follows (McLaughlin and Bennett, 2014):

> *A set of properties A supervenes upon another set B if and only if no two things can differ with respect to A-properties without also differing with respect to their B-properties. In slogan form, "there cannot be an A-difference without a B-difference".*

Or, conversely (c.f. Davidson, 1970; Vogelstein et al. 2011):

> *A set of properties A supervenes upon another set of properties B just in case two things which are identical to each other in terms of B properties, are also identical to each other in terms of A properties.*

The reason for the above two definitions being identical is that if, as in the first definition, a change in property A necessitates a change in property B, then if two individuals share

property B, they must also share property A. Otherwise the case would arise that there was a change in A-property, without a change in B-property.

As it relates to the issue under discussion here, supervenience presents us with the following picture: Let's say that **set A contains as its members all mental properties**, and **set B contains as its members all physical properties**. Supervenience, in this simple form, states that two individuals who exhibit the same physical property must also be exhibiting the same mental property, but two individuals can possess different physical properties while still possessing the same mental property. This allows physiological differences to exist between organisms.

Even on this brief definition, a number of questions can be raised. In order to determine whether supervenience is tenable, it is necessary to have a clear idea of what you mean by a mental or physical property, and what has to be the case in order for two individuals to possess identical properties of the relevant sort. Another set of questions arise from the modality of the statement.[5] The aspects of this vagueness can be stated as follows: does this sort of determination **hold solely for the way the actual world is**, or is it a **metaphysical necessity holding over all possible worlds?** Additionally, we can ask whether *supervenience* **acts over individuals**, or does it **hold for properties of the world has a whole**? The need for these clarifications was identified by Kim (1984/1993; 1987/ 1993), and as such he proposed a typology of various kinds of *supervenience*, and showed the different commitments that follow from them.

The first distinction is between *weak* and *strong* supervenience. The conceptual distinction here is whether supervenience should hold only in the state of the world *in actuality*, or whether it would hold in all possibilities. The informal and formal definitions for weak supervenience are as follows (Kim, 1984/1993: 64):

> (WS_{IF}): A *weakly supervenes* on B if and only if necessarily for any property F in A, if an object x has F, then there exists a property G in B such that x has G, and if any y has G it has F.

[5] A note of caution, this is modality in the sense of philosophy and formal semantics and is linked to the idea of possible worlds. It is not necessarily the same as the notion of modality in systemic functional linguistics, in which modality is linked to the interpersonal metafunction at the lexico-grammatical stratum, c.f. §3.1.2.

$$(WS_F): \Box\forall x\, \forall F \in A[F(x) \rightarrow \exists G \in B(G(x) \wedge \forall y(G(y) \rightarrow F(y))]$$

In contrast, strong supervenience can be stated as follows (*ibid:* 65):

> (SS_{IF}): A *strongly supervenes* on B just in case, necessarily, for each [object] x and each property F in A, if x has F, then there is a property G in B such that x has G, and *necessarily* if any [object] y has G, it has F.

$$(SS_F): \Box\forall x\, \forall F \in A[F(x) \rightarrow \exists G \in B(G(x) \wedge \Box\forall y(G(y) \rightarrow F(y))]$$

Sets A and B are defined as above. The crucial difference between the two is that (WS_F) contains only one necessary modal operator, while (SS_F) contains two. On the weaker view the dependency between mental and physical properties is contingent on the current state of our own possible world, but need not hold in others. The stronger version of supervenience, on the other hand, states that this particular dependency relation does not purely hold over the way the world actually is but also over other possible worlds, with certain conditions. The other kinds of possible worlds that are usually considered include those that are logically coherent (i.e. only those possible worlds that contain logical consistent and non-contradictory facts) and/or other nomologically possible worlds (other possible worlds in which the laws of physics as we know them also hold, c.f. Cranston-Paull and Sider, 1992: 834). It is for this reason that Jackson (2000: 9-10) dubs weak supervenience *intra-world* supervenience, and strong supervenience *inter-world* supervenience.

As well as the different forms of supervenience in terms of modal strength, another means of distinguishing different varieties of supervenience is whether they are aimed to be characteristic for entire worlds, or for individuals. The latter of these is dubbed *global* supervenience, after Kim(1984/1993), whereas the former can be dubbed *individual* or *local supervenience*. The definitions of weak and strong supervenience above were couched in local terms; they were defined as a dependency relation between objects, and how the properties of those objects were dependent over possible worlds. *Global supervenience*, however, deals not with individuals but mental and physical properties of worlds as a whole (Lewis, 1986: 14-15; Cranston-Paul and Sider, 1992).

Supervenience, then, is a thesis that considers cognition in a two-level framework: the *higher level* of mental properties, and the *lower level* of physical properties. The account gives a dependency relation between them that is argued to be constrained enough to link the two together, while allowing variation to exist between the realization of mental properties by physical properties. Many of the debates surrounding supervenience concern the metaphysical necessity of such relations, and whether they should hold for individuals, or for physical properties of entire worlds.

2.2.1.3. Supervenience and Neuroscience

The above outline of *supervenience* and its variants can only provide a partial account of the wide-spread views on the topics, and debates are on-going. Further discussions need not concern us here, as some of the issues raised at the end of this discussion will pertain to *supervenience* in all its forms. More interesting for the purposes of this thesis is that there have been some attempts to link metaphysical discussions about the dependency of the mental on the physical with neuroscientific accounts. For instance, Haykin (2008) and Ripley (2008) argue that work carried out in the connectionist framework and artificial neural networks (c.f. §2.4) are both based upon an implicit assumption of a local supervenience thesis between the cognitive/mental and a particular physical implementation.

Vogelstein et al. (2011), attempt to bridge the gap between philosophical considerations and neuroscientific empirical study by proposing a probabilistic version of *supervenience*, based upon Bayesian probability theory. The authors describe their version of *supervenience* as operating over two probability distributions. The first, $M = \{m_1, m_2, \dots\}$, is the probability space of all possible minds, where each m_n is a "possible collection of thoughts, memories, beliefs, etc." (*ibid:* 1). The second, $B = \{b_1, b_2, \dots\}$, is the probability space of all possible brains, where each b_n is a "possible collection for each possible position and momentum of all subatomic particles within the skull" (*ibid)*. This appeal to brain states as defined at the subatomic level is unusual for discussions on mind-brain supervenience, and the authors themselves are not consistent on this matter. In later examples, they present a thought experiment where the supervenience base is defined in terms of neural connectomes, that is, graph-structures consisting of neurons and

synapses (*ibid*: 3). The proposal in this work is that it is possible to take the conceptual idea of supervenience and from it derive a statistical hypothesis to determine whether it holds by looking at empirical data. The basic proposal is that statistical supervenience holds just in case the probability of two mental states being distinct, given that fact that the corresponding two brain states being identical, is 0. Their remaining proposal highlights the lack of a solid statistical test for determining whether the above can ever be shown or not, given the amount of empirical data that would be needed to test such a claim. However, such extensions are still based upon the basic definitions of the two probabilistic distributions mentioned above. Below in §2.2.2, I will argue that one of the problems with supervenience stems from the lack of consistent definitions of these terms. As a result, such extensions need not be discussed here.

2.2.2. Levels and Supervenience

To sum up the above account, the supervenience thesis provides us with a two-level view for the study of cognition and provides a characterization of the relation between these levels in terms of a dependency relation.

The relation of supervenience, and the levels involved, can be diagrammed as follows:

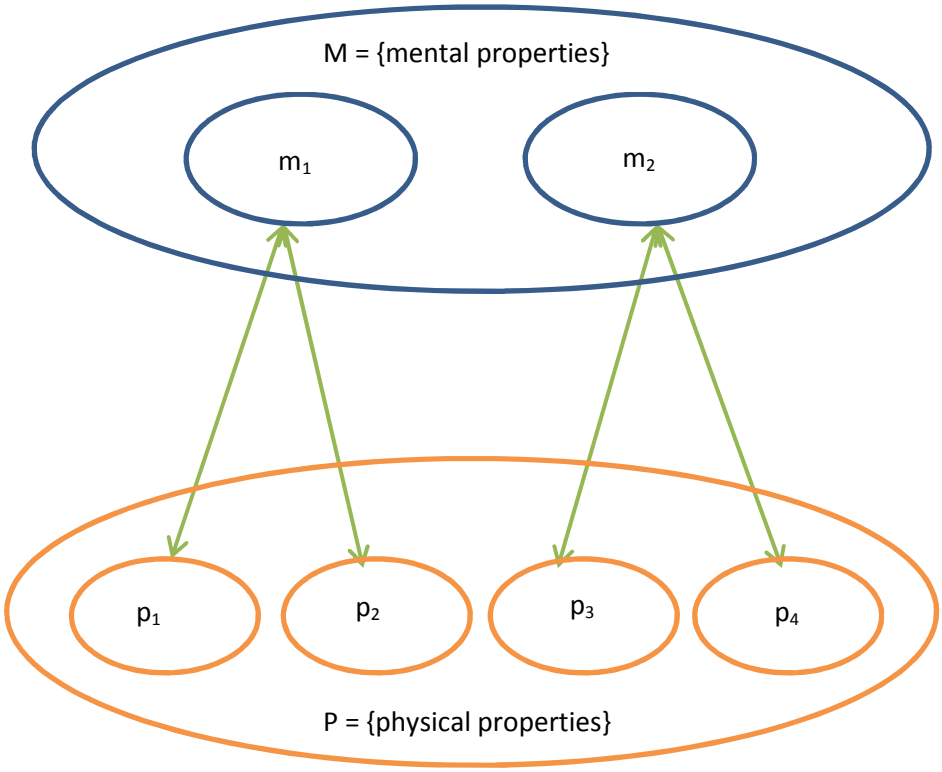


Figure 2.1: Supervenience Outline

As should be clear from the previous discussion, the above figure is a simplification and omits many of the nuances already discussed. Despite these omissions, the above sketch does illustrate the main features of supervenience: the conceptual framework consists of two levels, construed as sets of different properties with restrictions affecting the relationship between different members of these sets.

The pertinent question is what kinds of levels are we dealing with? Two answers to this question present themselves: either we are dealing with *levels of interpretation,* or *levels of ontology*. The first view would amount to the picture that the difference between mental and physical properties or states would be solely a difference in perspective on the part of the investigator. On the face of it, this would seem to be viable take on the difference between mental and physical properties in this case. The reasons for this can be related back to the reasons for taking a non-reductive physicalist approach that were discussed earlier. To recall, the major motivating reason was the lack of an available bridge-law, or a scientific law linking the higher-level to the lower-level properties. This could be considered to be akin to the links between different academic disciplines. It is not currently possible to fully reduce biology to chemistry to physics, but this could be

due to the nature of our understanding, instead of any metaphysical differences between biological, chemical and physical facts. The move to the higher level on the framework, therefore, would amount to taking a certain stance in our investigations, and the claim that this stance is irreducible to the alternative means of investigation of the system in physical terms.

There are reasons to doubt whether this is the sense the framework is meant to convey. For example, Kim (1996: 5-6) discusses the difference between the mental and the physical in the following way:

> *Having a mind [is having] a property, capacity, or characteristic that humans and some higher animals possess in contrast with things like pencils and rocks […] Mentality is a broad and complex category. Some properties are physical, like having a certain mass or temperature, being 1 meter long, and being heavier than. Some things – in particular, persons and certain biological organisms – can also instantiate mental properties, like being in pain and liking the taste of avocado.*

The philosophical thesis that Kim is espousing here is known as *property dualism*. It is the idea that although there is only one kind of *substance* in the world, physical objects can possess two kinds of *properties* that are non-identical. This position is clear in the above passage, where mental properties are directly contrasted with physical ones, with the implication that mental properties are not physical properties. It should be noted that not all proponents of supervenience based theses are committed to the notion of property dualism, but as it stands this metaphysical view is still compatible with it (Robinson, 2012).

The second interpretation amounts to the picture that the properties on the two different *levels* within the framework are of fundamentally different kinds. That is, there is a set of *physical properties*, of which the members of the lower physiological set would count as members, but then there is another set of *mental properties*, that are inherently *non-physical*. Searle (2002: 58-62) provides a succinct overview and critical discussion of

property dualism as it relates to the philosophy of mind. If this meaning is intended, then we are not dealing with *levels of interpretation*, since the difference does not depend upon the perspective of the observer, but something fundamental about the nature of certain properties. Instead, then, we are dealing with *levels of ontology*.

2.2.3. Problems with Supervenience as an informative model

Although supervenience may provide us with a notion of how this dependency might look, it can be argued that it is a rather weak thesis, and is of limit use to those who are actively engaged in developing cognitive theories.

The first problem can be termed as the *explanatory gap*. This is the idea that although supervenience can put certain restrictions on the kind of mental characteristics a certain physical system can take, it provides no explanation about **why** this should hold (Blackburn, 1984; Kim, 1998; McGinn, 1993). Whether or not supervenience holds does not depend upon any understanding of why it should be the case, simply that any description of the relation between the mind and the physical should show a certain type of dependence. Its usefulness as the basis for an explanatory model, as a result, seems limited.

Secondly, the dependence relation seems to be rather weak. All it requires is that objects that share physical properties must (to varying degrees of necessity) share mental properties, and that if two object differ in mental properties, they must also differ in physical properties. It gives no indication concerning whether there is any systematic link between the differences in physical properties and the difference in mental properties. For instance, the framework is compatible with the view that two individuals could differ physically in a very small detail, for instance the presence of a single firing neuron in one case, and its absence in another, and yet be in massively different mental states, say *seeing the colour red* in the first instance, and *being in pain* in the second. This is compatible with supervenience, since (i) the two objects are not physically identical, and therefore are not necessarily mentally identical, and (ii) there is a mental difference between the two individuals, and there is also a corresponding physical difference. This point is discussed by Kim (1989; 1993), who discussed the possibility of a tiny difference in

the world (say the presence of one extra hydrogen atom) could have massive consequences for mental properties in that world. Paull and Sider (1992) showed that the same arguments hold for small differences within individuals' brains, although they go on to argue that such a picture is not problematic. Certain solutions have been offered to combat this problem, but these solutions (Stalnacker, 1996; Paull and Sider, 1992), I would argue, would do little to help in our quest to understand the link between the mind and the brain. The crux behind these solutions is that there is no *a priori* reason to believe that a small difference in physical differences should not result in a large difference in mental properties. Regardless of our intuitions on the matter, we must allow for the possibility of this being the case. Whether or not this is true is indeed an empirical point, but steps towards discovering that fact will not be found due to any predictive power concerning the supervenience framework, since it makes no predictions either way.

Furthermore, and related to both of the previous points, there is the issue of defining **what exactly counts as a mental or physical state**. The problems of defining a mental state are well acknowledged, due to the fact that they are characteristically not publicly accessible (McKinsey, 1991; Allen and May, 2014). Also termed *privileged access,* the issue is that although we can know whether we ourselves are in possession of a certain mental property, it is difficult to know whether others are, or if their mental property is identical to our own. For instance, I can know when I am in pain, even in cases where I show no outward indication that I am feeling as such. This knowledge cannot be known in the same way for other people; I can be reasonably sure that someone else is in pain, but cannot achieve the same amount of certainty than I can for my own mental life. Furthermore, it seems difficult to ascertain whether my property of *being in pain* is the same as someone else's. On what criteria would such an issue be resolved? The answer not just to be a case of observable behaviour, since an outward expression of pain seems to be neither a sufficient nor necessary condition for having the mental property of *being in pain*. For instance, one can be in pain with no expression or such, or one might alternatively produce an exclamation of pain without the mental property being present (i.e. lie).

The situation seems not to improve when we come to look at the property of being a certain physical state. To show this, consider the classical example about the mental

property of *being in pain* supervening upon some physiological property, usually given as *C-fibres firing*, but again the exact details are less important (Davidson, 1970). Specifying the physical property seems relatively simple in this case; if the c-fibres of someone are firing, then they are exhibiting that particular property. However, the question is: how far do you go? Does the particular firing pattern in such nerves matter? And to how much detail? The particular spatial and temporal patternings of action potentials in neuronal systems are considered to be important in the processing of information (Rieke et al. 1999; VanRullen and Thorpe, 2001). Even more worrying is the ascription of a physical state as defined as the position and momentum of every individual particle that was seen above (Volgelstein et al., 2011), since it is currently outside the realms of possibility to know such a state in an organ as complex as the brain. The issue is not that these are not insurmountable concerns, but that they have a real effect on whether two properties are considered identical or not. Without further specification, it is difficult to assess the appropriateness of the supervenience position.

This section has outlined the tradition in the philosophy of mind known as supervenience. Its major characteristics have been discussed, as well as some of the problems associated with basing an explanatory model of cognitive science upon it. In the next sections, we will turn our discussion away from solely metaphysical philosophy, to discuss some frameworks that are more closely associated with 20th Century cognitive science.

2.3 Classical approaches

Since the advent of the cognitive revolution in the mid-20th Century, one of the most prominent themes in cognitive science has been theorizing about cognition using the vocabulary of computation (Gigerenzer and Goldstein, 1996). A quick survey of some titles in the field confirms this: Pylyshyn (1984: xiii), in his book *Computation and Cognition* states that "my proposal amounts to the claim that cognition is a type of computation", Churchland and Sejonwski (1992: 6), in *The Computational Brain*, claim that "nervous systems and probably parts of nervous systems are themselves naturally evolved computers", while Edelman (2008: 13), in *Computing the Mind* , argues that "minds are best explained in computational terms". The above are only a few examples,

but it is clear from both the titles of the works and the views expressed within make the link between cognition and computation explicit.

Before we turn to examining this notion of computation in more detail, it is important to introduce the highly influential conceptual framework introduced by Marr and Poggio (1977) and Marr (1982). This proposal introduces a three-levelled framework for theorizing about the nature of cognition as based upon a computational view. Marr (1982: 19-20) argues that considering certain phenomena at multiple levels is critical since it can act as a epistemological tool that allows researchers to posit different sets of explanations for different elements of a given system, without needing to initially posit the ideal circumstance of all aspects of a system being derivable from one set of theoretical laws (c.f. McClamrock, 1991; Bechtel, 1994). The framework that Marr and colleagues proposed is as follows (Marr, 1982: 25):

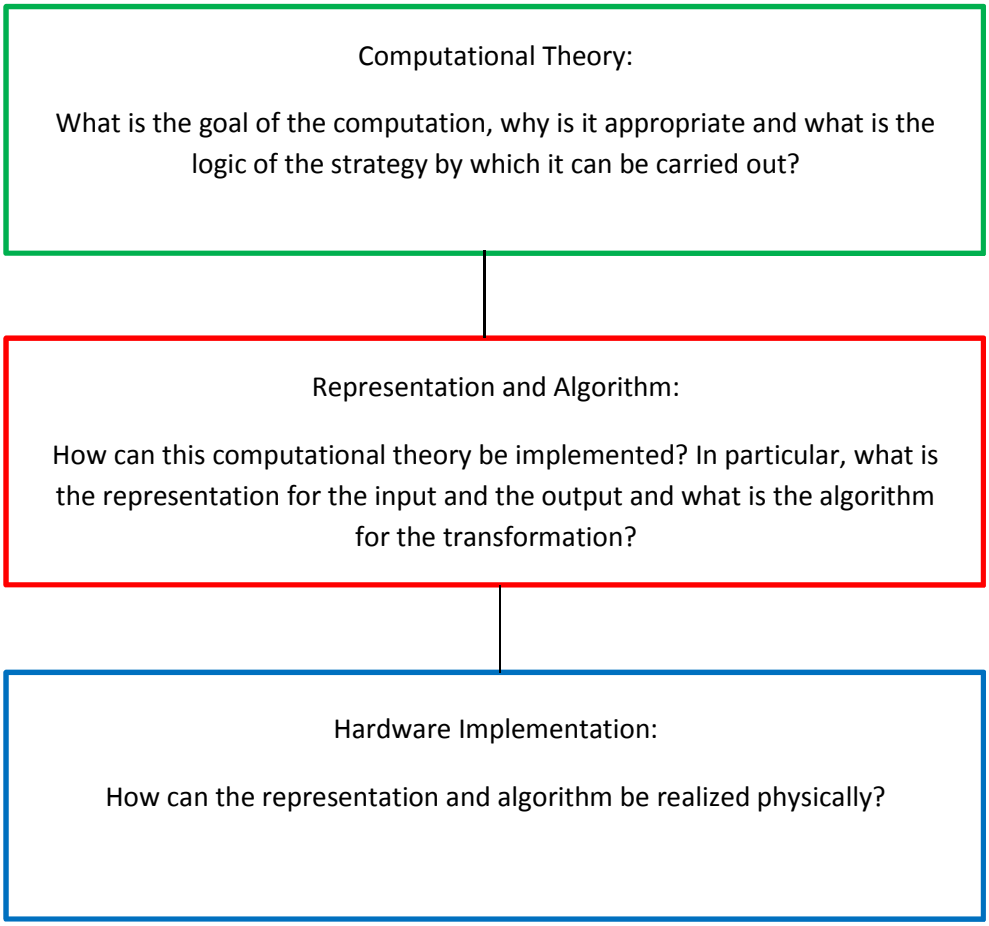


Figure 2.2: Marr's Three-Level Framework

In order to introduce how exactly this framework should be utilized, Marr gives the example of a shop checkout till (*ibid*: 22). In order to gain a full understanding of how this machine operates, one needs to consider the system at all of the levels given above. Starting from the highest, most abstract level, you need to understand **what** the machine is attempting to do, and **why** it is doing it. In the case of a checkout till, it needs to take the prices of any number of goods and produce the total price as the sum of all the individual goods. Addition is thus needed, and this computational strategy is appropriate because it satisfies the needed constraints of *commutativity*, *associativity*, *inverses* and that a price of zero should not change the previous sub-total. Moving down to the level of *representations and algorithms*, the system needs some way of representing entities that are crucial to the computation, and perform algorithms on those representations in order to achieve the stated goal. In this example, the inputs to the computation could be represented by Arabic numerals, and the output representation, being the sum of the inputs, could take the same representational form. An algorithm is a step-by-step procedure needed to get from the input to the output, and in our case would have to include instructions for 'carrying the one' when a given numeral exceeds 9. For Marr, the crucial point here is that there could be a wide number of representations or algorithms that could be used to solve any computational theory. For instance, one could use Roman instead of Arabic numerals, and this would necessarily change the nature of the corresponding algorithm. Finally, the lowest level would be a theory of how the above process is to be implemented in a physical system. In this, this would involve understanding how a physical system would have to be (for instance an electronic computer or a human brain) in order to run the computation, possess the relevant representations and run the algorithms defined by the higher two levels. As before, there will be more than one physical system that can implement a given set of representations and algorithms.

Certain aspects of Marr's conceptual framework should be noted. Firstly, it relies heavily on the computational metaphor: the idea that in order to understand cognition we need to view it as we would artefacts such as digital computers. Linked is the prominent role that algorithms play in the model, which implies that human (and other animal) cognition should be partly understood as operating via a step-by-step process. Both of these facts

will become clearer when we come to investigate the notions of representation and computation in more detail below, and issues with both of these construals will be highlighted.

With this outline in mind, we can ask ourselves what kind of *levels of analysis* this framework utilizes. The nature of the relationship between the three levels in Marr's model is a degree of perspective. This can be seen in Marr's views on understanding complex systems (Marr, 1982: 19-20):

> *Almost never can a complex system of any kind be understood as a simple extrapolation from the properties of its elementary components. [….] If one hopes to achieve a full understanding of a system as complicated as nervous system […] then one must be prepared to* ***contemplate different kinds of explanation at different levels of description***.

The levels are not linked by any of the standard notions of constituency, or hyponymy, and nor do they refer to different metaphysical realms of substances, as we saw in the case of supervenience. Instead, each particular level in the theory relates to a different perspective one can take upon a cognitive system: either it is an implementation of some computational theory, or it is a system with a set of representations and a set of algorithms that it performs upon them, or it is a physical system that implements the above two aspects. In fact, on this view a cognitive system is all of these things at once. With each level comes a different set of theoretical notions, thereby eliminating the need for solely one theoretical account. Instead of Marr's term *level of description,* this thesis will use the term *level of interpretation*.

It is worth highlighting the nature of the links between these levels. As just mentioned, each member of the top two levels can be linked to more than one members of the level immediately below it. In other words, more than one set of representations and algorithms can implement any given computational theory, and any one set of representations and algorithms can be implemented in numerous different physical systems. Much has been made of this many-to-one mapping from the lower to the higher

levels on Marr's framework, and the ramifications of this view will be discussed below in the discussion of the notion of *multiple realizability*[6].

Before this discussion, however, it is worth investigating the appeal to the computational metaphor in Marr's and other's viewpoints on cognition, as in many ways these have come to define much of the work in 20th Century cognitive science. The critical notions here are those of *representation* and *computation.*

2.3.1 Representation

For those who adhere to the classical view of cognitive science, the idea of *representations* holds a central position. Bechtel (2002: 332, emphasis added) indicates the importance of this notion in the following way:

> *One of the hallmarks of cognitive explanations of behaviour is that they appeal to* ***mental representations*** *and* ***operations over them****.*

Similar statements are made by Grush (2002: 349):

> *There is no more notion more crucial to the study of thought and cognition than representation. It is the fact that cognitive systems traffic in representations that sets them apart from merely complex and interesting, but non-cognitive, systems like elms, oceans and microwave ovens.*

The importance of representations on this framework in cognitive science is also given by Churchland and Sejonowski (1992: Ch. 4), Churchland (1986: 9-10), Fodor (1975: 27), Edelman (2008: 13), and Dretske (1981; 1996; 2003).

It should also be noted, that in the context of the framework under discussion here, the term *symbol* is often used interchangeably with the term *representation*. This stems from early work in computational cognitive science, specifically the *physical symbol system hypothesis* proposed by Newell and Simon (1976). This hypothesis states that: "a physical symbol system has the necessary and sufficient means for general intelligent action" (*ibid*:

[6] It should be noted that the mention of 'realizability' here, and the discussion in§2.3.3., is not intended to also refer to the notion of 'realization' in SFL. Although the same term is being used in both cases, for the moment I leave open the question concerning whether they are referring to the same or similar ideas.

116). For this reason, the computational framework has been dubbed “symbolicism” by some (Eliasmith, 2003; 2013).

The pertinent question is: what exactly is a representation? The most common theme in discussing this term is that a representation is something that **stands in for something else**. Just as in the checkout till example above, where a symbolic representation within the computational system stood in for a particular price of an object in the external world, symbols in the mind/brain are taken to stand in for something else. This is one of the senses in which cognitive systems are thought to have **meaning**, in a similar way that the referential view in linguistic semantics considers means as a certain link between linguistic units and external objects (c.f. Heim and Kratzer, 1997: 21-22). It is worth noting that mere correlation between some aspect X and another Y does not necessarily entail that X *represents* Y or vice-versa (Haugeland, 1991; van Gelder, 1995). Doing so would weaken the notion of representation to such an extent as to result in it losing all of its claimed explanatory power. Many things can be said to correlate with something else, but if one wanting to say that every instance of correlation was also an instance of representation, then many more things would become cognitive than seems sensible to desire.

If representations are argued to have any value in cognitive science, then, there must be something more to the story than simply correlation between cause and effect. This extra element is often said to be that mental representations have *semantic content* or *intensionality*: mental representations **stand in** for things because there are **about** things. For instance, on seeing an apple, we do so by means of being in possession of a mental representation that is *about* the apple, and this representation ‘re-presents’ the apple to us. (c.f. Edelman, 2008). How exactly mental representations come to have semantic content is an on-going issue within the philosophy of cognitive science. Shortly, we’ll briefly survey some of the possibilities that have been proposed. This overall view, however, on the classical framework of cognitive science is that cognition, and ultimately behaviour, is mediated by the existence of symbolic mental representations, that have semantic content and are therefore meaningful for any organism that possesses them.

One common misinterpretation concerning representations needs to be addressed, namely that an appeal to representations entails the existence of some conscious observer that the symbol is meaningful for. Such a view would potentially posit the presence of some kind of *homunculus*, that is, a 'little person' inside the head who 'reads' all the symbols as they are processed and imbues them with meaning. As well as postulating a rather preposterous ontology, it also runs the risk of being defeated by a *reductio ad infinitum* argument: the existence of each *homunculus* needs the postulation of an additional one that sits in *its head,* and so on with no logical end point (Dennett, 1991). This misunderstanding potentially stems from an over application work in semiotics, where the meaning of a semiotic symbol is dependent upon the existence of an *object* (the thing being represented), a *sign* (the representation), and an *interpretant* (an entity to interpret the sign).

Language is a well-studied example of a semiotic system, and indeed language has been considered to consists of a series of symbols, constructed in particular ways, that stand in for something else, although the precise nature of what exactly words stand in for is a matter for some debate (mental concepts, real world objects, etc.). Of course, in the case of language, a word has a particular meaning (whatever the correct characterization of meaning ends up being) in part because it is used by a particular speech community, which is made up of cognitive agents. It is perhaps natural, therefore, that any talk of symbols or representations to hint towards the existence of cognitive perceivers who "interpret" those symbols. Classical cognitive theorists, however, would not adhere to such a view.

A further reason to doubt the intuition that the postulation of representations implies the existence of a conscious observer comes from the nature of uncontroversial physical symbol systems: digital computers, since they are the paradigm of symbol manipulation systems. On the classical view, it is the nature of the computation itself that generates qualitative experience, and the operation of this computation does not itself need cognitive control. Pylyshyn maintains this view (1984: 40):

> *Many people are familiar with this idea [symbolic representation - JW], though it remains fairly controversial outside of cognitive*

> *science. The uneasiness this causes in some people is absent among those who work with computer models, because this story is precisely, and literally, what happens in computers – with one difference.*

This "one difference" is that, whereas in standard computing systems, meanings of symbols are set by the programmer, there is no external programmer who sets up the meanings of mental symbols in the case of biological cognitive agents. This approach needs some account, then, of how meaning in representations arises in cognitive systems.

The means by which mental representations come to have semantic content has been the subject of much debate. They will be briefly mentioned here, but the criticisms that follow will pertain to both of them, and space does not allow me to give a detailed description. The first group of ideas is known as *causal-informational theories* (Adams and Aizawa, 2008; Rupert, 2008). Examples of such ideas include proposals by Dretske (1981) who says that semantic content can be thought to arise in terms couched in information theory. Information passed from a transmission source (e.g. an object), to a receiver (i.e. an organism) creates a mental representation with the appropriate semantic content. Fodor (1987) believes that a special kind of causal operation creates semantic content, and these causes are a fundamental part of nature.

A second group is known as *teleosemantic theories* (McDonald and Papineau, 2006). One such idea is given by Millikan (1984; 1989; 2004) proposes that the semantics of a mental representation is determined by the function, or affect, that the representation has in the system. Papineau (1998) on the other hand, believes that the semantic content is not determined by an actual affect, but by the affect the possessor expects or believes it to have.

One of the main issues with attempts to provide an account of how the content of a representation gets fixed, is that they all suffer from problems (Neander, 2004). This is known as the *symbol grounding problem* (Harnard, 1990). For instance, Fodor's view seems to simple raise the question of why certain causal laws should have this special status. Teleosemantic theories, one the other hand, have been criticized for lack of explicitness, and the idea that more than one function may be relevant in any given task (Fodor, 1990).

This poses a problem for how the content of a mental representation is to be fixed when there is more than one possible option. One might be tempted to think that difficulties in providing such an explanation provides evidence that there is no such account to be given. More general criticisms concerning the nature of both *representations* and *computations* are given below in the following sections.

2.3.2 Computation

"Computations are, by definition, operations defined upon structured objects; in particular, on objects that have constituent structure" (Fodor, 2008: x). This quote sums up the basic assumptions behind the computational view: a computation must act upon something, in this context usually symbols/ representations, and these objects can themselves be made up of smaller units, i.e. subsymbols.

As in the case of representations, it is important to fully understand what is meant by the term as it can take a quite general meaning that can affect our assessment about its validity. Perhaps the most famous example of a symbolic computational framework is Jerry Fodor's *Language of Thought* hypothesis (LOT, Fodor, 1975; 2008). According to this view, acts such as thinking or believing are based upon the existence of a "mental language", or *mentalese*. It is argued to have properties that are also found in human languages, mainly that they have a combinatorial syntax and semantics. The LOT is primarily discussed in terms of propositional attitudes of the form *S thinks that P*, or *S believes P*, where *S* is any person, and *P* refers to a proposition like *John likes Mary*. The ability to have this thought is dependent upon the proposition being expressed in units of mentalese that can then be combined via a compositional syntactic computation to produce the larger proposition. So, on the above sentence, the proposition *John likes Mary* can be decomposed into its basic units, each of which is a symbol in a language of thought with semantic value: namely the individuals *John* and *Mary*, as well as the two-place predicate *likes*. As the LOT is based upon a propositional analysis, propositions in mentalese are taken to be semantically interpretable, i.e. have reference, truth-conditions, and truth-values. Some of the arguments that are given for something akin to a LOT being necessary for cognition are as follows (Fodor, 1975; 1987; 2008; Fodor and Pylyshyn, 1988):

1. *Productivity*: There are theoretically an unbounded number of thoughts any one person can potentially have. If you think *P* and *Q* independently, then it is also possible to think *P&Q* as a single thought. It would also be possible to think *P&Q&R&…* This capacity is based, however, on a finite number of individual units. This entails some kind of combinatorial computation.
2. *Systematicity:* The ability to think certain thoughts is intrinsically linked to the ability to think others. If, for instance, you can think the thought *John loves the girl,* you can also have the thought that *the girl loves John*. It would be surprising if there was no link between these two thoughts, and the ability to think the latter was completely independent on the ability to think the former.
3. *Compositionality*: The ability for, say, *John loves the girl* and *John loves the cake* are both about the same person is because the part of the thought corresponding to *John* has the same meaning in both cases. Thought works by means of a combinatorial syntactic and semantic structure, in the same way that language is thought to operate.
4. *Inferential Coherence*: If you have a belief like *John is at the supermarket and Mary is at the gym*, then necessarily you also hold the constituent belief that *John is at the supermarket*. It is a logical necessity that this is the case, and the only way that this would work is if the larger belief had an internal syntactic structure containing the two smaller beliefs.

A further argument that has been proposed has been dubbed the "what else could it be?" (Haugeland, 1978), or the "only game in town" argument (Fodor, 1975; 1987). The force of this argument was that, at the time of writing, no other coherent account of how to think about cognition was available. However, with the development of contrasting views of cognition, this argument has lost its force.

What should be clear from the above arguments is that they are heavily influenced by Chomsky's view of language as proposed in his generative framework (Chomsky, 1955; 1980; 1995). Indeed, Fodor is explicit in the influence that Chomsky has had on his thinking, and many of the examples and justifications are directly drawn from the human ability of language.

Basing his arguments on some claimed properties of human language does raise an interesting question: either there is some big difference between the ways that human cognition and the cognition of other animals operate, or if not, why would you ground a basis for cognition *in general* on properties that are, if anything, atypical of general cognitive capacities? The first would be surprising given that the majority of animal cognition is based upon the same set of basic features (neurons, action potentials, etc.), although the complexity of the neuronal structures may be different in each case. The second option seems like an odd strategy, to define the broad properties of the discipline based upon what only occurs in the minority of cases.

In order to highlight that the computational theory of cognition is not identical to the LOT hypothesis, Marr's (1982) theory of vision is also based upon computational principles, without attributing a propositional like structure to thought. According to Marr's theory, the computational process of vision involves a step-by-step computation whereby a representation is transformed by a number of set stages. The input of the computation is light activating retina cells, while the output is a representation for a 3-D image of whatever is being observed. The first step of the computation involves light hitting the retina, resulting in a representation of the observed object consisting of just the intensity of light at the various points of the visual field. This initial, basic representation then goes through various transformations that change it in various ways. The first transformation produces a *primal sketch*. This representation using information about the broad intensities, and how they change in the overall picture, to produce a rough idea of edges, curves, lines, blobs, and so on. Later transformations produce a *2.5-D sketch*, which adds extra information such as basic depth, surface orientation and the spatial contours of the object seen. The final transformation takes this representation as an input and produces a full *3D sketch*, which provides details about how the various lines, shapes and depths fit in hierarchically to produce one coherent three-dimensional picture of the object observed. Although Marr provided some neurobiological evidence for his proposals, his main concern was to provide a computational story of how vision could occur. As such, many of these hypotheses, especially the idea of the *2.5D* image were not motivated by biological facts. Its computational nature can also easily be seen; the process of vision is characterized to operate upon information-carrying representations that are transformed

in a serial manner, with each transformation adding incrementally more detail, resulting in the desired outcome.

Some commentators have questioned the notion of cognitive processes acting like computations. One of these criticisms comes from Searle (1980), who proposed the *Chinese Room* argument in an attempt to raise some of the problems with this approach. The argument attempts to show that it is possible to have a fully functioning program that performs a task perfectly, but in which there is a lack of intentionality. This lack of intentionality is particularly problematic for the classical approach, as it is a major hallmark of representations as discussed above.

His argument proceeds as follows. Imagine a person (Searle imagines himself) locked in a room with a stack of Chinese symbols. Furthermore, let's imagine that Searle neither understands Chinese to any extent, nor can even confidently judge the symbols of instances of Chinese. In other words, the Chinese symbols are absolutely meaningless for him. Next, another set of Chinese characters is presented to him, along with a written set of instructions, in English, that links the first set of characters to the second. Some outside agent then slides Chinese symbols to Searle under the door, and following the written instructions he matches this symbol with ones on the first list, finds the corresponding symbol(s) on the second list and gives this as an output. The important fact is that to an outside observer, it seems as though Searle is a fluent Chinese speaker as his actions correspond completely with how a fluent Chinese speaker would behave. The crux of the argument is that although a system can follow a computational strategy according to a set of algorithms (the English rules in this case), it can take place in a situation where there is no meaning. In doing so, Searle was arguing against computational based approaches to cognition, by arguing that a purely syntactical, algorithmic based computation is never enough to achieve intentionality, which is an important part of the symbolic account. It should be clear that Searle's arguments are closely related to the *symbol grounding problem* that was discussed above in §2.3.2, as they both concern how exactly symbolic structures in the mind achieve their meaning. Searle's argument is that, assuming that the computationalist view is correct (which he doesn't, c.f. Searle 2004), it cannot be down to syntactic rules operating over potentially meaningless symbols.

Another criticism of the idea of symbolic computation as the means for understanding cognition is that the conceptualization of linear algorithms is too far removed from what actually takes place in the brain. For instance, both operations on a language of thought, via some kind of mental combinatorial syntax, or the transformation of visual representations, both operate in such a manner. These criticisms amount to saying that the symbolic approach has idealised cognition to such an extent that they fail to appreciate biological facts about organisms, where information seems to be processed in a parallel, and not linear, fashion. Many of these criticisms originally came from those who developed the connectionist framework (c.f. §2.4 below, and the references cited there), who went on to attempt to capture the nature of cognition in terms of developing artificial neural networks that processed incoming information by parallel and not linear means.

This section has given an account of what is meant by the notion of *computation* in classical accounts of cognition, as well as some of the problems that have been raised against it. Next, we turn back to the notion of *levels* within this view, by discussing the notion of *multiple realizability* in classical work.

2.3.3 Multiple Realizability

In the discussion of Marr's three-level framework above, it was noted that the links between the lower and higher levels are to be considered as a many-to-one mapping. That is, many different physical systems can implement a certain set of representations and algorithms. Additionally, many different sets of representations and algorithms can realize any given computational theory. This general form of many-to-one mappings between levels is known as *multiple realizability* (MR).

Although holding such a position seems intuitively plausible, it has led some proponents of the classical position to propose a more controversial argument. Simply put, due to MR a full understanding of cognition need not take into account any of the facts about the physical system that is implementing the higher levels. In other words, you don't need to consider facts about the brain as a neurophysiological system in order to study cognitive processes.

Fodor (1974; 1975; 1997) is closely associated with this view. He argues that, even in principle, the "special sciences" (e.g. psychology and sociology) cannot be reduced to more basic sciences such as physics because of the vast number of disjoint ways in they can be linked. More formally, he argues that the "bridge-law" between special sciences and physics would involve linking one, say, psychological law with a possibly unknown large disjunct of physical predicates, be they neural, electronic, or potentially of many different sorts

One immediate consequence of such an approach is that as soon as one considers MR to be true, then the importance of the physical implementation of any representational/ computational system quickly diminishes. It creates the situation whereby the study of cognition is no longer as dependent upon the facts of the physical system. Wilson and Clark (2008: 56) sum up the situation as follows:

> *[T]he neurosciences were often viewed, at least until the 1980s, as a kind of junior partner in the venture to understand an explain cognition. Thus, while the neurosciences could tell us much about the* ***realization*** *of cognitive processes and structures in certain kinds of cognizers, their findings could not usurp the investigations of psychologists, linguists and computer scientists, who were exploring the functional structure of mind*.

The reason for this is: if higher level descriptions can be realized by an extremely large number (possibly infinite) of physical descriptions, then the only way the nature of the lowest level in the model could inform the higher levels of representation and computational theory would be if there was some common feature among the myriad of possible realizations, besides the fact that they all can realize the representational and computational characterization. However, this commonality is doubted by MR adherents. Fodor (1975: 20), for instance, argues that the possible set of physical realizers of any predicate in the special sciences "may be […] heterogeneous and unsystematic." The intuitions that are usually drawn upon include both octopi and humans realizing the same psychological kind *pain* in neural systems with different physical make-ups (Putnam, 1967), or whether it is possible for robots, humans and Martians with completely different physical descriptions can all be described in terms of the same psychological type. If the only

property common to all of these realizations is their representational character, then one can ignore large aspects of physical level, and simply treat cognition in abstract, computational terms.

This account has been subject to criticism. Bechtel and Mundale (1999) argue that just because no systematic link between psychological and physical states has yet to be found, it does not necessarily mean that one does not exist. The concepts currently utilized in the cognitive sciences may not reveal any, but a change in our conceptualization of the issue in the future may reveal systematic relations that are currently invisible to us.

Shapiro (2000) notes the absence of a consistent diagnostic for determining whether two given physical systems are the realization of the same kind, leading to the situation whereby "philosophers are much too quick to claim that a given kind is multiply realizable" (*Ibid*: 636). As an example, he considers a thought experiment whereby a person's brain is replaced piece-by-piece with silicon chips. This replacement takes place in such a way that the overall organization of the brain is maintained with each step, until finally we end up with a complete silicon brain that operates the same way as biological brains. He rejects the conclusion this thought experiment is meant to show, namely that psychological notions are multiply realizable due to the different implementations in biological and silicon brains. Instead, he argues that the two cases are not distinct realizers, since they both exhibit the same relevant physical properties, namely the structural organization and the ability to receive and pass on electrical signals. All the other physical differences, the materials, etc. are in this case unimportant or uninteresting. In the ways that matter for the task at hand, the two systems are physically identical, and therefore MR does not hold. In raising this point, Shapiro is echoing Adams (1979: 158), who says "just because two [physical-chemical] systems are different kinds of stuff does not mean that they do not share some identical property-kinds".

Polger (2002) takes the opposite tactic in his criticism. He states that in order for MR to carry the force it does, it needs to show that two creatures (say humans and aliens) could have radically different physical make-ups and still be in *exactly the same* psychological state. He doubts this is the case, or at least argues that it may not be as obvious as it has been claimed. Drawing upon Nagel's (1974) famous argument about conscious states of bats, and

the similar case of dolphins, he argues that due to the physiological differences in these animals we are less likely to attribute to them exactly the same psychological states as we do to ourselves. In the case of aliens, or robots, there is no *a priori* reason to claim that they have identical psychological states to our own, and as a result MR, while not refuted, is in need of further backing (*Ibid*: 153):

> *It may seem […] that Martians could have sensations like ours. But it does not seem to me that even avian or piscine sensations are quite like ours, although I am inclined to believe that birds and fish have sensations of some kinds. This might be wrong. But if it is we will have to do much work to discover that other creatures have sensations emphatically like our own. It is certainly not a platitude of common-sense psychological reasoning that all conscious terrestrial animals have experiences that are exactly, empathetically of the same kinds. And if they do not, then multiple realizability never gets a foothold.*

I agree with many of the above criticisms about the relationship between the higher and lower levels in classical theories of cognition and Marr's hierarchy. Arguments for the relative unimportance of the physical, implementational level have so far been argued for from intuition, and not backed up with any empirical support, even though even some of its adherents claim it is fundamentally an empirical thesis (Putnam, 1975; Fodor, 1975).

A point of clarification: the arguments presented and reviewed here should not be read as stating that it is impossible that other systems apart from biological organisms with neural systems can be cognitive. That possibility still remains. However, one should not assume that it is the case, and solely based upon intuitions one should not argue for the detachment of the study of cognition from neurobiological facts. As far as we know, the only systems that are capable of cognitive behaviour are biological ones, and this needs to be kept in mind.

2.3.4. What do you study on a classical view?

So far, this section has reviewed the following aspects of the classical framework in cognitive science:

- Marr's three-level hierarchy, and the nature of levels on this framework.
- The notion of representations in terms of symbols.
- The prominent role of computation in this view.
- The idea of multiple realizability, and the purported detachment of cognitive science from neurobiology.

To conclude and summarize this discussion, it is worth briefly discussing what aspects of the world are taken as important for studying cognition from a classical perspective. Marr's three-level hierarchy is again useful for this discussion. In the first instance, you need to consider the type of computation that an entity has to achieve, be it counting or perceiving an object, producing speech, etc. Next you need to work out the linear, step-by-step algorithm that can achieve this, as well as deciding the kinds of symbolic representation that the entity needs to be in possession of in order to achieve this goal. Finally (if so desired) an account needs to be given concerning how this computation would be implemented in a physical system, be it a neuronal system or a digital computer.

It is also useful to consider what you *need not* consider on adopting such a view, as taking a classical perspective on the nature of cognition means that many details about the daily life of entities are ignored. Firstly, you do not need to consider many facts about the external environment. Of course the existence of something outside of the cognizing agent (permitting the possibility of artificial intelligence) is required, say, for setting the content of mental representations or to provide input for the computation, but it has a distinctly passive role. The main explanatory target is to give an account of the transformational or computational strategy that the holder of mental representation has to go through. Furthermore, the role of action is back-grounded on this view. Many cognitive phenomena that the classical framework concerns itself with are of the more 'passive' variety. As was seen above, many of the discussion on this issue have centred around aspects of our cognitive life like perception (Marr) and knowledge of states of affairs (Fodor). Although, action would be intended to be explained, it would only be upon the viewpoint that it is led by such passive cognitive abilities. They are, therefore, kept separate; on this view the goal of cognitive science is to explain the knowledge and belief capacities that go on between sensing and action.

It should be noted, that I do not agree with this detachment from environmental concerns. My biggest concern is that the appeals to notions such as multiple realizability enable one to dismiss many potentially important facts about cognition. When one stops considering the lowest level on Marr's hierarchy to be important, facts about the organism cease to be relevant. Furthermore, the appeal to computational notions, as well as possessing their own problems, mean that cognition is in danger of being considered as a passive, purely information processing act, with less emphasis on features such as acting, behaving and existing in a particular environment. In §2.5, the importance of these features will be emphasized, making these criticisms clearer.

Next we will consider a competing view that, although one which still involves the idea of *computation*, took issue with some facets of the classical accounts.

2.4. The Connectionist Framework

In the last 30 years, there has emerged a different view of cognition that intentionally set itself up against the classical view of cognition as symbol processing computational systems. It has become known as *connectionism* and was entered into the discussion via the works of James McCelland and David Rumelhart (McCelland et al., 1986; Rumelhart et al., 1986). Also known as the *Parallel Distributed Processing* approach, it took issue with the classical conception outlined above, particularly with its dependence on symbol-manipulation, step-by-step algorithmic explanations and appeals to computational structures such as von Neumann architectures.

The motivations they give for their approach are that: (1) a classical, computational framework posits that cognition involves *linear* computation, whereas much of the computations in the brain are performed in parallel (2) their alternative is much closer to neurological reality, with the nature of the models inspired directly by neural facts, and (3) it focusses more on the kinds of tasks that cognitive agents are actually engaged in, rather than focusing on more passive phenomena like perception of objects or an explanation of belief based upon logical connections between symbols in a language of thought. (McCelland, Rumelhart and Hinton, 1986).

The major factors in a connectionist model are as follows (Rumelhart, Hinton and McCelland, 1986: 46):

- A *set of processing units*
- A *state of activation*
- An *output function* for each unit
- A *pattern of connectivity* among units
- A *propagation rule* for propagating patterns of activity throughout the network
- An *activation rule* for combining inputs impinging on a unit with the current state of that unit to produce a new level of activation for the unit
- A *learning rule* whereby patterns of connectivity are modified by experience
- An *environment* within which the system operates.

The clearest way to show all of the above, is by providing a simple sketch of one simple kind of connectionist architecture: the simple feedforward network (McCelland and Rumelhart, 1986; image from Jain et al, 1996: 38)

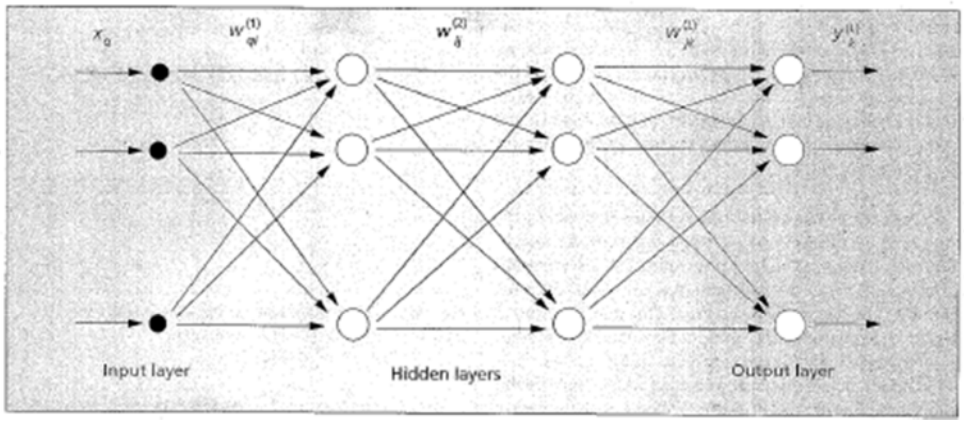


Figure 2.3.: A simple feedforward network

Each of the nodes in the above figure are members of the set of processing units, and the arrows between nodes represent the connectivity pattern in the network. In many networks, nodes are hierarchically arranged into at least three distinct *levels of organization*. The input layer contains those units which directly receive signals from the environment, and the output layer describes how the activity in the whole network interacts with the wider environment, or can describe the result of the network computation. In between these two levels containing number of *hidden* units. They are

considered *hidden* because they are not directly connected to anything not itself in the network. The hidden units play an important role in connectionist architectures, and are the locus of many of the learning algorithms in the framework. These factors will be discussed shortly. At any specific point in time, each node in the node will have a particular state of activation. Usually, this can either be discrete (each node is either on or off), or continuous (each node can exhibit an activation level taken from a pre-determined set of possible values). The output function determines how the signal emitted from a node affects the other nodes it is connected to.

The directions of the arrows indicate direction of signal propagation. In this toy example, they all operate from left-to-right. Taken together, the arrows provide the pattern on connectivity within the network. In this example, as with many others, the network exhibits widespread connectivity: each node is connected to all other nodes in the next organizational level. The arrows on the far left indicate input originating from a source external to the stated network. Similarly, the arrows emanating from the far right represent output signals from the network having an effect on an external source. Taken together, these give an indication of the environment in which the network is operating. Furthermore, the connections between nodes are *weighted*, i.e. they are defined as a particular function that specifies how strong or weak the connection is. Stronger connections will have a larger effect on the next node, while weaker connections will have less effect. These weights are variable, so as the network is exposed to a certain training task, useful connections will be strengthened, while others will be weakened. In this way, networks can adapt themselves to training data. The means by which weighted connections alter is given by the learning algorithm, which is set by the experimenter in advance.

Above, it was mentioned that the hidden level of units are prominent in the learning algorithms of connectionist networks. A connectionist network is of little use if there is no means for it to adapt itself so as to confirm with the output desired by the designer of the network. The most common method of learning within connectionist frameworks is known as *backpropagation* (Rumelhart, Hinton and Williams; 1986). Crucial to this type of learning algorithm is that each of the connections between units at different levels within a network has a certain weight attached to it; this corresponds with how “strong” or “weak” a

connection is. Furthermore, these connections are not fixed over the course of a network simulation, but can be altered. This will have an effect on the outputs that a given network will produce, with units at the output layer with predominately strong connections being more likely to be activated given appropriate input than others. The mathematics of backpropagation relies upon the knowledge of two factors: the *actual* output of the network, and the *desired* output of the network. With knowledge of both of these facts, the computational system can determine how far away the actual output is from the output that is needed in order to successfully accomplish the cognitive task, in other words the *error* of the system. A calculation known as the *delta rule* calculates the total error, as follows:

$$Error = \frac{1}{2}\sum_{u}(d_u - a_u)^2$$

In the above equation, u refers to units at the output layer, d_u is the desired output of a certain output node and a_u denotes the actual output seen. The backpropagation rule uses this total error to work out the connections in the network that contribute highly to this overall error, and alter the weight of that connection in such a way to decrease the error observed. It begins by altering the weights from the hidden layer to the output layer, using the following calculation:

$$\Delta weight_{uh} = lrate(d_u - a_u)a_u(1 - a_u)a_h$$

$\Delta weight_{uh}$ denotes the change in weight from the hidden layer to the output layer, $lrate$ refers to a learning rate is a constant that determines how fast the connections change. All other symbols are defined as above. It can be seen above that the same error calculation is crucial to determining how much the connections should change. *Backpropagation*, however, can also use the same error calculation to affect connections deeper in the network:

$$\Delta weight_{hi} = lrate\ a_h(1 - a_h)\sum_{u}(delta_u\ weight_{uh})a_i$$

All symbols are as above; with the added meanings that i denotes an input note, $weight_{hi}$ refers to a connection from the input to the hidden layer, and $delta_u$is the inverse of the

change in error given how the input to the network changes. In sum, given knowledge of actual output of the nodes in a network, the desired output of the network and some constants such as the learning rate, connectionist networks can adapt themselves to achieve certain tasks.

So far, we have seen the basic features of connectionist networks: they consist of hierarchically structured nodes that are linked by alterable weighted connections. They also receive signals from, and output to, some outside source. One last feature of artificial neural networks needs to be discussed: how the nodes in a connectionist network come to represent certain features that are seen to be important for determining the success of a given network. Unlike in the classical framework, where the semantic value of the symbolic representation is determined by some causal or teleological factor (although the details are not certain), within connectionist networks nodes are inherently meaningless, and have to be given a semantic value by the network builder. There are two means of representing meaningful features within a connectionist network: either by means of ***local*** or ***distributed*** representations (Hansen and Burr, 1991). The major difference is that in local representation, each node is used to represent a single concept, word, etc., while in distributed representations, the representation of, say, a word is not associated with only one node, but with many. In this sense, the representation is *distributed* across more than one unit. Networks using local representations are easier to represent, but come at the cost of increased neural implausibility, and with a danger of pre-made ideas of how concepts should be encoded entering into the simulation (Bechtel and Abrahamsen, 2002: 41). It should be emphasized that, in this approach, we are still dealing with **computations**, as these networks are computationally simulated, and **representations**, as units within the networks are set to 'stand-in' for other things, either by local or distributed means.

This section has outlined the basic features of the connectionist framework within cognitive science. To reinforce these points, the next section will briefly look at some implementations of the perspective.

2.4.1. Applications of the Connectionist Framework

Perhaps more so than with other conceptual frameworks, an appreciation of the value of the connectionist framework comes from seeing the range of cognitive tasks that have

undergone connectionist modelling. Connectionist models have been proposed to explain a range of cognitive behaviour, from face recognition (DeMers and Cottrell, 1993), to visual and audio-visual attention effects, and their role in development and learning (e.g. Harrison and de Kamps, 2011), amongst others. However, given part of the subject matter of this thesis, and since it remains an active area of research in the connectionist movement, we will turn now to overview how connectionist networks have been applied to language-related issues.

An example comes in the form of the NETtalk model that aimed to provide a means of converting orthogonal input to phonetic output in order to achieve a mapping between written and spoken language (Sejnowski and Rosenberg, 1987). The architecture is sketched as follows:

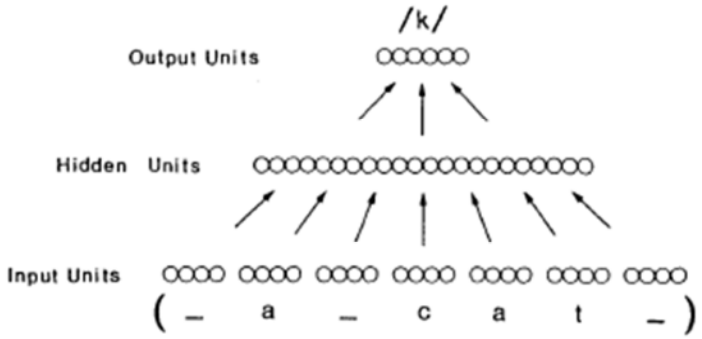


Figure 2.4: NETtalk model (Sejnowski and Rosenberg, 1987: 147)

Above, certain units at different organizational levels are associated with different interpretations. The input layer was divided into seven groups of nodes, with each group representing one orthographical letter. Each input group was made up of 29 units, one for each letter in the English alphabet, plus extras for punctuation. For example, if the input node was representing the letter <c>, only the node representing <c> would be activated. In this way, the network was presented with a seven letter sequence with each input, including gaps, full stops and other punctuation. The output layer consisted of one group that represented the phonetic output. The task of the network was to match the letter in the middle of the seven character string (<c> on the above picture) and to provide the correct phonetic output. The other letters in the input string provided informational context, to reflect the fact that the production of a particular phoneme can be affected by

the other phonemes in the same phonetic environment. After the middle latter was matched with an output, the input string was moved in a step-by-step manner, so that a new orthographic input occupied the middle position. By similar means, the output group consisted of 26 units, each one either representing articulatory features such as voicing, manner of articulation, etc. or representing other features such as intonation or stress. The hidden layer are not given a direct interpretation, but instead provide a **distributed** internal representation linking the input and output layers. Initially, the network was trained on a corpus of 1024 words from the continuous speech of children in the first grade of American elementary school. The authors showed that the network 95% accuracy after 50 passes through the entire corpus. To test the performance of the network, the researchers took a network trained on the initial corpus and exposed it to a novel passage of continuous speech containing 439 words from the same speaker. In this case, the network showed 78% accuracy. The network also showed partial resistance to damage, and the ability to rapidly re-learn after output was interrupted by intentional damage.

It was mentioned in §1, a relatively new set of *deep learning neural networks* have been proposed. These exhibit a larger processing power than the simpler example just given. Below is an example from Collobert and Weston (2008), which has been designed in the context of *natural language processing* (*ibid*, 163):

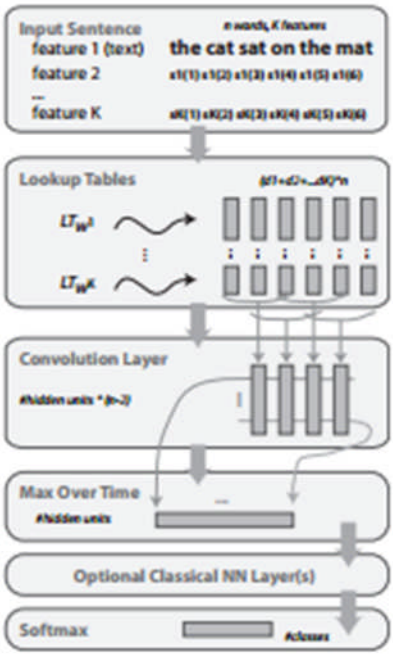


Figure 2.5: Deep learning neural architecture for natural language processing

The details are more complicated in this example, but it is enough for this work to appreciate that each of the grey areas in the above figure represents a given neural network. These have been organized into a series of *levels of organization*, with different tasks being undertaken at each stage. Setting these architectures enables researchers to achieve certain cognitive-like output, and on the motivations of the approach, provide a hypothesis into how biological systems may achieve similar tasks. It is this link to the brain that this discussion now turns.

2.4.2. Connectionism and the Brain

Without doubt, there exist today numerous impressive connectionist accounts of psychological phenomena. However, one controversial question remains: how much can connectionist networks tell us about functioning brains? Certainly, one of the appeals of connectionist networks is the inspiration it took from neural architectures and the use of single units with connections of variable weights, although the instigators of this framework have stated that this was of secondary importance in contrast to the empirical results gained (McCelland, Rumelhart and Hinton, 1986: 10). In order to frame this discussion about the links between consciousness and the brain, it is useful to introduce

the *levels of interpretation* discussed by Smolensky (1986) in the context of the connectionist framework. His description can be sketched as follows:

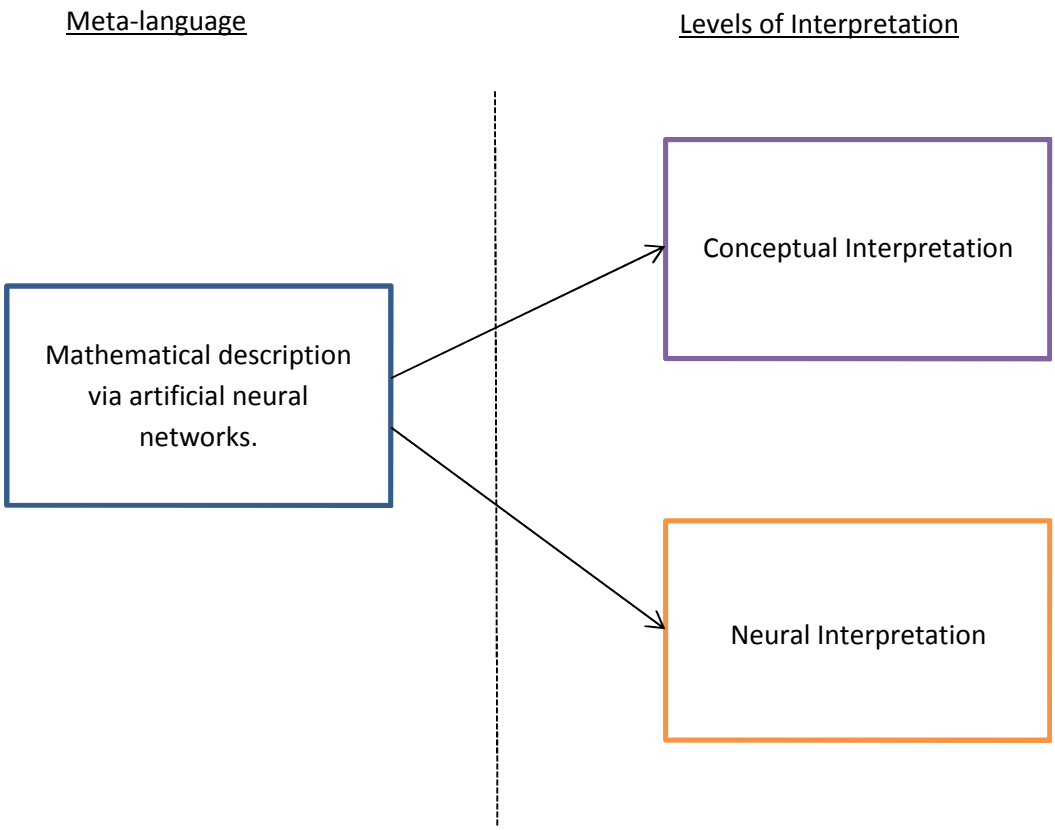


Figure 2.6: Metalanguage and levels of interpretation in connectionist networks (based on Smolensky, 1986: 392-393)

The intuition between the two interpretative levels in the above sketch is that one can consider a connectionist network in one of two ways: either it is a mathematical model of real neural systems (the lower interpretative level above), or it can be considered to be a mathematical model of conceptual or behavioural output: the ability to recognize a particular face as male or female, or the ability to correctly visually identify a particular letter in a word. The best case scenario would be if the mathematical description given, via the simulation of a given connectionist network, provided simultaneously a model of both of these interpretations. In other words, ideally a connectionist description provides a homomorphism between the conceptual and neural interpretations (*ibid*).

Some of the links between a connectionist model and the two interpretation levels are summarized as below (*ibid*: 393):

Neural	Mathematical	Conceptual
Neurons	*Units*	*Hypotheses*
Spiking Frequency	*Activation*	*Degree of Confidence*
Spread of Depolarization	*Spread of Activation*	*Propagation of Confidence: Inference*
Synaptic Contact	*Connection*	*Conceptual-Inferential-Interrelations*
Excitation/ Inhibition	*Positive/ Negative Weight*	*Positive/negative inferential relations*

Table 2.1: Relations between the neural, mathematical and conceptual levels in a connectionist framework.

Details about the neurobiological aspects above are given in §3.8.2.2.1. Although it has always been inferred that the hope was connectionism could provide the means of linking these levels of interpretation, in practice they have operated more-or-less independently. On the subject of neural interpretation, some of the early work within the connectionist framework attempted to provide a more realistic neural basis for connectionist models (Zipser, 1986; Munro, 1986). It is also interesting to note that in a recent survey of connectionist approaches to language learning, Joanisse and McCelland (2015: 1) state that connectionist approaches to this issue are beneficial because they propose "a model of the mind that closely reflects the functioning of the brain".

However, the biological plausibility of connectionist models has frequently been called into question (Green, 1998; 1999; Bechtel and Abrahamson, 2002; Eliasmith 2003; 2013). Green (1998; 1999) argues that nothing analogous to the hidden level of units, which is crucial in many connectionist approaches, has been discovered neurologically. Furthermore, he argues that connectionist networks are not constrained enough, i.e. with the correct setting of number of nodes and weighting functions you can design networks to achieve almost *anything*. These concerns were also part of Fodor and Pylyshyn's (1988) critique of the connectionist enterprise. The implication is that the success of connectionist

frameworks may be down to the skill of the programmers, over the fact that such models have biological relevancy.

Bechtel and Abrahamson (2002: §10.3.1) outline a number of biologically implausible aspects of connectionist networks. These include: (1) the fact that units, once activated, emit a single signal output, whereas neurons can emit a train of action potentials, (2) units have one means of summing incoming signals to determine whether to activate or not, whereas signal propagation in biological neurons happens via a wide range of chemical neurotransmitters, each of which have different properties, (3) connectionist neurons treat all inputs homogenously, whereas in biological neurons the specific location of a synapse on a neural dendrite can be relevant, (4) in classical connectionist frameworks there is full connectivity between units at adjacent organizational levels, while biologically connectivity patterns differ depending on the location and type of neuron, and (5) the major means of learning in connectionist networks is by the process of *backpropagation*, and there is no evidence that any analogous mechanism is found neurologically.

Other proponents of connectionist models concede this point. Houghton (2005: 3-4) says the following:

> *Connectionist modes of cognitive function are said to be "neurally inspired". However, it is fairly rare to come across models containing features specifically based on data regarding brain mechanisms involved in the particular capacity being modelled. In most cases this is simply because nothing very specific is known about these mechanisms. […] However, it is necessary to start somewhere, and current work has shown that very interesting and empirically successful models can be constructed using a handful of basic ideas.*

Many connectionist modellers, it seems, are interested primarily in the conceptual interpretation of their models, without focussing overly on the neurological correlates they may have. Considering the applications of connectionist models above, it is relatively easy to see some of the assumptions they are making. On both localist and distributed networks, each node is given a semantic interpretation by the programmer, and it is

unclear currently whether such attribution of semantic meaning to units has any correlates in the brain.

It should be noted, that I do not believe that it is *impossible* that connectionist models can provide useful means for studying cognition, simply that its current state is a long way from realizing the idealistic linking between the two interpretative levels discussed above. The only means to know whether connectionist models of a particular psychological phenomenon held neurologically viability would be to discover more about how such processes work in the brain. However, on such a discovery, connectionist networks lose some of their explanatory value. Let's assume that a future theory of neuroscience discovers that the differences in the types of neurotransmitter involved in signal propagation in biological neural network has some important explanatory value in a theory of cognition. I'm sure that on the discovery of this fact, this could be incorporated into artificial neural networks, so that not all inputs are summated in an equal manner. But this is not a fact that can be gained from neural networks **before** the fact is discovered neuro-biologically. Connectionist architectures can provide hypotheses about whether or not such factors are relevant based upon their simulations, but whether or not these or other factors are important in the only system that uncontroversially produces cognitive behaviour remains within the domain of modern neuroscience.

To summarize, connectionist networks have the advantage of being more biologically plausible than the classical framework, which generally sees the business of neuroscience as of secondary importance to a description of cognition at an algorithmic and representational level. However, although they are neurologically inspired there are still major differences between the operation of connectionist networks and the operation of biological neural networks. Whether or not these differences are important is an unanswered question, but it can only be answered with further neuroscientific inquiry. Given that organisms with biological neural networks are currently the only entities that without doubt exhibit cognitive behaviour, this is an important shortcoming of the connectionist enterprise.

2.4.3. Studying Cognition in a Connectionist Framework

The aspects of connectionism discussed in this section were:

- An overview of artificial neural networks, and the important role they play in connectionist theorizing.
- The notion of representations and computations on this approach. The distinction between *localist* and *distributed* representations was discussed, as well as the computational nature of the framework.
- Some of the results that have been gained by those using such an approach.
- The links between these models, the brain and cognition, via. Smolensky's conceptual framework.

Seeking to provide an explanation within a connectionist framework, therefore, amounts to developing artificial neural networks that can achieve the same behavioural or cognitive output that one would expect from (generally) humans. Much of this work involves setting an appropriate training set of inputs, determining the number of units in the network, how they are to be arranged into hierarchical levels of organization, determining the connections between them, and finding the correct learning algorithm.

Interpreting connectionist networks involves looking at how the network adapts itself over time to the set task, and by looking at how information about this input is distributed across the network. By doing this, connectionist networks can provide hypotheses about how conceptual information may be represented in analogous biological neural nets. They can also be interpreted as providing hypotheses about how such processes may work neurologically, although there are many differences between artificial and biological neural networks that make this interpretative link unlikely as things stand.

As with the classical framework, the role of the environment plays a limited role on many connectionist models. Usually, the environment includes a training set of static items that are used to provide stable input into the network. On the models mentioned above, these took the form of lexical items that were presented to the network. The output also is distinctively static: the production of a set of static phonetic symbols, or the correct recognition of a face. Therefore, it is primarily concerned with aspects of cognition such as

knowing or *perceiving*, and less interested in how these aspects of cognition drive behaviour to achieve goals in a constantly changing environment. Although there is a connected external source of information, it is much more impoverished than the rich environments that organisms interact with in daily life.

So far, this chapter has given an overview of the thesis of supervenience, the classical framework and connectionism. In the main, all of these approaches can be characterized as 'internalist'. That is, the main focus is on factors internal to the brain/ mind, with less focus on other aspects of the organism and the world. The next set of approaches aim to "broaden" the study of cognition to include such factors.

2.5. 4E approaches

The subject matter of cognitive science is the mind, and so far most of the previously discussed conceptual frameworks have addressed the nature of this phenomenon by focussing on the brain, or matters closely linked to it. For instance, philosophers usually talk about supervenience in terms of the link between mental properties and activation of certain neurophysiological structures firing. Classical approaches consider a mental state to be defined as a certain computation that could be implemented in the 'hardwire' of the brain. Connectionist theorists, on the other hand, study cognitive functions by creating computational models that are inspired by biological structures in the form of artificial neural networks. In the main, all these ways of conceiving the mind do so while ignoring or underemphasising the role the organism's body and the external environment may have on the cognitive life of organisms.

Recently, many researchers have argued against these overly **brain-centred** frameworks of cognition and have since argued for the important role that the body and the environment have in the study and nature of cognition (c.f. Lynott et al, 2014 for recent studies embracing this view). The collective name for approaches to cognition that more overtly include the body and the environment has been dubbed *embodied cognition* (Wilson, 2002; Cowart, 2005; Shapiro, 2007; 2010). Rowlands (2010: 3) dubs this group of related but non-identical viewpoints *the 4E* or *Non-Cartesian* approaches, since technically *embodied cognition* is only one of the possible options. He gives them as follows:

1. Embedded Cognition
2. Embodied Cognition
3. Extended Cognition
4. Enacted Cognition

In this work, we will continue to use the term *4E approaches* to designate the above collection of views. In the rest of this section, each of the above ideas will be discussed in more detail.

2.5.1. Embedded Cognition

One way to widen the study of cognition outwards away from simply focussing upon internal processing or connectionist models that take inspiration from neuronal structures is to make more explicit the idea that brains do not work in isolation. Every organism that possesses a brain or some other type of neuronal structure also has other parts that make up his/ her body, and this organism is constantly interacting with the wider environment. Although this may seem readily evident, we have already encountered a number of theories of cognition which underplay this fact. One means of overcoming this is to consider cognition to be embedded, or situated, and that the nature of this embedding has real effects upon the nature of cognitive life.

One of the first attempts in modern cognitive science came in the form of *dynamic cognition* (van Gelder and Port, 1995; van Gelder, 1995; 1998a.; 1998b.; Beer, 2000; Spivey, 2007). Beer (2000: 91) describes the enterprise as "[seeking] to understand the unfolding of [a cognitive] processes over time and the multiple internal and external influences whose interplay shapes this unfolding." Researchers who adhere to the dynamic view of cognition argue that previous attempts to understanding cognition, especially symbolic theorists, considered the phenomenon to be too static. Van Gelder and Port (1995: 1-2) explain their misgivings with the computational approach as follows:

> *The computational approach makes a series of […] assumptions. Representations are transformations from one static symbol structure to the next. These transformations are discrete, effectively instantaneous, and sequential. […] Many impressive models of*

> *cognitive processes have been developed within the computational framework, yet none of these models are wholly successful even in their own terms, and they completely sidestep numerous critical issues. […]* ***The heart of the problem is time. Cognitive processes and their context unfold continuously and simultaneously in real time.***

In order to fully encapsulate this intuition of cognitive processes "unfolding in real time", Dynamic approaches to cognition rely heavily upon the mathematical tools utilized in the wider field of dynamic systems theory. Previously these tools have been used to model and explain natural phenomena such as weather systems and the kinematics of human physiology (Fusella, 2013). Some of the basic notions critical to an understanding of dynamic systems are as follows:

a) *System.* A system is defined as a set of aspects of the world that are changing. Usually, different aspects of the world are considered to be parts of the same system if they interact. That is, if a change in one aspect of the world is dependent upon how another aspect changes, then those two aspects should be considered to be part of the same system.
b) *State*. A state of a given system is the description of the way all the different aspects of the system are at a single moment in time.
c) *State-space*. A state-space is a multi-dimensional mathematical structure that gives all the possible states a certain system can be in.
d) *Behaviour*. The behaviour of a given system is a description of the way the system changes over time. In other words, the behaviour of a system can be thought of as the movement from one state-point to the next, and so on, within the overall state-space of the system.

An explanation in terms of dynamic cognition, then, is to investigate how the behaviour of a system changes over time. Its embedded nature comes from looking at how the behavioural dynamics of a neural system is influenced by external factors. For instance, one can investigate how the nature of the environment affects the dynamics of the neural system, by biasing it to behave in certain ways and not in others. Dynamic approaches have been applied to a number of cognitive phenomena, such as decision and reaching tasks in

young children (Thelen et al. 2001), analysis of a walking automaton (Beer, 1995) and a description where an agent needed to catch moving targets and avoid others (Beer, 2000). Beer (*ibid*) has also proposed a multi-level framework for the dynamic account, but it also involves incorporating an embodied view of cognition. As such, this will be discussed below in the next section. However, it is worth at this stage emphasizing his views on the notion of *embeddedness* in cognitive life (Beer, 1995: 182)

> *I cannot overemphasize the fundamental role the feedback plays in this relationship. Any action that an agent takes affects its environment in some way […], which in turn affects the agent itself through the feedback it receives from its environment. […]. Thus, each of these two dynamical systems is continuously deforming the flow of the other, […] and therefore influencing its subsequent trajectory.*

I agree with the above view. Simply treating the environment as a stable factor providing more or less stable input to the cognitive system ignores a potentially important fact about the mental life of all organisms, namely that **what we do is both affected by what is around us and in turn we have an effect on it.**

A further example of a framework that should be identified as embedded is proposed by Morton and Frith (1995), discussed further in Frith (1995) and Frith and Happé (2001). Based upon their studies into the various factors that can lead to psychopathological conditions, such as autism and dyslexia, the authors argued that one can only appreciate the whole explanatory story by considering causal factors at various levels: the biological, cognitive and behavioural. Furthermore, the authors emphasise that any model of behavioural deficits should include an account of environmental factors (c.f. also Johnson and Morton, 1991). We can sketch their conceptual framework as follows (Morton and Frith, 1995: 371):

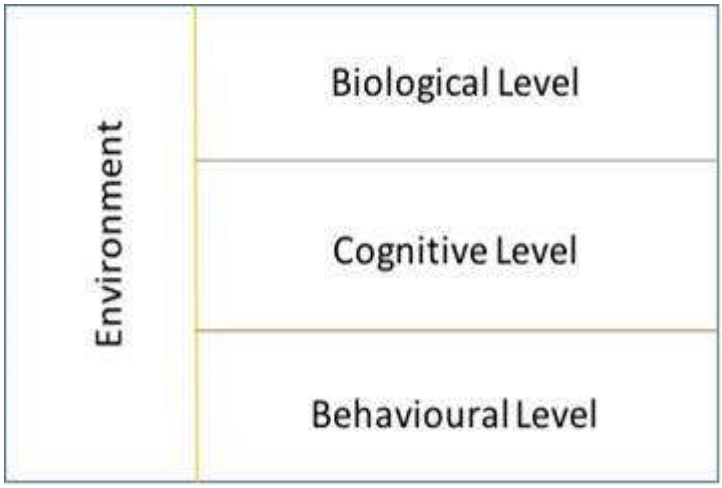


Figure 2.7: Causal Modelling Framework

The lowest level aims to map causes for conditions such as dyslexia or autism that are solely biological in origin, this can be genetic as well as neural. The *behavioural level* amounts to a description of observable facts about the individual, including producing spoken language, or attempting a reading task. The intermediate *cognitive level* provides a description of those aspects of the system that "can only be distantly observed" (Frith, 1995: 7) and involves factors such as emotion, knowledge of phonological categories and orthography-phonology mappings and attention deficits (*ibid*: 7-12).

It should be noted that the above framework was developed with a particular aim in mind: the study of certain developmental disorders. Therefore, it was not proposed as framework for general cognition, unlike all the other accounts that are under discussion in this chapter. However, if possible it would be beneficial if a particular framework of cognition could be utilized both in cases of typical and atypical cognitive cases. The major issue with the applying the above framework in a more general sense comes from the causal nature between the levels that constitute the framework. Taking the framework at face value, then, would lead one to the position whereby biological facts about the individual *cause* cognitive and behavioural ones. This is problematic, given the generally held position that an attribution of causation makes certain implications regarding the temporal ordering of the cause and the effect (Papineau, 2013: 127-128). Taking a causal account between the biological and cognitive/ behavioural levels in this regard, would amount to the picture whereby a biological state of an organism would exist **before** the corresponding cognitive

state, thereby raising questions concerning the ontology of the picture of cognition presented.

On the face of it, there seems little reason to object to the notion of embedded cognition, since it seems obvious that cognition could not take place without the existence of some outer world. In some of the cases we have seen in this section, especially the dynamic framework and the idea of environmental scaffolding, this conception of the link between the environment and the mind is more than simply one existing to provide representations of stimuli for the latter. One of the prominent themes has been that of *coupling*, the idea that the external environment takes a more active role in the study of cognition.

This thesis will take the view that the role of the environment is important in understanding cognitive systems. As will be shown in §3, this is in line with the SFL approach to language, and those in this field will probably take the embedded view as given. It is important to realize though, that such considerations are relatively recent in modern cognitive science.

2.5.2. Embodied Cognition

A related, but necessarily distinct thesis is the idea of *embodied cognition*. In general terms, it states that nature of cognition is the way it is because of certain facts to do with the organism's body. Usually, embodiment theses also contain a tacit approval of embedded views on cognition (Beer, 1999; 2000). Although computationalists would agree that the body is important for biological cognition in order for the neural system to receive external stimuli, it plays a more passive role in these construals in contrast to the views we will review below (Leitan and Chaffey, 2014). The embodied view, then, is that the nature of cognition for a particular organism or groups of organisms will be dependent upon certain facts about how that organism's body is organized. Embodied accounts, however, argue that the nature of cognition itself is highly dependent upon the bodily properties. Many of the reasons for motivating such a view come from empirical studies from psychology and neuroscience, which are discussed below.

Lakoff and Johnson (1980; 1999) provide examples to motivate their argument for an embodied view of cognitive science. In the first instance, Lakoff and Johnson (1980) is an

in-depth look at the role that metaphors have in linguistic semantics. In this work, the authors highlight the prominent role that metaphor plays in many aspects of natural language. Some of the metaphorical mappings that they point to include: states as locations (I'm *in* love; We're *far from* safety), changes as movements (He *went* crazy; he *fell into* depression), causation as forced movement (the home run *threw* the crowd into a frenzy; the news *propelled* the stock market to record heights), and difficulties as impediments to movement (Harry *got over* his divorce; we ran *into a brick wall*). However, all metaphors must be based upon some, non-metaphorical basis, otherwise there would be no account of how meaning could possibly arise. This non-metaphorical basis comes from meanings that are in the semantic field of spatial terms, such as *up, down, in, out, around, towards, away from,* etc. The influence of these spatial terms on the above examples is evident. Furthermore, they argue that our understanding of these basic, non-metaphorical terms can only be accounted for by understanding humans as embodied creatures (Lakoff and Johnson, 1980: 476)[7]:

> *Imagine a spherical being living outside of any gravitational field, with no knowledge or imagination of any other kind of experience. What could UP possibly mean to such a being?*

And later in Lakoff and Johnson (1999:230-231):

> *The study of human characterization has revealed that our conceptual system is organized around basic-level concepts, concepts that are defined relative to our ability to function optimally in our environment, given our bodies. Concepts of human agency […] are among the basic level anchors of our conceptual system in general and our system of causal concepts in particular.*

Lakoff and Johnson are correct in pointing out that metaphor is indeed a frequent phenomenon in certain aspects of natural language, especially when it comes to grounding

[7] In the below quote, UP in capital font designates the concept of "up".

our understanding of abstract nouns such as emotions, hardships, etc. It is also feasible that our understanding of these basic terms could be based upon factors concerning humans as embodied, and embedded, organisms.

Additional evidence that has been put forward to motivate an embodied view comes from work done on sentence processing. Glenberg and Kaschak (2002) proposed the *action-sentence compatibility effect* (ACE). The ACE states that the processing of a sentence is mediated by a mental simulation of the actions described by the sentence. To back up their claim, they investigated whether there would be interference between the processing of a word that implied a certain direction of movement, with actual bodily movements that involved movement in a different direction. In these experiments, a group of participants was presented with target sentences such as the following:

(1) Open the drawer
(2) Put your finger under the faucet
(3) Boil the air

In response, the participants had to judge whether the sentences were sensical or not. With respect to the above three sentences, the first two had the target response of *sensical* and (3) would be non-sensical. In order to record their responses, the participants had to press either a button marked 'yes' (for sensical), or 'no' (for non-sensical). A third, 'neutral' button was also present, which participants pressed in order to bring up the next sentence in the trial. The buttons were arranged in particular patterns, with one button being close to the participant, another was further away from the participant, and the third was between the two. In all cases, the middle button was the 'neutral' button. The functionality of the other two buttons was treated as an independent variable in the experiments: on one condition the "far" button was pressed to signal 'yes' (the yes-is-far condition), while on the other, the near button performed this function (the yes-is-near button). On the yes-is-far condition, the participants had to move their index finger away from them to register a sensical response; in the yes-is-near condition, they had to move their index finger towards them to signal the same response. Crucially, the sensical sentences implied some kind of action-movement relative to the speaker: sentence (1), for

example, implies movement towards, while sentence (2) implies movement away. The results showed that participants took significantly longer to respond when the movement implied by the target sentence contradicted the movement required to answer 'yes'. For instance, participants took longer to respond 'yes' to sentence (1) in the 'yes-is-far' condition, than on the 'yes-is-near' condition. They argued that this could only be explained if one accepted that actions that take place at the level of the entire organism are part of processing the meanings of terms that was traditionally seen to only be due to disembodied symbolic manipulation.

Other researchers have argued for similar claims on the basis of neuroimaging data. They argue that certain aspects of perception utilize the same parts of neural architecture that are also used in the processing of active responses. Bucinno et al. (2001) showed that on perceiving an action with a certain type of effector quality, the same areas of the motor cortex were activated that would also be activated on the actual performance of the action. For instance, upon seeing another human eating an apple, areas of the frontal and parietal lobes exhibited increased activation, the same areas that you would also expect to be active if the participant performed the action themselves. This was also argued to occur with hand-related and foot-related actions. Similar effects have also been shown in the case of language comprehension. Hauk et al. (2004) showed a similar effect in the comprehension of certain lexical items. In their study, participants passively read 'action words' that were related to either the face, arm or leg, and on the basis of fMRI data they argued that during lexical comprehension there was co-activation of areas of the somatotopoically organized motor and pre-motor cortex. For instance, passively reading an 'action word' associated with the face correlated with increased activation of the areas of the motor cortex that are also active during movement of one's own face. The claimed theoretical importance behind these works is that language comprehension is not based upon some specialized modular neural system, but also incorporates parts of the brain that are utilized in bodily action. As a result, these researchers argue that language itself, especially the processing of linguistic meaning, is embodied in nature.

As mentioned in §1, there is a lack of specific proposals concerning the ideas of *levels of analysis* in these works. As just seen, many of the arguments come from specific cases.

However, a few can be found. For instance, Beer (2000) provides the following sketch for an *embodied*, *embedded* approach to cognition:

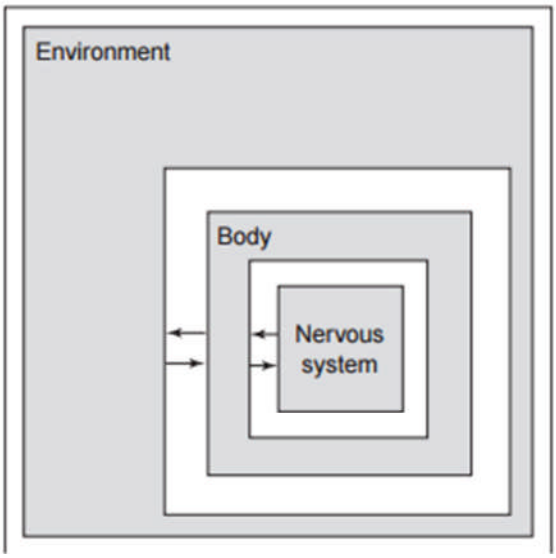


Fig. 2.8. Three levels in an embodied dynamic framework

The relevant features of the above framework, and how it contrasts with others we saw in this chapter are as follows: firstly, each level on this framework is considered to be a state-space in the terms defined in §2.4.1. Therefore, the lowest level can be considered to be the totality of states an organism's nervous system can be in, couched in terms of neural dynamics. Moving up, the next level is the state-space of extra-cranial bodily features, while the highest level is the totality of states one can use to describe the extra-organismal environment. The hierarchy is ordered in terms of constituency: the nervous system is clearly a part of the wider body in the same way the heart or an arm is, and arguably an individual is simply part of the wider environment, if one considers the environment to be something akin to the world. As already seen in previous frameworks, this use of constituency indicates that this particular hierarchy is an example of *levels of* organization. The differences between this characterization and others encountered in §§2.1 – 2.3 are: (i) the explicit inclusion of bodily and environmental factors in the conceptual framework of cognition, and (ii) the influence between each level is bidirectional, in stark contrast with the computational framework, that have largely utilized the arguments for multiple realizability.

The shortcomings with the above framework, however, stem from its lack of specificity. Although I believe the framework is correct in specifying the interlinkage between the levels

of neuronal detail, body and environment, in order for the framework to have any predicative power, or to be of any use in the practice of studying cognitive phenomena, hypotheses about the nature of these links need to be given. For instance, when does a difference in dynamic behaviour become meaningful in terms of any cognitive theory? Neural systems are highly complex, and could be described by a mathematical space containing an extremely large number of dimensions. Let's say that the average human nervous system contains around 100 billion neurons (Stein and Stoodley, 2006: 11), with the electrical properties of each neuron contributing one dimension to the description. This would mean that the neural dynamics is best described by a 10^{11} dimensional space. Obviously, there are an extremely high number of possible ways the behaviour of this system could unfold. What is unclear is whether all differences in these unfoldings are useful for understanding cognition. Also, what should one count in a description of an environmental system? Only objects, or more abstract ideas such as social relationships? The world is indeed constantly changing around us, but not all of these changes are down to our own actions, and not all changes seem to influence our behaviour in any way. Finally, what is the precise nature of coupling? If it remains a relatively broadly defined notion such as changes in one effects the other, does this add much to our understanding? The issues here, then, are that such a framework is not explicit enough in what it predicts how cognitive behaviour would unfold.

Embodied cognition, therefore, is the view that cognition can only be meaningfully studied by considering how aspects of the organism's body influence their cognitive behaviour. It is more, however, than the simple recognition that one needs to have a body to serve as the source of input to the internal computational, cognitive system. Instead, it argues that cognition is highly dependent about the way that the body is organized. They also point to ways in which the processing of motor and other sensory information can impact upon cognitive factors such as language processing.

2.5.3. Extended Cognition

Extended cognition states that external factors can, in certain cases, be constituents in mental processes. As a result, proponents of this view argue that the specific objects in the environment are as much a part of mental processes as are neurons firing in particular

parts of the brain, and a result the position that mental processes are *conceptually* to be defined solely in terms of brain processes is a mistaken view (Dennett, 1996; Hutchins, 1995, Clark, 2010a; Walter, 2010; Menary, 2010).

The most widely discussed arguments for such a view come from Clark and Chalmers (1998), who provide a thought experiment to back up their position. In this, the authors imagine two individuals: Igna and Otto. They are told that there is an interesting exhibition at the Museum of Modern Art. Having previously visited the museum, Igna remembers that it is on 53rd Street, and using this recalled knowledge, makes her way there. Otto, on the other hand, suffers from a form of Alzheimer's disease, and as a result is unable to form long-term memories. In order to counteract this, at all times Otto carries with him a notebook in which he notes down information that for most people would be stored in long-term memory. Having also previously been to the Museum in question, he looks up this information in his notebook and discovers that it is located on 53rd Street. With this information, he makes his way there.

The crucial part of this scenario is that Igna's biological memory is performing the same function for her as Otto's notebook does for him. The authors then investigate whether these two people differ in their belief concerning the location of the museum. They argue that in both cases each believes that the museum is on 53rd Street, even before Igna recalls this fact from her biological memory, or Otto looks up the information in his notebook (*ibid:* 13):

> *The alternative is to say that Otto has no belief about the matter until he consults his notebook; at best he believes the museum is located at the address in his notebook [… But] to say that beliefs disappear when the notebook is filed away seems to miss the big picture in just the same way as saying that Igna's beliefs disappear as soon as she is no longer conscious of them. In both cases the information is reliably there when needed, available to consciousness and available to guide actions, in just the way that we expect a belief to be.*

In Otto's case, the mental process involved in belief states extends to include as a proper part elements that are external to the individual. Therefore, in a conceptual sense, cognition can be considered to be *extended*. They go on give four reasons for considering Otto's belief process to be *extended* (*ibid*: 17):

> *First, the notebook is a constant in Otto's life – in cases where the information in the notebook would be relevant, he will rarely take action without consulting it. Second, the information is directly available without difficulty. Third, upon retrieving the information from the notebook he automatically endorses it. Fourth, the information in the notebook has been consciously endorsed at some point in the past, and indeed is these as a consequence of this endorsement.*

The last of these conditions is required otherwise the position argued for would lead to an over-extension of the amount of things any person believes. For instance, without this final caveat, the position would entail the counter-intuitive view that simply being in constant possession of the Yellow Pages means that you possess a belief about the correct telephone number of every entry found within it.

Note that this step in the argument is crucial to their argument. If one doubts the above, and instead says that Otto only believes in the location of the museum upon checking, and that Igna only believes the same on the act of remembering, then the appeal to externalism, in the sense the authors argue for, fails. The reason is that, if the belief only returns on Otto's reading of the notes in his diary, then the process itself can be said to no longer have an external part, even though, of course, the notebook is still important. This is the distinction, raised by Rupert (2004: 395-396), between *cognitive mereology*, and *epistemological dependence*. The former is the thesis under discussion here, namely that external objects form *proper parts* of the cognitive process, whereas the second is the claim that cognition can only be fully understood by considering the environment in which the organism is situated. Therefore, this latter notion is the view of those who advocate some form of embedded cognition, as discussed above. Crucially, the two are not identical. Something (broadly construed) can be indispensable for understanding a system, without that thing

itself being a part of that system. This point is also raised by Adams and Aizawa (2001: 56; 2008: 85), who dub such an argument an instance of the "coupling-constitution fallacy". They give examples such as the filtration process in the kidney. Although an important part of this process is that kidney enters into causal relationships with parts of the circulatory system, i.e. the fact that the kidney must receive blood from the heart, that the blood must pass along blood vessels of various sizes, etc., does not give motivation for the view that the filtration process occurs anywhere other than in the kidney. In sum: "a process P may actively interact with its environment, but this does not mean that P extends into its environment" (Adams and Aizawa, 2008: 56).

It should be obvious at this point that the thesis of extended cognition is stronger than that of either embedded or embodied cognition. In those cases, the processing that one might want to label 'cognitive' can still take place solely inside the head, although the nature of that processing is heavily dependent upon and can only be studied with reference to an external environment (embedded) and/or the organism's body (embodied). On the extended view, however, external items can literally enter into the processing, by means of performing processing tasks that could as well have taken place intra-cranially. I find myself unconvinced by the arguments for *extended* cognition, for the same reasons as the critics above.

2.5.4. Enacted Cognition

The final 4E approach to be discussed is known as "Enacted Cognition". It shares common links with the other approaches discussed so far, that of the increased role of environmental and bodily factors in cognition. However, the distinguishing feature between this approach and the others is that it emphasizes the *active* interaction of the whole organism and its environment. Hutto and Myin (2013: 4) describe the approach as follows:

> *Enactivism is inspired by the insight that the embedded and embodied activity of living beings provides the right model for understanding minds. To understand mentality, however complex and sophisticated it may be, it is necessary to appreciate how living beings* ***dynamically interact*** *with their environment: ultimately,*

> *there is no prospect of understanding minds* ***without reference to interactions between organisms and environments.***

A few facts are immediately obvious on the above account: its link to the ideas of embedded and embodied cognition, the appreciation of the dynamic aspect of cognition and the focus on constant interaction between an organism and the environment. Much of the work done in the enactive tradition stems from Varela et al. (1991), who focussed on how embodied cognition interacted with the wider biological, psychological and cultural contexts that organisms constantly find themselves in (c.f. Rowlands; 2010: 70-71).

Following this tradition, some authors (Hurley, 1998; O'Regan and Noë, 2001; 2002; Noë, 2004; 2009;) have proposed a version of enacted cognition dubbed "Sensorimotor Enactivism" (Noë, 2004; Hutto and Myin, 2013: 23). These approaches are based upon the notion of *sensorimotor contingencies*. This is the idea that many facts about cognition are reliant upon expectations that our actions will have similar effects in similar situations. The position is sometimes summed up by the oft-cited quote from Noë (2004: 227): "[perception] isn't something that happens in us, it is something we do." The oft-given example of this idea is the notion of haptic perception. To exemplify this, consider how someone who is blind comes to perceive an object, say a bottle (O'Regan and Noë, 2001: 945). Firstly, she has to grasp it but this is not enough to successfully identity the object correctly, for many things may feel similar upon first contact. On further exploration, she notices that the way that her perceptual input changes is contingent with what you would expect from a bottle: as she moves her hand up the object narrows, further down she feels something of a different material, consistent with the presence of a label. On this account, then, to perceive something is to be aware of the affects our actions will have in the presence of a certain object. Nor has this phenomenon been restricted to simply the tactile domain, with proponents of this view also arguing that similar considerations hold equally well for visual perception. To perceive the bottle visually would be to have knowledge that certain eye movements would lead to certain changes in the incoming visual stimuli, e.g. when one fixates at a point higher up the bottle, the area of the incoming stimuli that is rightly identified as the bottle narrow, etc.

There are a number of points to be raised concerning the above view. Firstly, it presupposes a vast amount of knowledge on the part of the organism. By depending upon the notion of contingent sensorimotor knowledge of expectations, the idea places a great epistemological burden on the part of the organism. That is, in order for something as simple as perceiving an object to take place, the cognitive organism but be in possession of a large body of knowledge of how (some) of its potential actions will affect its own incoming sensory stimulus. The theory would therefore need some account of how this information is extracted in the first place, and how tacit knowledge of this information is available for later use and stored within the system. Second, even if we were to grant that expectations could play a role in the perception for tactile and visual stimuli, it is less clear how they would be translated to the other sensory modalities, especially auditory and olfactory. In these cases, the issue is not that organisms do not use expectations in cognitive life, but that the sensorimotor enactive view states that it is necessarily dependent upon changes that the organism itself brings about. With vision, eye saccades and more intentional eye movements are necessary for visual perception; with tactile perception the expectations of conscious movements played an important role. What would be the analogous movements in the case of the other modalities? Certainly, we can move our head to alter the nature of incoming auditory stimuli, but it seems unlikely that such movements would be necessary for auditory perception.

A different take on the enacted framework is known as “radical enactivism”. It is radical in the sense that unlike the previous two version of this approach, it argues that the study of cognition should entirely reject the presence of representations (Hutto and Myin, 2013). In its place, it opts for a theory of *teleosemiotics*, which draws upon the versions of teleosemantics that has been argued for in classical views on cognition, but instead argues that the role of functions is not to provide content for mental representations. On this account, issues such as the intentions of organisms can still play an important role in the explanation of cognition, just not in setting the content of mental representations. One of their main reasons for forgoing the idea of representations is essentially along the lines of the *symbol grounding problem* encountered earlier, that they dub the “Hard Problem of Content” (*ibid*: §4). Building upon the idea that covariance or correlation is not enough to establish meaningful content, they argue that no means currently exists to show how

content **can arise out of essentially meaningless naturalistic laws of causation**. Their solution to the symbol grounding problem, then, is to simply say that there is no problem to solve. The fact that content is so difficult to define points to the fact that there is no such account to be given.

Unfortunately, the authors do not provide further elucidation of this idea other than characterizing it in terms of a negative account of the existence of mental representations (c.f. Matthen, 2014). Without such a positive account, their arguments seem less than convincing. Furthermore, they stop short of saying that all minds are content-less. The exceptions are language-using humans. The semantic content of linguistic communication points to the fact that, *for those minds*, there must be content (*ibid*: §7). However, this would pose a problem. It would necessarily entail that there would be a stark difference between language-using content-full minds, and language-less content-less minds, and that the explanation of cognition as applied to linguistic behaviour would differ from other cognitive practices. Not only does this entail a disjoint between human and non-human cognition, it would also say that if, as seems reasonable, parts of our cognition operates independently of language, than there would be two types of cognition in operation in the minds of humans. This would be a great claim, and not one that they provide any evidence for.

In sum then, enactive cognition is the view that actions play a large role in the determination of cognition. As with the *embedded* and *embodied* ideas discussed in this thesis, I am sympathetic to this view, although I appreciate the criticisms just mentioned. The hope is that some of the ideas presented later in this work can bring some extra clarity to this position.

2.5.5. 4E approaches and the absence of conceptual frameworks

The discussion so far has been on the conceptual basis behind the 4E approaches, and has come to the conclusion that some of these views are more plausible than others. The aspects of this viewpoint that were discussed are as follows:

- The increased importance of the environment and the body in the study of cognition

- Embedded Cognition, Embodied Cognition, Extended Cognition and Enacted Cognition.
- The implications that each of these views has for the study of cognition.
- A discussion of various conceptual frameworks proposed by researchers working within this viewpoint (when one is given).

However, unlike many of the previous approaches that have been discussed so far, the issue of *levels of interpretation* has not been frequently raised.Wilson and Clark (2008: 55) sum up the situation as follows:

> *Importantly, researchers who self-identify with the "situated movement" [read 4E - JW] in the cognitive sciences have not been paralyzed by a relative lack of attention to conceptual articulation of their paradigm*. (Wilson and Clark, 2008: 55).

It is, of course, an open question whether a "conceptual articulation" of a certain viewpoint is paralyzing. Some of the virtues of such articulated frameworks have been seen throughout this section. It allows empirical findings from a diverse range of sources and researchers to be incorporated and evaluated with regards to one overarching picture. It also enables new findings to not just be incorporated into one wider picture, but also to amend that picture when necessary. In short, it provides a largely agreed upon "vocabulary" for those working in the same area. The absence of such raises the risk of unnecessary vagueness or lack of explicitness to enter the communication of ideas. Furthermore, being explicit in terms of conceptual frameworks allows predictions to be made more easily. As we saw above, although the frameworks of supervenience and classical, symbolic, cognition were argued to have shortcomings, one positive aspect of their approach was that they had explicit frameworks on which results and theories could be related to. It allowed predictions to be made, and links between the cognitive and physical were made definite.

Apart from the specific problems that face some of the particular varieties discussed in this section, then, the major problem facing these proposals on a general level is this lack of an over-arching framework. The exception was the framework provided by Morton and Frith (1993) and Frith (1995), but it was seen that attempting to take their proposal as a general

model of embedded cognition poses conceptual problems concerning treating inter-level relations in terms of causation. Also, the three-levelled framework of embodied dynamic cognition (Beer, 2000) was too weak as to enlighten the study further than just vague statements concerning the influence of extra-cranial factors on neuronal and cognitive processing. Indeed, Northoff (2014: 240) states the idea of the "extended brain", which on his interpretation is "characterized by an intrinsic relationship between the brain and both the body and the environment", as one of the areas of future research. The idea is well set, but more work needs to be done to make this idea into a fully-fledged empirical hypothesis, set in terms of an explicit conceptual framework.

As is hopefully clear from the above discussion about the 4E approaches, and also the criticisms that were raised to the "inter-cranial" approaches dealt with in the first part of this chapter, this thesis takes the position the former are generally along the right path. The evidence that is stemming out of cognitive science is that facts about the wider organism, not just its brain but also its entire body, and also the world with which it is in constant interaction, have a real effect upon the nature of cognition. Furthermore, I would argue, a full understanding about the nature of cognition can only be gained by considering all of these factors. The main failure of these approaches is that no convincing or powerful overarching framework currently exists in order to bring together the myriad of experimental work, and to generate hypotheses that can be tested and verified or disproved. The model that is introduced in the remaining chapters of this thesis is an attempt to do just that.

2.6. Summary

The following table summarizes the approaches to cognition that have been discussed in this chapter. In each case, the following points will be discussed:

- Conceptual assumptions of the framework
- Theoretical Levels highlighted in the framework. Includes the types of levels involved and the nature of their connection.
- Whether it can be broadly considered to be *internalist* or *externalist*.
- Advantages of the proposal
- Disadvantages of the proposal.

Framework	Conceptual assumptions	Important levels in the framework	Internalist or Externalist?	Advantages	Disadvantages/ problems
Supervenience	Cognition as a thesis of non-reductive physicalism	"mental" and "physical" levels of different kinds of states, properties, etc. One mental property can be linked to more than one physical property, and no physical property can be linked to more than one mental property.	Internalist	Provides an account of linking physical and mental properties in a way that allows for variation. Makes certain predictions concerning changes in mental and physical properties	Weak empirical conclusions Primarily a metaphysical treatise Low predictive power
Classical/ Symbolic	Cognition as computation over symbolic representations	Marr's levels of computation, algorithms and implementation.	Internalist	Understanding complex systems simultaneously at different levels of interpretation	Cognition as linear processing

		Interpretative levels linked by a top-down relation of realization. One higher level unit realized by multiple lower level units		Partial appreciation of cognition taking place in a wider context	Problems concerning determining content of mental representations Intuitions concerning multiple realizability and the divorcing of cognition from neuroscience Problems concerning explaining content of symbolic representiations.
Connectionist	Cognition as distributed activity over artificial neural networks	Levels of organization within neural networks: input, output and hidden layers.	Internalist	Distributed processing Neurologically Inspired Variety of Empirical Results at the Cognitive Level of Interpretation.	Differences between artificial neural networks and biological neural systems Learning system dependent upon prior knowledge.

		Levels of interpretation: the cognitive and neural levels.			Reliance upon computational metaphor
Embedded	Cognition highly dependent upon the wider environment	Biological, cognitive, and behavioural levels. Linked in terms of causation (Firth; Morton)	Externalist	Dynamic models focussing on the importance of cognition unfolding over time. Explicit linking of cognition to the wider environment	Unclear exactly *how* the environment affects cognitive life. The nature of the environment itself is underspecified
Embodied	Cognition dependent upon facts about the organism's body	Neural, Body and Environment levels. Bidirectionally linked (Beer – also embedded)	Externalist	Neurobiological evidence for the potential importance of embodiment. Allows the organism as a whole to be considered	Conceptual framework lacking in predictive power

Extended	Objects or processes external to the organism can enter into cognitive processing	None given	Externalist		Differences between constitution and coupling. Lack of conceptual framework
Enacted	Cognition construed in terms of constant action in an environment	None given	Externalist	Role of action, and the interaction between individual and environment is emphasized	Lack of specifics in ideas Difference between language-using organisms and others

Table 2.2: Overview of Conceptual Frameworks in Cognitive Science

The purpose of this chapter has been to critically survey the current literature on conceptual frameworks in cognitive science. In sum, classical and connectionist approaches were criticised, amongst other things, for constraining the study of cognition to purely intra-cranial features and ignoring or under-emphasising the role of the whole organism and the wider environment. Although this issue has been corrected in the *4E* approaches, in these approaches there was a general problem of lack of specificity that has been highlighted by various authors. This weakness inhibits these proposals, and their current empirical results, from being linked into an overarching picture of cognition. This shortcoming is striking in comparison with the other accounts of cognition, which are explicit about the picture of cognitive life that they provide.

The following chapters of this thesis go on to examine the possibility of a conceptual framework for cognitive science that is based upon theoretical notions that stem from Systemic Functional Linguistics. By comparing the picture of cognition that emerges from this discussion with the proposals discussed in this section, this thesis can then determine how the proposal in this work intersects with ideas that have come before, and can assess whether by taking the approach I am advocating here, some of the shortcomings discussed here can be overcome.

Chapter 3

Thoughts on a Stratificational Model of Cognitive and Behavioural Processes

> *"Christopher Robin was sitting outside his door, putting on his Big Boots. As soon as he saw the Big Boots, Pooh knew that an Adventure was going to happen, and he brushed the honey off his nose with the back of his paw, and spruced himself up as well as he could, so as to look Ready for Anything"*
>
> *A.A.Milne. Winnie-the-Pooh*

So far, the major research aims of this thesis have been stated, and past proposals from the literature has been reviewed. This purpose of this chapter is to begin our investigation whether the theoretical heuristics of Systemic Functional Linguistics (SFL) can be usefully applied to the subject matter of cognitive science more generally. In doing so, it will investigate the theoretical dimensions of *stratification*, *instantiation, delicacy* and *rank*, as these all pertain to the issue of *levels of analysis*. In §1.3, various initial motivations for this approach were given. They are repeated again below:

- *Cognitive systems are semiotic, and SFL draws its inspiration from this field.*
- *SFL has shown that it can be applied to non-human forms of semiotics, as well as non-language based human semiotic systems.*
- *Both language and cognition can be described as dynamic, open systems*
- *SFL provides a number of different levels of analysis for studying semiotic systems of this particular type. Some of these have already been shown to be important in studying cognitive systems.*

However, these motivations only mean that the inter-disciplinary approach taken by this work is feasible, not that it will be successful. It will only have been shown to be successful if there are advantages to be gained by taking such an approach. The purpose of this chapter is to begin our investigation into whether any advantages can be gained.

Before we begin, however, a major factor omitted from this chapter needs to be highlighted. Within SFL, the relationship between language and the wider eco-social context is of utmost importance (Halliday, 1979; Lemke, 1990; Thibault, 2004). This relationship will

form an extremely important part of the arguments developed in this thesis, but, in order to give them the space they deserve, such a discussion will be postponed until §4. The current chapter will focus instead on the nature of this particular linguistic theory as it relates to 'language-internal' factors. It should be kept in mind, however, that these ideas were always meant to extend beyond language and into the semiotic context in which it operates.

The structure of this chapter is as follows: the basic assumptions and tenets of SFL will first be laid out. This includes primarily describing how they are both *systemic* and *functional*, and explicitly uncovering what these terms designate. Secondly, the conceptual framework of SFL will be presented. In this section the theoretical dimensions of *stratification*, *delicacy, rank,* and *instantiation* need to be fully explicated. From this point, an overall sketch of a stratificational framework of cognition and behaviour will be presented, by applying the notions that have been discussed in a linguistic context to the context of cognitive science. This will involve investigating whether the proposal finds any support in the psychological and neurological literature. It will also compare and contrast the framework developed in this chapter to those discussed previously. In doing so, the novel aspects of this framework can be highlighted, as well as advantages it may have. It will also allow us to investigate, at this stage, any short-comings or areas of investigation that can be gained from this work.

3.1 Foundations of Systemic Functional Linguistics

SFL is a linguistic theory that highlights two aspects of language (Halliday, 1979; Eggins, 2004; Matthiessen, 2007a; Halliday and Matthiessen, 2014): (1) the idea of paradigmatic distinctions in linguistics categories, that has its roots in Saussurean structuralism (Saussure, 1916), and (2) that linguists should focus upon active language use in the development of linguistic theory, and less upon providing context-independent linguistic computations that all humans are said to have (unconscious) 'knowledge' of (Halliday, 1979; Eggins, 2004; Matthiessen, 2007a; 2007b.; Halliday and Matthiessen, 2014). It is these two aspects that make language, respectively, *systemic* on one hand, and *functional* on the other.

3.1.1. Language as Systemic

The focus of the paradigmatic relation between linguistic categories stems from the highlighting of one of the two crucial sign relations that was proposed, amongst much else,

by Ferdinand de Saussure in his influential *Course in General Linguistics* (*ibid*). In this work, Saussure showed that in order to fully understand semiotic values, one needs to appreciate two kinds of relationships between one sign and another. Firstly, the *syntagmatic* relationship between signs and others that appear in the same structural context. The other is the *paradigmatic* relationship, which held between a sign and all other signs that could be legitimately substituted in its place. For instance, consider the following:

(1) The angry dog scared the postman

Syntagmatic relationships would hold between linguistic signs as they appear in linear order. For instance, there exists a relationship between *dog* and *angry*, or between *dog* and *scared* in (1). Syntagmatic analysis would then provide a description of the syntagmatic relationships that exist in a particular language, as well as the construction of rules that determine what kinds of syntagmatic relationships are valid. Paradigmatic relationships describe legitimate substitutions that can be made between linguistic units. For instance, compare the following with (1) above:

(2) The angry cat scared the postman

(3) ??The angry green scared the postman[8]

In both (2) and (3) *dog* has been replaced with an alternative linguistic unit. However, only one of these substitutions is legitimate; the substitution of *dog* with *green* in (3) producing a seemingly odd sentence.[9]

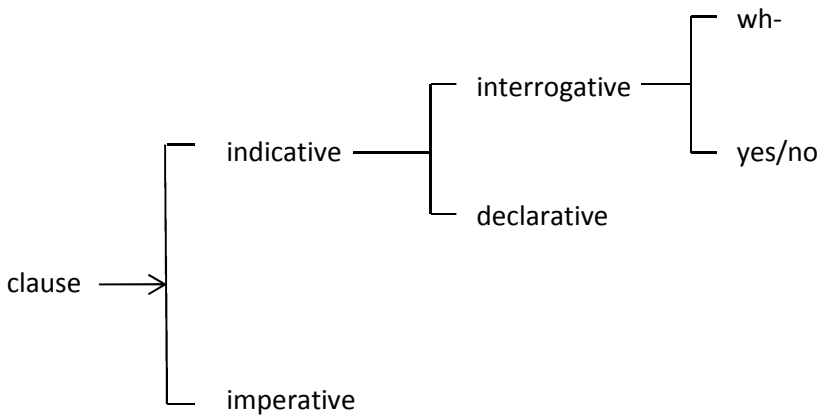
SFL highlights the role of paradigmatic relations in language, over syntagmatic ones. In doing so, it provides descriptions of the various "options" that language provides the users in the production of meaningful linguistic behaviour. This is in contrast to many other linguistic theories that instead focus upon syntagmatic explanations of language. This is prominent in the variety of *generativist* theories, including Chomsky's theories, Generalized Phrase Structure Grammar, and Lexical Functional Grammar. All of these theories, regardless of the

[8] Question marks (??) indicate semantic anomaly/ inappropriateness of piece of linguistic data.

[9] Unless, of course, *green* is used here as a metonym for "member of the Green party". However, this would be a different linguistic sign to *green* as an adjective denoting a certain colour, differing in terms of their content. Members of left-wing, environmental political parties are also a friendly bunch, raising questions of the validity of this sentence even in this other sense!

major differences between them, all attempt to provide an account of the ordering of linguistic units that distinguishes well-formed from ill-formed linguistic sentences.

In emphasizing the paradigmatic dimension in language, one of the main modelling tools in SFL is the creation of *system networks* (Halliday, 1966; 2009; Halliday and Matthiessen, 2014: 22). System networks are set up to describe the various options that language provides the language-user at the strata of phonology, lexicogrammar, and semantics (c.f. §3.2). For example, consider the following simplified system network for the *lexicogrammatical* system of MOOD[10] (adapted from Matthiessen and Bateman, 1991: 92) :



Network 1: Simplified network showing paradigmatic relations for the clause system

The above network gives the (simplified) range of paradigmatic relations that hold between different kinds of major clauses. In the most general terms, there are two types of major clauses in English: indicative clauses that contain both a Subject and Finite element, and imperative clauses, which lack both of these elements. Indicative elements can either be interrogative, where the Finite element precedes the Subject, or declarative, in this case this ordering is reversed. Examples of these types of clauses are given below:

(4a.) Imperative: *Wash the dishes!*

(b.) Declarative: *I wash the dishes*

[10] Unless indicated otherwise, use of SMALL CAPS will indicate the name of a particular system network.

(c.) Interrogative: *Am I washing the dishes?*

Fuller details concerning system networks and their full implications for the study of language will be given below in §3.3. For now, however, it is important to highlight that by creating such descriptions of the "options" that are visible in a language, one is modelling language along the paradigmatic axis. Syntagmatic relations are then derived from the paradigmatic relations that hold between particular linguistic units. Again, this will be discussed in more detail below.

A further point should be raised at this juncture. In the description just given, I spoke about "options", and used scare quotes to qualify my use of this term. The reason is because this notion need not always involve a conscious act on the part of the speaker, whereby a speaker consciously chooses between various options (Matthiessen et al. 2010: 69). However, it should be noted that this is still a controversial and open issue for those working in the SFL framework (c.f. papers in Fontaine et al. 2013 for different perspectives on the notion of 'choice' in SFL). Unless made explicit, I do not intend for the use of 'option' or 'choice' to entail or imply any involvement of a conscious 'chooser'.

3.1.2. Language as function

As well as seeing language as systemic in the sense that it can be modelled in terms of the various options available to the language user, SFL is also a **functional** linguistic theory. In this, it follows a similar path to other theories in linguistics, including Dik's Functional Grammar and Role and Reference Grammar (c.f. Butler, 2003 for a comparison of these approaches). These theories are often contrasted with so called formalist theories of language (Newmeyer, 1998). Matthiessen and Bateman (1991: 4) give the difference between SFL and formal linguistic theories as follows: "systemic linguists interpret and represent language not as a rule-system for generating structures, but as a *resource* for expressing and making meanings" (emphasis in the original). Halliday himself has stressed many times that the main aim for linguistic theory should be the characterization of language as it is used by speakers in real life situations (Halliday, 1961; 1979; 1996; 2003). As such, Eggins (2004: 3) gives the main aims of SFL as follows:

1. To investigate how many different sorts of meaning we use language to make
2. To investigate how language is organized to make meanings.

At first glance, then, the notion of *function* in SFL seems to be akin to the role that language has for its users. Language is an essential part of humans' eco-social life; it is used for organizing our world around us into different categories, which stand in certain taxonomic relations to one another (Halliday and Matthiessen, 1999). At the same time, however, we use language to bring about certain events in the world; we use language to build social relationships, to manipulate the world to bring about a certain state of affairs.

In order to more explicitly distinguish between the kinds of functions that language can have, one of the central notions in SFL is *metafunction*. Metafunctions are generalizations over the wide range of functions that language can have, and groups them according to semantic similarity.[11] The metafunctions in SFL theory are as follows (Halliday, 1969; Halliday and Matthiessen, 1999: 511-513):

1. *Ideational* – language is the means by which we interpret our experience. Language gives names to things that we experience, constructs relationships between objects and distinguishes between different kinds of processes that we encounter.
2. *Interpersonal* – language is used to construct, alter and maintain social relations between language users. By using language, we signal and locate ourselves within the complex social structures that are evident in our lives.
3. *Textual* – Since language must be used in order to signal information to others, it must be ordered in a certain way in order to achieve this. The textual metafunction describes the processes in language that lead to the creation of coherent texts.

The first of these, the ideational metafunction, can be divided into two separate components:

1a. Experiential – language as a theory of experience; it represents the processes and individuals/ objects we encounter

1b. Logical – language as a representation of the relations between those processes.

[11] The number of metafunctions given here is based upon the 'Sydney School' of SFL, as exemplified by Halliday, Matthiessen and Hasan, amongst many others who have used their ideas. The number of metafunctions differs in other versions of SFL. For instance, the variation of SFL known as the 'Cardiff Grammar' contains 8 metafunctions (Fawcett, 2008).

Certain system networks are associated with each of the above metafunctions. At the stratum of lexico-grammar, for example, the ideational metafunction is associated with the system of TRANSITIVITY that distinguishes between the different kinds of processes our construed in our language. Examples would be the distinctions between *mental*, *behavioural*, *material*, *relational*, *verbal*, and *existential* processes. The subsystem of VOICE, which includes the traditional distinction between active and passive clauses, is also included in this system. The interpersonal metafunction is associated with the lexico-grammatical system of MOOD, which was seen above and also includes the subsystem of MODALITY. Finally, the textual metafunction is linked to the system of THEME. This system deals with the ordering of themes and rhemes in text, and has connections to the ordering of Given and New information in texts.

One of the important aspects of the proposed metafunctional diversification in language is that a text is meant to achieve all of these functions simultaneously (Hasan, 2009).[12] This means that, for instance, a given clause will not only serve an ideational function, but also an interpersonal and textual one. As an example, we can provide simultaneous analysis for the following clause (adapted from Halliday and Matthiessen, 2014: 143):

(5) The Duke gave the teapot away

In terms of the textual metafunction, it is possible to analyse the above sentence in terms of **Theme** and **Rheme**. The **Theme** element of the clause describes the part that "serves as the point of departure of the message; it is that which locates and orients the clause within its context" (*ibid*: 89). **Rheme**, on the other hand, is the rest of the clause that is not offered this distinct status. The textual analysis for (5), then, can be given as follows:

The Duke	gave the teapot away
Theme	Rheme

In terms of the interpersonal metafunction, we can in the first instance distinguish between the **Subject**, **Finite**, and the various categories, which together form the **Residue** elements. The **Subject** in this case is defined in functional terms; it is the part of the clause that is

[12] The exception would be the logical metafunction, which is only relevant in the analysis of clause complexes. As a result, single clauses cannot be analysed in terms of this metafunction.

referenced to when assessing the affirmation of denial of the proposition (Halliday, 1984). The **Finite** element is the element of the clause that is marked with the interpersonal meanings of tense and modality. Together, the **Subject** and **Finite** elements form the **Mood** element of the clause. This contrasts with the **Residue** and, as the analysis below will show, there are various different categories that constitute it. The interpersonal analysis of (5) can be given as follows:

<table>
<tr><td>The Duke</td><td colspan="2">gave</td><td>the teapot</td><td>away</td></tr>
<tr><td>Subject</td><td>Finite</td><td>Predicator</td><td>Complement</td><td>Adjunct</td></tr>
<tr><td colspan="2">Mood</td><td colspan="3">Residue</td></tr>
</table>

Finally, (5) can be analysed in terms of the ideational metafunction. The important notion in this case is the *transitivity* of the clause, which distinguishes between the **Process**, **Participants** and **Circumstance** of the clause. The first of these describes whether the clause has construed the aspect of experience described by the clause in terms of *material, mental, relational,* etc. processes. The **Participants** are those members that refer to entities that are engaged in the **Process**, while the **Circumstance** elements give information about the situation of the event. The ideational analysis of (5) in these terms can be given as follows:

The Duke	gave	the teapot	away
Participant: Actor	Process: material	Participant: goal	Process: material

In sum, SFL as theory envisions language as both *systemic* and *functional*. It is systemic because SFL highlights the options that are made available through language, and focusses primarily on the paradigmatic relation between linguistic units. It is *functional* because it sees the main goal of linguistic theory as explaining the role language plays in our eco-social environment. The different kinds of functions that language can fulfil are grouped in terms of three or four general *metafunctions*. For the moment, this characterization of the notion of function will suffice for the goals of setting out the central ideas and theoretical tools in SFL.

With this outline in mind, we can now turn to the theoretical dimensions that are integral to the **architecture** of systemic theory. Each of these will be discussed in term, with efforts made wherever possible to make their characterization as explicit as possible, in order to aid the discussion in the second part of this chapter.

3.2. Stratification

The first central dimension for studying language from a systemic viewpoint is that of *stratification*. This is the appreciation of a central fact about language that has been noted since antiquity, namely that an explanation of language must give some account of how linguistic form is related to linguistic meaning. Most, if not all, modern linguistic theories capture this need to view language at a number of different *theoretical levels*: generative grammar has its core syntactic module interacting with other cognitive systems: the conceptual-intentional system that gives language its meaning, and the articulatory-perceptual system that provides language with its sounding. The *theoretical levels* of Logical Form and Phonological Form describe where the core syntactic system interacts with these other modules (c.f. Chomsky, 1995). In Lexical Functional Grammar, the *theoretical levels* of *c-structure* and *f-structure* describe language in terms of, respectively, it's syntactic and semantic organization, and a major part of the theory is to determine how these two descriptions are linked (Kaplan and Bresnan, 1982). The idea that language can be studied in more than one way is also evident in the way that many university linguistic departments are organized, and the specialities that practising theoretical linguists identify as belonging to. Such specializations include *phonology* (the study of sound systems), *morpho-syntax* (roughly the study of word and sentence structure), *semantics* (the study of context-independent, usually truth-conditional, meaning), and *pragmatics* (the study of context dependent meaning, or the meaning of language in use). Although the theoretical architecture of SFL does not divide up the study of language in precisely this way, it is inspired by these traditional distinctions.

The stratificational framework breaks the phenomenon of language up along these lines. Following Martin and Matthiessen (1991), it is usually diagrammed as follows:

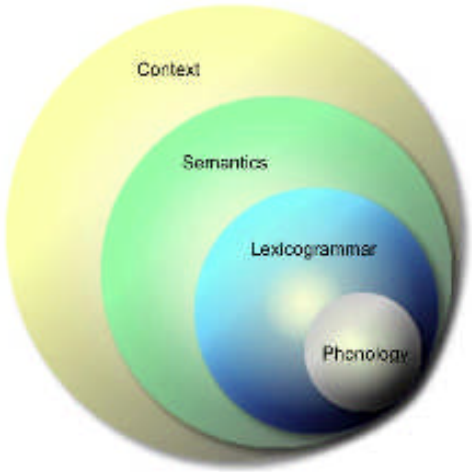


Figure 3.1. The stratificational framework (image taken from Taverniers (2011: 1101)

A few aspects of the above figure need to be highlighted: (1) there is no distinction between *morphology* as the study of grammar below the constituency level of the word, and *syntax* as the study of grammar as it applies at the level of the sentence (c.f. §3.3 below). In SFL both of these traditional disciplines are unified into a single stratum of *lexicogrammar*. The link between these two aspects of language is not according to the dimension of stratification, but instead is understood in terms of *delicacy*, or specificity. (2) No distinction is made between the linguistic levels of semantics and pragmatics. The notion of *semantics* in SFL is markedly different from the ways it is used in other linguistic theories, especially formal, truth conditional theories of meaning. This is due to SFL focussing on semantics from an inter-organism perspective, which includes its contextualized and social nature (Butler, 2003: 192). The notion of semantics as context-independent meaning therefore carries no force in the study of language from this theoretical viewpoint.

The first motivation behind the stratificational framework within SFL stems from the need to consider language in different terms, or as it is put by Halliday and Matthiessen (2014: 24) the need to "talk about language under different headings". On this reading, the strata amount to the ability to view a linguistic text in terms of "soundings" (phonology), "wordings" (lexicogrammar), or "meanings" (semantics). On this reading then, it seems that stratification should be best interpreted as a hierarchy of *levels of interpretation*, in a similar way as was seen for Marr's proposal in §2.3. There is only one phenomenon under study but in order to understand it fully we need to view it in different ways, in order to aid our study of such a complex system.

In order to fully appreciate the nature of the stratificational framework, it is useful to refer back to its inspirational roots. The two major influences on the development of the stratificational framework come from Lamb's stratificational linguistics (Lamb, 1966), and the work of the Danish linguist Louis Hjelmslev (e.g. Hjelmslev, 1943/1961). A detailed discussion of the work of Hjelmslev is given by Taverniers (2008), and its inspirations for SFL theory are discussed by Bache (2010), Taverniers (2011) and Martin (2014). Some of these ideas, specifically the idea of *content and expression* planes, will be outlined below.

The distinction between *content* and *expression* stems from the Saussurean distinction between *signifier* and *signified* in relation to the structural unit of the sign. The first, the *signified*, roughly corresponds to a concept or thought. The second, the *signifier*, is an account of a mental representation, or "sound-image" of how the language expresses this concept. Figure 3.2 presents a famous example from Saussure (1916 [1983]), where the 'higher' level *signified* represents the concept or idea of a *tree*, while the 'lower' level *signifier* is the linguistic form that expresses this concept, in this case ARBOR:

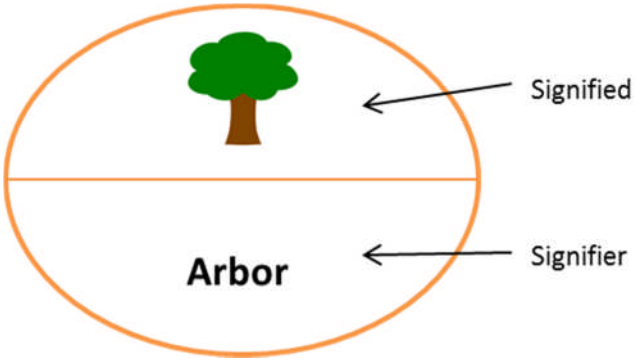


Figure 3.2. The Linguistic sign in structural linguistics

Hjelmslev continued the ideas of Saussure by relating the above distinction not only to the basic unit of linguistic analysis, but also to the overall architecture of the study of language itself (Taverniers, 2008: 369). On this view, all semiotic *systems,* and not purely the units within those systems, could be split along similar lines. Therefore semiotic systems can be divided into two *planes*, the *content plane* on one hand, and the *expression plane* on the other. The former can (roughly) be considered to be the collection of concepts, ideas or meanings expressible in a particular semiotic system, and the expression plane is the collection of means the system has for making those meanings evident.

Returning to the stratificational framework, then, a distinction can be made in the first instance between an expression plane and a content plane. Traditionally, the content plane contains the two strata of lexico-grammar and semantics, while the expression plane contains the stratum of phonology (and phonetics)[13]. This results in the following situation:

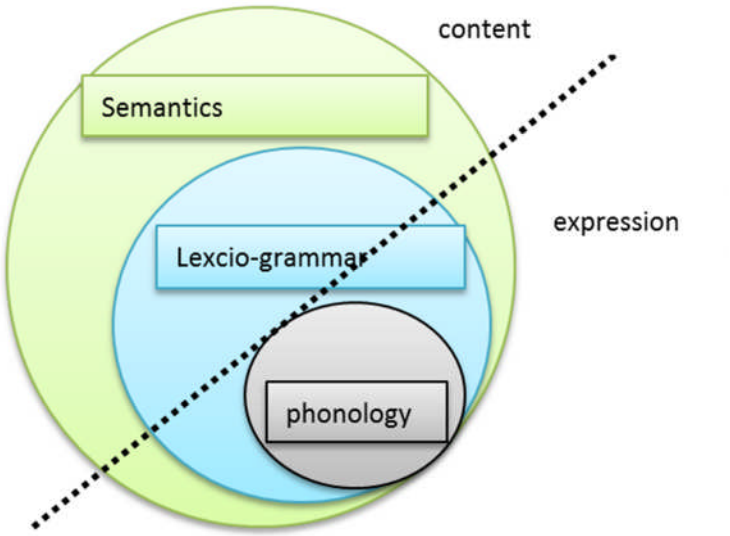


Figure 3.3: Content and Expression in the stratificational framework.

The nature of the division between *semantics* and *lexicogrammar* needs to be explained, as these are two *levels of interpretation* both found within the content plane of language. Within SFL, this division is termed as the *internal stratification* of a semiotic plane, in this case the content plane. This indicates that although we have two different strata in this case, the divide between them does not cross a semiotic boundary between content and expression.

A question can be asked at this point: the idea of lexicogrammar in SFL is markedly different to how it is considered in other linguistic theories, so what is the evidence for this "internal stratification" of the content plane of language? The motivation behind this stratal distinction comes from Halliday's research into, and discussion of, the development of language in children (c.f. especially Halliday, 1975; 1984; 1993). Halliday argued that when a child is developing their language abilities can indeed be described as a linguistic system

[13] It should be noted that in many theoretical discussions in SFL the stratum of *phonetics* is, rightly or wrongly, omitted from discussions. This can be clearly seen in some of the theoretical discussions on the stratificational framework. I believe there may be reasons for thinking that the link between the two is in terms of instantiation, instead of stratification. For this reason, I will not discuss phonetics in this work, and phonology will remain as the lowest level for the discussions here.

containing only two strata. Their language is characterized by a direct mapping between elements of expression (sound) and elements of content (meaning), without evidence for any intermediate level. He dubbed this stage of linguistic development *protolanguage* (Halliday and Matthiessen, 2014: 24):

> *In infants' protolanguage, which has yet no grammar in it, there are simple signs; for example, a meaning 'give me that!' is expressed directly by a sound, like 'nananana', or maybe by a gesture of some kind. Here we have just two strata, a stratum of content, and a stratum of expression.*

More examples of *protolanguage* are given below (Halliday, 1993: 95)

(6a.) [grasps object and release]. 'I want to hold that'

(b.) [touch object lightly, momentarily] 'I don't want that'

(c.) [touch object firmly for a measurable time] 'Go on doing (what you were doing) with that (e.g. throwing it up in the air).

Before a child begins to develop some kind of grammatical ability they communicate their wants and desires using a distinct form of expression, which can be either vocal or gestural. One important fact about *protolanguage* is that they vocal and gestural means cannot be combined in any way: a new form of expression must be created for each new meaning to be communicated. For a time, this strategy can be relatively successful, since the child only interacts with a relatively small number of people and objects, and only needs to express a limited number of meanings. However, as time goes on this direct mapping between expression and content reaches breaking point. As the child needs to communicate an increasing amount of meanings in an increasing number of situations, the protolanguage becomes inadequate. At this point, the child needs to develop the ability to combine forms of expression to communicate a larger number of meanings. In other words, the child develops a grammar. Without the intermediate level of *lexicogrammar*, there would be no means for any human to combine existing forms of expression and language would consist of a simple one-to-one mapping between the content and expression planes.

We can sum up the differences that would hold between semiotic systems with and without an internal stratification in the content plane as follows. A simple semiotic system with only one stratum in the content plane and another in the expression plane would simply be a collection of one-to-one mappings between expressions and meanings. Internal stratification of the content plane is then be required when this basic mapping is no longer sufficient, and a user needs a means of combining different semiotic signs into more complex units.

One crucial notion behind the idea of *strata* needs to be emphasised. Each of the strata of *phonology*, *lexicogrammar* and *semantics* are **systems**. In other words, they describe different aspects of the **meaning-making potential** of language. This idea of potential stems directly from considering language in systemic terms, i.e. highlighting the paradigmatic organization of semiotic systems. Considering network 3.1 again, in setting out the various options in this clause system, the **potential of possible MOOD values** is given. In this way, the phonological stratum represents the **potential soundings** a particular language has, the lexicogrammatical stratum provides a characterization of the **potential wordings** and the stratum of semantics states the **potential meanings** in the language variety under study.

The strata are necessarily separate, with each one focussing on different aspects of a text. They are, however, linked. The relationship between strata on the stratificational framework is known as *realization*. We have encountered this term before, when the classical, symbolic cognitive theory was discussed in chapter 2. However, the term is used in a different way within systemic theory, and it has different properties. The most important of these is that, whereas *realization* in the work of Marr and Fodor was considered to operate purely from a higher level to a lower level, whereby a particular system of algorithms and transformations or computations was realised by a certain number of physical systems, within SFL the relation can be interpreted as operating in both directions, with each of these inter-relations altering how you view the relation between two adjacent strata.

Firstly, the relation between a higher and a lower stratum within the framework can be interpreted in terms of construal or encoding, depending on the semiotic relationship between the strata involved (Cloran, 1992: 58; Halliday and Matthiessen, 2014: 26). The general picture can be summed up as follows (*ibid*):

[from environment to] meaning	Interfacing, via receptors	**Semantics**
[from meaning to] wording	Internal organization	**Lexicogrammar**
[from wording to] composing	Internal organization	**Phonology**
[from composing] to sounding	Interfacing, via motors	**Phonetics**

Table 3.1. Realization as a top-down relationship

The difference between construal and encoding comes from the nature of the relationship between two adjacent strata: if one is crossing the semiotic boundary between a content plane and an expression plane then the term *encoding* is more appropriate, and if one is instead dealing with an internal stratification of a content or expression plane then the relationship is instead interpreted in terms of construal. This major difference stems from the arbitrary relationship between a content and expression plane on one hand, and the 'natural' or non-arbitrary relationship between a plane-internal stratification on the other. This is most often described in terms of the relation between semantics and lexicogrammar, as mentioned above. The reverse perspective on realization, one that holds from a lower-stratum to a higher one, is interpreted in terms of "activation" (Cloran, 1994: 54). From this perspective, although lower strata may semiotically encode or construe the information at the higher level, in reverse the higher level can be thought to activate features, making them more likely to occur.

A further important property of the notion of realization is that it is inherently *non-causal*. The claim is not that semantic features cause lexicogrammatical ones to be present. This point is clearly made by Halliday and Matthiessen (1999: 25):

> *In any stratal system (i.e. any system where there are two strata such that one sis the realization of another) there is no temporal or causal ordering between the strata. It makes no sense to ask which came first or which causes which. That would be like taking an expression such as x = 2 and asking which existed first, the x or the 2, or which caused the other to come into being.*

This lack of causal relations between the strata in the framework is, I think, an important one. It fully encapsulates the idea that this theoretical dimension is best considered to be *levels of interpretation*. The discussion at the beginning of this section of treating the distinctions made by the stratificational model in terms of "talking about language under different headings" seems apt here. Language is a highly complex phenomenon. In order to bring some sense to this complexity and make the job of linguists more manageable, it seems sensible to divide the study of it into more manageable "chunks" by focussing on different aspects of language itself. This relationship can be viewed in terms of stratification providing an "epistemological sieve" through which a text (or series of texts) passes. A text is divided into strata that represent these different aspects, and then theoretical work can be done on each of these aspects independently (although in SFL an eye is always kept on the other strata, in line with the holistic picture of language it hopes to give)

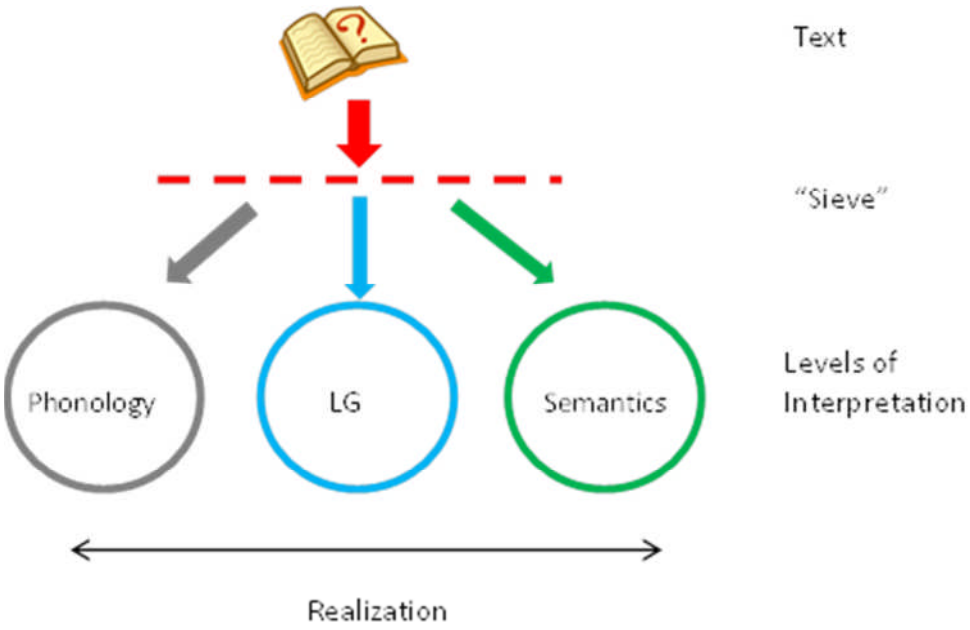
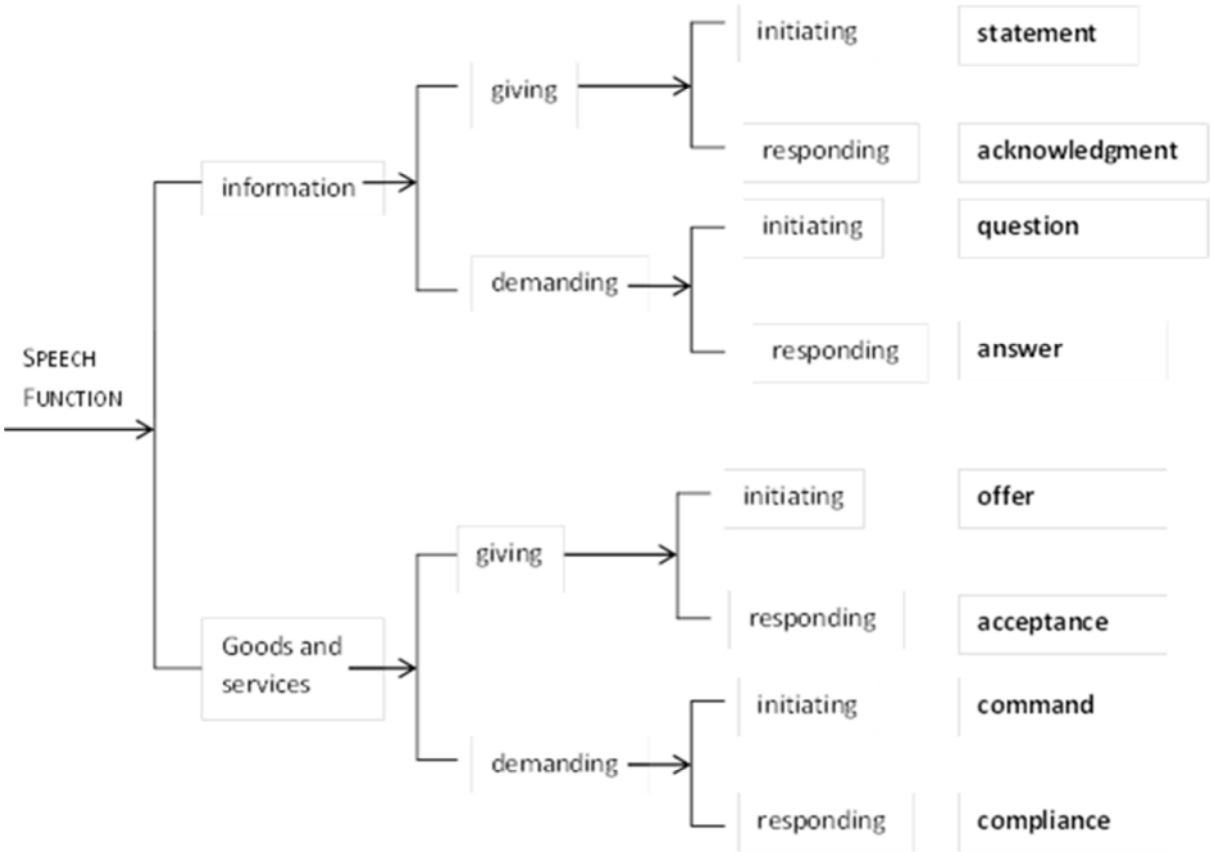


Figure 3.4: Stratification as levels of interpretation

The view presented by interpreting strata in terms of levels of interpretation is that it postulates only one object: the text. The strata then arise from viewing this text in different ways: as an instance of 'sounding', 'wording', or 'meaning' in turn. Semantics is not different to the other strata in the sense that it deals with essentially different ontological objects, instead they define different epistemological categories imposed by theoreticians and researchers. It also accounts for the non-causal relationship stated above: one stratum cannot cause another if a text is described by all the strata simultaneously.

In order to best bring out this relationship, let's consider two descriptions of language, one at the semantic level and the other at the lexicogrammatical level that are claimed to operate over the relationship of realization. This is the link between the semantic system of SPEECH FUNCTION and the lexicogrammatical relation of MOOD, both of which describe meanings in the content plane in terms of the interpersonal metafunction. The system of SPEECH FUNCTION is given as below (adapted from Martin and Rose, 2007: 226):



Network 3.2: SPEECH FUNCTION system network

This system network represents the various options available at the semantic stratum in terms of the kinds of negotiation that utterances can have in social settings. It distinguishes between those that involve a negotiation of information from those that concern goods and services, as well as within these two categories, those that are to do with giving the respective types of "thing" under negotiation, from those that demand it. In the following discussion only the initiating features in the network above will be considered. Some examples are given below:

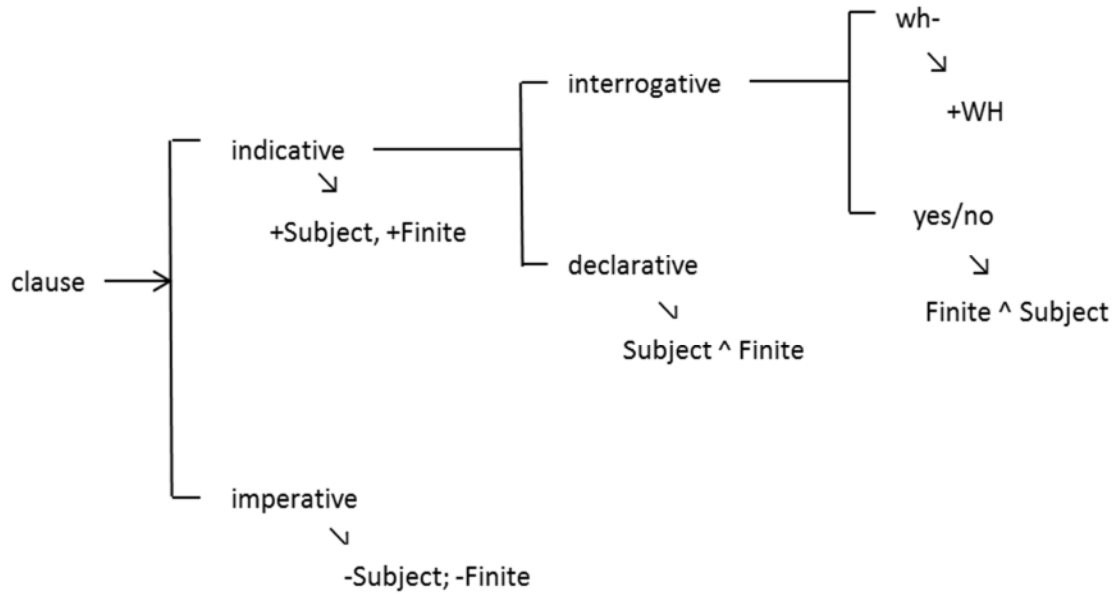
(7a.) The key is under the mat. **[statement]**

(b.) Is my coat at the dry cleaners? **[question]**

(c.) Do you want a cup of tea? **[offer]**

(d.) Give me the hairdryer! **[command]**

The corresponding lexicogrammatical system of MOOD is given as below (adapted from Matthiessen and Bateman, 1991: 92):



Network 3.3: Lexicogrammatical MOOD system

Again, some examples of the above system features are given below:

(8a.) Go away! **[imperative]**

(b.) We have to leave now .**[declarative]**

(c.) Where has the cat gone? **[Wh-interrogative]**

(d.) Have you seen my wallet? **[yes/no interrogative]**

It is clear that these two sets of system features are related, but what may be less clear is what distinguishes them as in both cases we seem to be dealing with questions, statements, etc. The difference stems from the criteria that are used to determine classification in each case. This can be most easily seen by considering the realization statements (see §3.3 below) in network 3.2 just presented. For instance, the criterium for being a declarative clause is the presence of both a Subject and a Finite element (denoted by +Subject, +Finite), and the Subject element preceding the Finite element (denoted by Subject ^ Finite). The criterium for classification in terms of SPEECH FUNCTION, on the other hand, is along the lines outlined above, namely by means of the type of "thing" under negotiation, etc.

These two networks stand in a *realization* relationship to one another, with the lexicogrammatical system of MOOD, realizing the semantic system of SPEECH FUNCTION. There are clearly links between some of the features on the two system networks. For instance, declarative clauses are more often associated with statements, interrogative clauses are frequently linked to questions and offers, and imperative clauses are regularly correlated with commands. These links would count as the *unmarked* or *congruent* means of realizing speech function through the mood system. However, this picture does not always hold, as pointed out by Martin and Rose (2007: 227). For instance, although a *command* is usually realized by an imperative clause type, the same SPEECH FUNCTION can also be realized by the other major MOOD types at the lexicogrammatical stratum. For example, consider the sentences below:

(9a.) There are dirty dishes in the sink

(b.) Would you mind washing up those dirty dishes?

On their *unmarked* realizations, the sentence in (9a.) would grammatically be considered to be a realization of a *statement*, while (9b.) would be the realization of a *question*. However, by considering how one would respond if the *unmarked* realization was interpreted, it can be seen that these frequent realizations are not always intended:

(10a.) A: There are dirty dishes in the sink
B: So there are.

(b.) A: Would you mind washing up those dirty dishes?
B: Yes, I would mind

Although the above responses by B in each case are not necessarily wrong, it can be seen that in some cases A's intended meaning has not been correctly understood. In (10a.), B has interpreted A's utterance as a *statement* and so has responded with the corresponding responding SPEECH FUNCTION of *acknowledgment* (c.f. network 3.1 above). Similarly in (10b.) a literal interpretation on B's part has caused him to respond with the SPEECH FUNCTION of *answer*. It is not difficult to imagine situations, however, where these unmarked realizations are not intended. In (10a.) instead of giving information, A is instead issuing a command on B's part to wash the dishes, but for certain social reasons (most likely relating to the *tenor* of the context of situation, c.f. §4.2.3.) A has decided to mediate their command in the form of a declarative clause. A similar situation can be seen in the case of (10b.), where A's intention is to convey a *command* and not a *question*.

Each stratum in the stratification framework, then, describes a certain aspects of the linguistic system. At each point, the theoretical description of the system of language will describe the patterning of language behaviour at each of these levels of interpretation. The relation of *realization* will then provide the links between these different groups of patterned activity. In most cases, certain linguistic types will often correlate, but the relationship is more fluid in that it also captures the fact that unmarked, or less frequent, correlations also occur. It is possible to diagram the relationship between two strata as follows:

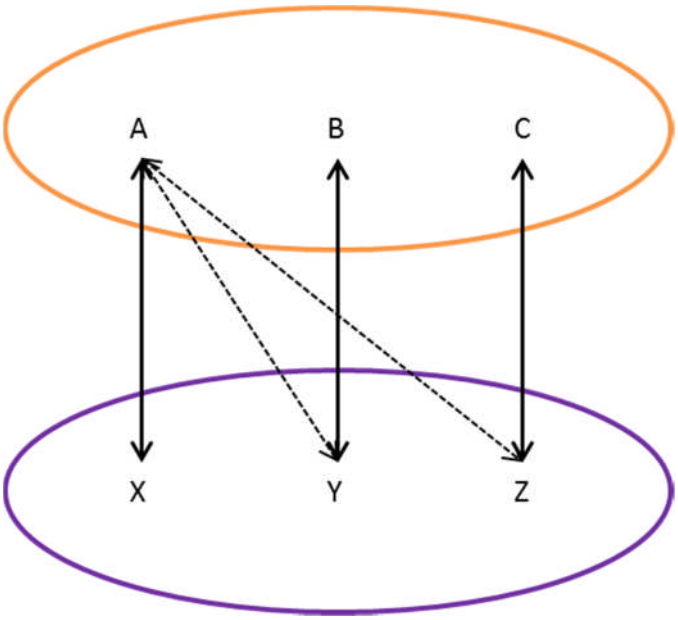


Figure 3.5: Realization: A blueprint

In figure 3.5, the two circles represent two strata, with one described as being 'higher' than the other. The letters represent linguistic types at each of these strata; for instance, the linguistic types at the higher stratum could be *command*, *question*, *statement*, and at the lower stratum they could represent *interrogative*, *declarative* and *indicative* clauses. The double headed arrows represent the realization relationship. The bold arrows stand for the *unmarked* relationship, while the dashed arrows stand for the less frequent correlations. For illustrative purposes, on the above figure the *marked* links have only been included for one of the higher level types, but it should be noted that the same could potentially hold for all of them.

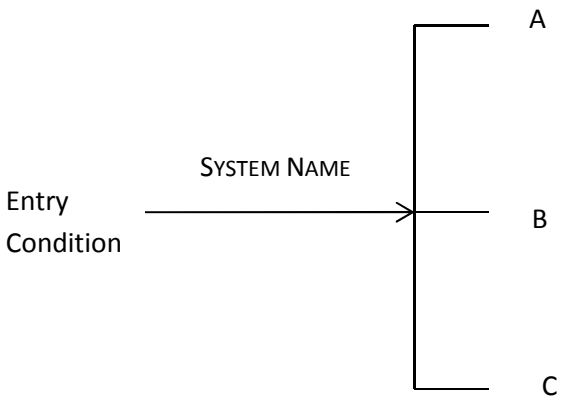
In sum, the dimension of stratification can be considered to be an ordered collection of differing *levels of interpretation*. Each of the strata within the framework views **the meaning-making potential** of language in different ways by focussing on only certain aspects of the phenomenon: treating language either as a **system** of sounding, wording or meaning. These *levels of interpretation* are linked by the semiotic relationship of *realization*, which correlates patterned activity at all of these points based upon the frequent and infrequent co-occurrences. This relationship is properly considered to be **bidirectional**, in that the lower systems **construe** or **encode** the systems above, and the higher systems **activate** by those further down.

3.3. Delicacy

Language, then, can be described at different levels of interpretation. In doing so one looks at language in different ways and accounts for different aspects of a text in each case. The main modelling resource in SFL for describing the potential of language at the various strata is through the use of system networks. These were briefly introduced in §3.1.1. Further details regarding these networks will now be given, and their links to the dimension of *delicacy*.

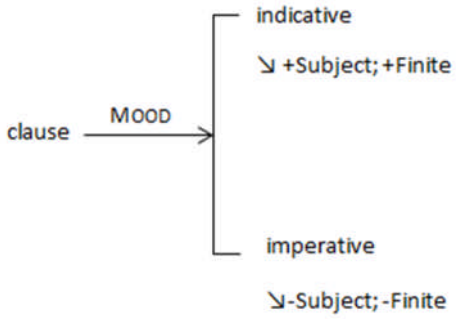
Within SFL, the notion of delicacy is linked to the idea of the *grammarian's dream* (Halliday, 1961; Hasan, 1987). According to this hypothesis, *lexis* and *grammar* are not to be considered to be two separate phenomena within the linguistic system. Instead, the relationship between the two should correctly be considered to operate over the scale of specificity, or in SFL terms, *delicacy*. *Lexis* is thus construed as the "most delicate grammar", as indicated in the title of Hasan's (1987) piece. The proposal is that in the *lexicogrammatical* system, the system of 'wordings', we start with distinctions that are traditionally covered by the term *grammar*, but as the description becomes more specific in the distinctions it makes, a move is made towards what is traditionally called *lexis*. In other words, at a general point on the scale of delicacy, the majority of the distinctions are between different types of clauses, but ideally the scale would end with the choices made between individual lexical types.

A system network describes the notion of potential in language by laying out the various options that are available at each *level of interpretation*. At the most general end, a system network can be described by an *entry condition,* which specifies the nature of the phenomenon under discussion, and a number of options that define the different types of that phenomenon that the available in the system. Additionally, system networks are often given a descriptive name:



Network 3.4: Basic system network to one degree of delicacy

The logic behind networks of this form can be stated as follows: If [entry condition], then [in terms of SYSTEM NAME] either A, B, or C. In the rest of this discussion I will omit the slightly awkward phrase "[in terms of SYSTEM NAME]" for ease of elucidation. As a concrete example, below is the lexicogrammatical system of MOOD, which we will be returning to at various points in this thesis. In this case, the system only operates at a low point on the scale of delicacy:

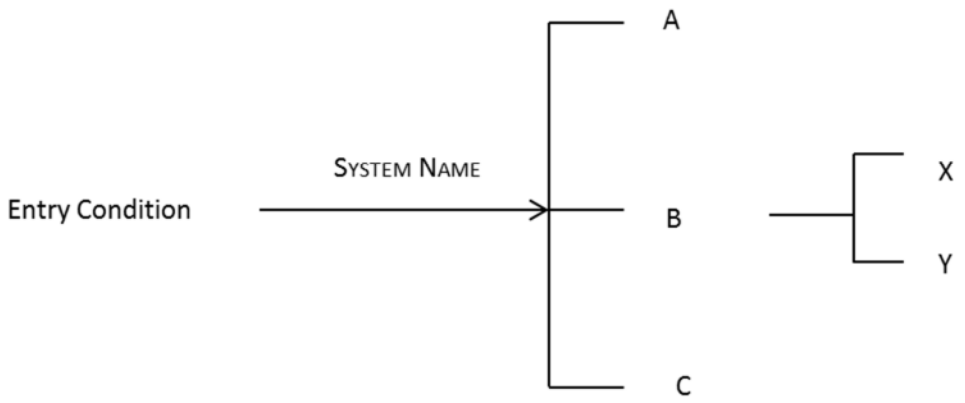


Network 3.5: system of MOOD at low level of delicacy

The above network specifies that in terms of this particular system, all major clauses in English can be divided into two groups: those that are *indicative*, and so contain both a *Subject* and *Finite* element, and those that are *imperative*, and so generally lack both these elements. No other option is available, so all major clauses will fit into one of these two groups. Indicated above are also the *realization statements* for the above two options. That

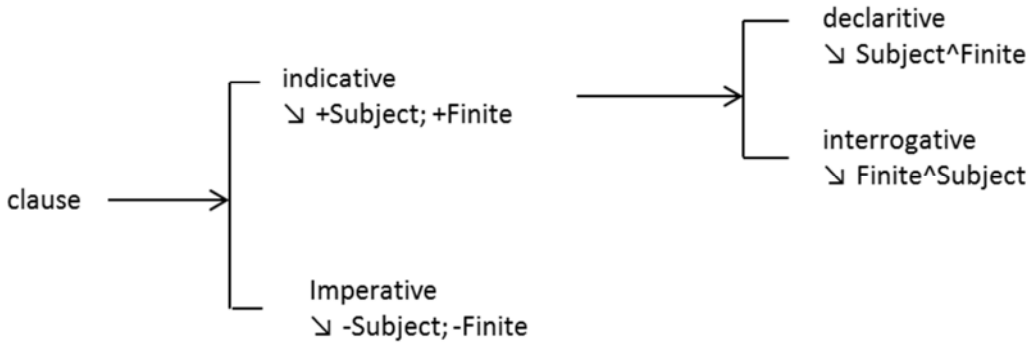
is, if a choice of *indicative* is made, then the clause will have both a Subject and Finite element. This is not the case, if *imperative* is instead relevant.

Such basic system networks can be increased in terms of delicacy. That is to say, after an initial "choice" has been made, further options become available corresponding to more specific divisions that can be made in that system network. Extending the general system network above in terms of delicacy results in the following:



Network 3.6: Basic system network to two degrees of delicacy

The logic behind the above network is as follows: If [entry condition] **and** if B, then either X or Y. Increasing our linguistic example above gives the following:

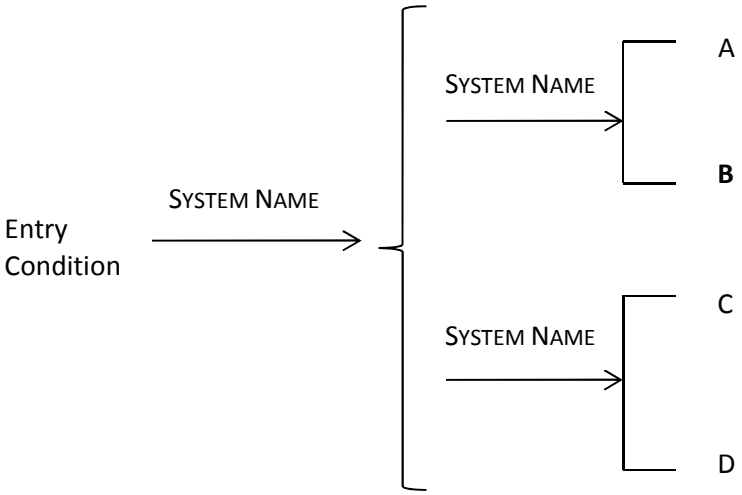


Network 3.7: system of MOOD with increased delicacy

The logic described above can now be seen. All major clauses are still either *indicative* or *imperative*, but now we can make the distinctions between different kinds of indicative clauses: *declaratives*, where the *Subject* precedes to *Finite* element, and *interrogatives*,

which describe the cases where this ordering is reversed. Again, a particular choice in delicacy will result in different structural features being seen in the clauses.[14]

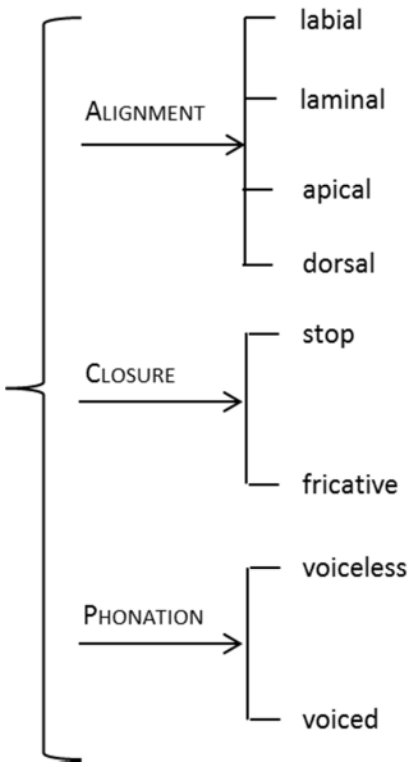
Many system networks presented in the literature follow the above patterns. However, when more detailed patterns in language are observed, it is necessary to introduce more complicated types of system network. One such network type is a simultaneous network:



Network 3.8: A simultaneous system network

The logic behind this can be stated as: If [entry condition], then either A or B, **and** either C or D. An example of a simultaneous system is as follows (Cléirigh, 1998:2):

[14] For instance, the term Subject^Finite refers to an ordering. In declarative clauses the subject precedes the finite element.



Network 3.8: Consonantal potential

The above system is for consonantal phoneme types. At this stage, it is only necessary that all consonantal phonemes will take a certain value at all of these systems (ALIGNMENT, CLOSURE, and PHONATION) simultaneously.

In sum, then, the dimension of delicacy divides the potential for meaning-making in language at each *level of interpretation* in terms of *levels of specificity*. That is, it makes increasingly fine-grained distinctions between the linguistic categories that are thought to be important for describing the meaning-making nature of language.

3.4. Rank

The dimension of *rank* is included in the architecture of SFL theory in order to account for the various constituency relationships that are evident in human language. Whereas stratification has been seen to be a hierarchy of *levels of interpretation*, and delicacy is best considered as containing various *levels of specificity*, the dimension of *rank* is intended to account for the various constituency relations that are relevant for the study of language.

Constituency is a structural relationship that refers to the groupings of linguistic units into larger ones. As a result, the overall guiding principle is one of *mereology*, or 'part-whole'

relations. So far in the theory of SFL, two rank scales have been identified; one of these operates at the phonological stratum, and the other at the lexicogrammatical one (Halliday and Matthiessen, 2014: 21-22; Eggins, 2004: 121-131):

(6) ***Phonological rank scale:***

phoneme → syllable → foot → tone group

(7) ***Lexicogrammatical rank scale***

morpheme → word→ phrase/group → clause

In both (6) and (7) each unit to the right of the arrow is made up of a collection of units to the left of it. In this way, the dimension of rank acknowledges the constituency relations that are included in many linguistic theories; although it should be pointed out that the details of rank within SFL are relatively simple in comparison with other linguistic theories.

As an example of a rank analysis at the stratum of lexicogrammar, consider the following sentence (Eggins, 2004: 127)

(8) All students must satisfy all assessment requirements

In terms of rank, example (8) would be analysed as follows (*ibid*):

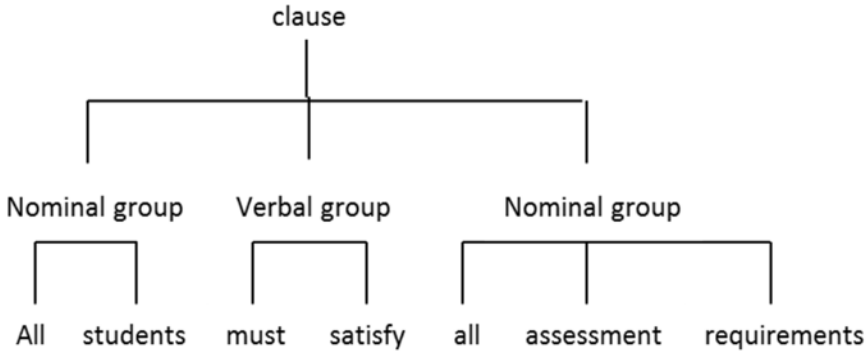


Figure 3.6: Rank scale analysis for (8)

Strictly, the example only gives a partial analysis of the sentence in (8) in terms of rank; it provides an analysis from the word rank to the clause rank. The units *students* and *requirements* could be divided further into a conjunction of morphemes (*student + s*;

requirement + s). However, the above picture is sufficient to capture the kind of relation that the rank scale is intended to convey.

The hierarchy of rank theoretically operates at every stratum in the stratificational framework. That is to say, with each new stratum comes a new set of units ordered by the constituency relation. So far, the rank scales for the strata of phonology and lexicogrammar are well established, but if this picture is true then one should also be evident at the strata of semantics (Martin and Matthiessen; 1991; Martin, 1992). For instance, Martin (1992: 20-21, emphasis mine) writes: ”within strata, description is further organized through layering (simultaneous metafunctions) *and constituency (rank)*”. Without further clarification, it seems as though the notion of rank must also apply at the (discourse) semantic stratum. This is further corroborated by Martin’s assertion that discourse semantics deals with “larger units than lexicogrammar” (ibid: 19). If the units at this higher stratum are indeed larger than those at the level of grammar, than this implies that there is a different set, or at the least further sets, of constituency relations that are relevant at this level of interpretation. However, as of yet there has been no widely accepted account of the kinds of constituency relationships one would expect to find. As discussed by Martin (1992: 46-47), attempts have been made, for instance, by discourse analysts (but not necessarily directly working within the SFL framework). Sinclair and Coulthard (1975) and Coulthard and Montgomery (1981) posit the following rank scale at the level of discourse, or in Martin’s terms discourse semantic, consisting of the higher rank of *exchange* and a lower rank of *move*. In this sense an exchange in discourse is made up of one or more moves. This was extended by Burton (1980) and Butler (1985) to include the lowest rank of *act*. This presents us with the picture as follows:

Discourse Semantic Rank Scale (Burton/ Butler)

act → move → exchange

However, in his discussion Martin (1992: 56-57) seems sceptical about a rank analysis of discourse semantic structure in this way. He argues that such accounts can already be captured at the levels of lexicogrammar, and possibly phonology. The example he gives is that the textual metafunction at the level of lexicogrammar, that is analysing clauses in

terms of various kinds of topics, can already provide the relevant account of move sequences. Also, he points out that at these lower strata, there is only one rank scale that is relevant for all metafunctions: the clause is a relevant unit of analysis for all three metafunctions: the experiential, interpersonal and textual. In this case of exchanges, moves and acts, this seems to only be relevant for the interpersonal metafunction, and not for the other three. This point is also raised by Halliday and Matthiessen (2014: 661), who question whether a unified rank scale can be given at the stratum of semantics. If so, then this would be a major difference between semantics and they leave this question open, commenting that more research on the nature of the semantic stratum is needed before this question can be answered. Even so, they postulate the basic units of rank at a semantic level, although the nature of these units differs depending on one's metafunctional perspective:

> ***Semantic Rank Units (Halliday and Matthiessen, c.f. Halliday and Matthiessen, 1999; 2014)***
>
> Ideational: figure [quantum of change of flow of events]
>
> Interpersonal: move (proposition/ proposal) [quantum of interaction]
>
> Textual: message [quantum of information]

Regardless which one, if any, of these proposals are along the right lines, it is interesting to note that as one comes to consider the higher levels of interpretation within the framework, it becomes harder to identify a well-defined set of constituents that enter into a hierarchy of rank. With regards to the highest stratum of context, there seem to be no proposals as to what a rank scale at this level might be, even though if it is to be properly considered as such, one should be evident. I believe that many of the problems can be linked to the nature of the units themselves as one moves upwards along the stratificational model. This will be discussed in detail in §3.8.3.2, and therefore this discussion will be put on hold until that point.

In sum, the proposal in SFL is that with each new stratum, or as they have been interpreted in this chapter, with each new level of interpretation comes a new hierarchy of constituency. So far within the theoretical literature, the hierarchies of constituencies at the strata of

phonology and lexicogrammar have been presented, although they are hypothesized to exist at the semantic stratum as well.

3.5. Instantiation

One of the major controversies of modern linguistics concerns the relationship between abstract descriptions of language, and language use. The modern source of this controversy was introduced by Saussaure's (1916) distinction between *langue* and *parole*. For Saussure, *langue* was the study of the abstract nature of language, which he considered to be social in nature, whereas *parole* was the study of the implementation of this system in actual use. For Saussure, linguistics should be considered with studying *langue*. A similar dichotomy has also placed a major role in Chomsky's generative linguistics. The major division is this case is between *competence* and *performance*, or the similar distinction between *I-language* and *E-language*. This is the claimed difference between knowledge of grammaticality, or the ability to generate all those and only those sentences that are grammatically well-formed (*competence*), and the actual production of linguistic utterances (*performance*). The object of study in generative linguistics is firmly within the domain of *competence*, and furthermore, the *competence* of an ideal speaker-hearer (Chomsky, 1986; c.f. Jones, 2005; Ludlow, 2011, c.f. §5.3.1).

SFL, on the other hand, dismisses this divorce of knowledge of language, or language as a system, with the production of linguistic texts. In place of such a strict distinction, the theory instead conceptualizes the relationship between the underlying system of language and the actual production of texts as a cline. This *cline of instantiation* can be diagrammed as below:

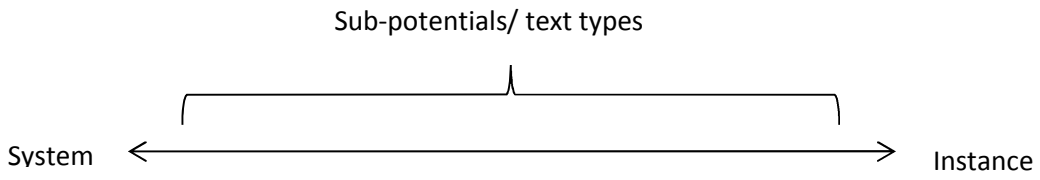


Figure 3.7: The cline of instantiation

In between these extremes there lies a number of intermediate points that correspond firstly to generalizations over a number of texts, and secondly, as sub-potentials for the creation of new texts and new meanings. Further detail about the nature of these mid-

points on the cline will be given below, after a few more overall details about its nature have been presented.

SFL practitioners often liken this relationship to that of weather and climate. This analogy is discussed by Halliday and Matthiessen (1999: 328-329; 2014: 27), Halliday (2003), and Matthiessen (2007a.). Halliday and Matthiessen (1999: 328) discuss the analogy as follows:

> *The terms **climate** and **weather** are related to each other as **langue** to **parole**: as the system of language is to its text. Weather is the instantiation of climate, climate is the system 'behind' the weather. As with language, there is, of course, only one set of phenomena here not two; when we refer to climate we are construing general principles and tendencies that 'explain' the multidimensional microvariation that is what we actually have to live with when gardening or planning an outing*.

This analogy aims to highlight two aspects about the link between the system and an instance. Firstly, they operate over different time scales. A text unfolds in a relatively short amount of time, whereas the overall system evolves over a much more drawn out process. The analogy to meteorology here is apt, since any change to the climate of a particular geographical region is going to come about much more slowly in comparison to how the weather changes, say, over the course of one day. To stay with this comparison for the moment, it also captures the point that although in one sense the ascription of climate is a statistical generalization over what has come before, it can also be used to predict what is likely to happen in the future. It is, of course, not infallible, but it does provide a description of what is likely to happen on any particular day.

To elucidate this idea of a mid-point on the cline of instantiation further, presented below is the example provided by Halliday and Matthiessen (1999: §8.3) of the domain model of cooking recipes. The term *domain model* is the term provided by the authors to describe what has been labelled as *sub-potential* in Figure 3.7 above. This seems to be only a notational difference, so I will continue to use the term *sub-potential*. In this work, the authors noted, on the basis of a corpus of recipe texts, that the following differences could be seen for this particular sub-potential/ text-type:

- Participant roles will be exclusively filled by concrete “things”; there is an absence of “metathings” in the ideational semantics of cooking recipes. So, for instance, there is a lack of facts, ideas and locutions in participant roles in figures.
- Even within the category of concrete “things”, only a few sub-types are readily found in certain participant roles: Agents are predominately persons, while objects, animals, etc. occupy the Goal/ Patient role.
- Certain types of figure are preferred: the majority of figures are of the “doing” type, and specifically “doing to/with” at a further level of delicacy. Less frequent are figures of sensing and figures of being, while figures of saying are extremely rare.
- In terms of sequences: generally there exists more coarse-grained distinctions between projecting sequences and expanding sequences. In the sub-potential/ text type of recipes, only expanding types are attested.

The general picture of the nature of these mid-points along the cline of instantiation can be shown by considering the third characteristic above. As stated above, certain types of figure are found more frequently when one looks at recipes, with other types of figure seemingly missing altogether. When talking about relative terms, such as *more frequently*, this obviously implies a comparison to something else. This ‘something else’ is the consideration of language from the system end of the cline, i.e. from the consideration of the regularities in language as seen in all actual and potential contexts. In other words, a mid-point in instantiatal terms can be thought of as a variant of the overall system of language, but one in which the probabilities attached to the systems at each level of interpretation are ‘re-set’ or differ (Halliday and Matthiessen, 1999: 554). A toy example will suffice here. Consider the system network of figures, and further assume that at the system pole of instantiation there is an equal distribution of each figure in texts. This is almost certainly not true, but it will help illustration of this point (figure adapted from *ibid*, 362):

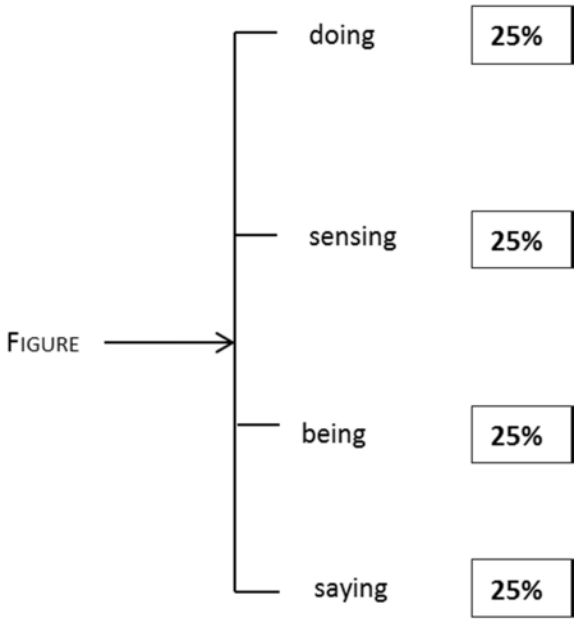


Figure 3.8: Variant of the system of FIGURE with equal probabilities

Describing a sub-potential for the ideational semantics of figures for recipes would then amount to showing if and how this group of probabilities is skewed in those particular examples. For example, the same system in the domain model of recipes could result in a situation similar to that given in figure 3.9 below. Again, although the general picture is sketched out in the work mentioned, precise values for probabilities are not given, so these values given below should only be seen as an indicator:

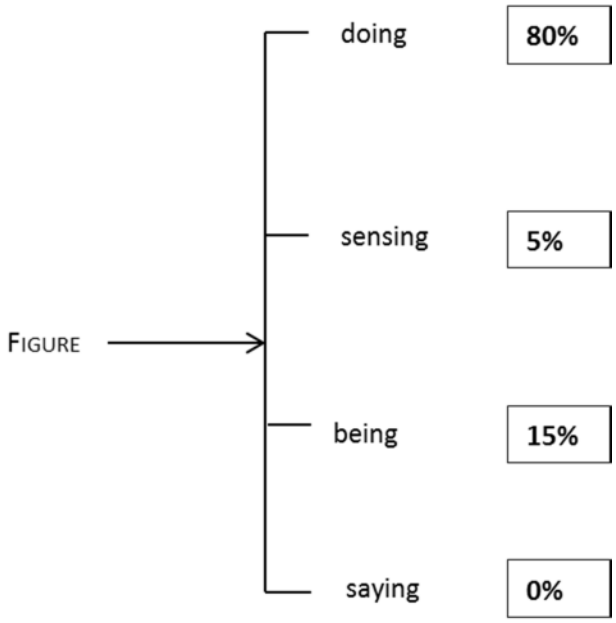


Figure 3.9: Illustrative example of the system of FIGURE in the sub-potential of cooking recipes

The mid-point on the cline of instantiation, then, does not refer to a different system being in operation in specific contexts. Instead, it is to be considered as a **version of the same system where the probability of certain features occurring has changed.**

This discussion can also illuminate one way in which the instance and the system are supposed to interact, and why they should not be considered in terms of a dichotomy on the SFL view. The probabilistic variants that one can describe at various points of instantiation (including at the system end of the cline) do not come from no-where, and neither are they fixed. They stem, ultimately, from the texts themselves. Since we are constantly producing and interpreting new texts that would fit into various text types/ sub-potentials along the cline, these probabilities are constantly in flux. How much they change, however, more likely than not depends upon how 'close' to the text or system end we position our mid-point. The utterance of a single sentence is unlikely to have any major effect on the overall system of language. Even a single utterance of, say, a figure of being in the text type of recipes is also not going to alter the overall probabilistic distribution in any meaningful way, although it will by a greater amount than the extreme system end, simply because the mid-point covers a restricted number of texts in comparison to the overall

system of language. *Repeated* occurrences of a previously dis-preferred option, however, does have the potential to alter the overall shape of the intermediate instantiatial area.

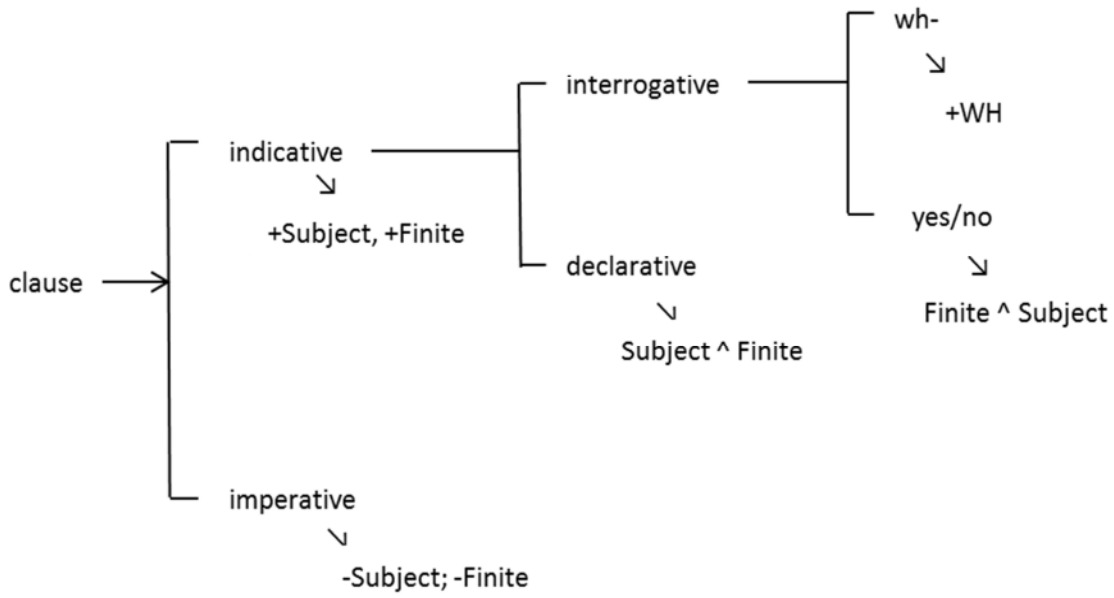
Instantiation, then, is a dimension that operates over various *levels of regularity*. It provides the means of highlighting what regularities may exist in linguistic texts in particular contexts. It also provides the means of moving beyond the strict divorce of the system from actual texts. By conceptualizing the dimension as a cline, it is possible to conceptualize of points in between these two extremes. These can be interpreted either as *sub-potentials* of the overall system of language, or generalization of texts with common properties.

3.6. Relations between these dimensions

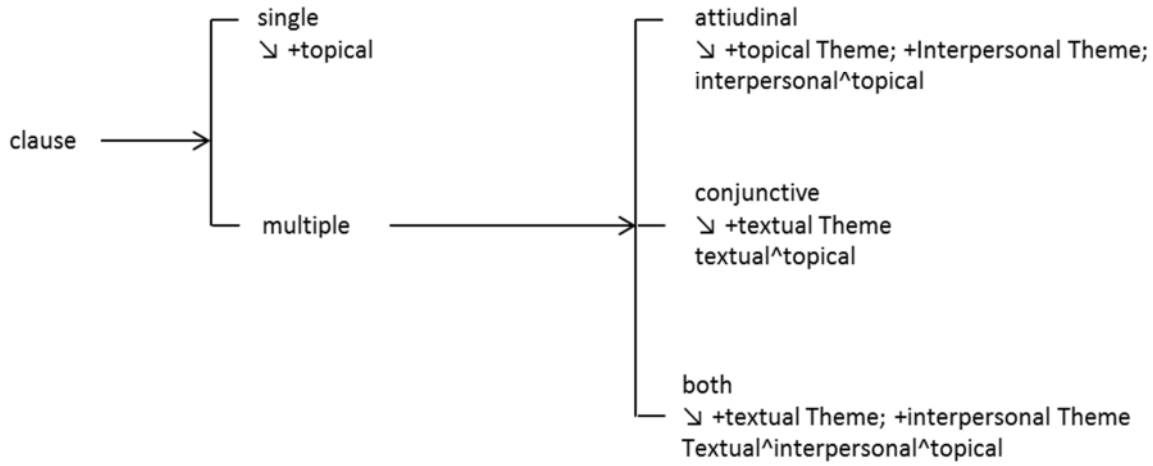
Although the above dimensions are stated to be independent of one another, and should be kept separate, they nevertheless interact with one another. This final section on the theoretical notions in SFL will describe the ways in which this is so. The major interactions in SFL are between: (1) the dimensions of rank and delicacy, and (2) between stratification and instantiation.

3.6.1. Rank and Delicacy

The relationship between rank and delicacy can be exemplified by considering many of the system networks that have been given for the stratum of lexico-grammar. On each of them, the entry condition for the network, which can then be described at various points along the dimension of delicacy, is a unit from the rank scale. For instance, consider the following system networks (repeated from the previous discussions):



Network 3.10: System of grammatical MOOD



Network 3.11: System of THEME

These system networks capture paradigmatic choices at the lexicogrammatical level. In both of these cases it can be noted that each system network begins with the statement of a certain structural unit at which the following choices are meant to hold for.

It is thought that there is a principle unit at each stratum that is the main locus of the meaning-making potential of language. Martin (1985: 249), for instance, states the following:

> The fundamental unit at each stratum (where the action is in terms of the choices made) is the phoneme [for phonology – JW], the clause [for lexicogrammar – JW] and the text [for discourse semantics – JW].

Clearly, many of the system networks at the same level of interpretation seem to focus upon the same level of structure. For instance, the system networks given above at the lexicogrammatical level describe the meaning-making potential associated with the interpersonal and textual metafunctions as they relate to the structural level of the clause. Furthermore, many of the potentials discussed in the discourse semantics associated with Martin (NEGOTIATION, APPRAISAL, PERIODICITY, etc. see Martin and Rose, 2007), are associated with the structural level of the text.

One of the general principles that arise from the interaction of the *rank scale* and *meaning-making potential* described in terms of *delicacy* and system networks is that **certain phenomena only make sense at a certain structural level.** Put in these terms, the interaction between these two dimensions can be stated as follows: rank provides us with a way of modelling the structural levels at which certain meanings can be made, and delicacy provides the means of describing what the possible range of meanings associated with that particular structural level are, to varying degrees of specificity

It is also worth returning to the idea of the *grammarian's dream*, that was encountered in §3.3 above. To recall, the idea behind this proposal was that grammar and lexis are not considered to be separate systems, but instead they are encapsulated within the same lexicogrammatical system and are linked by the scale of delicacy. To repeat the oft-quoted phrase above, lexis is to be considered "most delicate grammar". This view, if correct, tells us something additional about the links between delicacy and rank on the SFL framework: **as one moves along the scale of delicacy, one also moves down the scale of rank**. This conclusion (at least for the stratum of lexicogrammar, but see below) seems to be inevitable, if one is assuming the grammarian's dream to be true. For, as was just seen, the starting point for most scales of delicacy is the ascription of which *level of constituency* the following paradigmatic options are applicable to. In nearly all grammatical systems that have been presented so far, and the majority of those in the literature (but c.f. Tucker,

2007 for a proposal that sits in the middle of grammar and lexis) the relevant grammatical unit has been given as the clause. If one could hypothetically create a full lexicogrammatical system that operated over both grammar and lexis, then the hypothesis is that the most delicate points on the system will give the paradigmatic relations that hold between individual lexical items. Since lexical items occupy a lower position on the rank scale than clauses do, at some point as the description of language moves along the scale of delicacy there must be a descent down the scale of rank.

In sum the relation between rank and delicacy can be stated as follows: rank specifies the structural unit which certain paradigmatic options, as described by delicacy, operate over, and as one moves to more delicate options in the system, at some point a descent is made to lower structural units.

3.6.2. Stratification and Instantiation

The relationship between the notion of stratification and instantiation is usually given in terms of the *stratificational-instantiation matrix* (c.f. Halliday, 2002: 19)[15]:

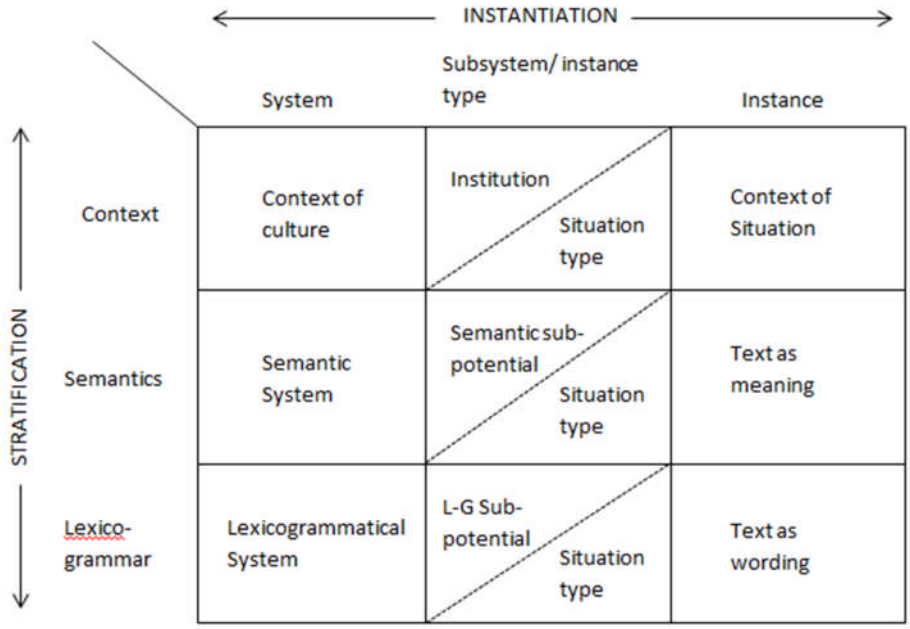


Figure 3.10: The stratification-instantiation matrix

[15] Figure 3.10 has been slightly altered from the original. In the original, what I have called (semantic/ lexicogrammatical) sub-potential is instead given the label 'register'. Register can take multiple means in SFL. This change in label aims to make the intended meaning clear. See §4.1 of this thesis for further discussion

This table highlights the fact that although it is possible to discuss the cline of instantiation as a single "thing", in systemic theory it is thought to operate over all of the strata of language simultaneously. Another way of considering it is that there is a different version of the cline at every level of analysis. The difference between them is that in each case you are discussing different elements of a text. The elements that are considered in each case are determined by the nature of the level of analysis the practitioner is currently working in. So, on this interpretation, any particular text is *simultaneously* an instance of sounding if considered phonologically, a system of wording if described in lexico-grammatical terms, and so on. This reinforces the fact stressed in the discussion above, that when we move from one stratum to another, we are still dealing with the same thing in each case, the difference is in terms of the *perspective* we as researchers take upon the object, and the kinds of relations we are looking for.

Another factor concerning figure 3.10 needs to be highlighted, although it does impinge upon discussions in the next chapter. When discussing the ideational semantic *sub-potential* of cooking recipes in instantiation, no discussion was given into **why** such regularities happen. This comes from the interaction between the context and the linguistic system. Cooking recipes often function in similar semiotic environments, involving similar types of objects, processes, etc. This similarity in context will be captured by the texts that operate in said context. In this way, stratification influences instantiation by means of the highest stratum of context motivating the groupings we see lower down in terms of instantiation. We say that such sub-potentials are **contextually motivated**.

3.7. The different types of levels in Systemic Functional Linguistics

In sum, then, the preceding discussion of the systemic literature has highlighted that systemic linguistics consider the following kinds of *levels* to be important in the study of language from a functional viewpoint:

Name	Type of Level	Example	Further Details
Stratification	Levels of Interpretation	Phonetics – Phonology- Lexicogrammar – Semantics - Context	Makes epistemological, but not necessarily ontological distinctions in the study of language. Strata are linked by the bi-directional relation of *realization* that maps probabilistic links between features at different strata.
Delicacy	Levels of Specificity	Clause – indicative – declarative or interrogative	Can be described from a typological and topological perspective. Each level of interpretation contains its own dimension of delicacy.
Rank	Levels of Constituency	Phoneme – syllable – foot – tone group. (Phonology) Morpheme – word – group/ phrase – clause. (Lexicogrammar)	Hypothesized to exist at every level of interpretation. However, rank scales at the higher strata of semantics and context are not universally accepted. Interacts with the scale of delicacy to determine the structural units at which the paradigmatic oppositions hold.

Instantiation	Levels of Regularity	System – Sub-potentials/ Text Types - Text	Conceptualized as a cline. The intermediate points between the two extremes can be interpreted in two ways – firstly as a text type or a generalization over a corpora of texts, and secondly as a variant of the overall potential for meaning-making. Each level of interpretation has a different cline of instantiation, or alternatively, it is conceived to operate over all strata simultaneously.

Table 3.2: The variety of levels in Systemic Functional Linguistics

3.8. Towards a Stratificational framework of cognitive and behavioural processes

As summed up above, SFL identifies a range of levelled hierarchies that have proved useful in the study of language. This section moves the discussion to the wider field of cognitive science, and investigates whether treating cognitive systems as a semiotic can aid theorizing in this field. In doing so, this exploration will keep the following questions in mind:

1. Is there any evidence of similar hierarchies currently existing in wider cognitive theory?
    a. If not, is there any evidence of data or findings that could be interpreted as operating over these hierarchies?
2. If so, what are the gains of postulating that these types of hierarchies exist?

Before turning to the dimensions discussed so far in this chapter in more detail, however, it will be useful to discuss previous accounts in SFL that have considered cognitive science, in order to locate this work alongside previous ideas.

3.8.1. Previous accounts of the relations between SFL and cognitive science

In the introduction to this thesis, it was claimed that one of the motivations for undertaking the work found in these pages was the lack of attention that SFL theorists and practitioners have given to humans as biological individuals, in comparison to the species in social terms Although I stand by this characterization of the current work being undertaken in this field, as well as the arguments that a general linguistic theory should account for both the 'inner' and 'outer' aspects of the 'languaging' human, there have been some previous work that has attempted to breach this gap. Before we come to the assessment of the potential use of a cross-disciplinary analogy between SFL and cognitive science, these proposals need to be described in greater detail.

Close to the perspective taken in this thesis is the work presented by Thibault (2004). In this work, the author "is concerned with the relations between the brain-body system, meaning-making activity, and the ecosocial environment in which the former are embedded" (*ibid*: 3). These are much the same considerations that occupy these pages, and just as in the perspective taken in this work, Thibault is also inspired by Halliday's views on the social-ecological nature of language. His main proposal in this work is to outline an approach to

cognition that incorporates the notions of *dynamic open systems*, *contextualization*, *embeddedness*, and *embodiment*. His main goal is to argue that by taking this approach it is possible to collapse many of the dichotomies that he argues are detrimental to the study of cognition. These include the distinctions between: "biology/ society, behaviour/cognition, brain/mind, communication/ cognition, form/meaning, mind/body, meaning/ matter, nature/nurture, organism/ environment, perception/ semiosis, signifier/ signified and structure/ function (*ibid*: 314). It is important to make clear, then, at this stage how the work presented below either differs from or builds upon the work above. This is important to maintain the novelty of the ideas presented in these pages.

In the first instance, this work can be seen to be a continuation of some aspects of Thibault's work, as well as the comments about the brain as a bio-semiotic system as made by Halliday and Matthiessen (1999: 607), in the sense that whereas previously the work has mainly concerned the link at a conceptual level, this work does this and further more concretely discusses the dimensions of *stratification*, *delicacy*, *rank* and *instantiation* in cognitive systems. Although Thibault does discuss the notion of *stratification* in terms of its conceptual semiotic roots (*ibid*: 64), and makes a claim that the expression plane of language can be further extended to contain the lower level of the neurophysiological capacities of the organism, his task is not directly to apply the ideas of SFL to create a semi-formalized conceptual framework of cognition incorporating many different types of *levels of analysis*. In this sense, the work considered below can be considered to broadly follow the overarching spirit of Thibault's work (but see below), but extends his work in terms of proposing a more practical application of SFL inspired-themes to cognitive science.

Another way in which this work can be considered to be a continuation of the above work is the range of cognitive and behavioural phenomena that Thibault discusses. Although he talks about the importance of sensory-motor activity in terms of meaning-making, much of his discussion considers that actions of humans in terms of articulatory movements (c.f. §§2.4 – 2.7). As has been mentioned, the aim here is more general and the application I intend to defend should hold for all cognitive systems and all organisms which are claimed to instantiate these systems. In this sense this work claims novelty in extending the general idea behind this previous work in this way.

There are, however, a number of differences in focus between his proposals and the approach taken here. Firstly, the major idea from SFL applied in this work, apart from stratification as indicated above, is that of *metafunction* (c.f. §3.1.2.). Many of his proposals include an ascription of proposed correlations between the metafunctional organization of language and certain neurological phenomena. For example, he discusses the case of mirror neurons, and argues that this has links to the interpersonal metafunction of language. Mirror neurons are those that fire both when an organisms performs a certain action, and when the organism perceives the same action being performed by another individual (Rizzolatti and Arbib, 1998). He claims that this is essential for "enacting interpersonal relationships between addresser and addressee" (Thibault, 2004: 268), not just giving a more 'ideational' representation of the external world.

Metafunctions do not feature heavily in the discussion below, since the main focus of this thesis is the different types of *levels of analysis* that are important in the study of cognitive systems. Although metafunctions are of paramount importance in SFL theory, they do not easily match on this idea as they are the three simultaneous types of meaning that themselves hold at each *level of interpretation*. They are not ordered in any way, as they are all operate at the same time (Matthiessen, 2007a) and therefore seem different to the other dimensions discussed in this chapter, where such implicit organization is required.

In sum, this work should be considered an extension of Thibault's work in the following ways: (1) the focus is not mainly on human cognitive systems, and the links between linguistic acts and neuro-cognitive facts are not our sole concern, and (2) whereas Thibault focusses mainly on the notions of *stratification* and *metafunction*, in this work we will discuss *instantiation, rank, delicacy* and also *stratification*.

With this justification of the novelty of this thesis given, this chapter now turns to the application of SFL heuristics and ideas to cognitive science.

3.8.2. Stratification

To reiterate, stratification describes language as operating over various different *levels of interpretation* simultaneously. This type of level has already been argued to be important in

various frameworks in cognitive science, but this section will investigate whether this proposal differs in detail from previous ones.

3.8.2.1. Towards a more general view on stratification

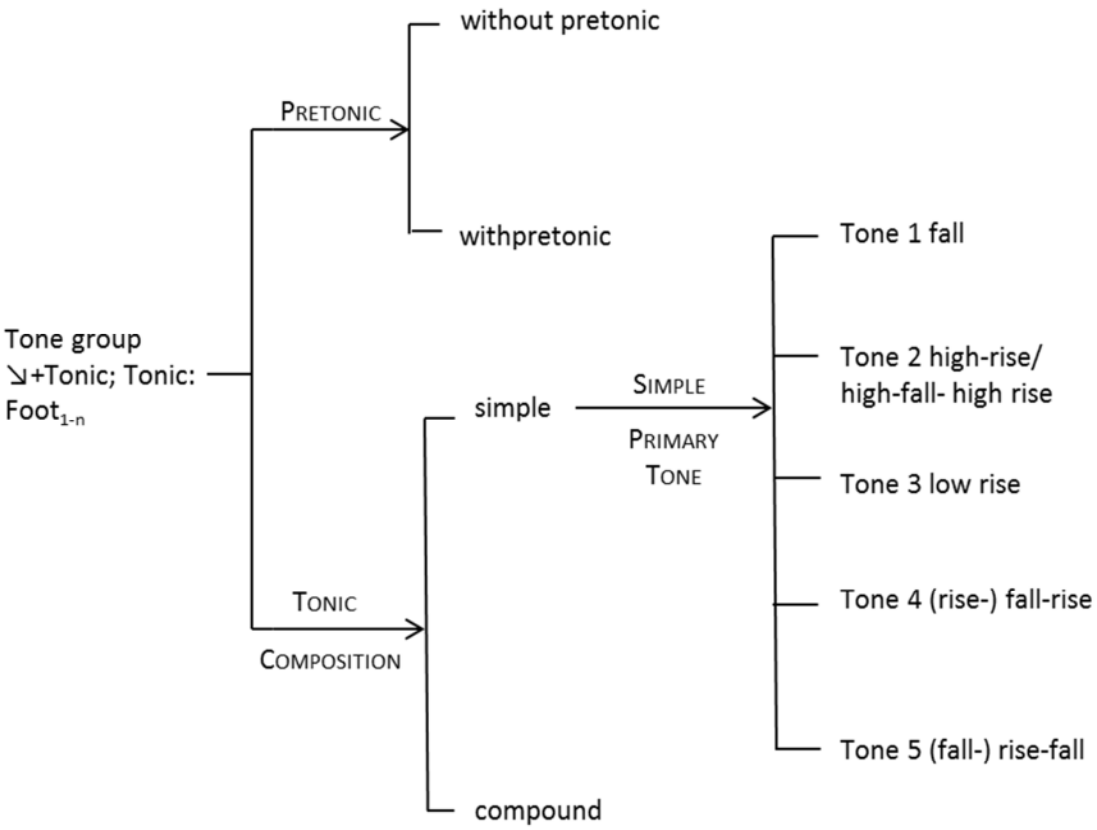
In §3.2 stratification was described as it has been developed in the context of linguistic theory. However, there are certain problems that present themselves when one comes to attempt to apply it to a different setting. Firstly, in the context of language, many of the linguistic categories relevant at the various levels of interpretation are already well established. For example, categories such as SUBJECT, OBJECT or FINITE have an established history in traditional grammars. Similarly for the structural units such as MORPHEME, GROUP or CLAUSE, although the precise nature of these units may differ according to a linguists theoretical perspective, what some of the units relevant to the study of language **actually are** can be considered to be more or less pre-given. Secondly, the different levels of interpretation themselves have a rich history. The idea that linguistics has the aim of linking sounds and meanings can be traced back at least to the Stoics early investigations on language (Halliday, 1977a. [2003]: 112). However, this raises an issue, since in this application it is less obvious whether there are similar well-established notions in cognitive science in general.

In order to overcome these problems, it is necessary to attempt to provide a more general idea about the nature of the stratificational framework. This should move beyond seeing the framework in its linguistic context, and instead bring out some of its more general properties. This section presents a novel interpretation of the nature of the stratificational framework, based upon my own interpretation of the theoretical discussion within SFL. This discussion will take the form of defending the following hypothesis, which I believe captures some of the essence of the idea of stratification:

> ***The Spatial-Temporal Tangibility Hypothesis of Stratification***
>
> As one moves from a lower to a higher level within the stratificational framework, one is investigating relations between units that form an increasingly spatially and temporally disbursed set.

Before fleshing out the above proposal in more depth, it is worth looking at some of the initial motivating factors for the hypothesis. Firstly, it provides a more explicitly stated means of differentiating which phenomena in language belong at each level of interpretation. Secondly, it captures the idea of stratification given by Halliday (1981/ 2005). In this piece, he discusses the similarities between a clause (a structural unit at the level of lexicogrammar) and a text (a structural unit at the level of semantics). He argues that although a clause is typically bigger than a text, this is not the whole story. It also involves considering the unit in **more abstract terms**. Although the notion of abstraction will be taken up in detail in §5, as a foreshadowing to that discussion, it is claimed that the above hypothesis is at the core of the term *abstraction* as it applies to the dimension of stratification. See §5.4.5 for a detailed discussion of this point. Furthermore, keeping in line with the rest of the discussion so far, this section will omit the level of interpretation of *phonetics* from our discussion (c.f. figure 3.1. above).

We will start with the phonological stratum. To begin, consider the below intonational system, which has been argued to apply at the phonological stratum since the beginnings of SFL (Halliday, 1963). The below network is simplified from Greaves (2007: 982):



Network 3.12: Intonation system in English

The question to consider in this case is how do you identify a particular tonal group as being, say a **falling tone**, or a **low rise tone**? We are dealing here with linguistic types, but the characterization of such types is based upon the phonetic aspects of the speech signal in terms of phonological contours (Greaves, 2005: 986). For example, consider the below example (*ibid*):

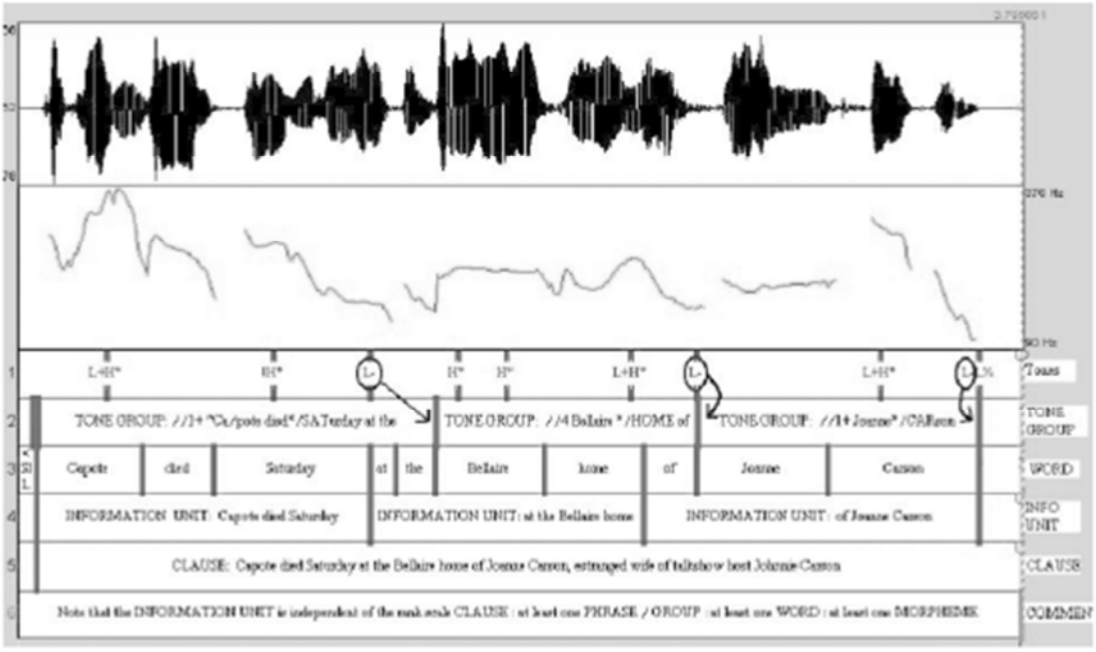
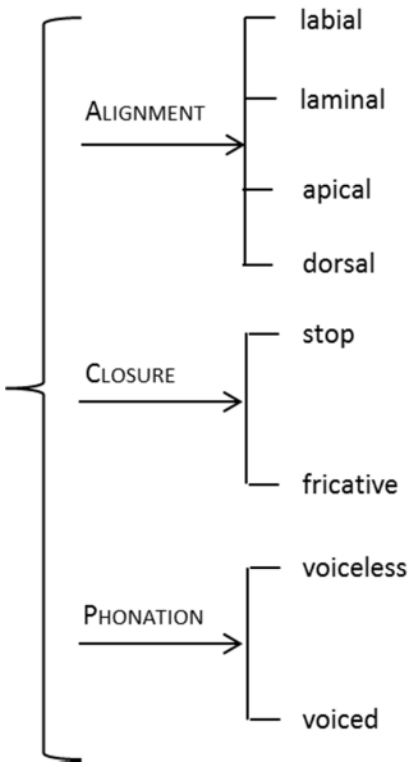


Figure 3.11: Tone Group analysis for "Capote died Saturday at the Bellaire home of Joanne Carson"

The different tones are given in the third row in the above figure. The second row shows the frequency of the speech signal, while the third row describes the tones as low (L) and high (H), with combinations representing rising and falling tones (c.f. Wells, 2006: 261-262). Taking the first example, on the third row we have the notation L+H*, referring to a rising tone (type 3). The means of identifying this unit as such can be found within the frequency trace itself.

A second feature can be noted here that is important for the purposes of this section. The grouping of signals that make up a particular tone in each case forms a **coherent, unbroken whole** that are more-or-less identifiable boundaries (although it would be mistaken to think that there are completely strict boundaries in all cases). This is a property that will differ when we come to look at the *lexicogrammatical* stratum shortly.

As another example, consider the following network that described the various features that make up the consonantal potential of language (Cléirigh, 1998: 2):



Network 3.13: Consonantal potential

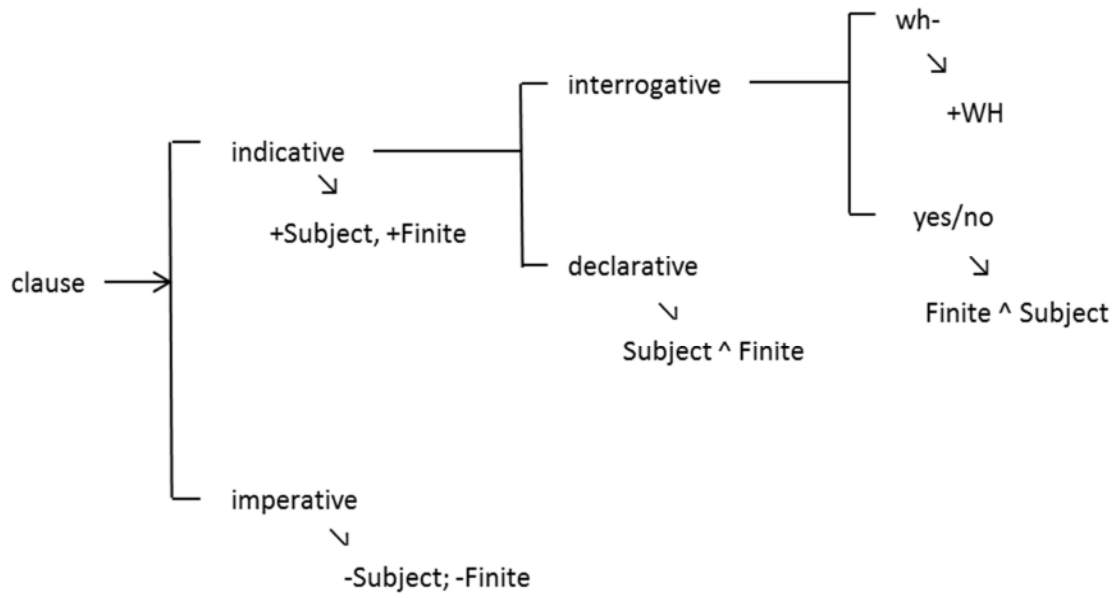
For example, on the above system network, the phonological consonant /p/ would be characterized as the choice of feature *labial*, *stop* and *voiceless* respectively on the above three simultaneous systems.

It is possible to proceed along the lines just attempted, by looking at the nature of the features given above and the relations that hold between units in terms of phonological features. The first question that should be asked is: on what criteria does one determine which of the options above a certain consonantal phoneme is represented by? For instance, above it was noted that the phoneme /p/ could be characterized in terms of the system features *labial*, *stop* and *voiceless*. These features, I would argue, are intrinsic to the nature of the linguistic unit itself: the very reason why we consider a particular phoneme to be, for instance, a devoiced bilabial plosive is due to features that are intrinsic to the phoneme itself.

Having presented some system networks that operate at the lower linguistic stratum of *phonology*, we now consider the stratum of *lexicogrammar*. By interrogating these systems as we did above, and by comparing the discussion with that for the *phonological* systems

just presented, it will be possible to determine whether there is any evidence to support the above hypothesis.

A number of lexicogrammatical systems have already been presented in this chapter. Below is the lexicogrammatical system of MOOD, repeated from §3.2 above (Network 3.2.) along with the relevant realization statements:



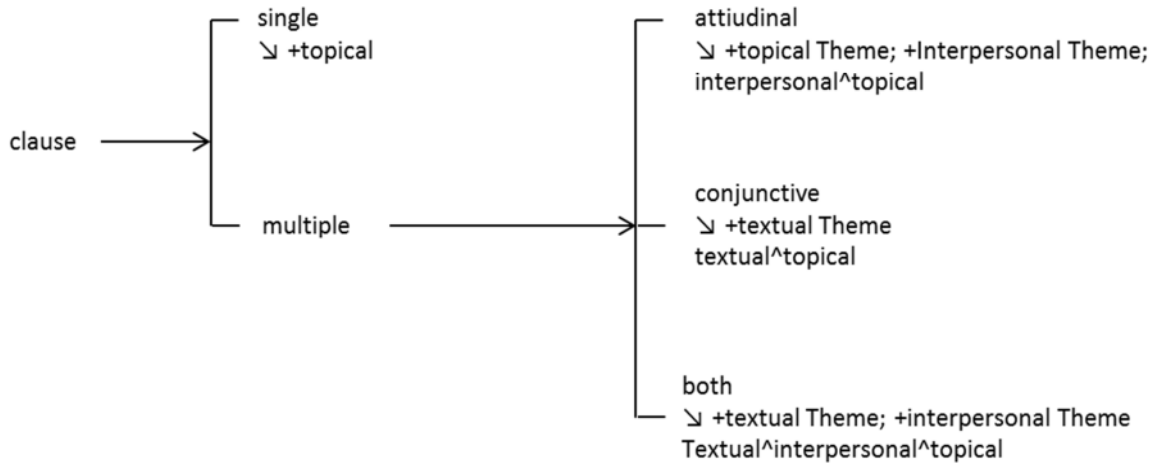
Network 3.14: Lexicogrammatical system of MOOD

As previously, this discussion begins by considering the means by which we come to know how a certain clause should be described according to the system of MOOD. As was noted in the discussion in §3.2, the main means of doing so is by the relative ordering of grammatical categories such as *Subject* or *Finite*. Our question, then, is as follows: how is it we know the status of *these* elements? Is it by factors internal to the linguistic units themselves, as was argued for in the phonological examples given above? Or is it by some other means?

The choice to base part of this discussion on the MOOD system was not completely random, as Halliday does provide a discussion on the nature of both the *Subject* and *Finite* elements in the literature (Halliday, 1982/ 2002; Halliday and Matthiessen, 2014: §§2.6, 4.2.2.). A discussion of the ideas found in these works will therefore be presented. As noted in the above works, giving a precise notion of what exactly a *Subject* **is** can prove difficult, as it can be associated with three different notions: the *grammatical subject*, which is the element of

the predication and is the concord of verbal inflectional morphology, the *logical subject*: the 'doer' of the action, and the *psychological subject*, roughly equivalent to that element that the clause is about. Important for our purpose is that in every case the definitions are not in terms of any intrinsic fact about the words themselves, but instead in terms of some relation between units. In the first case, this is a grammatical relation between the *predicator* and the *subject*, in the second case as a relation between things, and in the third case, a relation between a word and some idea in the mind of the speaker (Halliday and Matthiessen, 2014: 79).

As a further example, presented below is the system of THEME associated with the textual metafunction (c.f. §3.1.2 above). The below network diagram is adapted from Eggins (2004: 299):



Network 3.15: System of THEME

Regarding the status of these categories, a similar situation to that just argued for can be seen. As discussed in Halliday (1984) and Halliday and Matthiessen (2014: 88) the theme is the part of the clause that is given special status in terms of the clause as a message; it is "the element that serves as the point of departure of the message" (Halliday and Matthiessen, 2014: 89). For example, below is repeated the previous example given in §3.1:

The Duke	gave the teapot away
Theme	Rheme

The question now remains: is there something intrinsic to the nature of the words <the Duke> in the above example that guarantees its status as topical theme in this case? The answer is clearly in the negative, as the same words can appear in a Rheme situation, as in:

The jury	indicted the Duke
Theme	Rheme

In this case, then, we have the same lexical items, but their status within the thematic system has changed, they are now part of the Rheme element. This pattern holds for all cases of this *lexicogrammatical system*. The general picture is that, again, these grammatical notions are not defined by any intrinsic value inherent to the linguistic units themselves, but instead they are defined in terms of relations between the relevant linguistic units and others in the same linguistic context, in this case the clause. However, in these examples at least, the relations that are used to define such grammatical notions are bound within the nature of the clause itself. This fact will be important when we come next to look at the highest linguistic stratum of *semantics*.

Finally, it was seen previously that prosodic units form a coherent, unbroken whole. This is not always the case at the *lexicogrammatical* stratum. In §3.1, we saw the ideational description of the clause *the duke gave the teapot away.* This is repeated below:

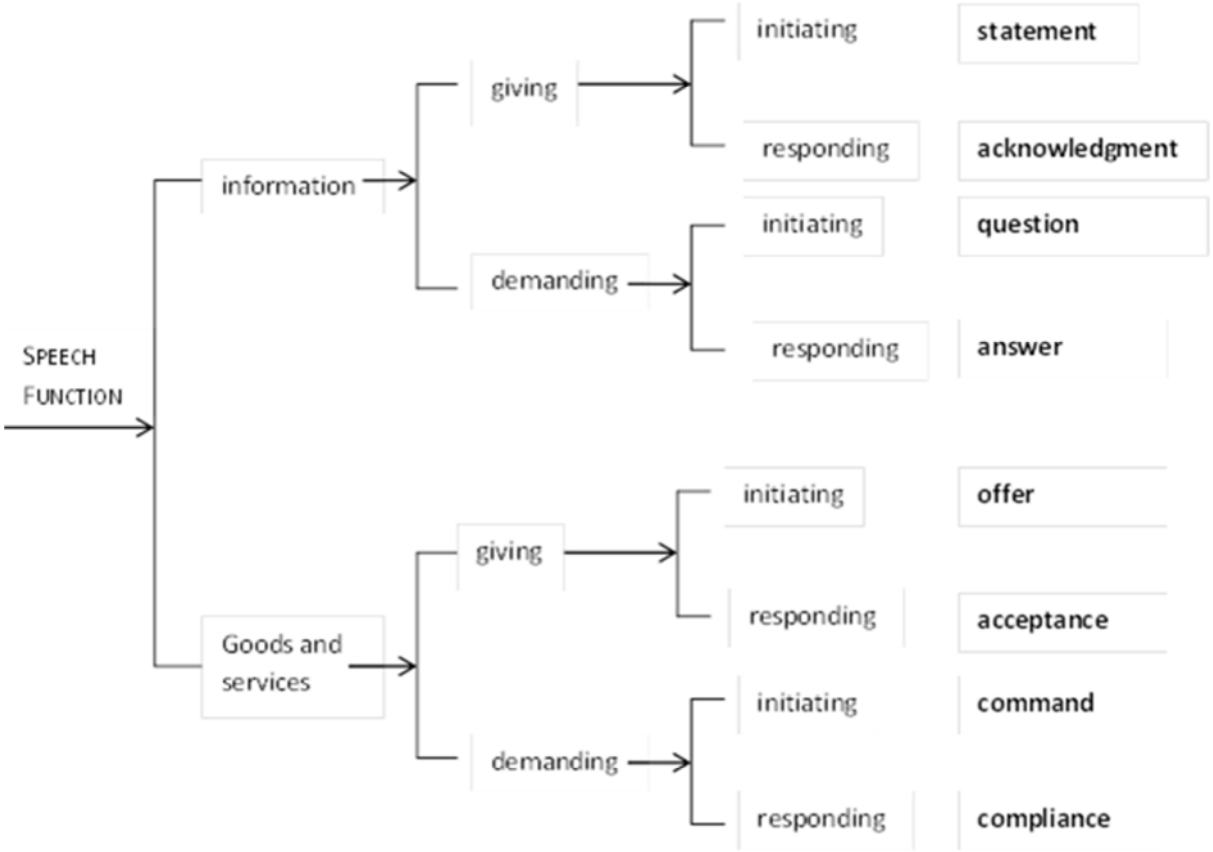
The Duke	gave	the teapot	away
Participant: Actor	Process: material	Participant: goal	Process: material

In this case, the feature of *material process* is distributed in two places along the clause. In such a way, some lexicogrammatical units can be dispersed throughout the relevant structural unit, in a way that did not seem possible at the lower stratum.

The final *level of interpretation* that needs to be discussed it the stratum of (discourse) semantics. To frame the discussion, given below are two system networks that have been claimed to operate at that point on the stratificational framework. The clearest way to note the difference between these two strata is by comparing, again, the system networks of SPEECH FUNCTION at the stratum of *semantics* and the network of MOOD at the

lexicogrammatical stratum. This comparison will prove useful, since these two have been linked explicitly in terms of the realization relationship, and their placement on adjacent strata allows a useful comparison to be made.

The system network for MOOD has just been given in network 3.13, the system network for SPEECH FUNCTION is repeated below:



Network 3.16: System network of SPEECH FUNCTION

How do we know what particular speech function a certain clause is construing. The answer cannot be in the text itself, for in this case we are not dealing with grammatical relations that exist between linguistic units, but instead **those aspects of the social setting of language that the text is realizing**. We are not talking about features such as *Subject* or *Finite*, but instead in terms of the kinds of activity that are taking place. To reiterate, as was shown in §3.2, a speech function is not identical to the grammatical mood of the clause, since the same speech function has *unmarked* realizations in terms of the types of mood that it can be linked to. Neither, however, is it wholly outside the text itself, since it is still a property of this linguistic unit. Semantics, then, seems to be at the interface between the

linguistic system and the context, as indicated in table 3.1 above. It has this increased spatial scope, by being a feature of the text that construes those aspects that are outside it in the wider social setting.

The above examples have all been presented to motivate the hypothesis given at the beginning of this section, namely that we can consider the move upwards along the stratificational framework in terms of increasing our viewing of language in terms of increasing spatial and temporal tangibility. We can sum up the proposal diagrammatically as follows:

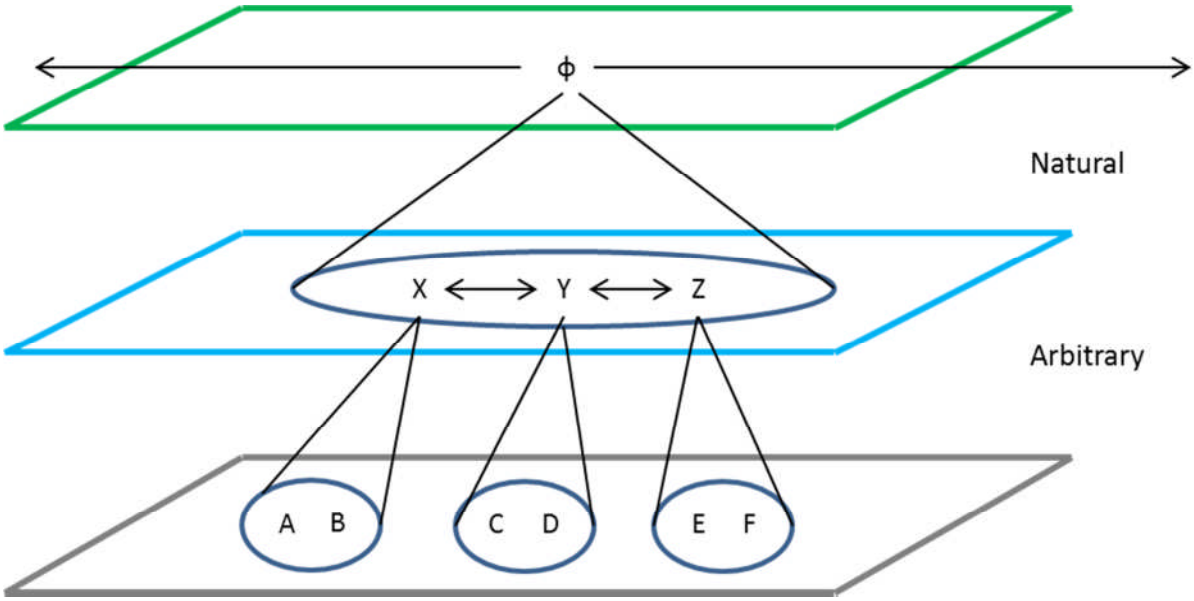


Figure 3.11: Spatial-Temporal Hypothesis of Stratification

This general idea presented by this figure is as follows: lower points in a stratification framework deal with units and features that have a restricted spatial and temporal scope, and form coherent units. As one moves higher, these relations become more far-reaching. In the linguistic example, *lexicogrammatical units* have no inherent value but are defined in terms of relations within the clause or clause complex. Semantics takes this one stage further, by acting as the interface of the linguistic system and the wider social environment.

3.8.2.2. Towards Identifying the Strata

The question, then, is as follows: can the more general proposal for the idea of stratification given above provide the means of hypothesizing about the nature of the "higher-level"

strata in a cognitive framework? It is the proposal in this section that the idea of *stratification* as applied to cognitive science more generally involves identifying units and patterns that increase in their spatial and temporal extent, in a similar way that was just argued for in the previous section in the case of linguistics.

Instead of three-levels as has been stated in SFL, this thesis will restrict itself to a two level framework (with the addition of the highest level of *semiotic environment*, to be discussed and defined in the next chapter). These two *levels of interpretation* will be stated as the **neurophysiological stratum** and the **psychological stratum**. They will both be discussed in turn:

3.8.2.2.1. Neurophysiological stratum

The lowest *level of interpretation* on this framework consist of interpreting the cognitive system in terms of *neuronal events*, that is in terms of the neurological signals that are found in neuronal circuits and are the object of study in neuroscience and, in a certain sense, cognitive neuroscience. This is the "brain" approach, but crucially this is not the only perspective that makes up the framework. In this regard, it shares properties with Marr's lowest *level of interpretation*.

What exactly is meant by neurophysiological signals, though, and why are they thought to be important for the study of cognitive systems? Neuronal systems consist of neurons, which have identifiable parts. The *dendrite* is the part of the neuron that receives incoming electrical signals. The *soma* is the cell body, and the *axon* passes on electrical signals to further neurons in a circuit. Where the *soma* and *axon* meet is known as the *axon hillock* and the end of an *axon* is the *terminal bouton*. Within neuronal systems, the most important signal is known as the Action Potential (AP). This electrical signal is thought to be the basis of information processing and biological cognition in organisms (Rieke et al., 1997; Stein and Stoodly, 2006). An AP is created by the electrical properties of a neuron rapidly rises and falls during an excitatory event. Specifically, at any given time a neuron will have a particular value for its *membrane potential*: the electrical potential across the cell membrane. This membrane potential is caused by particular concentrations of electrically charged *ions* being found both inside and outside the cell. The important types of *ions* in this

case are Potassium (K^+), Sodium (Na^+) and Calcium (Cl^-). The distribution of these ions result in a typical neuron having a membrane potential of around -65mV at rest.

During stimulation, either from a source outside the nervous system, or from another neuron, channels will open on the dendritic spines of the neuron. These channels are selective for Sodium ions. This will cause an influx of Na^+ ions **into the cell**. Since these ions are positively charged, the membrane potential of the cell will **increase**. Given enough stimulation, the membrane potential of a neuron will reach a threshold level (typically around -50mV), after which an Action Potential will result, due to a further increase of sodium ions. The AP is formed at the *axon hillock*, that is the area where the soma meets the axon. After around 1ms, channels on the soma will open, that allow Potassium ions to **exit the cell**, since there is a greater concentration of them internally then externally. Since these are also positively charged, this movement will cause the membrane potential to **fall.** This gives the AP its characteristic "spike" shape, as shown below (Rieke et al. 1997: 3):

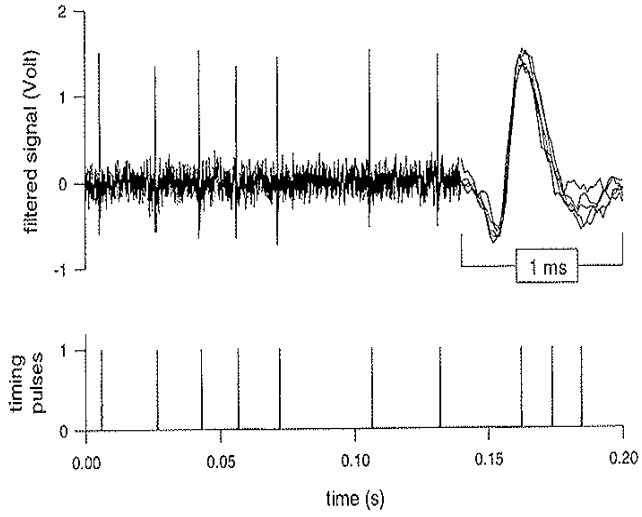


Figure 3.12: the Action Potential

APs are passed between neurons at *synapses*. This is where the axon of a *pre-synaptic neuron* comes into close proximity with the *dendrites* of a *post-synaptic neuron*. When the AP travels down the *axon* and reaches the *terminal bouton*, it causes the release of *neurotransmitters* into the *synaptic cleft*: the gap between the two neurons. These *neurotransmitters* bind to receptors on the *post-synaptic neuron*, causing sodium channels to open, and beginning the process described above again.

APs are considered to be important for the processing of sensory information, and nearly all other cognitive processes. Most work in cognitive neuroscience is involved in finding the neurological correlates of certain cognitive processes. This may be through specific recordings from a small number of neurons, or larger activity as evidenced through neuroimaging techniques such as functional magnetic resonance imaging (fMRI), electroencephalography (EEG), or magnetoencephalography (MEG). All these techniques are, indirectly or directly, measuring the electrical activity of neuronal system, typified through the Action Potential.

It should again be highlighted, that on the view presented here we are dealing with a **potential**, i.e. a **system of neurological processes**. Construed as such, the lowest level of a stratificational cognitive framework can be considered initially to be the **set of possible neuro-physiological states** that a certain organism can be in. The nature of this potential will, of course, depend upon the physiological facts of the organism in question: the number of neurons, how they are arranged either in terms of circuits or more macroscopic regions such as brain lobes. In §3.8.3.1 below an example of a system network in terms of *neural oscillations* will be given, to give an indication of the ways in the beginnings of such a potential may look. The general idea, though, is for any given system, there will be a **resource** for neurophysiological behaviour, characterized by the possible **neurophysiological states** the individual can be in.

3.8.2.2.2. Psychological Stratum

The second, higher, *level of interpretation* on this framework involves interpreting the cognitive system not in terms of physiological signals being transmitted within a neuronal network, but instead in terms of interfaces between the organism and the external environment. In the first instance, this will involve observable actions, such as moving, eating, speaking and communicating, etc.

However, also on this view I mean to include certain cognitive capabilities that are generally considered to be 'internal', that is unobservable. In this I intend to refer to processes such as *emotion*, *belief*, *thoughts*, etc. To recall the framework provided Morton and Frith (1995), this proposal contained two levels labelled *cognitive*, for those that are not directly observable, and *behavioural* for those that involves externalized actions. On the framework

I am putting forward here, this distinction no longer applies. On the face of it, this would seem like a major omission in the proposal, but I believe that when one comes properly to consider cognition in an *embedded*, *enacted* way, as this framework is intended, then this conceptualization is, I think, unavoidable. Take an emotion like *fear*. This emotion has a rich history, that is, it extends in time beyond the present. My fear of snakes, for instance, has a rich history, and cannot be solely located in the present time, although I may currently be feeling it. Its nature has a history stretching back to my past interactions and experiences; in this sense it is neither internal to me, nor solely located in the present. One downside to the *computational frameworks* seen earlier in this thesis is that they bias us to consider such facts as wholly internal to us, omitting the role that past action and experience has had in shaping the nature of such psychological states.

It is also important at this juncture to highlight the *embodied* nature of this *level of interpretation*. This is crucial for linking this proposal to previous ideas seen in §2.5. The psychological stratum, as I intend to use the term here, refers to the **set of possible cognitive acts** an individual can exhibit. This includes both 'internalized' and 'externalized' acts. It would be an omission is such a characterization was not heavily influenced by the nature of the organism's general physiology. That is, although the notion of the *body* does not have an explicit place on this framework, it permeates through this entire account. This will range from seemingly trivial cases such as the presence of wings in order to fly, or the possession of a gall bladder to achieve buoyancy control underwater, to more subtle cases such as features of the sensory organs enabling a different range of external stimuli to be potentially meaningful for the organism (c.f. §4).

I will concede that the above characterization in this section is brief and remains highly preliminary. In §3.8.3.2 some reasons for this will be given, and they will be repeated in the conclusions under avenues for future work. However, in the section just referenced, an example of *memory* will be stated at various degrees of *delicacy*. The general idea, however, is that for any given individual there will be a system of **potential cognitive acts** that it can perform. Considering this as a system enables the researcher to ask why a certain act was performed instead of another, thereby maintaining the paradigmatic picture that SFL provides us with.

3.8.2.3. Realization: comparison with previous uses

In any stratificational framework adjacent *levels of* interpretation form a *realization* relationship. As seen in §2.3.3, the term *realization* is already part of the terminological currency in cognitive science. It has been used to describe the relationship between higher and lower levels in hierarchies containing different levels of analysis. The most obvious example of this was the idea of *multiple realization* as found in the work of Marr and Fodor. As a reminder, this was the idea that one feature at a higher level could be linked to a large number, on some accounts potentially infinite, number of lower level instances. It was taken by some to motivate the position that cognitive theory could relegate the study of neuroscience to the periphery, with the main focus being on the subjects of psychology, artificial intelligence and linguistics. It was also highlighted in the same section the problems that have been raised against this view, that this thesis agrees with (c.f. §2.3.3. for all of these points).

It is important to make clear that, although the same term in used in SFL theory, the nature of the relation in both cases is markedly different, and this affects how one conceives of the relationship between different levels in conceptual frameworks. The first major difference between the terms as used in the different frameworks is that **while realization is the classical sense was strictly a top-down relationship, within SFL realization is bi-directional.** In §3.2 this was discussed in terms of lower levels construing information in higher ones, and alternatively higher-level features activating aspects of the stratum below. For our interpretative purposes, this creates the situation that a description of any cognitive fact must take into account the description in all levels. This multi-perspective evaluation of terms in the framework is known as the 'trinocular' perspective in SFL, whereby the meaning of any term can only be properly understood in terms of the features it construes (from above), the features it activates (from below) and its relationship with other features at the same level (from around). Halliday and Matthiessen (1999: 504) discuss this idea.

This difference has a major consequence on how cognitive systems are conceived between the approaches. To show this, it is necessary to return to the discussion in §2.3.3. In this section, doubts were raised concerning the trend for certain cognitive scientists to intuitively argue that the same higher level phenomenon could be realized in a multitude of

different lower level systems. This perspective only comes if the features of the lower level are ignored, as was noted. On the idea of **bi-directional realization** this intuitive argument no longer holds: the higher level psychological system can **only be** considered in terms of the lower level system that realizes it. The linguistic analogy will prove useful here. SFL, unlike formal semantic theories, does not assume that the same meanings are expressed in every language. That is, they are more cautious in ascribing to the view that grammar is simply a passive realizer of a universal set of meanings. Instead, the bi-directional nature of *realization* means that lexicogrammar has an important role to play in the nature and types of meanings that a particular linguistic variety can express. For instance, see Halliday and Matthiessen (1999) for a discussion of Chinese, Teich (2003) for German and the papers found in Caffrel et al. (2004) for a discussion of the interpersonal metafunction in a variety of languages. As far as cognitive science is concerned, this is an important point. On this view, you cannot divorce the study of cognition from the nature of the lower-level features that are in the same stratificational framework. In most cases, this means cognition can only be understood in terms of the brain.

Having noted this distinction, the properties of the *realization* relationship itself can be stated. In SFL, it is an intrinsically **probabilistic** relation. This is necessary since SFL embraces the 'fuzzy' nature of language (Halliday, 1991., 1992b., 1993). On this picture, each member of the psychological system, each type of act, will have an *unmarked* realizer. In the context of the memory system outlined in §3.8.3.2.1, the prediction is that each feature will correlate strongly with a certain area within the overall neurophysiological potential. One use of such a framework is that each member of the higher-level system must have a different unmarked realizer at the system below. However, being a probabilistic notion, this view permits variations both within the same individual, and in comparison with others.

3.8.3. Delicacy

Earlier in this chapter, the scale of delicacy was presented as an ordered series of different *levels of specificity*. Within SFL, its purpose is to enable researchers to divide up the meaning-making potential of language in more or less general ways. It is also important to reiterate the links between delicacy and the scale of rank. The relation between these two is that the scale at rank determines at what structural unit the specified meaningful

oppositions hold at. Both of these above points will be important in the following discussions.

3.8.3.1. Delicacy at lower strata

In §3.8.2, it was argued that the key behind the stratificational framework is that a move towards a higher stratum indicated that the observer is attending to units and relations that exhibit a wider scope in space and time. It was also argued in that section that at the lowest level of the stratificational framework of cognition being developed here, the lowest stratum referred to signals and relations that held within the neuronal system itself. In broad terms, this refers to what have previously been called *brain-states* (c.f. §2.2.1.), but it can also be characterized as patterned neuronal activity. It should be remembered, however, that SFL considers each stratum within a stratificational framework to be a potential. It is worth providing the details of this as they apply to, what will be called, the *neurophysiological level interpretation*.

One issue for the current proposal is that the set of possible neural activity that any organism can exhibit is extremely large. In the previous section, the *neurophysiological system* was described in terms of the firing of individual neurons. Each neuron in a system can be considered to be either firing or not at a given moment in time, and the combined possibilities of firings provide you with the overall system potential. However, this would be describing such potential at an extremely specific point in delicacy. In the rest of this section a less delicate idea of neuronal potential will be considered, to give an idea of how delicacy operates at this point in the framework.

Within conscious, behaving organisms it is difficult to acquire data concerning the firing patterns from many single neurons simultaneously, and as just mentioned many of the ideas concerning the encoding of information in such cases is based upon cellular recordings of a few at a time. However, there are a number of studies that have looked at neuronal signalling at a higher structural unit. This is the idea of *neural oscillations* that operate at the higher structural units within neuronal systems, namely neural ensembles and larger brain regions (Ward, 2003; Buzsáki and Draguhn, 2004, c.f. §3.8.4.). Important for our purposes here is that this kind of synchronized activity between neurons and groups of neurons has been linked to a variety of cognition functions, indicating that they are likely to be important

for the semiotic activity of neuronal systems. Some of these claimed links will be discussed below.

Neural oscillations can be described as rhythmic or repetitive neuronal activity. Outside of a neural context, one of the simplest examples of an oscillation is a swinging pendulum, which repetitively moves from one point to another. Such repetitive activity can be described mathematically as different types of waves: usually sine or cosine waves. The below figure provides an example of a neural oscillation taken from both an electroencephalogram and from a computational neural simulation (Ward, 2003: 554):

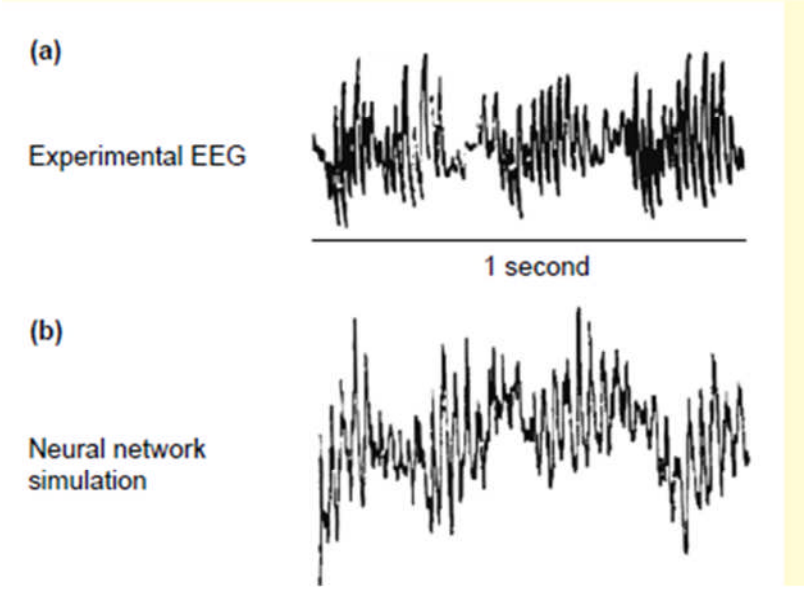


Figure 3.14: Oscillations in neural systems

Neural oscillations can be defined in two ways, firstly by their frequency, and secondly by their amplitude. In terms of frequency, the standard procedure is to group similar frequencies into bands denoted by the terms *alpha*, *beta*, *gamma*, *theta* and *delta* frequencies. The below figure diagrams these different frequencies, and others, along with the numerical values that are included in each band (adapted from Buzsáki and Draguhn, 2004: 1926):

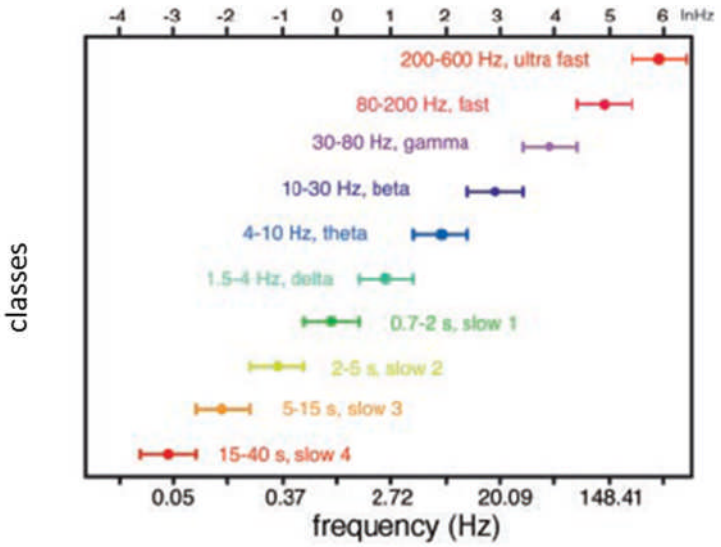


Figure 3.14: frequency classes for neural oscillations

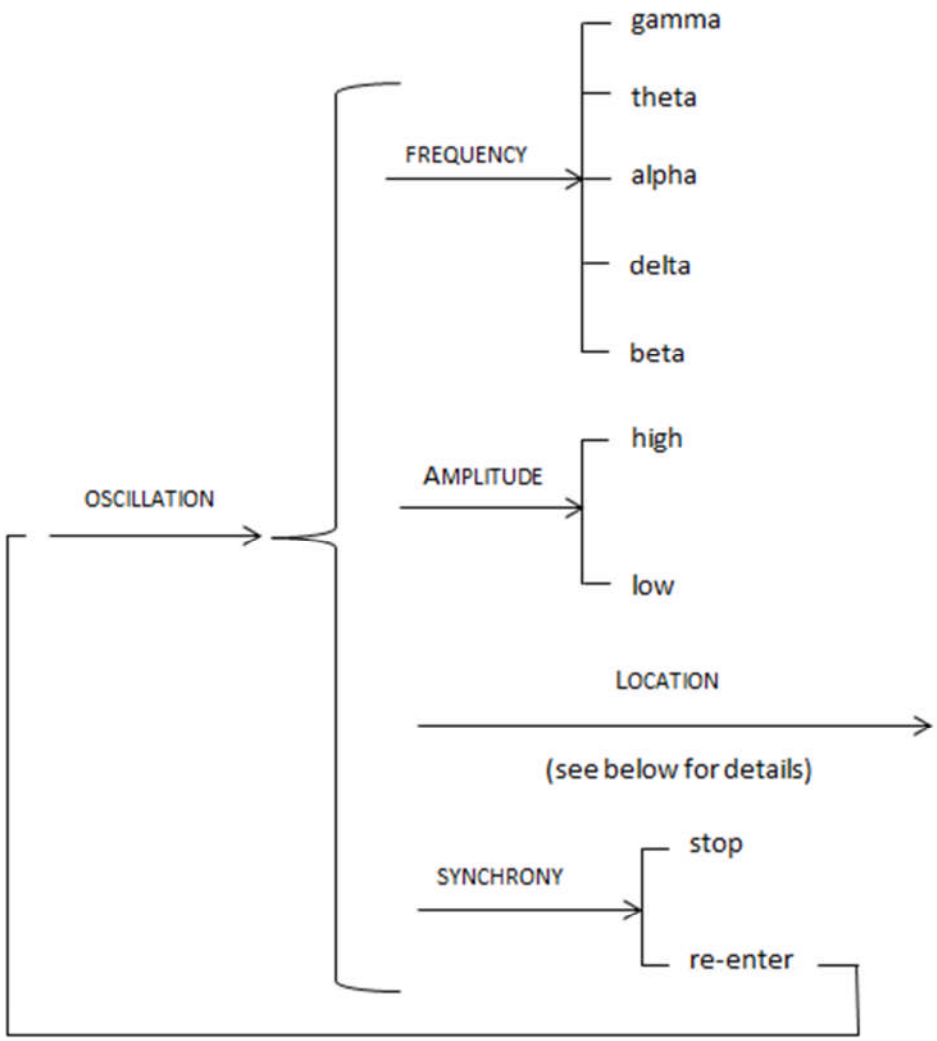
Different cognitive tasks have been associated with various types of neuronal oscillation, especially memory (Hsieh and Ranganath, 2014; Kahana et al., 2001) and attention studies (e.g. Cooper et al., 2003), and more recently to issues such as sentence comprehension (Meyer et al,. 2015) and problem solving (Cao et al., 2015). In these cases, certain types of signalling at this higher neuronal rank contribute to the meaning-making potential of cognition.

In terms of accounts of memory, neural oscillations in the *theta* category (4-10Hz) have been associated with spatial memory tasks in rodents, as shown with experiments using maze protocols. Specifically, it has been shown that theta oscillations are observed in the hippocampal place cells of rodents during exploration of the maze. These particular cells in the hippocampus have been argued to be important in the forming of spatial memories (O'Keefe and Reece, 1993; Skaggs et al., 1996). Within humans, difficulties in the measurement of such signals can be overcome by the use of intracranial electroencephalogram (iEEG), which allows the recording of electrical activity within the skull. This allows for a more accurate recording than can be gained from extracranial EEGs, and can by-pass artefacts created by muscle movements, for example eye movements. In the experiment outlined by Kahana et al. (1999), such methodology was used to investigate the neural correlations of a spatial memory task in a virtual maze environment on epileptic patients. Three participants took part in the study, and had 171 iEEG electrodes inserted into a variety of places in the cortex. For the spatial memory task, the participants were

asked to navigate their way through a computer-generated maze, using keyboard inputs to navigate their way to a pre-set goal area. For each case, the participants were given a number of learning trials, where signposting was presented towards the goal area. In the test condition, this signposting was removed, forcing the subjects to recall the route learned during the training trials. For both the training and test trials, the researchers observed theta oscillations in areas of the frontal and temporal lobes (see *ibid*: 783, fg. 3). However, it was observed that the strength of the theta oscillation was higher in test trials than in training trials.

To reiterate, the exact proposals put forward in these studies is not crucial for the view put forward here. They are only included as an indication of some of the cognitive roles that neural oscillations have been claimed to have in the previous literature. What is significant, however, is that part of the potential states that would make up the neurological potential for meaning making would be the presence of oscillating electrical behaviour at a number of different frequencies, and at a variety of different locations in the brain. Furthermore, the specific frequency of these oscillations seems to play a role in certain cases, indicating that this fact is an important one in understanding the lower-level realization of cognitive behaviour.

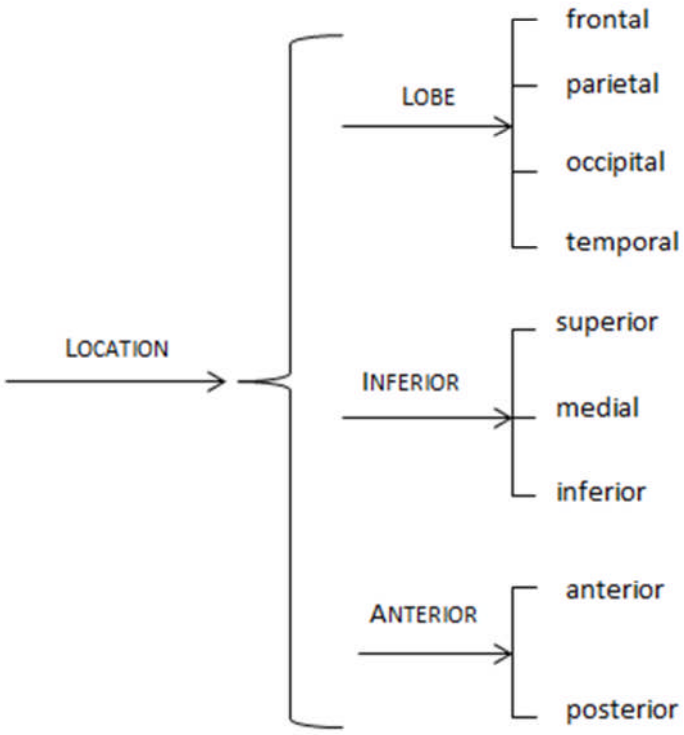
Summing up the previous experiments, which it should be noted only constitute a small number of the work undertaken in this field (c.f. Hsieh and Ranganath, 2014 for a recent review), the following factors of neural oscillations have been argued to be important in cognitive systems: firstly, the frequency of the oscillation, as grouped in various classes, secondly, the area in which the oscillating neural structure is found, and thirdly, the strength of the oscillation, with differing factors leading to higher amplitude oscillations in certain cases. It is also necessary to include the possibility that synchronous oscillations over different areas are involved in a certain cognitive phenomenon. It is possible to re-interpret these facts along systemic lines, by outlining the certain 'options' that are evident in cognitive systems:



Network 3.17: Oscillation network

The above network has been conceived of as containing four simultaneous systems. The first of these, FREQUENCY, captures the different groups of frequency that are usually the basis of discussion in the literature. The second, AMPLITUDE, has been given the values *high* and *low* to appreciate that sometimes this difference can have cognitive effects. This system could be made more specific as more information about its importance becomes known. The third, LOCATION, specifies the area where the particular neural structure is located, and for reasons of space this subsystem is given below in network 3.18. It is a simultaneous system containing three sub-systems. The first identifies the unit as located in either the *frontal*, *parietal*, *temporal* or *occipital* lobe. The second refers to its vertical orientation, in

terms of whether it is located in the *superior*, *medial* or *inferior* areas of these lobes. The final system refers to whether it can be located in the *anterior* or *posterior* regions. The subsystem of LOCATION is given below:



Network 3.18: Subsystem of LOCATION

The final system, SYNCHRONY, allows for the description to capture cases where more than one neural unit may be involved in a particular cognitive task. The crucial option here is *re-enter*, which allows the observer to also describe that an additional oscillating signal may be simultaneously involved.

The above elucidation provides just one example of how delicacy may operate at this stratum. I believe this is enough to illustrate the general point being made, but is clearly not the whole story.

3.8.3.2. Delicacy at the higher stratum

Just as the level of neurological signals can be thought of as a set of possible options, representing the resources an individual organism has for meaningful cognitive behaviour, the higher stratum can be thought to contain the potential an organism has for cognitive behaviour. That is, in certain situations, an organism exhibits some form of such behaviour out of a range of possibilities. This focus on **potentials** is central to the ideas being developed here.

When we come to look at delicacy, however, there is a problem as far as this work is conceived. Whereas the nature of the *neurophysiological* stratum is relatively uncontroversial, a similar situation does not hold at the higher level. For instance, we can think of different types of action: moving, eating, sleeping, thinking, believing, etc. One may think that this would form a good basis for describing the different types of psychological act that would make up the details of such as system. This remains a possibility. However, doubts have been raised concerning our use of 'everyday' notions of psychological predicates as the basis of such a theory. This is the debate surrounding *folk psychology* (Baker, 1999, Roth, 2013). Its supporters point to the fact that since we do linguistically carve up the range of actions in a certain way indicates that there is likely some important difference between them (Fodor, 1987), and that this is as good a candidate as any for talking about psychological notions. Its detractors, on the other hand, argue that it has failed as a useful theoretical tool, and instead argue that a new 'vocabulary' is needed in this context (Churchland, 1981).

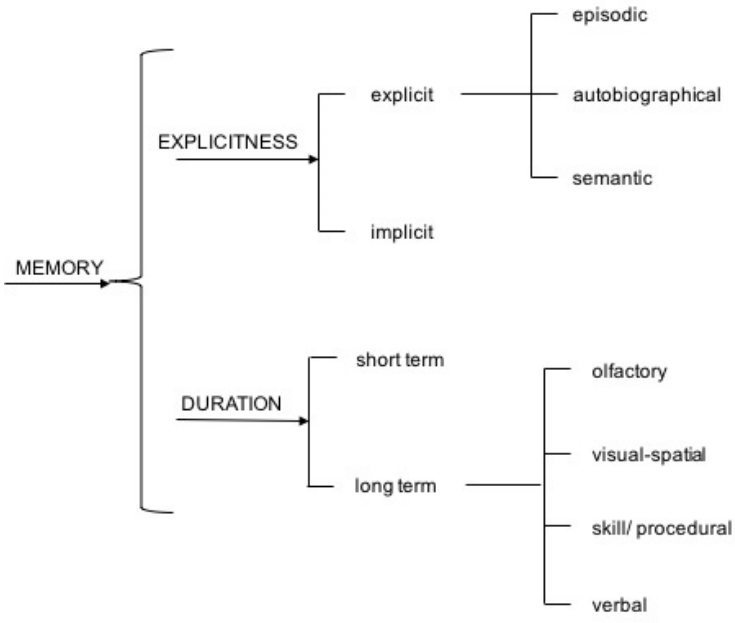
This thesis remains neutral on this debate. In this section, however, the example of *memory types* will be used to illustrate how some of the ideas here may be of use in resolving this issue.

3.8.3.2.1. An example: memory

In order to show how some heuristics from SFL can be usefully applied to this ongoing discussion within cognitive theory, I will now give an account of neuroscientific accounts of memory. One reason for choosing this aspect of cognition is that a range of different "types" of memory have been presented in the literature. Some such types include

emotional memory, implicit or non-declarative memory, and explicit or declarative memory, amongst others that will be discussed below (c.f. Nalbantian, 2011). Perhaps the record for the number of different types of memory is the 256 outlined by Turving (2007)! Also, the subject of memory has found interest from the various sub-disciplines that make up cognitive science, and as such the subject has been studied from both a neuroscientific and psychological perspective. In the below discussion, the different types of memory as outlined by Roediger et al. (2008) will be used as a basis.

The detailed typology of memory given by the authors just cited can be outlined as follows. One of the first broad distinctions is between *explicit* and *implicit* memory (Graf and Schacter, 1985). The first of these is elicited when subjects are explicitly asked to remember a certain fact, whereas the latter is used in cases when the memory is retrieved in more indirect ways. Different **types** of explicit memory have been identified. These include *episodic memory*: memories concerning particular events in time and space, *autobiographical* memory: memories concerning one's own life and *semantic memory*: general knowledge of the world. (Roediger et al. *op cit.:* 20). A frequent distinction is also made between *short* and *long term* memory: defined in terms of how long the memory persists. The authors also distinguish between different **types** of long term memory: *olfactory* memory; memory of odours, *visual-spatial memory*: memories of places and spatial relationships, *procedural memory*: memory of how to perform certain skills, and *verbal memory*: remembering events in terms of stored language (*ibid*). These are only some of the vast number of memory types that have been proposed. Their classification of different types, however, enables a delicacy analysis, as follows:



Network 3.19: System of MEMORY types

In terms of the concern surrounding *folk psychology*, systematizing our everyday notions along the lines above allows certain predictions to be made. Each one of these system features should have an identifiable *unmarked* neurophysiological realization at the lower stratum. Although, as it is inherently probabilistic, it should be emphasized that some cross-over is permitted, but only in *marked* cases. If this is found not to be the case, then there is evidence for abandoning our everyday notions of psychological terms. This would hold for, not just memory, but a range of psychological phenomena.

Delicacy in both strata involves asking: "what kinds of [neurological/ psychological] events are there?" In this way, it divides up the overall potential for meaningful cognitive or neurological activity into more specific sub-types, in the same way that different types of clauses or meanings are identified in SFL. At this stage of the development of this work, only a few examples can be given, but hopefully these are enough to elucidate the general idea.

3.8.4. Rank

If the concept of rank is to have any use in cognitive science, then there should be evidence of a hierarchy that operates over a part-whole dimension. Moreover, the framework presented here postulates that there should be (at least) one such scale for each new level of interpretation. This section will outline the kind of constituency structures that have been

argued to be relevant for cognitive science, and how these fit into to the overall framework presented here.

3.8.4.1. Rank at the neurophysiological stratum

Neuronal units, specifically the neuron, are not found in neuronal systems in some kind of unstructured mess. Physiological studies into the structure and function of neuronal systems have uncovered that there are organizational units that seem to be highly important in the study of cognitive science. One such hierarchy of constituency is as follows (Churchland and Sejnowski, 1992: 20):

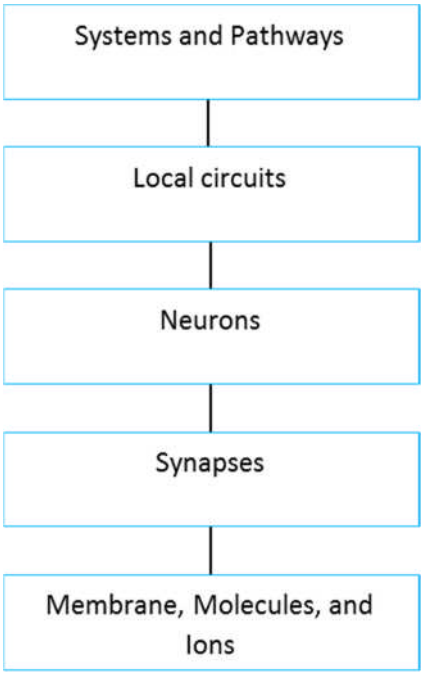


Figure 3.15: Structural levels of organization in neural systems

The lowest level on this hierarchy refers to various sub-components of the synapse, and includes elements such as the channels on the membrane. Synapses, themselves, are a constituent part of neurons themselves, as well as the other structural elements. Neurons combine to form local circuits, which themselves constitute larger systems and pathways within the brain (*ibid*).

A similar notion of constituent based considerations in cognitive neuroscience comes from Craver (2007: 166) in his discussion of spatial memory:

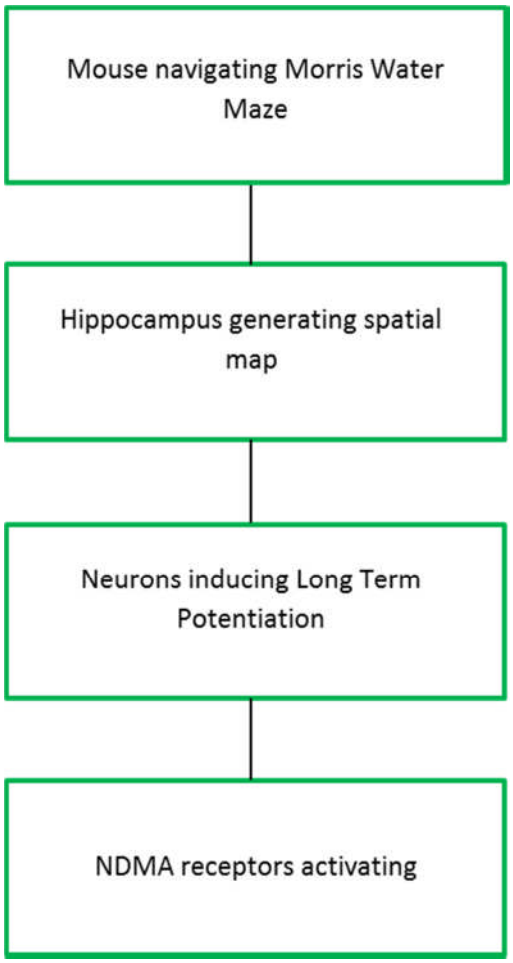


Figure 3.16: Levels of Organization in Spatial memory, based on Craver (2007: 166)

It can be seen here that we are operating over a hierarchy of constituency. The mouse is constituated by the hippocampus, as well as other parts of the body and brain. The hippocampus is made up of individual neurons, which have as some of its parts receptors including NDMA receptors. An additional aspect of this picture is that it not only identifies what structural levels may be important, but also the kinds of processes that occur at each level.

Within neuronal systems there is a well understood series of constituency relations that hold between parts and wholes. For the purposes of this framework under development, the following rank scale will hold at the lowest *level of interpretation*:

Rank at a neurophysiological level

Neuron > Neural Ensemble > Local Circuits > Brain Region

As a reminder, the importance of specifying this dimension in the framework of cognitive systems is that it allows **the identification of the relevant structural level at which each meaningful activity is taking place**. As an example, in the previous section an initial system for *neural oscillations* was stated. This was seen to operate at the structural level of *neural ensembles* or in some cases larger brain regions. Any 'choice' within the overall system of *neurophysiological* events would have to be stated in the same terms, i.e. at a given structural level.

3.8.4.2. Rank at the psychological stratum?

In terms of the lowest neurophysiological stratum there is a readily available and well-known rank scale in operation. When we come to look at the higher stratum, however, the situation is less clear.

Although it is less clear, it might be possible to make a few initial observations. As was seen earlier in this chapter, there is a similar situation in operation at the higher *levels of interpretation* within the linguistic stratification framework: the higher strata have less well-defined units of rank at which to operate. However, this has not stopped at least one initial unit of rank being identified: that of the *text*. It was this unit which was described as the predominant constituent level for meaning making at the semantic stratum (Martin, 1985). In the work that has been previously done on the structuring of human behaviour, the unit that is usually taken to be important is that of the **act** (c.f. Pastra and Aloimonos, 2012: §2). I will propose that in much the same as the text is the main unit at the semantic level, the **act** holds a similar position in this proposal. An act in this sense is characterized as any instance of cognitive behaviour, be it externalized (as in movement), or not (as in beliefs or emotions).

In any given act, however, there may be a number of smaller 'sub-acts' that can be described as combining to create it. For instance, consider a movement act like grabbing an apple. In the successful completion of this movement, there will be a number of constituent movements that are necessary to perform it: tiny motions of the arm, the opening of the

fingers, visual movements for spatial judgement, and the grasping itself. This example, and similar ones, has led some to propose sub-units known as motor primitives (Flash and Hochner, 2005) or 'movemes' (Del Vecchio et al., 2003).

Potentially, any cognitive act can be thought of as consisting of a number of smaller acts. This is especially true for cases of outward movement and motion, although less clear for cognitive processes such as emoting or thought. Past this intuitive idea, however, there is little more that can be given at this stage, and I leave the idea of a constituent scale operating at the psychological level an open question. This is not necessarily a problem, as a similar case is found in SFL, indicating that structural relations may be harder to identify at higher *levels of interpretation* within semiotic systems.

3.8.5. Instantiation

As detailed above, the crux of the notion of instantiation was the system of language as viewed at different levels of interpretation is not divorced from the output that the system gives in the form of texts. As well as conceptualizing this link between the system and the instance, the cline of instantiation hypothesizes mid-points on this dimension that capture regularities observes over various instances, as well as viewing the same points as sub-potentials that model the resources used to create new texts in novel situations. In concluding section §3.7, this dimension was observed to operate over many *levels of regularity*.

This section will proceed as follows: firstly, we will see if such a hierarchy has already been proposed, and if so to see whether it is the same in detail to the term in SFL. If not, then it is necessary to see if there is any evidence that such a scale would be useful for cognitive science, and propose a further scale into the discussions of cognition.

On this first point, some of the previous accounts in cognitive science that were discussed in the last chapter do contain the foundational idea behind the *cline of instantiation*, namely that the nature of the system is affected by the output of said system. This was true in the case of connectionist networks, whereby the changing weighted network connections were a result of the activity of the network itself and how well it compared to the desired output.

However, there was no explicit *level of analysis* in the conceptual frameworks based upon this approach. The discussion in this section provides such a theoretical tool.

We can start our motivation for the inclusion of a series of *levels of regularity* by considering the role of *neuroplasticity* at the lowest of the *levels of interpretation*. *Neuroplasticity* can be defined as the ability for neural systems to change their circuitry, so that future neural responses are more likely to occur in certain situations (Malenka, 2002). In terms of synaptic plasticity, this involves the strengthening or weakening of synaptic connections, so that certain pathways become more or less likely to occur in subsequent firings.

There are two major types of synaptic plasticity that are thought to be important for the cognitive behaviour of organisms, Long-Term Potentiation (LTP) and Long-Term Depression (LTD). The first of these occurs when repeated stimulation causes an increase in the efficacy in the synaptic connection, making is 'stronger' in the sense that the particular connection is more likely to be active over a time period that can last from minutes to days (Dayan and Abbott, 2001: 282). LTD on the other hand, has the opposite effect. The strengthening of certain connections causes certain pathways to be **less likely** to fire over a similar period of time. The difference between the two is down to particular mechanistic effects in each case (Malenka, 2002), although the precise details in each case are still under investigation. Important for this discussion is that both of these processes are **activity driven**, that is they are caused by the activity of neural system itself, an idea first postulated by Hebb (1949). Shortly, an interpretation of this in terms of SFL will be given.

Certain types of neural activity are thought to trigger either LTP or LTD. In experimental studies LTP and LTD effects can be induced by *tectonic* stimulation: repeated stimulation of a particular neuronal pathway. More recently, a proposal known as *Spike-Time-Dependent-Plasticity* has been proposed (Bi and Poo, 1999). In this case the relative orderings of the pre-and post-synaptic neurons are crucial. If the pre-synaptic neuron consistently fires **before** the post-synaptic neuron then the synapse is potentially induces LTP. If the reverse happens, then LTD is witnessed.

Furthermore, *neuroplasticity* is thought to be important for a number of cognitive processes. The most frequently discussed cognitive correlates are memory (c.f. Neves, 2008

for a review) and learning mechanisms (c.f. Caroni et al. 2012 for a review). This link between plasticity and cognitive function is crucial because it indicates how a neurological phenomenon can be *meaningful* for an individual.

In terms of an interpretation along the lines of SFL, the process of *neuroplasiticty* highlights two factors that fit in well with an *instatiation*-type analysis. Firstly, the process is **activity-driven**, i.e. it occurs based upon the activity of the neuronal system itself. Along the lines of the picture being posited here, this can be considered in terms of an **instance of the neurophysiological potential** having further effects upon the way **the system itself is organized and operates.** In other words, there is evidence at this lowest stratum for the interplay between the **system** (i.e. the potential number of neurological states an individual can exhibit) and the **instance** (i.e. actual occurrences of neurophysiological activity). The instance affects the system by causing future events to be more or less likely to occur, and the system influences the instance, since any changes to the system will influence the nature of any later instantiations of said system. These factors are exactly what the *cline of instantiation* has been designed to capture.

With this discussion in mind, I postulate the existence of an *instantiation cline* operating at the *neurophysiological stratum*. It is sketched below:

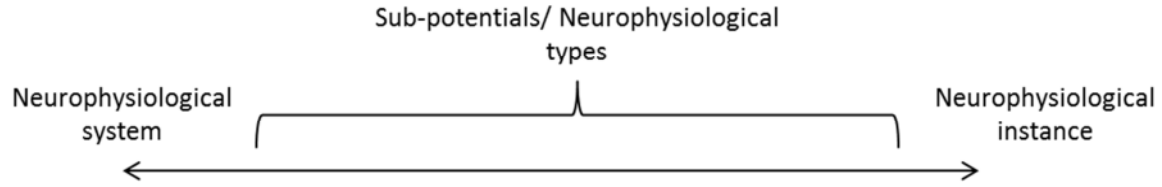


Figure 3.17: Cline of instantiation at the neurophysiological stratum

At one extreme lies the **neurophysiological system** and at the other lies **an instance of that system** as would be seen in experimental recordings. In between, we can hypothesize the existence of a number of **sub-systems/ neuronal types**, These are interpreted in the way they have been throughout this chapter: firstly, as a variant of the overall system, but where certain features are more or less likely to occur, and secondly as a generalization over instances, grouped together because they share certain properties.

At the higher *level of interpretation* in order to motivate the inclusion of a cline of instantiation in this case would be if there is evidence for any regularity in behaviour given certain situations. In a sense, in the case of humans, the discussions of regularities of linguistic patterns can be seen as evidence for this, since language is a sub-type of the possible behaviours a human individual can partake in.

Moving away from linguistic examples, the relevant question is whether it would be beneficial to consider of a similar dimension operating in terms of the **psychological system**, that is, the systems of cognitive **acts**, broadly construed as above?

Intuitively, it seems that such a dimension would be useful. After all, individuals do not seem to perform random behaviours in certain situations. Work on habitual behaviour has tried to show that, for certain cognitive and behavioural outputs at least. It is possible to see regular patterns in the ways the kinds of acts individuals produce. In the context of habitual research, Danner et al. opine that (2008: 245): "habits are formed when using the same type of behaviour frequently and consistently in a similar context for the same purpose." The appeal to context is important here. As was seen in the case of cooking recipes above, the driving factor behind any instantiational type is that it is contextually motivated. However, an in-depth look at the notions of *context* and *environment* is provided in §4.

One of the first influential studies into the importance of past behaviour on the likelihood of future behaviour came from Bentler and Speckhart (1979). In this study the authors investigated the consumption of alcohol and marijuana by university students. They argued that consumption of both of these substances becomes habitual over time, based upon the frequency of past behaviour, possibly even in the absence of reported intention to repeat the behaviour in the future. This approach has been applied to a range of habitual behaviour, including condom use, seatbelt wearing and drug use (c.f. Ouellette and Wood, 1998). In the studies in Danner et al. (2008) reported upon the possible causes of habitual drinking of alcohol, cycling, drinking milk and snacking. As a result of their findings, two important factors were identified: *context stability* and *frequency of past behaviour*. The first of these refers to similar situations occurring at the time, place, day, etc. The second is how frequently the individual has performed the behaviour previously. Interestingly, the

intention to perform the behaviour at a later time was not a significant predictor of occurrence of future behaviour.

These two factors, and the relative lack of importance of intention, fit in well with the notion of the *cline of instantiation*. The importance of *context similarity* reinforces the idea that any regular behaviour captured by this dimension is to be considered to be *contextually* driven (c.f. 3.6.2.). The factor of past behaviour presents us with a similar position to that just seen in the case of *neuroplasiticity*; the repeated instances a particular behaviour results in the same behaviour becoming more likely to occur at a later time. In the terms that have been discussed, repeated **instances** of the psychological **system** alters it in terms of likely events. This change to the overall potential of meaningful activity influences any future instances in similar contexts. Again, there is an example of the bidirectional relationship between the **system** and **instance** that can be captured by an *instantiation cline*.

The proposal is that there is a similar *level of analysis* operating at the **psychological stratum**:

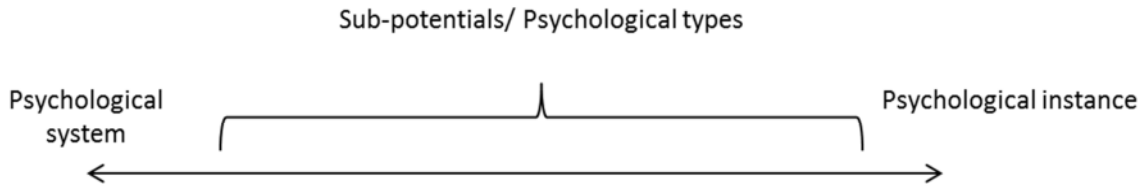


Figure 3.18: Instantiation at the Psychological stratum

Instead of running between the neurological system and instance, in this case the relevant terms are the **behavioural system** and the **behavioural instance**. The mid-points, **sub-systems** or **behavioural types** capture similar generalizations as were seen at the lower stratum.

This concludes the discussion about the application of *levels of analysis* from SFL to cognitive science more generally. This chapter concludes with an overview of the conceptualization of cognition that this provides us.

3.9. Summary

The purpose of this chapter has been to begin our investigation into whether it is useful to consider cognitive systems in terms of *dynamic, open, semiotic* systems along the lines that SFL considers language to be. It has involved a discussion of the *levels of analysis* that SFL theorists argue are important to study such semiotic systems. The *levels of analysis* discussed, and some of their properties, are given below:

Levels of analysis in cognitive science (from SFL)

1. *Stratification* – describing cognitive systems at various **levels of interpretation**. It was argued that this can be considered to be the identification of relations between units that increase in spatial and temporal scope. Levels of interpretation are linked by the relation of **realization**. This is a bi-directional probabilistic link, by which higher-level features are taken to "activate" features at the level below, and lower level features are considered to "construe/ encode" higher level ones. The *levels of interpretation* identified were the **neurophysiological stratum**, the system of possible brain states, and the **psychological stratum**, the system of possible cognitive acts.

2. *Rank* – describes the **levels of constituency** that are thought to be important at each level of interpretation. The hypothesis is that a relevant set of constituents is available at each level of interpretation. These interact with the dimensions of *stratification* and *delicacy* to posit the potential range of meaning-making structures important in each case. In the case of the neurophysiological system, the relevant scale is relatively clear. This is not the case for the higher stratum at the current moment.

3. *Delicacy* – describes each level of interpretation in terms of **levels of specificity**. The importance of this dimension is that it enables the modeller to divide the potential for meaningful cognitive activity at each stratum in

more or less general terms. The examples of *oscillations,* and *memory* were provided in this chapter.

4. *Instantiation* – sets up a cline consisting of **levels of regularity**. It provides the link between individual instances and the theoretical description of the system as a whole. Between these two extremes lie an indeterminable number of "mid-points". These can be considered in two ways: either as generalizations over a number of observed instances (viewing "from the instance end", or as a description of the potential for further meaning in similar situations (viewing "from the system end", mediated by constraining factors from context).

There are, however, some weaknesses in this proposal as it currently stands. Firstly, although much has been made about the potential for SFL heuristics in considering cognitive systems as coupled to the external environment, these ideas are yet to be developed. This will be the focus of the next chapter. Furthermore, the proposal is still very much at the conceptual stage, and it is as yet unclear how the development of such a framework may proceed. This will be the focus of the work in chapter 5, where the notion of *abstraction* in SFL will be discussed. As mentioned in the introductory chapter, given the holistic nature of SFL theory, the full proposal can only be considered once all the factors of cognitive systems and their environments have been considered. The next chapters will continue on this goal. Overall, however, it is true that the proposal is still at the initial stage, with many gaps and specifics to fit in. As discussed in §6, this is a necessary consequence of this thesis, but the discussion will be left until that point.

Chapter 4

Taking context seriously in cognitive science

""Now," said Rabbit, "this a Search, and I've Organized It-"
"Done what to it?" said Pooh
"Organized it. Which means – well it's what you do to a Search, when you don't all look in the same place at once."
A.A.Milne. The House at Pooh Corner

4.1 Introduction

It has been mentioned in both the previous chapter and the introduction to this work that one advantage of taking the approach advocated here is that it allows us to consider in a more explicit way cognitive life with reference to the environment that organisms interact with. It was argued that this is a crucial factor in thinking about cognition in an environmentally-driven way. Although only mentioned up to now, the details about how considerations in SFL concerning text-context interactions could help us towards a means of thinking about organism-environment interactions have not been given. The purpose of this chapter, then, is to do just that, by expanding our horizons to consider organism-environment interactions from an SFL-inspired viewpoint.

This chapter will be structured as follows: initially we will look at some of the ways that text-context interactions are thought to operate within SFL. This will include a discussion of the following:

1. The nature of context according to the SFL framework, incorporating the distinction made between the *material situation setting* and the *semiotic context*.
2. The contextual parameters of *field*, *tenor*, and *mode*, and the notion of a *contextual configuration*.
3. *Context* as a further stratum in the linguistic system, on par with the "language-internal" strata of *semantics*, *lexicogrammar*, and *phonology*.
4. The proposal for an "internal stratification" of the contextual strata into *genre* and *register*,

With this characterization of SFL in hand, this chapter will provide its potential application to cognitive science. Specifically, it will look at all of the above points, and consider whether

there is any use in applying them more widely. This will enable an example of a classical psychological experiment, the Morris Water maze, to be given in the concluding chapter (§6.2), which will bring together all of the work presented in this thesis.

Before the above discussion can commence, however, there is an issue of terminology that need to be made clear. This refers to the terms that refer to concepts such as *context* and *environment*. In this chapter, I will use a three way distinction between the *material situational setting*, *semiotic context*, and *semiotic environment*. A fuller discussion of all of these terms will be given in this chapter, but they will be used in the following ways:

a) The *Material Situational Setting* (c.f. Cloran, 1999) refers to the objective world not considered through either the 'lens' of language or the 'lens' of an organism's cognitive capacities. A closely-linked term is the 'mind-independent world'. I will use this term when discussing both SFL and cognitive science.
b) The *semiotic context* will be used to refer the notion of context as is relevant to systemic functional linguistic theorizing. This involves considering the wider contexts in which language and texts operate **through the perspective of the linguistic system itself**. In contrast to the *Material Situational Setting*, the *semiotic context* refers to the wider context considered in a language-dependent way. Furthermore, since the *semiotic context* is an explicit part of SFL's conceptual framework, it will also refer to that *level of interpretation* within the stratificational framework. This second meaning will be indicated in the text, when a possible ambiguity arises. See Table 4.1. for a fuller account of this idea
c) The *semiotic environment* will be used to refer in a similar way to (b), but in this case the focus is more broadly on cognitive science. It is the environment an organism is interacting with in a cognitive act, but seen through the lens of the organism's cognitive life. Again, I will use the term to talk about the phenomenon itself, and also the *level of interpretation* in the conceptual framework. I will mark differences between these meanings in the text.

One issue in using such terms is that they have been used more-or-less interchangeably in certain aspects of the literature. This is especially true for the terms *context* and *environment*. Bartlett (2013) provides a detailed discussion of this point. I am not claiming

that the use of these terms as defined above is wholly unproblematic. For instance, in the case of humans, the difference between the *semiotic context* and the *semiotic environment* may not exist on the above definitions, or be entirely clear. However, in applying these ideas in a novel way, I feel it is more important to make clear whether I am talking about similar ideas either in functional linguistics or in cognitive science. I accept, however, that these definitions may blur certain important factors, and may need revisiting in the future.

4.2 Semiotic Context in Systemic Functional Linguistics

The idea that language can only be properly understood by always thinking of it as contextualized forms one of the cornerstones of SFL. This is summed up by Halliday (1999: 6) as follows:

> *[A] theory of language in* [it's semiotic - JW] *context is not just a theory about how people **use** language, important though that is. It is a theory about the **nature** and **evolution** of language, explaining why the system works the way it does; but with the explanation making reference to its use.*

The above quote also highlights why this idea of context is so important: the nature of texts can only be understood by its function in a particular setting, and if the nature of the system is built from studying texts, then the study of the system itself must consider the coupling between text and context.

What, though, is semiotic context? The discussion begins by considering this question.

4.2.1. The material situational setting versus the semiotic context[16]

Context itself can be difficult to give a precise definition for, although all linguists would say that at some point it must be considered in the study of language. Within many linguistic theories, the study of linguistic meaning in context is given the name *pragmatics*. This study includes aspects of language such as presuppositions (Stalnacker, 1973; 1998; Karttunen, 1974; Heim, 1983), Gricean maxims and conversational implicatures (Grice, 1989), indexical

[16] It is noted that the term *semiotic context* is also used by Bartlett (2013). I chose this term because I felt it emphasized the difference between the non-semiotic notion of the mind/language independent world and the context as it is important in a semiotic system. Therefore, I do not necessarily use the term in the same way as the just cited work.

and demonstrative terms, and anaphora (c.f. Potts, 2009). All of the work in this tradition is closely associated with work in logic and the philosophy of language, and they attempt to investigate the role that context can have to linguistic meaning by means that cannot be derived from solely considering the nature of the linguistic sentence. However, their means of conceptualizing language is mainly in informational or epistemological terms. That is, they consider the context to be, largely, the information that is known to both interlocutors at the time of a particular linguistic utterance. This idea is at the heart of many pragmatic theories of presupposition, via means of the "common ground" that models the known information shared by both parties (c.f. Stalnacker, 2002).

The picture in SFL is different. In the first instance, an important distinction made is between the *material situational setting* (MSS) and the *semiotic context*. The former refers to the environmental context in which language operates in a **language-independent way**, or one could say, in a **mind-independent** way (Cloran, 1999: 178-179). It has been described as the "actual physical space containing actual physical elements" (Cloran, 1999: 178), while Hasan (2009: 177) states that "the material situational setting […] consists simply of material objects, person(s) – but not personalities, which always form part of the relevant context[17] – and their attributes." One possible way to think about it is to consider the MSS as the external environment as it would exist without the influence of language dividing it up in such-and-such as way; it is the "pre-lingusitic" setting in this sense. It would contain the aspects of the world **as they are *in actuality*** without any biases that may come from our linguistic capabilities.

The *semiotic context* on the other hand, is the context as understood *through the perspective of language*. Being part of the semiotic system, it is the aspects of the environment that are involved in the meaning-making act themselves. This idea of a semiotic context has its roots in Firth (1957), as discussed by Hasan (2009: 176). According to Firth, the context for linguistic study should contain the following aspects:

[17] Here, Hasan talks about the notion of "relevant" context. As I am using the terms here, I am not making any distinction between "relevant context" and "semiotic context", as part of the semiotic context will include objects made relevant through language (c.f. Table 4.1 below). This will also be clear in the discussion of cognitive science, where included in the notion of "semiotic environment" will be objects made relevant for the organism, based upon certain properties of that environment and perceptual capabilities.

A: the relevant feature of the participants: persons and personalities

(i) the verbal action of the participants

(ii) the non-verbal aspects of the participants

B: the relevant objects

C. the effect of the verbal action.

As pointed out by Hasan (*ibid*), there is sometimes difficulty in specifying exactly what would count as **relevant** on the above exposition of context. The partitioning of the *semiotic context* into the parameters of *field*, *tenor*, and *mode* (see §4.2.2. for description) enables those working on context to track those aspects of the context that are deemed important for any given linguistic event.

The second aspect of the *semiotic context* is that, in contrast to the way that the material situational setting was considered in language-independent terms as simply physical objects, the semiotic context can be considered to contain not only the relevant aspects of the context, but also how a text can help shape its own context. Hasan (*ibid:* 177, emphasis added) describes the situation as follows:

> *Relevant context* [read: semiotic context – JW] *is that which is based on the interpretation of the language of the text, then, clearly, it is something* ***'made by (worded) meanings.'***

One unfortunate turn of phrase in the above characterizations needs to be briefly discussed. In discussing the material situational setting in terms of "actual physical elements", the implication seems to be that the type of context usually assumed to be important for language, what is being called here the *semiotic context*, is somehow not physical. Any fears along these lines can be safely dismissed. This can be shown by considering the *orders of nature* as discussed by Matthiessen (2007a.: 545-546, c.f. Halliday, 1996; Halliday and Matthiessen 1999):

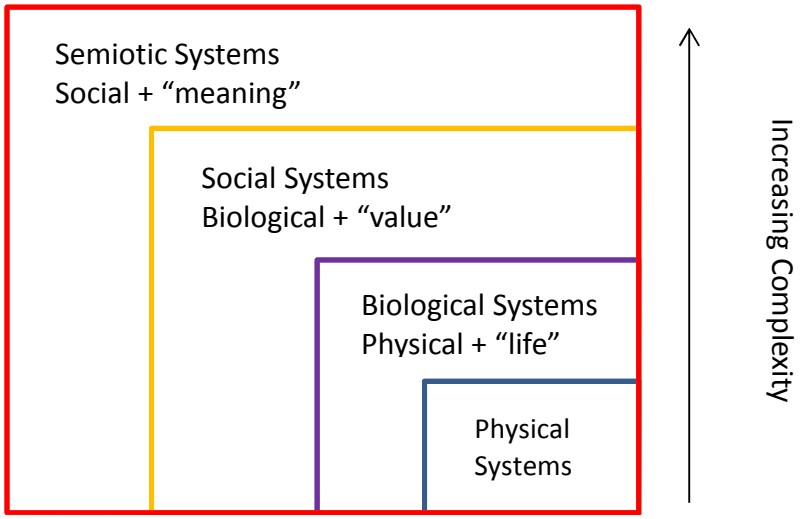


Figure 4.1: Matthiessen's systems of nature.

For our purposes here, it is important to note that each instance of a higher system is also an instance of the systems lower down in the hierarchy. As such, the groupings above should be conceptualized according to a ordering of sub-/supersets (Halliday and Matthiessen, 1999: 508):[18]

> *Physical systems are just physical systems. Biological systems, however, are not just biological systems; they are at once both biological and physical. Social systems are all three: social, biological and physical.*

The general picture presented above can alleviate the worries one might have about semiotics being "non-physical". The context is one part of the semiotic system of language. As such, it is also social, biological and physical. In other words, it is simply part of a physical system of a certain kind. The appeal to specifically stating the 'physical' nature of the MSS, then, is most likely to make more precise the idea that it is to be considered to "exist" lower down on Matthiessen's hierarchy of systems than the *semiotic context*.

In sum then, the differences between these two conceptions of context can be summed up as follows:

[18] Unfortunately, the diagram as given by Matthiessen seems to suggest that there are more examples of the higher systems, in comparison to the lower one. It should be clear from the accompanying discussion here, that the opposite is true.

Material Situational Setting	Semiotic Context
Independent of Language	Viewed through the perspective of language
Material, non-meaningful	Semiotic, therefore material and also meaningful
Contains all entities, properties and relations	Contains only relevant entities, properties, relations

Table 4.1: Differences between the *Material Situational Setting* and the Semiotic Context

SFL, then, conceives the context as appropriate for linguistic theorizing as a semiotic construct: it is meaningful and cannot be considered as separate from the linguistic system to which it is linked. In the next section, the status of this semiotic notion of context as a theoretical entity will be discussed further.

4.2.2. Contextual parameters: Field, Tenor and Mode

We have seen, then, how SFL makes the distinction between the environment as it exists 'language-independently', and the *semiotic context* as the context both created and influencing the linguistic system. The relationship between the two can be considered to be one of constrainment: the nature of the semiotic context is partly constrained what appears in the Material Situational Setting. Even with this distinction, however, the notion of context (I will drop the modifier *semiotic* from now on) still remains a vast and heterogeneous notion. In order to bring some systematicity to the notion of context, SFL proposes that the context can be metafunctionally diversified in much the same way that the 'inner-linguistic' strata can be understood in terms of the *ideational*, *interpersonal* and *textual* metafunctions. The elements of this metafunctional diversification are given as follows (Halliday and Hasan, 1985: 12):

> **Field** refers to what is happening, to the nature of the social action that is taking place: what it is that the participants are engaged in, in which language figures as some essential component.
>
> **Tenor** refers to who is taking part, to the nature of the participants, their statuses and roles: what kinds of role

> relationship obtain, including permanent and temporary relationships of one kind or another, both the types of speech roles they are taking on in the dialogue and the whole cluster of socially significant meanings in which they are involved.
>
> **Mode** refers to what part language is playing, what it is that the participants are expecting language to do for them in the situation: the symbolic organisation of the text, the status it has, and its function in the context.

As may be clear from the discussion of metafunctions in §3.1, each of the three contextual parameters given above matches to ones previously described. *Field* corresponds to *ideational*, *tenor* corresponds to *interpersonal*, while *mode* corresponds to *textual*. The reasons for these groupings is that both the *ideational* metafunction and *field* concern aspects of the action taking place, the *interpersonal* metafunction and *tenor* both deal with the kinds of social roles important to understanding the text, and *textual* metafunction and *mode* both cover the role that language is having in the act of text production. This link between certain contextual and linguistic metafunctions has become known as the *context-metafunctional hook-up hypothesis* (Halliday, 1977b.; Hasan, 1995; Clarke, 2013). The relationship between the two can be considered to operate over the dimension of *realization*, already discussed in §3.2. Taken together, values for these three contextual parameters provide a description of a particular semiotic context. In the remainder of this thesis, I will use the term *contextual configuration* to refer to a particular combination of *field, tenor* and *mode* values.[19]

In order to clarify the aspects of the linguistic context that these contextual variables are intended to capture, below are a few examples of the aspects that are covered by each of the above parameters.

[19] Again, there are some terminological issues with my choice of terms when talking about *field, tenor* and *mode*, as they relate to the notion if *register* as used by different theorists in SFL. In the work of Jim Martin and colleagues (c.f. §4.2.4) the *system* of register is defined in terms of these contextual parameters. However, the same term is used by Halliday and those who follow his writings more closely to refer to functional variation, particularly at the stratum of semantics (for example, see Halliday, 1997/2005: 256). I have chosen the more neutral *contextual configuration* in this context to avoid any misinterpretation of my ideas.

Examples of the features of the semiotic context covered by the parameter of *field* include: the type of action taking place, whether it is on-going or restricted to a particular place and time, and whether the action is part of "general life" or only in certain specialized contexts. As an example, Bowcher (2014) claims that there are three major strands of meaning-making contextual information relating to the *field* of discourse, or the nature of the on-going social action. The first of these describes the type of action taking place: is it practical in the sense of, for example, the "physical" exchange of items, or conceptual in the sense of debating and other purely verbal forms of task. The second major strand describes some aspects of the situation in which the action is taking place: whether, for instance, the nature of the task is highly specialized and so would only be meaningful to a small group of interlocutors, or whether it is more general in orientation, or whether the nature of the task is defined according to institutionalized practices or not. Finally, the system of PERFORMANCE OF ACTION captures details of the social action such as whether the action takes place in a well-defined singular place and time, and whether the goal is immediately achievable, or operating over a longer period of time. The more delicate options on this part of the description have been omitted for sake of space, but the reader can refer to the original citation for more details.

The *tenor* of context concerns itself with the social relations between participants that play an important role in the kinds of choices that are made within the linguistic system. Butt (2004) provides the following features of tenor: **social hierarchy**, **agentive role** and **social distance**. The first of these describes the social structure of the particular situation: who has the social prestige, and who does not. Related to this is the idea of social distance: given this social hierarchy, where are the participants situated and what are the social dynamics between them? Mapping these kinds of social, interactive elements of the context is deemed important from the social semiotic viewpoint, because they can greatly influence the semantic, lexicogrammatical and phonological features of texts in different situations. Agentive role describes the role that each interlocutor is playing within the social event in which language is playing a part.

Finally, the *mode* of context deals with the role that language itself is playing in any one particular setting. Again, according to Butt (2004: 42) one the major features is conceptualised in terms of the role that language is playing: either it is **constitutive**,

supported, or **ancillary**. On the first of these language plays a major role in the social event; the social event is based and negotiated primarily through language. On the **supported** option, language is not the primary medium, but is still an obligatory part of the event. Finally, **ancillary** role relates to situations where language is not necessary; any language that is used is entirely secondary to other kinds of social exchange (c.f. Hasan, 2014: 18 for discussion). Within the subsystem of **ancillary role,** sub-options include **exchange**, where the event contains more than one act, and **act**, when only a single act takes place.

It is important to note that the description a particular *contextual configuration* in terms of field, tenor and mode acts as a constraining factor on the types of language produced in a particular semiotic context (Halliday, 1985: 12). If it is possible to understand the nature of the context in which language is operating, it is possible to state, with some degree of accuracy, the relevant semantic and lexicogrammatical features the resulting texts will contain. In this, the systemic functional approach to language can be said to be *contextually driven*. Furthermore, if texts are the only way to get to system, then the system of contexts are also important in understanding the system of language.

Context, then, is described in SFL according to three parameters: *field*, *tenor* and *mode*. These parameters allow the researcher to bring some structure to the wide range of entities in the semiotic context, by grouping contextual features that relate to (i.) the nature of the social action, (ii.) the social relationships relevant in the situation and (iii.) the role that language is playing.

The next section contains a discussion of the stratum of context as a **fully-fledged theoretical stratum**. This will be important when this chapter moves to cognitive science, as the view in SFL is that the context is itself a **system**, and can be described along the same dimensions as the other linguistic strata.

4.2.3. The status of context as a theoretical stratum

The notion of the *semiotic context*, then, is given a particular meaning in SFL that is markedly different from the "world out there" approach that may constitute a theoretically naïve perspective of the term. It is defined as that aspect of the wider world that is both created through and effects language use, and maps the parts of the context that are

linguistically relevant for whatever the particular task at hand is. It is important to highlight, however, that within SFL this perspective on context is considered to be a stratum within the overall model, **which should be amenable to description in the same terms as semantics, lexicogrammar and phonology.**

The importance of this is that the theoretical dimensions that were discussed in the previous section also apply to accounts of the context. That is, it is a theoretical necessity that context should be described in terms of *instantiation*, *delicacy* and *rank*, and also its links to the other strata need to be stated. This latter notion will mainly be discussed along with an account of *genre* and the *register system*, as found in the work of Jim Martin, in §4.2.4.[20] As noted in the discussion of rank in the previous chapter, to my knowledge there are no accounts of rank at the contextual level of interpretation. This section, therefore, will discuss the dimensions of *instantiation* and *delicacy*.

To recall, *instantiation* is a hierarchy of levels of analysis understood to operate over a scale of regularity. It is the means of linking individual texts with the overall system, and, in doing so, attempts to bridge the gap between linguistic *competence* and *performance* that is maintained in other linguistic theories. By looking for regularities, this dimension allows researchers to account for certain re-occurring patterns that are found in language, whereas still maintaining that the system itself has the potential for undergoing change. The following diagram, repeated from the previous chapter, will be useful for guiding our discussion here:

[20] As noted in ff19, there is some variability of the use of the term *register*. When referring to the notion of *register* as a contextual stratum, I will use the term *registerial system*. Again, this is not meant to include the use of the term as originally used by Halliday.

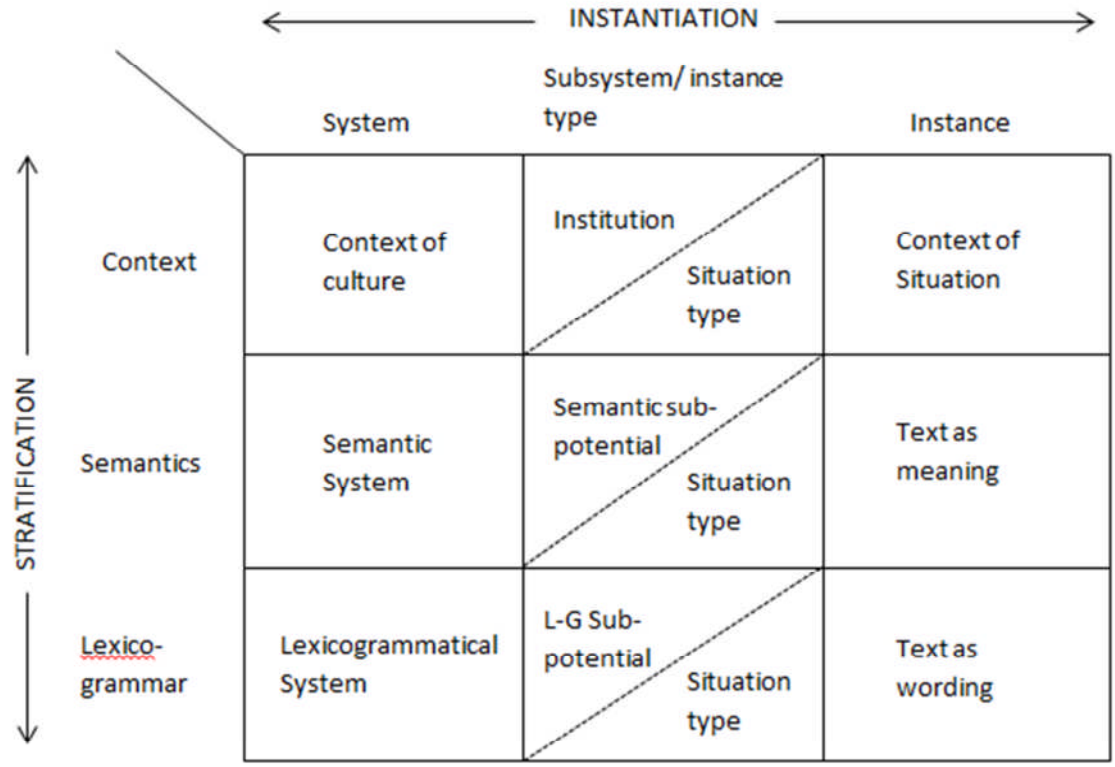
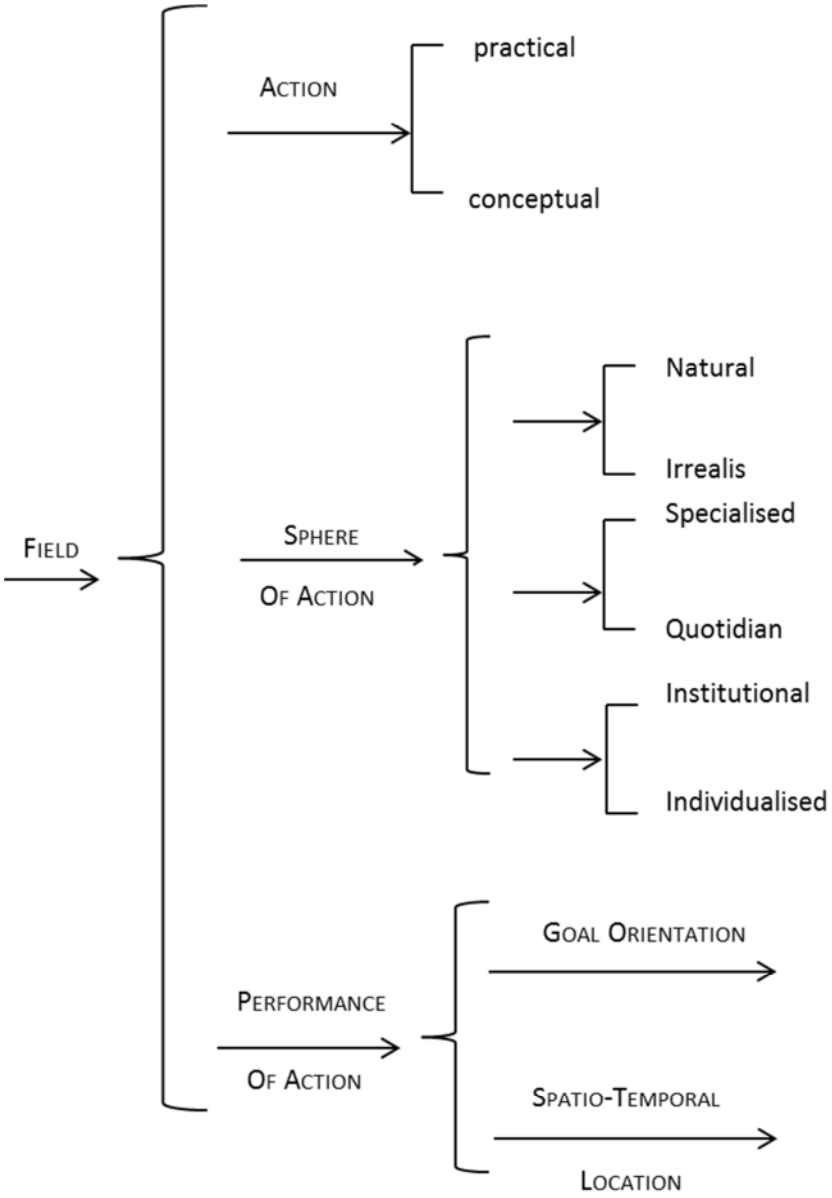


Figure 4.2: Instantiation/ Stratification matrix

Previously, instantiation has been discussed as it pertains to the bottom two rows on the above figure. The main difference in the case of context is that the cline of instantiation as applied to the levels of interpretation of *semantics* and *lexicogrammar* take *text* as the relevant unit at the instance end, but when we are dealing with context this is no longer applicable. As we are dealing here with the semiotic contexts that language occurs in, we are instead taking the instance end to consist of a particular *semiotic context*. In the terms used in the above diagram, this is indicated as a *context of situation,* although I will not use that term in order to avoid confusion with the notion of the Material Situational Setting. As detailed in §4.2.2 this can be described in terms of the three contextual parameters of *field*, *tenor* and *mode*. On the other end of this cline we have all the possible situations that language can occur in, or the potential number of combination of *field*, *tenor* and *mode* variables. This carries the name *context of culture*. Between these two, fall the intermediate points that can be described as either *institutions* or *situation types*. The former of these refers to thinking about the generalizations covered by the cline of

instantiation as a description of possible future situations, while *situation types* denotes generalizations based upon a certain set of described contexts that are alike in some way.

In terms of delicacy, there have been attempts to provide system network descriptions for the level of semiotic context as a theoretical stratum (Bowcher, 2007; 2014; Hasan, 1999; 2009; 2014). For example, in the last section, Bowcher's (2014) features of *field* were outlined. In terms of delicacy, these can be described in a system diagram as follows:



Network 4.1: Field network (adapted from Bowcher, 2014)

This system network sets out the features identified previously in terms of delicacy. For instance, the features that describe the various options in the system of SPHERE OF ACTION

have been organized into three simultaneous systems. Other systems have been given for the options within the contextual parameter of *mode* (Butt, 2004: 42). In each of these cases, various features within the stratum of context have been organized along the dimension of delicacy, motivating this stratum as a fully-fledged semiotic *level of interpretation*.

A picture of how the semiotic context fits into the overall picture of language via the dimensions of stratification and instantiation is given below (taken from Hasan, 2009:173):

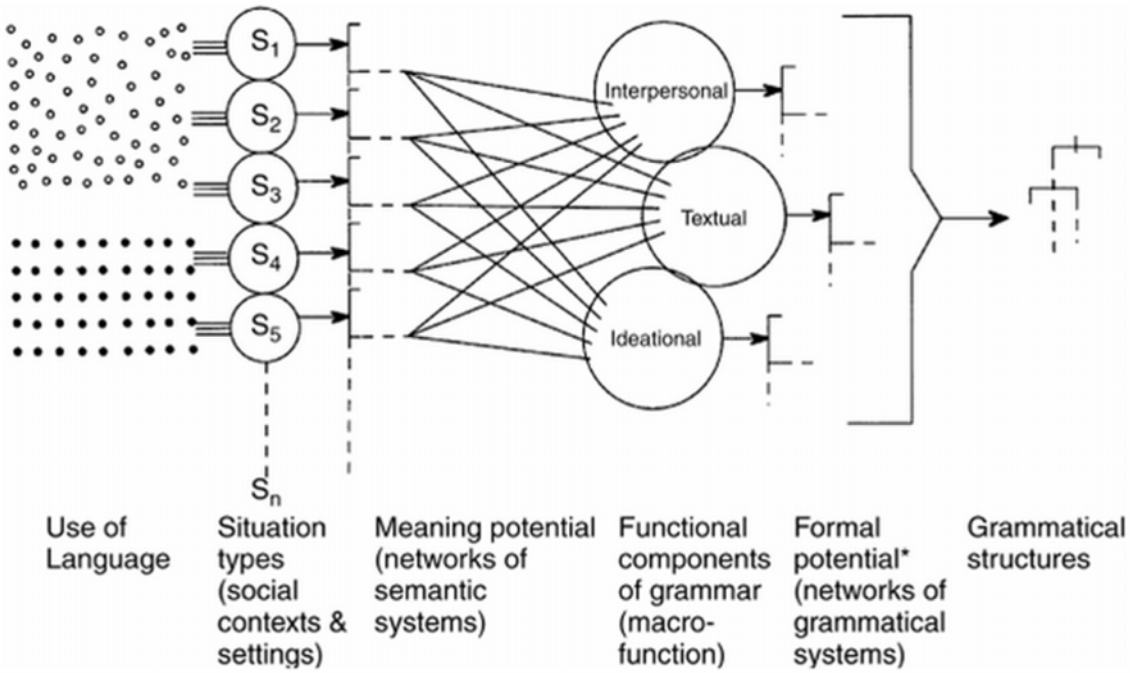


Figure 4.3: Trajectory from use, through context, to language (Hasan, 2009: 173)

Firstly, the initial point of departure for the theory is *language use*, seen through various instances of text. It is this *use-first* strategy that differentiates it from many linguistic theories. Secondly, once you have some data in terms of text, the next major aspect to consider is the context that the texts are occurring in. Of course, it is possible to consider the context on a case-by-case basis, but it is more useful to group similar contexts along the dimension of instantiation, as this enables one to see how similar contexts can affect language production in relatively similar ways. From here, you look at how aspects of this context correlate with certain regularities, or probability skewing, in the linguistic levels. Again, this should operate from a top-down perspective, moving from semantics, to grammar and finally to phonology. This is to maintain the perspective of SFL as a 'meaning-

motivated' theory of language (Halliday, 1979). Of course, this perspective requires a large amount of data to get the whole process running, and often this is difficult to acquire. Having said that, this methodological tax comes with the benefit of having a more empirically driven and motivated theory of language.

We have seen that the idea of a semiotic context can be discussed in terms of *instantiation* and *delicacy*. What has not been discussed is the possibility of a rank scale at this level of interpretation. Unfortunately, as with the case of the semantic stratum, I know of no such proposal. Although, if context should be rightly considered as a stratum, then one should exist.

4.2.4. Stratification of the contextual stratum: Genre and Register

As the content plane of language has been divided into the separate strata of *lexicogrammar* and semantics, some scholars also divide the contextual stratum into (at least) two strata: *genre* and, what I am calling, the *registerial system*. This view is most closely associated with James Martin and his associates (Martin and Rothery, 1980; 1981; Martin, 1985; 1997; Martin and Rose, 2007; 2008). As noted by Christie (2002), this view was influenced by the work of Gregory (c.f. Gregory and Carrol, 1978). In this work, instead of the three contextual parameters given above, Gregory's work included four: the idea of *tenor* was divided into *personal tenor* and *functional tenor*, although all of them operated as part of a single contextual *level of interpretation*. The first of these deals with the social relations between participants in the language event, while the latter deals with the purpose, or the goal, that the text is aiming to achieve. Martin takes this idea of *functional tenor* but instead places it at a level of interpretation higher than the other three and uses it as the basis of his definition of *genre*. *Genre*, then, becomes defined as "a staged, goal-oriented social process" (Martin and Rose, 2007: 8).[21]

Why is this discussion important? In the second part of this chapter, when we move away from linguistics proper and turn instead to cognitive neuroscientific studies, **goal** will play an important role. As just indicated, this notion is an integral part of *genre*. Therefore, this

[21] Actually, Martin has extended his contextual layer even further to contain a further layer of "ideology" at the top of his variant of the stratificational framework. Issues of space delimit me from discussing this here, but see Martin (1992) for the proposal and Hasan (1995) for critical discussion.

section lays the foundation for my arguments about where to place such considerations within the picture of cognition under development here.

An example of a *genre* is a type of story genre known as *exemplum*: a moral tale that is exemplified by fables and parables (Martin and Rose, 2007: 9). In order to typify a genre, two things need to be set out: (1) the overall "goal" of texts included in this genre, and (2) the stages that are usually found in texts that are members. It should be noted, however, that these stages should not be seen to be akin to necessary and sufficient conditions for inclusion in a particular genre. Keeping in line with one of the overall themes of SFL, it should be interpreted as the *likely* means of structuring a text. These two aspects for this particular genre are as follows (*ibid*):

Characteristics of a *exemplum story genre*.

(1) Goal – to present, interpret and highlight a problematic incident to an audience.

(2) Stages

(i.) Orientation – Provides the background setting for the story in terms of a time and place.

(ii.) Incident – Detailing some events that carry some moral content for the audience. There can be more than one of these, and they will usually be ordered in terms of chronology.

(iii.) Interpretation – the narrator comments upon the significance and/ or importance of these events as they relate to the overall purpose of the text.

(iv.) Coda – A summing up of the lesson of the story.

Within each of these stages, there may be identifiable "sub-stages", although these are much more variable across texts than these more general framing stages. Examples of other genres described in this work include *exposition argument genre*, and *legislation genre.* Rose (2006) gives further examples of *genre*, including *news story, anecdote, narrative* and *discussion*.

In §4.2.3 above it was seen that a context is defined by a certain configuration of the contextual parameters of field, tenor and mode. In the vocabulary of Martin and colleagues' proposal, this description is named instead *register*. It is important then, to look into the relationship between *genre* on one hand, and the *registerial system* on the other. Within Martin's work, the division between the two is in terms of *stratification* and *realization*:

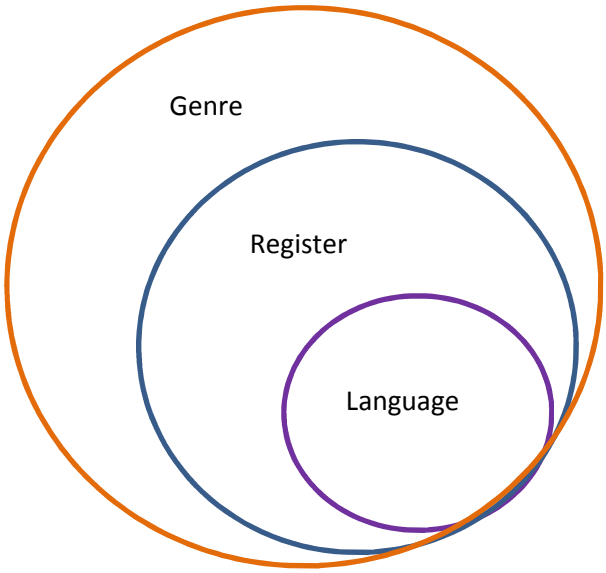


Figure 4.4: Language, Register and Genre. Following Martin and Rose (2007: 309)

Genre then stands at a higher level of interpretation that the *registerial system*, or what was simply called 'context' in the previous section. This means genre stands in a realization relationship to the contextual parameters of field tenor and mode. The fourth parameter of functional tenor proposed by Gregory and Carroll (1978) has been 'raised' above the other three to a higher level of interpretation. In the discussion provided in the previous chapter, it was seen that there are two kinds of divisions that can be applied when considering the relationship between two adjacent strata: either they cross a "semiotic boundary", i.e. we are moving from an expression to a content plane, or we have an instance of "internal stratification" whereby the inclusion of a new stratum happens within one semiotic plane. According to Martin (1999: 30) the first of these possibilities is considered to be applicable here, with genre acting as a content plane with register as its expression plane:

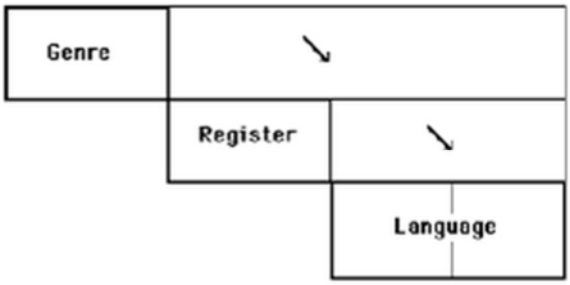


Figure 4.5: Realization relations between genre, register and language.

In order to justify this stratification of context into these two aspects, the authors provide the following arguments:

1. The *registerial system* is metafunctionally organized in terms of field, tenor and mode; genre is not. (Martin and Rose, 2007: 308).
2. A genre is not the same as the *registerial system* because more than one set of contextual parameters can be linked to one genre. Also, more than one genre can be linked to one set of contextual parameters. (*ibid*: 309; Martin, 1999).
3. A genre can be considered to be a generalization over the contextual parameters of field, tenor and mode. (*ibid*).

I believe that there are reasons to doubt the validity of all of the above arguments, by showing that they do not lead to the conclusions that Martin and Rose propose.

In short, my arguments are as follows:

1’ The “simpler” notion of the stratum of genre in comparison to the more “complex” nature of the rest of the strata is a concern for the division between (Martin’s) register and genre, given the fact that theoretical dimensions in SFL are taken to hold at all levels of interpretation.

2’ Martin is right that one goal can be linked to more than one set of contextual parameters, but mistaken that this trend works in reverse. This stems from the fact that two goals can be linked by a dimension of constituency, which is different in important respects to the many-to-many mappings in realization.

3’ The ability to generalize over genres using features of register to is in contrast to generalizations at the other strata. Typically, generalizations of features are done by means of the features that are found **also at that level of interpretation**. Furthermore, this ability to generalize is covered by the theoretical dimension of *instantiation*.

Having laid out the arguments in general terms, let’s now look at each one in more detail. This will be important for the purposes of this thesis, because clarifications concerning the nature of these terms in SFL will aid us in our attempts to apply them to the novel context of a general cognitive framework.

Firstly then, the first argument that Martin and Rose put forward for the diversification of the contextual stratum is that it is not metafunctionally diversified in the way that the *registerial system* is. From this conclusion, they argue that they are of a different kind and therefore should take place at a higher level of interpretation in the framework. This, however, would be a surprising claim, because one of the hallmarks of the idea of the stratificational model is that all of the strata in the framework are described using the same set of theoretical dimensions. Lexico-grammar and semantics have been frequently discussed in terms of metafunctional variety, as should be clear from much of the discussion in this thesis. There have also been accounts of phonology being metafunctionally diversified (Cléirigh, 1998). The *registerial system* itself is split into the metafunctional parameters of field, tenor and mode. So, if one of the characteristics of the strata in human language is the fact that it can be divided according to different strands of meaning, this omission at the level of genre is an anomaly. Its omission, at the very least, demands explanation, and other motivating factors need to be provided for the argument that one needs to add more levels of interpretation to the framework.

Secondly, they argue that genre is not the same as the registerial system because one genre, i.e. one definition of a staged, goal-oriented process can be associated with more than one set of contextual parameters. If true, this would be a convincing argument in favour of treating genre to act at a higher level of representation, because, as was seen in the discussion of stratification in the previous chapter, this kind of complex many-to-many mappings between levels of interpretation was a hallmark of the realization relationship. In

order to justify this point, Martin (1999: 34) provides the example of the children's book *The Very Hungry Caterpillar*. He argues that this one text, described using one set of contextual parameters can be linked to more than one genre, or more than one staged, goal oriented process. The first of these would be to provide the child with a basic understanding of the process of metamorphosis, the other is to tell a story about a caterpillar who happens to eat a lot. The first is informative, the second is for entertainment. We have two goals here, both of which are linked (in Martin's words *realized*) through the same contextual configuration. However, on closer examination, the situation in this case is not the same as that in classical arguments for stratification.

The argument for a split between the strata of lexicogrammar and semantics was that it enables language-using humans to move away from a simple one-to-one mapping between sound-types and meaning-types. If this was the case, then our linguistic resources for communicating meanings would quickly "run out", as a completely new means of expression would have to be created to deal with each new meaning, however subtle, would have to be linked to a new form of expression. One of the most frequently given pieces of evidence for this comes from a comparison of the grammatical system of MOOD and the discourse semantic system of SPEECH FUNCTION. To reiterate, the grammatical types of interrogative, imperative and declarative can all be associated (with influence from the semiotic context) with any of the semantic types of question, command, request and statement. That is to say, although in any one linguistic text a statement, say, will be realized by a interrogative (for example with rhetorical questions), when one comes to look at a collection a texts in order to describe the system, the mappings are more complex. One of the mappings will be the most frequent, or the unmarked mapping.

The key is to realize that this is *not* what is happening when Martin is saying that one set of contextual variables (say those that would describe the situation surrounding the *Very Hungry Caterpillar)* can be linked to more than one staged, goal orientated processes. In this situation, I believe, what is happening that there is only *one goal* but that it is *complex*, by which I mean it has two constituent parts. The goal of this particular book is *both* to educate and entertain *simultaneously*. Using this observation to postulate the existence of a higher level of interpretation is based upon the confusion between the dimension of *constituency,* and the *abstraction* relation that links higher and lower strata. If we have a dimension of

constituency between two features, then this is a link along the dimension of *rank* not *stratification* (c.f. table 3.2).

Thirdly, they argue that we need to differentiate between *genre* and the *registerial system* because one can use contextual parameters to generalize over genres. This point is, of course, linked to the previous one, where they have argued that one goal can be linked to more than one set of contextual parameters. It does raise an interesting point though, again one which seems on the face of it to be at odds with the situation at other points in the stratification framework. The question is: does generalization over one set of linguistic features using another set necessarily mean that we are dealing with linguistic types at two different levels of interpretation? The answer is, I think, in the negative. We will deal with this issue in much further depth in the next chapter, when the notion of *abstraction* in SFL will be discussed, but a short pre-emption of future ideas is warranted here. The notion of generalization in scientific practice amounts to making statistical inductions of common or likely properties of a set of instances, and in doing so one is able to make predictions about what is likely to occur in the future. As will be shown in the next chapter, this notion of generalization applies to the inter-level dimensions of delicacy and instantiation, and not to the inter-stratal dimension of stratification. Therefore, this is not a valid reason for postulating a higher level of representation.

So far, this chapter has examined the following aspects of the notion of context, and context-text interactions: the difference between the *material situational setting* and *semiotic context*, the status of the contextual stratum as a theoretical entity on par with the "language-internal" strata of semantics, lexicogrammar and phonology, the metafunctional diversification of context, and the proposed distinction between the registerial system and genre. This chapter now turns to look at the possible applications of these notions in cognitive science.

4.3. Organism-Environment relations through the lens of Systemic Functional Linguistics

In §2.5, the various members of the **4E approach** were discussed. A common factor in these proposals is that they all attempt to incorporate aspects external to the brain in their cognitive theories. From this discussion, however, the conclusion was reached that although there is a growing appreciation of the role of the environment in neurocognitive theory,

there is (1) a lack of specific recommendations or a framework on how to incorporate the environment into conceptual frameworks of cognition, and (2) the environment has not, so far, been dealt with on a systematic basis, with arguments about the role of a contextual factors on a certain cognitive function being considered on a case-by-case basis.

The position taken here is that by taking some lessons from SFL, steps can be made to rectify the above criticisms. Again, it should be noted that the purpose of this discussion, as with the discussion in the last chapter, is not to present a fully formed theory of environment-organism interactions, but to show that the principles described above can start us on the path to that eventual goal. A theory can only be reached on the basis of analysis of data, and as with the rest of this work in this thesis, no new data is being presented here.

The discussion begins with the first feature of contextual considerations in SFL, namely the distinction between the *material situational setting* and the *semiotic environment.*

4.3.1. Material vs. Semiotic Cognitive Context

The first distinction that we saw in SFL was between the material situational setting (MSS) and the semiotic context, the latter of which enters into the framework of stratification as a higher level of interpretation.[22] The former of these equates to the "mind/ language-independent" setting, while the latter corresponds to those factors of the environment that are (1) relevant for the linguistic event under discussion, and (2) represents the semiotic context that is both created by and simultaneously influences language itself.

Translating this perspective to cognitive science, we can therefore ask ourselves the following question about the environment, broadly construed: should it refer to: (1) the external world that is independent of the individual, that is the world as would be described by some ideal physical theory that can capture the world as it really is?, and (2) does the context that is relevant for the study of cognition somehow arise through the activity of the neuronal system itself? It will be argued in this section that there is room for both of these perspectives in the study of cognition, but the second of these should be dominant in our theorizing. In order to arrive at this conclusion, it is necessary to show that the 'real'

[22] It should be recalled that some proposals have the strata of *genre* and *ideology* higher than the stratum of *register*. However, given the discussion earlier in this chapter, it should be clear that I doubt the need/ validity of proposing such higher levels of interpretation.

environment, the material situational setting, and the *semiotic* environment **can be different at the same point in time**. The following sections present evidence from neurological and behavioural experiments to support this idea.

4.3.1.1. Evidence 1: Sensory Illusions.

Evidence for the distinction under discussion here comes from some of the vast amount of studies that have attempted to uncover the means by which certain sensory illusions arise. The number of studies is far too vast to even attempt a cursory overview, but I will outline a few of them below. To recall, the current purpose is to highlight how the perceived world, the semiotic environment may differ from the material situational setting, to the best of our understanding. It should also be pointed out that the arguments here are not novel, but merely point to a link between concepts discussed in SFL and the wider field of cognitive science.

One early example of a visual illusion where by the perceptive experience of an object differs from its objective characteristics are Mach bands, first discussed by Ernst Mach in 1865. An example of this phenomenon is given below (taken from Eagleman, 2001):

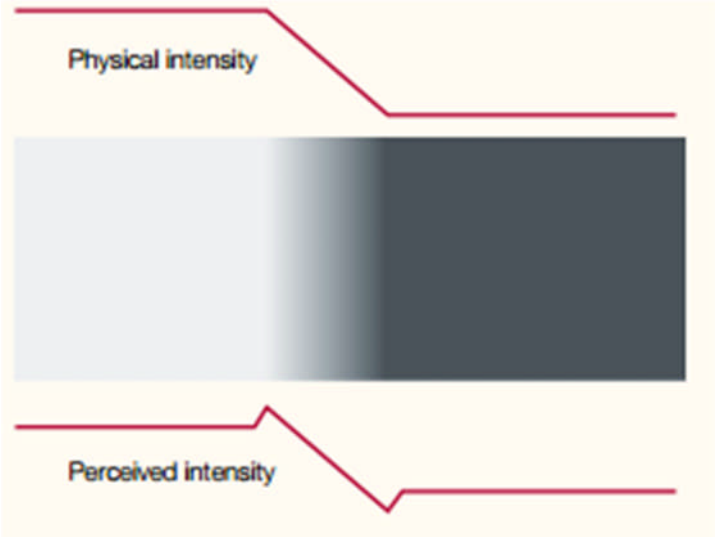


Figure 4.6: Mach band effect

This example exemplifies the phenomena of differing "realities" well. The observed illusion is that in the areas immediately preceding and following the gradient seem to be of a different

intensity to the shades that precede it (in the case of the white side), or after (in the case of the black side). That is, the white area immediately preceding the gradient seems brighter than what comes before, and the black area immediately following the gradient seems darker than that which follows. In truth though, there is no such difference, as indicated on the above figure. What this early experiment shows is that our perception of certain properties can differ from how they can be objectively described.

Such differences are not restricted to the visual domain, however. Within linguistics, there are many examples of auditory perceptual illusions that have become well-known in the literature. Two of the most famous of these are the phonemic restoration effect (Warren, 1970; Samuel, 1996), and the McGurk effect (McGurk and McDonald, 1976). In the former of these experiments, the auditory stimulus of a word was mechanically altered so that one of the phonemes was replaced by a distractor noise, either a cough or an auditory bleep. Even though the phoneme in question had been omitted from the stimulus, participants still reported that they perceived the phoneme, although they also reported that the phoneme had been overlaid by another sound. The McGurk effect is an instance of a cross-modality illusion, whereby different inputs in more than one sensory modality interact to produce an illusion effect. For instance, when a subject is presented with an auditory /k/ stimulus, and a visual /p/, they will report hearing the phoneme intermediate to these two, i.e. a /t/. This is the case even though the reported case is neither of the two presented stimuli.

Such examples are not new, they have occupied the thoughts of cognitive scientists, especially psychologists and philosophers, for a long time (Fish, 2010). These phenomena, amongst others, have been used in various arguments, especially those for the idea of *mental representations*. Searle (1999) calls them “arguments from illusion.” They are given here just to motivate a distinction made in SFL to cognitive science more generally.

4.3.1.2. Evidence 2: Cross-species differences.

The second piece of evidence for the distinction under discussion here is that although the material situational setting may remain the same, the semiotic environment may differ between organisms (or even within the same organism at different times). It is well known that when one looks at the differences between organisms, the range of sensory stimuli they are sensitive to differs, in terms of the types or the range of values of certain sensory inputs.

The importance of highlighting these aspects is as follows: it further shows that two individuals can be in exactly the same material setting, but yet their semiotic environments differ. This is further evidence that the two need to be differentiated from one another.

The first of these two differences can be discussed more briefly, as it is relatively simple. When comparing humans with other species in the natural world, we share the ability to sense and react to many of the same stimulus types. However, there are certain types of sensory stimulation that seem to have no analogous type in human cognition: for example echolocation and electrolocation (Rieke et al. 1997: 231). The case of electrolocation is particularly interesting for our purposes here, as it seems to involve the creation of a semiotic "reality" by means that have no analogous counterpart in human sensory systems. Certain species of African fish have the ability to produce weak electrical signals, using a specialized "electrical organ", in order to locate objects in their immediate environment, and furthermore have the ability to determine the size, shape and distance of such objects (c.f. von der Emde, 1999 for discussion). Electroreceptor cells in these specialized organs react, not to wavelengths of light or acoustic signals, but instead to changes in the voltage across the skin as these signals are emitted and then projected back to the organism. The relevant issue is to attempt to determine whether such a different modality of sensing results in a different conscious, meaningful experience that would be inaccessible to organisms without this modality of sensation. Intuitively, it seems as though they do, but answering such questions can be notoriously difficult (Nagel, 1974). This argument rests upon the premise that the means of sensing external stimuli, be it by sight, smell, sound, electrical signals, etc., results in a different qualitative experience, which then goes on to shape behaviour.

Secondly, even in cases where the sensory abilities of different species are more-or-less analogous to each other, the range within a certain type of stimulus that are meaningful will differ across species. Everyone who has ever had a pet dog, for instance, can attest to their reactions to various sensory stimuli that are "invisible" to humans. A widely studied sensory system is the visual system of the fly. Studies carried out on the fruit fly *Dropsophila melanogaster* have shown that it is sensitive to a range of the visual spectrum different to that of humans. Humans are sensitive to the range of frequencies from around 400-700nm (the "visual spectrum"), whereas for the fruit fly this range is closer to 100-510nm, covering

the area of the electromagnetic spectrum from ultra-violet light to the area associated with our perception of green (Zhu, 2013).

We must be careful, however, when talking about neuronal responses to different ranges of stimuli and immediately inferring that this is evidence that such a difference amounts to such a difference being meaningful for the organism. As noted by Rieke et al. (1997: 204), if the properties of auditory neuron responses closely match the behavioural ability in humans to distinguish between different signals, then our precision in auditory discrimination should be an order of magnitude better than it seems to be from what is observed behaviourally. Our actual ability to discriminate between auditory stimuli is within the range of a 1-3Hz change in frequency, when the overall frequency is around 10^3 Hz (*ibid*: 205). Our understanding of the properties of neural responses however, would predict that we should be able to notice differences of around 0.1-0.3Hz. To determine whether a statistically significant distinction in neuronal responses requires the joint effort of both neuroscientists studying the properties of neurons and neuronal structures, and psychologists who study behavioural-level phenomena. That is, we need to consider the neurological signals in the context of the higher-level features they realize.

The purpose of highlighting these examples is to show that two organisms can be in exactly the same material situational setting, but yet be in different semiotic ones. The material setting, as a reminder, would be the world as it really is, and if there is a difference in signal then this amounts to a change in the material situational setting, regardless of whether this change can have any meaningful consequences for the individual themselves. To take a simple hypothetical example: imagine that a simple behavioural experiment was set up that required an individual to perform a certain action in response to the area on the electromagnetic spectrum corresponding to ultra-violet range. Now consider the success of this task for two different organisms: a fly, which was seen to have neural mechanisms sensitive to this particular frequency, and another organism, say a human, which is not. It is relatively easy to predict which of the two is going to perform better on the task. It is the fly,

simply due to the fact it is sensitive to the relative aspects of the environment that pertain to the task.[23]

Trivial as this simple example may seem, the important point is that **they both are operating within the same material situational setting.** In other words, considered as a mind-independent "reality" of the situation, they both are operating in near identical settings.[24] But only one of our two test subjects have the neurological means of turning this material (and therefore not meaningful) constituent of the situation and make it semiotic, i.e. is aware of it and therefore has the potential to enter into cognitive and behavioural activity. We have a case, then, of an identical material setting leading to two different semiotic environments, motivating the distinction made in SFL to a cognitive context.

4.3.2. Contextual Parameters in Cognitive Science

In §4.2.2 it was seen that SFL attempts to systematize the nature of semiotic context by grouping together various parameters that are thought to be important in considering language as a social semiotic. The purpose of this section is to ask whether such as approach be taken to cognitive science more generally.

4.3.2.1. What to include?

If ideas about contextual parameters and their influence on language can help us bring some order to the vast messiness of the environment, we need to have some ideas about the kinds of things that would constitute elements of this contextual level of interpretation. The guiding question throughout this section, then, will be: what kinds of factors about the environment are necessary and/or sufficient for cognitive behaviour?

The position taken in this thesis is that the features of the semiotic environment for cognitive behaviour can be grouped along the following lines:

[23] Humans, incidentally, are extremely skilled at turning the "unknowable" to the "knowable" through technological advances. We will gloss over this possibility here, but even such considerations should not affect the validity of this discussion, as it simply involves the transformation of a non-semiotic stimulus into a semiotic one.

[24] Clearly they will not be completely identical, since the temperature, time, humidity may differ, but concerning the stimulus relevant for the task the material setting would be identical in this respect.

1. The objects in the semiotic environment and their properties, and relations between them
2. Relations between entities in the semiotic environment and the individual itself.
3. The set of goals that the organism currently has: the possible situation that it wants to bring about.

It should be noted that although the idea of giving contextual parameters for cognitive science finds its inspiration from the work on contextual modelling in SFL, these three categories are not directly analogous to *field*, *tenor* and *mode*. This is because, although many cognitive abilities may indeed by dependent upon the *social* situation, this does not seem directly applicable in every case. Also, often cognitive behavioural experiments do not take place in a setting that would be 'social' in the terms conceived of in SFL. This means, that if we are to link the work presented here to much of the work already performed in cognitive psychology and cognitive neuroscience, we must downplay this social aspect for the moment.

Although, the above set of parameters are not direct translations of the terms found in SFL, they are drawn from an, admittedly loose, interpretation of the three contextual parameters are taken to model. To show this, consider the broad definitions of *field*, *tenor* and *mode* again (Halliday and Hasan, 1985: 12, repeated from §4.2.3 above):

> **Field** refers to what is happening, to the nature of the social action that is taking place: what it is that the participants are engaged in, in which language figures as some essential component.
>
> **Tenor** refers to who is taking part, to the nature of the participants, their statuses and roles: what kinds of role relationship obtain, including permanent and temporary relationships of one kind or another, both the types of speech roles they are taking on in the dialogue and the whole cluster of socially significant meanings in which they are involved.

> **Mode** refers to what part language is playing, what it is that the participants are expecting language to do for them in the situation: the symbolic organisation of the text, the status it has, and its function in the context.

In general, then, *tenor* seems to be operating over relations, albeit only relations of a certain kind. *Mode* maps aspects along the thematic lines of *function* and *goal* as indicated by the expectations that speakers have that language will do something for them. *Field*, on the other hand, seems too complicated to give a simple label to. By denoting the properties of the action itself, it would consist of the various objects and entities in the environment, but it also includes relations that are non-social in nature. In the initial hypothesis for environmental parameters given above, all relations, be they social or otherwise, were separated into their own parameter. The reasons for this will be given below when a set of experiments from the literature are discussed.

Postulating some further theoretical distinctions into one's work needs motivation; specifically its needs to be shown that there is some practical use in doing so. So the question that will be considered now is: assuming that some identification and grouping of environmental parameters can be given, why should we do it? What can be gained by doing so? As well as highlighting that some kind of grouping of environmental parameters is possible, then, the discussion of the following group of cognitive phenomena and studies will also need to focus upon what reasons we have for thinking that such groupings would be *useful*.

In order to determine whether the initial attempt at identifying the relevant sets of contextual parameters is correct or sufficient, the rest of this section will consider a group of experiments from the psychological and cognitive neuroscientific literature, and ask what factors of the environment seem to be useful in influencing success on the reported tasks. In this way, we can begin to see what factors seem to be relevant to cognition, and whether any identifiable groupings can be unearthed.

4.3.2.2. Case Study 1: Attention Studies

In recent times, neurocognitive research has broadened its horizons to study not just the "classical" cognitive processes of memory, thought, belief, vision and the other types of sensory processing, but also what may be called *metacognitive* processes, or those cognitive processes that have an effect on the success or efficiency of others. One such group of experiments, in both psychology and recently also in cognitive neuroscience, has looked into the role that attention has had on cognitive processes, and the external factors that lead to subjects exhibiting increased "attentional effort". This first case study of environment–organism interactions will look at some of these affects.

The reasons for choosing this particular set of experiments to form part of this discussion are as follows: (1) the phenomenon has been discussed with reference to the environment, as well as being studied from both a psychological and neurological perspective, enabling us to talk about it at many different levels of interpretation, and (2) a variety of different factors have been argued to be relevant for understanding the effects on behaviour. Some of these are properly seen as existing outside the organism at the environmental level, while others seem to exist more at the interface between the external and the internal. This will present us with many different kinds of environmental and individual based factors for attention, allowing for any broad categories to emerge, and allowing us to place these at the appropriate *level of interpretation*.

A sustained level of interest into attention studies has produced both an identification of the factors that can influence our attention levels, and also some general psychological theories about how best to conceive of this metacognitive phenomenon, and some studies about the potential neurological consequences of the phenomenon.

Initially then, a wide-ranging set of factors have been thought to influence attention and, as a result, are thought to influence cognitive processing and behavioural success. Some such factors include *task difficulty* (Kahneman, 1973; Sunaert et al. 2000; Washburn and Thompson-Putney, 2001; Huang and Pashler, 2005), *object novelty* (Berlyne, 1960; Daffner, Scinto et al., 2003; Daffner, Ryan et al., 2006, Wu and Huberman, 2007), and *distraction effects* (Sarter et al. 2006; Demeter et al. 2010). Considerations of space only permit me to

discuss the impact of the above factors briefly, but hopefully the following discussion should suffice to present the overall picture.

Task difficulty is known to have an important role for attention study. This has perhaps been best shown through the means of visual search tasks. The foundations of a visual search task are exemplified through the *Where's Wally* series of books; a participant is tasked with locating a given target object in a visual field containing a number of distractor elements. *Where's Wally* is actually a rather extreme version of a visual search task; the kinds of tasks used in psychological experiments are usually much simpler. One of the most well-discussed factors that can influence performance on such tasks is the number of features that distinguish the target object from the distractors. First discussed by Treisman and Gelade (1980), if the target object and the distractor object differ in only one feature, then the time taken to find the object does not significantly increase regardless of the number of distractor objects present. A case of this would be finding a green circle amongst any number of red circles of the same size. In cases, however, where the target object differs in more than one feature, the reaction times measured in the response significantly increase as the number of distractor objects increases (see Huang and Pashler, 2005 for discussion). An example of this would be attempting to locate the letter 'T' amongst a number of 'L' letters, which are at different angles of orientation. Here we have the target differing in two respects: the letter shape and the angle of orientation.

However, it was soon discovered that this simple division based upon the number of differing features proved to be too simple. Carter and Carter (1981), for instance discovered that a similar correlation between feature differences and reaction time for only one differing feature, but in cases where the nature of the difference was more subtle, no such affect was observed. Specifically, they argued that in the case of differences in colour, the same affect seen with multiple feature differences was also seen when the colour of the target object is more similar to the distractors. This would be, for instance, the difference between trying to find a pink circle amongst a number of red circles and trying to find the green circle as above. In the former, but not the latter, the reaction time of the participants increases in line with the size of the distractor set. Similar results concerning subtle differences were recorded by Huang and Pashler (2005) as they related to slight differences in size.

The purpose of highlighting these studies is to show that cognitive performance on these kinds of visual search tasks are influenced by factors in the semiotic environment of the individual. Specifically, it seems that in these cases it is not primarily the features of the objects themselves in the environment that are influencing the cognitive performance, but **instead the relations between the objects in terms of the features in which they differ.** It is not the features themselves that make the visual task more difficult to perform, but the way that the features of the target object differ in relation to the others in the scene.

Another factor that has been argued to affect attention in cognitive studies is *object novelty*. In humans, it seems that we have a predisposition to pay attention to, and explore visually or otherwise novel events that enter our sensory field (Berlyne, 1960; Sokolov, 1963). In the case of mammals, where the affect has been studied most extensively, and where common neurophysiological features of the frontal lobe to mediate attention to novel stimuli have been argued to exist (Mesulam, 1998), this factor enables animals to successfully adapt to a fast-changing environment, and increase chances of survival (Daffner et al., 1998). Attention to novel objects, in comparison to already sensed or experienced objects has been shown to play an important role in child development. Early experiments by Fantz (1964) showed that infants from the ages of 2-6 months old were more likely to focus their attention on presented novel stimuli, as indicated through eye-gazing. Their attention decreased on repeated showings of the same stimulus. Similar considerations have also been argued to be important in language development (Kersten and Smith, 2002).

In terms of helping to identify what kinds of environmental variables may be appropriate for this study, the presence of a novel object can again be considered to be a relation. However, in this case it is not a relation between objects in the semiotic environment, but instead this is a relation between an object and the individual themselves. Factors of novelty are not features that are inherent to any particular object *per se*, but **instead the feature arises out of a relationship between the individual and an object.**

The world that organisms interact with is usually not as directed towards a task as it is in carefully designed experimental situations. In many real-life cases, there is a whole wealth of events taking place in any one scene. A large part of any successful cognitive behaviour is the ability to focus upon the aspects of the semiotic environment that are relevant, or useful, to

the task at hand, and by the same token ignore any irrelevant aspects that may prove a distraction. In light of this, research has been carried out on the role that distracting elements can have on attentional studies. One such study is presented by Demeter et al (2010), where the authors aimed to uncover some of the neurological correlates associated with performing a simple visual attention task under distracting circumstances. Their experimental set-up was quite simple; the participants were tasked with paying attention to a display and determining whether a target object appeared on the screen or not. An auditory stimulus signalled to the participants to either answer (by button-press) whether they had seen the target object or not. The experimental set up is diagrammed below (taken from Demeter et al. (2010: 1520).

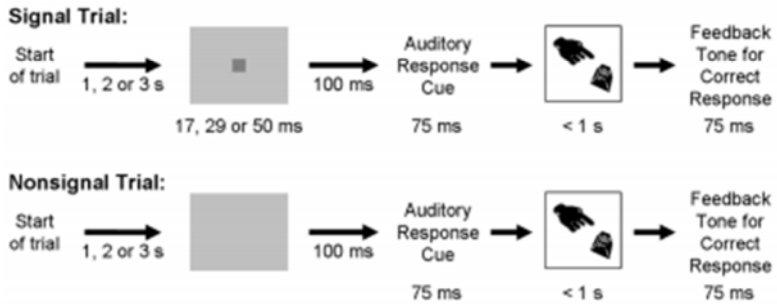


Figure 4.7: Experimental set up in Demeter et al. (2010). In each trial, after a varying length of time (1, 2 or 3 seconds) the target object (dark square) either appeared or did not appear. The duration of the object presence also varied, remaining on the screen in signal trials for either 17, 29, or 50 ms. After a constant delay, an auditory cue was played, signalling to participants to record their response to whether the target object was present or not.

In order to investigate the factors that any distractor effects may have, as well as having signal and non-signal trials, the investigators also incorporated distractor and non-distractor versions of the above two types. On the non-distractor versions, the target object was presented (or failed to appear) on a non-changing silver background. On the distractor versions, the background screen instead alternated between silver and black at a frequency of 10Hz. At a behavioural level, the authors discovered that the presence of the distracting background led to reduced performance on the behavioural task (*ibid*: 1523). This was in line with previous experiments conducted by Demeter et al. (2008). There was also a slight effect shown for the length of time the target object was presented for, but

only in the distractor conditions. That is, in cases where the object was presented, or not presented, on a stable background there were less errors in (non-) identification, but when a unstable background was presented, the shorter the target was presented for, the worse the performance on average.

The picture, then, is that there is a mass of factors that are thought to be important in determining this metacognitive ability. This, of course, is to be expected when considering cognitive capabilities as heavily influenced by both environmental and internal factors. The question facing us now is whether they can be grouped according to any discernible categories that could form the basis for any parameters. Two such groupings seem to arise out of the discussion of attention studies: firstly, we have factors that either seem to do with the properties of the external entities themselves[25], or the relationship between entities. Examples of these include the presence of distractor effects, whose influence on attention studies come from properties that are inherent to the entity itself, or the feature effects seen in the first studies discussed above. In these studies the factor influencing performance on the tasks were not factors that were intrinsic to the objects themselves, but instead they arose out of consideration of the relations between the various objects that constituted the semiotic environment. To recall an example from earlier, where there was an interaction between the number of features that differed between the target object and distractor objects in a visual search task determined whether the number of distracting objects influenced performance. This is not a feature that can be considered to be inherent to the target object itself, but can only be gained in comparison with other objects. However, it remains a relational effect between objects external to the individual. Let's call both of these factors as instances of a wider contextual parameter dubbed *objective* here.

The second grouping is again a relational one, but in this case it is not between external objects, but instead between objects and the individual themselves. The case discussed here was that of *object novelty*, which was also seen to be a factor influencing performance on attention studies. Let's call these factors *relational*.

[25]I am using the term "entities" in an extremely broad sense. In the way I intend to use it here, it could refer to objects, processes, properties, etc. It is a "catch-all" term for the various things that are usually the focus of arguments in philosophy. There is much to do in determining whether any one of these entities has precedent over the other, but I will leave that question open.

We have seen evidence for two contextual parameters, which will hopefully help experimenters and theorists treat environmental factors in cognitive experiments in a more systematic manner. In the next section a number of studies that have looked into the role of *tasks* or *goals* in cognitive performance.

4.3.2.3. Case Study 2: Goal-oriented studies

Behaviour and cognition is not completely random and unguided. In most, if not all, cases organisms behave in a certain way or undergo a certain set of cognitive processes with the aim of **achieving something**. In many respects, this section intersects with the previous one, since as will be shown the particular goal of a task can have an effect on attentional affects one sees in psychological experiments. However, given the prominent view that goals will have in the account presented here, it is worth discussing it independently.

Intuitively, the idea that **goals** would be important in the study of cognition (at whatever level of interpretation) seems highly plausible. As summed up by Brass and von Cramon (2004: 8847):

> *The human ability to selectively adapt to behaviour to different situations and tasks is a crucial requirement in our daily life […]. Our environment generally affords a variety of different behavioural options, only some of which are adaptive in any given situation. Therefore, individuals are required to select the contextual information that determines adaptive, appropriate behaviour. We refer to this control process as the selection of task-relevant information.*

In their study, the authors attempted to discover some of the neurobiological correlates that may underlie the task-relevant information selection procedure mentioned above. They argued that activation in the posterior inferior frontal sulcus, in line with previous results, and additionally in the inferior parietal sulcus was associated with selection of stimulus that was **related to the task set**, and not to features of the stimulus considered in isolation.

Apart from intuitive arguments giving a prominent role to *goals* as pertaining to one of the contextual parameters under discussion here, it would bolster the case if some experimental

evidence could be found whereby the nature of the task or goal of a study affects some aspect of cognitive processing. The rest of this section will survey such experimental considerations.

Simons and Chabris (1999) undertook a study into the ways in which the presence of a certain task can affect the attentional processes of participants. They performed a study into the phenomenon known as *inattentional blindness*, whereby focussing on one aspect of the environment results in co-occurring unexpected events being ignored. In this experiment, 228 participants were asked to watch a video of 75 seconds in length. In the video, two teams of three players, differentiated by either wearing a white or black shirt, passed a basketball between members of their own team. Passes between players could either be aerial (without bouncing), or bounced. The players were also not stationary, but instead moved around and made other motions while the basketball was being passed between them. There were a wide number of different independent variables in this experiment. One condition was whether the participant was asked to keep track of the white team, or the black team. A second condition was whether the participant was tasked with counting both kinds of passes together (easy task), or to simultaneously keep one count of aerial passes, and another of bounced passes (hard task). Thirdly, two different kinds of unexpected events were filmed. In the first of these, a woman holding an opened umbrella walked from left-to-right across the screen, passing between the basketball players. In the second case, a man dressed in a gorilla costume entered the screen in the same manner. Both unexpected events lasted for 10 seconds. The final condition pertained to the method of filming. In one case, each set of players, and the actor for the unexpected event were filmed separately, with these three separate films being combined later. In the second case, all actors were filmed at the same time. Each participant was, therefore, sorted into one of the sixteen groups that arise out of the combination of the above experimental parameters.

Instances of *inattentional blindness* were reported across all conditions: overall, 54% of participants included in the analysis (N=192) noticed the unexpected events, while 46% failed to notice them, confirming earlier results Neisser and Becklen (1975). In terms of task related effects, however, there were further conclusions that were drawn from the results of the study. In the ‘easy task’ condition, 64% of participants reported seeing the unexpected event, compared to 45% in the ‘hard task’ condition. A Chi-square analysis showed that this

was a statistically significant difference. The 'hard task' condition was also shown to be more difficult for the participants by the authors noting that the reported counts of passes in this condition exhibited a higher degree was variability than the 'easy' task condition. Furthermore, the type of unexpected event also had an effect on the participants' performance: the woman holding an umbrella was noticed far more often than the gorilla (65% versus 44%), and, again, this difference was deemed to be significant.

There are a few important aspects of the above study as they relate to the discussion here. Firstly, the significant effect of task type on *inattentional blindness* points to an important factor concerning the nature of the semiotic environment. In both of these cases, the participants were in the same MSS, that is, they were watching the same video, with the same actors, the same movement of the ball, and the same unexpected event. The only difference was the task that the two groups were given, and this resulted in a real psychological difference. This motivates the inclusion of the notion of *goal* in the description of the semiotic environment.

The second of the above cited results is also interesting for the purposes of the discussion here. The important factor in this case is not related to the idea of a *goal* but certain facts about the objects in the *semiotic environment* and their relationship to the individuals partaking in the study. The difference between the umbrella-carrying woman and the gorilla was argued to be due to expectations on the individuals' part about what is expected in such a situation. The more 'typical' event, that is seeing a person walking through a group of other people, resulted in a better performance on the task than a less 'typical' event. Here, again, we have an example of a *relational environmental parameter*; the relevant factor here is the relation of an entity and the individual involved in the cognitive act.

As well as affecting psychological factors of the individual, factors relating to the task at hand have been shown to have a number of neurophysiological consequences. An overview of one such study was discussed above, where certain brain regions have been argued to be involved in the selection of task-relevant information, and so their activity is partially driven by the nature of the task. In this case, different tasks would have different neurological consequences. Further neurobiological evidence for the inclusion of a task or goal based contextual parameter comes from the experiments undertaken by Scheich et al. (2007). In

one of these experiments, they measured neuronal activity, via fMRI, in the auditory cortex in a number of different auditory discrimination tasks. They noted a different response in the secondary auditory cortical areas (T3) of the *planum temporale*. The two tasks that were given to the human subjects were (1) to correctly identify a directional frequency modulated auditory signal as either "rising", or "falling", and (2) to identify whether the auditory signal could be categorized as "long" or "short". Crucially, both of these tasks were undertaken using the same set of stimulus material (*ibid*: 218). The results showed that there was a correlation between the participants undertaking the first of these tasks and an increased blood oxygenated level dependent (BOLD) signal in the T3 area in the right hemisphere, while on the second trial the increased BOLD signal was located in the T3 area of the left hemisphere. Not only were the differing responses found in different hemispheres, but also the nature of the responses differed. The more subtle features of the first task, i.e. when the "fall" or "rise" of the auditory sweep were less pronounced, correlated with a higher response in the right T3 areas, while the reverse was true for the duration discriminatory task.

Tasks or goals, then, seem to be an important contextual factor when considering the cognitive abilities of organisms. This can be seen by noted differences at both the psychological and neurological levels of interpretation. This motivates containing such a consideration in the parameterized description of the semiotic context under consideration here.

4.3.2.4. Environmental Parameters: Objective, Relational and Goal

The proposal, then, is that the study of the interaction between the organism and its environment can be structured according to (at least) three contextual parameters. The labels for these are: *objective*, *relational* and *goal*. Their details are summed up in the table below:

	Objective	Relational	Goal
Broad Question	What is in the environment? What are the relations between the objects in the environment?	What are the relations between the individual and the objects in the environment?	What is the purpose of the task?
Specifics	Any cognitive task takes place with reference to certain objects/ entities in the external world. Relations between these objects are also included	In certain tasks the relationship between the individual and objects the external environment and the individual will have an effect on success	Describing the nature of the situation that the organism would like to bring about.

The above groupings are supposed to capture the following aspects of the external environment in which organisms interact. Firstly, in any cognitive or behavioural task there will be certain **objects** that are important and relevant in the situation. These objects will have certain properties, and there will be a certain set of relations that can be stated between them. Secondly, there will be a certain set of **relations** that hold between the **individual** and those objects. In this discussion, an example of *novelty* was discussed but there are surely more. Finally, in most situations there will be an explicit or implicit **goal**, a certain event that the individual would like to bring about. All of these aspects will guide and shape both the psychological and neurological output of the lower systems, causing certain instances to be more or less likely to occur.

With these factors in mind, this chapter finally turns to treating the *semiotic environment* as a *system* by describing it along the same theoretical dimensions previously discussed.

4.3.3. The contextual stratum in relation to the others.

4.3.3.1. Theoretical dimensions at this level of interpretation

To recall, one of the claimed idiosyncrasies of the SFL approach to context was that it is amenable to the same theoretical framing as the other levels within language. Although modelling of the context in SFL should theoretically be possible along all of the theoretical dimensions discussed above, as was seen, progress in terms of *delicacy* and *rank* lags behind

that of *instantiation* and *stratification*. Therefore, this thesis will concentrate on only these two, as there is a risk of applying an under-developed idea in one area to another. Hopefully though, as theoretical progress continues in one of these disciplines, these gaps can be overcome.

Firstly, then, it is necessary to discuss the stratum *of semiotic* environment in its relation to the other lower strata discussed in the previous chapter that were argued to be relevant for the study of cognitive behaviour. In SFL, one of the more general thoughts about the relation between the text and the context is that the context places a constraint on the possible options that are available at the linguistic levels of interpretation. In a linguistic context, this was seen above in the discussion of the field-motivated ideational semantic choices that were found in the domain of cooking recipes. The presence of a certain contextual configuration skewed the probabilities of certain types of figures and sequences appearing, when compared to their distribution in the system as a whole. The fact that this picture can be transferred relatively easily over to the context under discussion here is not really surprising. Once one considers psychology and neuroscience as a system of "options" instead of **purely** a number of different structures (be they symbolic structures, sequences of action potentials, or whatever structure your particular theory of cognition presents you with), describing the various possible states the individual can be in, then this constraining picture becomes wholly natural. The various experiments described above all amount to this picture, when certain objective, relational or goal factors result in certain regularities at the lower cognitive levels, be they regular actions or common activation in certain cortical regions. However, this is also its weakness. If the claim only amounts to an ascription that certain environmental factors will result in regularities in psychological or neurological events, then this is hardly surprising. In fact, considered in this way, this is at the heart of much neuroimaging work currently taking place in modern cognitive neuroscience, in that the aim is to see how, or whether, certain cognitive functions can be localized to particular brain regions.

In order to make slightly stronger claims, then, it is important to realise that realization is not the only dimension in operation here. The other dimension of *instantiation* also needs to be considered, and it is possible that considering this other dimension, and how it interacts with the idea of constraining choices, can lead to any stronger predictions.

Instantiation may also prove to prove to play an important role in theorizing about the environment. Giving a description about the contextual parameters for one particular experiment may prove useful to some extent, but little can be gained from simply considering these descriptions alone, or in isolation from one another. The cline of instantiation provides us with one means of beginning to compare different environments, as well as providing the means to determining whether the contextual parameters postulated here have any practical use and as such deserve inclusion in any cognitive theory. To recall, the cline of instantiation determines *levels of regularity* at various points on the stratificational framework. In SFL, the two extreme ends on the scale are referred to as the *context of situation* (analogous to the instance/ text end at the language internal levels of interpretation), and the *context of culture*, which considers the system of all actual and possible situations that could influence the linguistic choices made in the linguistic system. Keeping in line with how instantiation is considered throughout the stratified system, in between these two extreme poles are an indeterminable number of mid-points. The situation in SFL is diagrammed below:

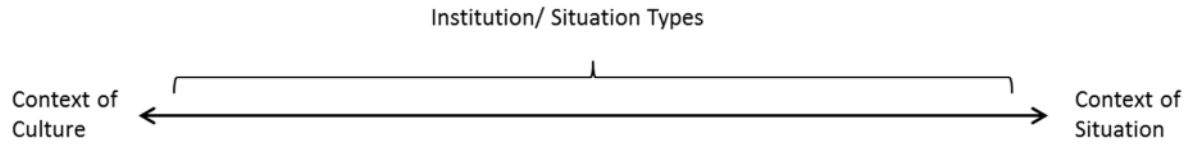


Figure 4.8: Cline of instantiation at the contextual level in SFL

Just as the context of situation for a particular text can be described as a combination of different values for the contextual parameters, the hypothesis here is that the same is possible, and also advantageous, in wider cognitive science. Instead of the parameters of *field*, *tenor* and *mode*, however, the hypothesis developed in the preceding sections has been one possible set of useful parameters are the *objective* the *relational* and the *goal.* With these three parameters in mind, it is therefore possible to conceive of an instantiational cline operating at the level of semiotic environment in cognitive science, in much the same way as one has been utilized in SFL.

On such a dimension, a few immediate points need to be identifiable. In the first instance, the nature of the two extreme poles on such a cline needs to be identified. The analogous "instance end" would account for a description of a particular semiotic environment along

the lines given above; that is, a description of the environment in terms of what is present and the relations between them, the relations between said objects and the individual and the goal or situation that the individual is trying to bring about. Let's call this end of the cline *environmental instance*. On the other extreme we have a theoretical construct corresponding to the system of all possible combinations of the various values for the three contextual parameters seen above. Let's call this end the *environmental system.* It is not possible, at this stage, to begin to envision what such a system may look like, just as the vision of a systematic description of a contextual system as pertaining to language is still far from coming to fruition, although as mentioned earlier in this chapter there have been some attempts to begin to flesh out the details. This picture will only be illuminated if (1) the environmental description postulated in this chapter is shown to be of use in describing the constraining force that different environmental factors may have on the many different possible options at both the psychological and neurological level in cognitive science, and, if this can be shown to be the case, (2) a large investiture in describing individual semiotic environments along these lines, so the relevant factors can be uncovered. Given the exploratory, and overall preliminary, nature of this thesis, this will not be attempted in these pages. However, the environmental system as described above remains a hypothetical theoretical construct. For this reason, our discussion will mainly focus upon the instance end of such a cline. In the next chapter, we will see ways in which such a system may be built up from experiments run and descriptions attempted.

With these two ends in mind, the postulated dimension of instantiation at this level of interpretation can be preliminary sketched as follows:

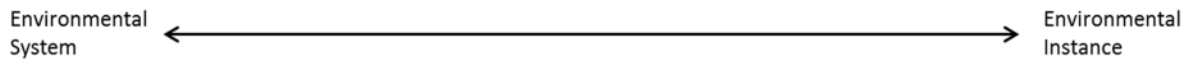


Figure 4.9: *Preliminary cline of instantiation at the stratum of semiotic environment.*

Of course, the overall point of instantiation is that one should not consider the instance and system ends to be a member of a dichotomy. It is a true scale, in the sense that mid-points exist between these two points. As mentioned, in this discussion we will only be focussing on these mid-points as viewed from the instance end of the cline, and these will be called *environmental types*. This gives us the picture as below:

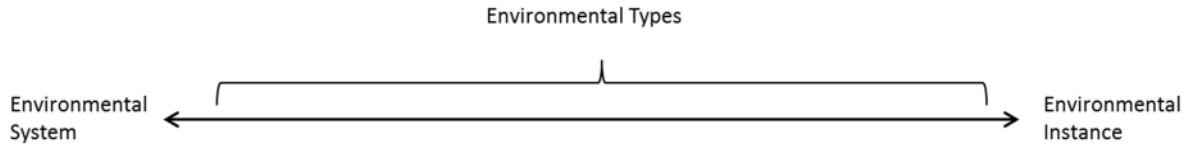


Figure 4.10: Cline of instantiation at the stratum of semiotic environment

What do these mid-points actually relate to? In the same way that mid-points at the other levels of interpretation refer to the relevant levels of regularity, these environmental types would refer to theoretical constructs created by grouping individual environmental instances in terms of common values for the environmental parameters. If any given situation can be described as a collection of these parameters, be they in the lab or descriptions of more "natural" events that can be observed without of these highly controlled situations, then it is possible to begin to look for common properties between said situations and to group situations accordingly. This is the essence of what the idea of *situation type* above is meant to convey, and it is in keeping with how the term is used with regards to the instantiation cline at the levels of interpretation within the stratificational framework.

The usefulness of considering an environmental cline of instantiation in any cognitive theory is that it allows one to move beyond the particular situation of any given experiment and fit it in to other results that may be gathered through other descriptions or experiments. Different environmental instances can be prepared, and on the basis of shared or common features form the theoretical notion of an environmental type. This collection of environments can then be treated as a single type, enabling investigation into the regular effects that can be found at the lower levels of interpretation. Furthermore, the ascription of contextual parameters allows one to not treat all entities in the environment as of the same kind, but by bringing some structure to what can initially look like a heterogeneous mess, provides the researcher with some framework on which to justify such groupings.

4.4. The story so far

This thesis began with the following argument for the approach taken in these pages: the brain, many people seem to agree, produces meaningful events, or as it has been considered here, is one part of a *meaning-making system*. If this is true, then it can therefore be modelled as a *semiotic* system, since all meaning-making systems are by definition also

semiotic. From this, the possibility arises of modelling cognition along the lines taken by SFL as applied to language, as the broad theoretical tools used therein have been hypothesized to hold not just for language but for semiotic systems more generally. The goal of this thesis was to take the idea of the brain as a semiotic system seriously, attempt to reconsider its study along the lines taken by SFL, and ultimately answer the question: **what picture of cognition do the principles of SFL present us with, what would such a theory look like, is it of any use to us, and how does it compare to previous discussions about the nature and study of cognition?**

Nearly three chapters later, we can begin to arrive at such an answer. Many of the building blocks towards a description of an SFL-inspired cognitive framework have now been laid down. To recap, this thesis has discussed the following possible applications. These can be split into two groups: (1) the levels of analysis that are potentially useful in theorizing about cognition, and (2) ways to think about environment-organism interactions:

Levels of analysis in cognitive science (from SFL)

1. *Stratification* – describing cognitive systems at various **levels of interpretation**. It was argued that this can be considered to be the identification of relations between units that increase in spatial and temporal scope. Levels of interpretation are linked by the relation of **realization**. This is a bi-directional probabilistic link, by which higher-level features are taken to "activate" features at the level below, and lower level features are considered to "construe/ encode" higher level ones

2. *Rank* – describes the **levels of constituency** that are thought to be important at each level of interpretation. The hypothesis is that a relevant set of constituents is available at each level of interpretation. These interact with the dimensions of *stratification* and *delicacy* to posit the potential range of meaning-making structures important in each case.

3. *Delicacy* – describes each level of interpretation in terms of **levels of specificity**. The importance of this dimension is that it enables the modeller

to state how much detail in a particular description can be ignored (due to "noise") while not overlooking potentially important information.

4. *Instantiation* – sets up a cline consisting of **levels of regularity**. It provides the link between individual instances and the theoretical description of the system as a whole. Between these two extremes lie an indeterminable number of "mid-points". These can be considered in two ways: either as generalizations over a number of observed instances (viewing "from the instance end", or as a description of the potential for further meaning in similar situations (viewing "from the system end", mediated by constraining factors from context).

The summary and discussion in §3 details how these proposals differ from previous accounts. It was seen that many, but not all, of these hierarchies are already present in previous approaches. However, the details of the hierarchies may differ in this case, and no framework so far enables one to bring all of these together into one coherent picture.

In terms of environment-organism interactions, the following aspects have been discussed in this chapter:

<u>Environment-Organism Interactions</u>

1. *Material vs. Semiotic Environment* – the need to make a distinction between the world "as it really is", and those aspects of the world that are important for the behaving organism. The former of these places a constraining (but not determining role) on what can be included in the second. The semiotic environment can also include entities are that are not present in the immediate material setting.

2. *Systematizing the environment in terms of environmental parameters* – this involves moving beyond the conception of the environment in terms of a heterogeneous "mess" of entities, and instead grouping the features of the semiotic environment in terms of certain themes. The "groupings" that were

claimed to be important for cognitive science were: (1) the objects that are in the semiotic environment and their properties. One potential property that was discussed was that of *prominence*, (2) the relations between the objects, and (3) the goal of the situation that has some beneficial effect for the organism.

3. *Theorizing over the environment* – One of the hallmarks of systemic linguistic theory is that the level of context is described in the same terms as the "language-internal" levels that are the usual object of description in linguistic theory. This means that the same theoretical dimensions previously discussed can also be applied to the context in which the behaviour occurs in. By this means, it is possible to start to bring some order to what can initially seem like a vast and complicated set of entities.

4. *Constraints on the neuronal meaning-making potential* – neuronal systems being as complicated as they are, the details of which we are only just starting to fully comprehend, it is sensible to assume that they are capable of a vast number of meaningful events. However, this vast potential does not instantiate itself in an unconstrained fashion. A large part of determining what output we are likely to see comes from the values of the environment placing constraints on possible outputs, by skewing probabilities on what will occur at the "lower" strata.

So far, the discussion that has led us to the identification of the above possible applications of systemic theory on cognitive science has operated at an abstract, conceptual level. The arguments for these positions have mainly taken the form of a philosophically-orientated discussion of some of the grounding assumptions of cognitive theory. It is useful, then to ground this far-reaching and lengthy discussion in the actual day-to-day work of cognitive neuroscientists. In the concluding chapter, this thesis will detail a particular example of a neuropsychological experiment, the Morris Water maze, which has been used extensively to study aspects of spatial navigation and memory. This will enable a concrete example of the above ideas to be stated.

Chapter 5

Putting the ideas to use: on abstraction in systemic functional linguistics and building the framework from the data.

"Piglet was so excited at the idea of being Useful that he forgot to be frightened any more, and when Rabbit went on to say that Kangas were only Fierce during the winter months, being at other times of an Affectionate Disposition, he could hardly sit still, he was so eager to begin being useful at once."

A.A. Milne. Winnie-the-Pooh.

5.1. Introduction

The previous chapters in this work have attempted to develop and defend a framework for cognitive science based upon the principles of SFL. So far, the focus has been detailing this proposal at a conceptual level, defending and motivating the ideas through psychological and neurological phenomena, with little focus upon how you might start developing the framework. This chapter aims to provide some means of doing so. It hopes to achieve this aim is by critically assessing the meaning and role of the theoretical term *abstraction* in SFL. This is important because, as we will see, this term has been associated with many of the theoretical dimensions that have taken centre stage so far in this work. This means of *abstraction*, and its role in theory creation, will set the groundwork for linking data gathered from descriptions and experiments to the overall model.

Much of this thesis has been concerning the theoretical dimensions of *stratification*, *instantiation*, *delicacy*, and *rank*. The first three of these dimensions have all, at various points in time, been described as operating over a dimension of *abstraction*[26]. Therefore, it would be worth investigating this notion in more depth, to see whether any information can be gained into how to start building the theoretical framework that has been discussed and

[26] As I am using the term in this chapter, I intend to convey the meaning of abstraction as a process. In some accounts abstraction is considered to be a process, in other times it has a nominal use. If I intend to convey the nominal meaning, I will use the term *abstract object*.

outlined in the previous chapters. Firstly, *stratification* is linked to the idea of abstraction by Halliday and Matthiessen (1999: 4, emphasis added), who state that "language […] is a resource organized into three strata according to *order of* **abstraction**". Out of all the theoretical dimensions mentioned above, stratification is the one that is most often linked to the idea of *abstraction*. Similar links can be found in the works of Martin (2013), Taverniers (2011), Matthiessen (2007a.), Halliday (1961; 1981), and Hasan (1995). The dimension of *instantiation* is also linked to this notion. As well as being one of the distinct orders of *abstraction* mentioned in Halliday's (1961) foundational work in SFL, Matthiessen (2012: 450) makes the link between i*nstantiation* and *abstraction* in a reinterpretation of Fariclough's Critical Discourse Analysis (Fairclough, 1995) along the lines of SFL. He clearly identifies the potential end of the cline of instantiation as more *abstract* then the instance end, which is correspondingly identified as *concrete*. Delicacy, as well, has occasionally been identified as operating over an order of abstraction, most evidently in Halliday (1961/ 2003: 41). Considering them all together, Halliday and Matthiessen (1999: 327, emphasis added) say the following:

> *It is important to make quite explicit the distinction among the three concepts of realization, delicacy and instantiation, since* ***each of these is a distinct scale of abstraction*****.**

Considering the above quote, there is an ambiguity that can be read into the identification of abstraction in these cases. It could mean any of the following:

A. Stratification, delicacy and instantiation are all scales that utilize the **same** process of abstraction, but in each case operating over **different** units.
B. Each of these scales utilizes **different** processes of abstraction, and it operates over **different** units
C. Each of these scales utilizes **different** processes of abstraction, and in each case operating over the **same** units.[27]

For the purposes of this discussion, we will focus less on the nature of the units in each case, and so the difference between B. and C. will not be focussed upon. However,

[27] The forth logical alternative, namely that each uses both the same kind of abstraction over the same units, is untenable on the assumption that we really do have three different scales in operation here.

immediately (3) seems problematic since delicacy and stratification operate over linguistic types, as presumably, while instantiationoperates over linguistic tokens. Therefore, it seems that the nature of the units in all cases is not the same. The choice between A. and B., though, still remains.

What reasons do we have for initially thinking that there might be more than one type of abstraction? After all, the same word is used in each case, so the presumption might be that the same process is being referred to. However, as we shall see in the course of this chapter, the situation surrounding the meaning of this term in the sciences and the humanities is not that simple. As a precursor to the discussion below, Barsalou (2003; 2005), in his work on psychology and symbolic cognitive science, has argued that there are six possible views on abstraction as pertaining to the formation of categories or concepts in the human mind. Additionally, Saitta and Zucker (2013: §2), in a survey of philosophy, art, computer science, psychology, linguistics and mathematics, argue that there were five possible uses of the term in these different disciplines. These were identified as:

- Information hiding
- Generalization (or a close variant of it)
- Taking a distance from the concrete world
- Keeping relevant information
- Reformulation or approximation

Most of the above points will be important in the discussion below. For the moment, though, it is enough to realize that although the same word has currency in a number of different cases, **it does not always mean exactly the same thing**. This can lead to confusion, since if one is unaware about the exact meaning of a theoretical term, then the use of said term can become compromised.

This goal of this chapter, therefore, is to interrogate the meaning of the term *abstraction* with respect to the three theoretical dimensions mentioned above, and which have been at the centre of the work in this thesis up until now. More specifically, this chapter aims to answer the following questions:

1. Do the dimensions of *stratification*, *instantiation* and *delicacy* all utilize the term *abstraction* in the same way?
2. If not, how do they differ? Is it possible to give a more precise meaning of the term in each case?
3. How does this affect the means of building the conceptual framework from actual linguistic data?

How exactly, though, does one start to enquire about the nature of a theoretical term? The method employed here is to ask what other researchers have had to say about the topic in other disciplines, as abstraction is a term (as has already been seen) that is important in many of them. Much of the ideas and thoughts discussed in this chapter come from the discipline of *philosophy of science*, since this is the discipline that has as its aim to enquire into the methods, status and consequences of scientific practices.[28] Part of their work is to enquire into subjects such as the development of scientific ideas, the status of law-like generalizations over observed phenomena, the robustness of induction, abduction and deduction in scientific accounts and the links between science and truth. In short, then, they are concerned with critically looking at the way scientists carry out their work, and highlighting potential problems or inconsistencies that can usually only be discovered through close reflection that is the hallmark of philosophical methodology (Rosenberg, 2012). Perhaps unsurprisingly, *abstraction* and related terms have attracted attention in this literature, and has occupied the thoughts of many of the major philosophical thinkers throughout history. By mining these thinkers and commentators for ideas on the subject, the hope is that a means of interrogating the term in SFL can be found.

The rest of this chapter, therefore, will be structured as follows. It begins with a short discussion building upon the methodology just outlined. After this, an overview of the relevant literature will be presented, with the aim of providing a range of meanings that have been attached to the term *abstraction* in the literature. Having identified some features of *abstraction*, the discussion will then turn to the theoretical dimensions in SFL

[28] The term "philosophy of science" is a bit of a misnomer in the way I intend to use it here. Although much work in philosophy of science concentrates on the "core" scientific studies of chemistry, physics, biology and mathematics, similar questions can be raised for disciplines such as history, art, etc. The term "philosophy of enquiry" may be more appropriate, but has no currency in the literature. It would also seem to entail that I am treating linguistics uncontroversially as a science, which is not my intention.

mentioned previously, and applies the lessons learned from the philosophically-orientated literature to tease apart the different meanings of this term in the linguistic literature under discussion. It will conclude with a *typology of abstraction types* in SFL, which aims to provide the different meanings the term can take.

5.2. Philosophy of Science and Systemic Functional Linguistics

Before this chapter turns to a discussion of abstraction, it is necessary to defend to methodology taken here, as I believe there may be some in SFL who would question its validity. The most explicit of these criticisms come from Halliday (1992c./ 2005). Halliday makes two potential criticisms that are pertinent to the present discussion. Firstly, he argues that the objects of study in the "core" scientific studies and in linguistics are fundamentally of a different kind. Secondly, he questions the viability of the philosophy of science because they operate within a highly *idealized* viewpoint, far removed from the daily practice of scientists themselves. He goes on the argue that if linguists want to learn from the methodology of scientists then this is best brought about by emulating the practices that scientists undertake, instead of "studying the models constructed in the name of philosophy of science" (Halliday, 1992c./ 2005: 200). Both of these views, if correct, would be disastrous for the perspective taken here. On the face of it, if the subject matter between the linguistic and the other sciences is radically different, then turning to discussion about the latter to inform the former seems like a non-starter. The second of these criticisms is directly the opposite of what I am proposing here, which is that we can learn from the ideas that have been constructed in the "name of philosophy of science". Initially, then, both of these points need to be addressed.

Although a full defence of the applicability of philosophy of science to linguistics (commonly referred to as the *philosophy of linguistics*, c.f. Scholtz et al., 2014, and the various papers in Kempson et al., 2012), and particularly with respect to SFL deserves far more attention that can be attempted in these pages, some initial remarks can be made. In my view, far from models being constructed in the name of philosophy of science, the models or other targets of the attention of those in this field are often the creation of the scientists themselves. Creating abstract models, rules and other types of generalizations are as much of the day-to-day practices of scientists as providing rigorous methodologies for the acquisition of new

data. Otherwise, scientists would not be able to link the results found in a particular lab on a particular day to the wider scientific situation, and therefore make their results relevant and useful for the wider scientific community. Those undertaking philosophy of science take these constructions and examine the processes that go into creating them and interrogating their validity. In short, it is not often the philosophers that create the idealized, abstract entities, but instead the scientists themselves. As an example, Cartwright (1983) provides a discussion of the difficulties in linking mathematical laws in physics with the real-world, by highlighting the problems of marrying the *ceteris paribus* foundations of physical enquiry with the multitude of causes that may act to produce a certain outcome *in actuality*.[29] Her main point is, briefly, that law-like statements in physics are only true in highly idealized situations, i.e. in other cases where all other things are equal, but that this is never found in reality. Interestingly, this is the exact same criticism that Halliday lays at the door of philosophers. From laws we create models, so in the end the laws of physics do not tell us truths about the actual phenomenon or the actual world, but instead tell us truths about the "world" that is created and described by the model themselves. The point here is not whether Cartwright is right or wrong in her assessment of physics, but simply to give a feel for the kinds of enquiries that philosophers of science partake in, and to show that the criticism made above is not warranted, since often the idealizations are not undertaken by philosophers but instead scientists themselves.

Additionally, although this work takes the view that although it cannot be assumed *a priori* that theoretical notions can transcend discipline boundaries, neither can it be immediately dismissed out of hand. One of the central premises of the work being undertaken in these pages is that such trans-disciplinary comparison and application of ideas is indeed possible, and its appropriateness rests on the use that can be made of it. Furthermore, as was hinted at in §5.1 and will be discussed in the next section, overviews of abstractions have already been undertaken, with common themes being identified, although this is not to also downplay the differences that do exist. The fact that something as wide-spread and ubiquitous as *abstraction* may have common currency across discipline boundaries has come from accounts that have based the academic process of abstraction on general properties of human cognition (Barsalou, 2003; 2005; Martinez and Huang, 2011). If, as

[29] The principle of *ceteris paribus* is usually stated as "other things being equal or held constant".

seems reasonable, at some level the same or similar cognitive capabilities are in operation by researchers in all disciplines then it is not beyond the realm of possibility that such a cross-disciplinary comparison would prove useful.

However, such a view does not directly lead to the conclusion that the practice of modelling phenomena will be the same in the various sciences and other disciplines. Halliday is correct in saying that there may be differences between language and the objects of study in the (other) sciences. If these differences amount to a difference of *type*, however, is a question that needs further elucidation and clarification. Firstly, what exactly would a difference in *type* between the subject matter of different disciplines amount to? As stated in §2.2, the perspective taken in this work is broadly of a *physicalist* orientation, that is there is only one type of "substance" in operation in the natural world, and therefore the different disciplines would all involve inquiring into different aspects of this one underlying reality. Therefore, any account of a different type between linguistics (or any of the human sciences) and the "core" sciences along the lines of a material-immaterial divide are, in my view, untenable for the broadly metaphysical reasons that have been given earlier in this thesis. Even without this consideration, any appeal to an immaterial-material divide as accounting for the type difference seems to fail once ones considers Matthiessen's (2012) orders of nature, which has been discussed at various points in this thesis (see figure 4.1, §4.2.1) In fact, in this paper Matthiessen describes the higher systems of nature as being "immaterial", which I believe is a mistake (*ibid*: 451). Consider the following argument:

(1) If system S_i is physical, then it is material.
(2) If system S_j is semiotic, then it is immaterial.
(3) If system S_m is semiotic, then it is also physical.

From (1)-(3) we are lead to the following conclusion:

(4) If system S_n is semiotic then it is both material and immaterial

The conclusion follows from a simple deduction from the three premises given. I will assume premise (1) to be true, as it seems reasonable to consider physical systems within nature to be material, if one wants to include some kind of material reality into their world-view.

Premise (2) is what one would need to propose if immateriality of language is to be drawn upon. Premise (3) has already been justified in §4.2.1, and it stems from Mathiessen's own writings on the subject. Given all of these, the conclusion is unavoidable, and it is immediately clear that this is a contradiction, as it amounts to saying that a particular system both has and does not have the same property. Therefore, we can conclude that the material-immaterial distinction cannot account for any purported difference in type between the linguistic and scientific systems under study. Without a clear idea of what a difference in type would entail, this criticism can be overcome, or at least postponed the nature of the criticism can be clarified.

At this point, there is a need to make a further clarification. The appeal to scientific practice is not intended to be an implicit claim for the superiority of those subjects over those who are usually categorized as part of the "humanities", or whose status as such is less secure. If it is discovered that SFL uses the term *abstraction* in a different way to that which is found in other subjects, then this is not meant to imply that the term is being used in an incorrect manner. Creators of theoretical categories are free to use terms as they so wish, but on the condition that its meaning is clarified if it differs from any standard usage. Nor does it prevent such uses from undergoing a process of close inspection and critiquing. The appeal to the philosophy of science here, therefore, is merely provides a means to instigate our discussion on the notion of *abstraction* in SFL. In the end, as always, the proof will be in the pudding; if the discussion can illuminate certain uncertainties or inconsistencies in SFL theory, then it will be justified.

5.3. Abstraction in the philosophy of science.

It was briefly outlined above in §5.1 that various meanings have been associated with the term *abstraction* in various disciplines. The purpose of this section is to provide a critical overview of this various ideas. In sum, the following aspects of the literature will be discussed:

1. Abstraction and the idea of "omission"
2. Abstraction and generalization
3. Abstraction and decontextualization

We will look at these three ideas in turn. As we do so, it will be seen that all of the features identified by Saitta and Zucker (2013) mentioned above will be covered.

5.3.1. Abstraction and "Omission"

Possibly one of the first influential ideas on abstraction came from Aristotle, who discussed abstraction in terms of the creation of the different scientific disciplines. For him, the various academic disciplines define their subject matter by not treating the world as a whole. Instead, they only attempt to give a description and explanation of a certain part of it. Part of the reason for why academic disciplines are abstract is, for Aristotle, because they **omit certain features** of the wholeness of reality in their efforts (Bäck, 2014: 12). As another historical example, consider the opinions of Frege (1884) on the nature of abstract objects in mathematics. For Frege, and other thinkers such as Kurt Gödel, abstract objects were to be defined as objects that **lack** certain properties that can be found in their more concrete counterparts. Dubbed the *way of negation* by Lewis (1986), Frege's view was that abstract objects had an independent ontology, but that they lacked spatial and temporal extension, as well as being devoid of causal powers.

The reason for raising these historical examples is because they highlight a common theme that has long been associated with *abstraction*: abstractions are often linked to the idea of **leaving things out**, or **lacking properties**. These ideas will be termed *omission*, and they have formed the basis of several more modern attempts to clarify the notion. For instance, Jones (2005) takes this idea of *omission* as the foundations of his theory of abstraction, and for him it is at the heart of being able to distinguish between two related terms: *abstraction* and *idealization*. On this account, both of these processes involve omitting certain features of a phenomenon or system, but the means by which they come to leave things out differs.[30]

Abstraction, on one hand, is characterized as an omission of a truth, or the model "simply [being] silent in [some] respects" (Jones, 2005: 184). Shortly we will see examples of

[30] In some cases, such as here, *abstraction* is defined in comparison to *idealization*. This chapter will have little to say about the latter of these two, and it will only be discussed in this comparative sense. This does not mean that *idealization* does not play an important part in linguistics (c.f. Ludlow, 2011; Stockhof and van Lambalgan, 2011 for discussion as it related to generative linguistics), but it is not our primary aim here.

abstractions in a simple physics example. On the other hand, *idealizations* arise when the model **specifically states** that a system lacks features that it is known to have; in this sense it is not an omission of a truth, but instead a statement of a falsehood (*ibid*: 182-183). Within linguistics, there is a famous example of idealization in Chomsky's proposal of the "ideal speaker-hearer" (Chomsky, 1965: 3-4). His characterization of such a speaker hearer is as follows (ibid):

> *Linguistic theory is concerned primarily with an ideal speaker-hearer, in a completely homogenous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance.*

It is clear from the above quote that Chomsky is stating certain that the speaker-hearer as relevant to linguistic theory lacks certain properties that we know they possess in reality. Note here that Chomsky is not remaining silent on, for instance, the restrictions placed upon the individual by short-term memory limitations; his proposal is instead a statement that the ideal individual here is not burdened by such limitations. On Jones' view, then, this is an obvious example of an *idealization*.

To make this distinction clearer, the author presents a relatively simple example which contains a number of *abstractions* and a number of *idealizations*. The example he considers (Jones, 2005: 181-184) is someone firing a ball from a cannon, and producing a model that can be used to predict where the ball will land. The model that Jones has in mind is as follows:

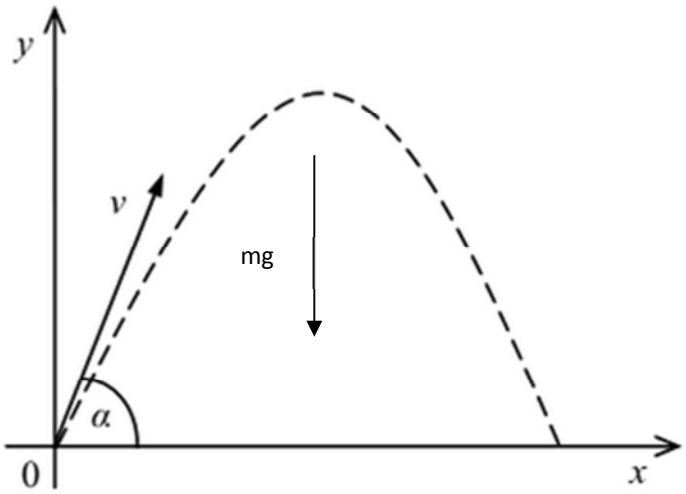


Figure 5.1: Model to predict the trajectory of a cannonball

(adapted from Jones, 2005: 181)

It is clear that there are many aspects of the actual-word situation that have been omitted in figure 5.1. In his proposal, the author divides this multitude of omissions or discrepancies into two groups. The first of these include: the force of gravity has the same direction and magnitude at all points in the ball's flight, the only force acting on the ball in gravity, the ground (represented by the *x-axis*) being perfectly flat, and the ball beginning at the same level as the ground. Summing these differences up, the author notes the following:

> *Here we have then, a number of discrepancies between the way the model (or sort of model) in question represents the modelled system as being, and the way the system really is. It is perfectly in keeping with at least some standard usage to call each of those ways in which the model misrepresents the system an idealization, and in each case there is some property the system has which the model represents it as not having* (*and, correlatively, some property the system does not have which the model represents it as having).* (Jones, 2005: 182-183)

There is also a much larger set of omissions that are included in the above figure. For instance, there is no mention of the material which the ball is made up of or its internal

structure, colour or temperature, the mechanism of the cannon itself that is propelling the ball forward, factors about the setting, i.e. where and when it takes place, and much else. Summing up this second group of omission, the author (*ibid: 184*) states the following:

> *The model is simply silent on all these respects. And according to the regular recommendation I wish to put forward, we should say that a model of a particular system involves an abstraction in a particular respect only when it omits some features of the model system without misrepresenting the system as lacking this feature*.

The difference that Jones is aiming to highlight is that *idealizations* involve **an omission with explicit misrepresentation of the system.** *Abstractions*, on the other hand, can be defined as **omissions without explicit misrepresentation of the system**. Possibly, one could question whether 'remaining silent' in the cases he calls *abstraction* should not also be called misrepresentations. In a sense, can 'remaining silent' also be termed a variety of falsehood, given general pragmatic principles? Views on the matter are mixed: some believe that all falsehoods must involve explicit commission. On the other hand if you believe that omission is as much an 'act' as explicit statements, then this could be doubted (c.f. Clarke, 2013 for discussion on omissions).

Regardless of this thought, however, which I do not think has been settled satisfactorily, I take the general point that Jones is trying to make: namely that there are cases where discrepancies between a model and a system involve a deliberate and explicit statement of counter-factual facts, while others are instead a non-committal to the existence of certain facts. Whether these two are difference in type, or whether what we have instead is a difference in degree is a question I will leave open for the moment. This idea of (mis)representation is also endorsed by Godfrey-Smith (2009) in his discussion of evolutionary biology. He differentiates between *abstraction* and *idealization* as follows (*ibid*: 47):

> *An abstract description of a system leaves a lot out. But it is not intended to say things that are literally false*. *An idealized description of a system is a description that fictionalizes in the service of simplification*.

This ascription raises another important point when discussing *abstraction* and *idealization*. On the definitions presented so far, *abstractions* are defined in terms of omitting information without explicitly making counter-factual statements, while *idealization* has been defined as explicit misrepresentation. However, these definitions cast the net too wide, for they force us to consider as *idealizations* and *abstractions* things it seems to make little sense to consider as such. This danger is more evident in the case of *idealizations*, as it is this process that requires an act of imagination. It is perfectly possible to state a great number of falsehoods about any particular system, but we do not want to consider them all to be *idealizations*:

> *Consider some specific episode of combustion and a model of the process according to which the burning object releases phlogiston into the air. The model represents the piece of wood, say, as initially containing phlogiston; this is clearly a misrepresentation, but just as certainly it would not ordinarily be called an idealization*. (Jones, 2005: 186).[31]

In short, we also need to differentiate those cases which are mistakes, from those that are idealizations. To do so, other criteria have been proposed that idealizations must also satisfy, but it should be noted that these do not constitute a set of necessary or sufficient conditions. These further criteria are: **approximation** to truth, a strive towards to **simplicity**, and misrepresenting **relevant** features. Since this chapter is not specifically about idealization, I will not discuss these features in any detail.

We can also ask similar questions about *abstraction*: namely, is abstraction simply leaving things out without explicit misrepresentation, and if not, then what else is there to the process? Considering the three points just raised above, Jones convincingly argues that if there is anything more to *abstraction* than leaving things out without explicit misrepresentation, then these factors can be neither a strive towards simplicity nor approximation. One of these additional factors cannot be a move towards simplicity, because any omission of features is necessarily also a simplification. In other words, if one

[31] Phlogiston was a hypothesized substance that a constituent of all combustible materials. IT was thought that the burning was caused by the release of this substance. It was shown to be unnecessary by Antoine Lavoisier in the late 18th Century after the processes of oxidation and reduction were discovered.

meaning of *abstraction* is just mere omission, then in this sense *abstraction* necessarily entails simplification, on the assumption that a model containing fewer features will be simpler than one containing more. Therefore, a motivation for simplicity cannot be used to distinguish 'correct' abstractions from 'incorrect' ones, since all omissions do. On the view that an abstraction is not an explicit statement of a falsehood, nor an explicit statement of truth, then they cannot approximate the truth.

More interesting is the question of *relevance*; in other words, if you omit relevant information, would this still count as an abstraction, or would it count as a mistake? Views upon this are mixed. Cartwright (1989: 187) and McMullin (1985: 258), in separate discussions on the topic, both argue that abstraction is only correctly applied when irrelevant aspects of the system under study, or at least features of the system considered to be irrelevant at any given time, are omitted. This view is also maintained by Saitta and Zucker (2013:12), who state as one of their features of abstraction that it involves the retention of relevant information. Jones himself (2005: 190) alters this perspective slightly, by arguing that the term *abstraction* can apply to the omission of relevant feature, but only in cases where an abstraction process occurs on some other idealized basis. In other words, an idealized model can make things seem irrelevant, when in actuality they actually are relevant for the specific phenomenon under study. The main difference between these two views is that the former views relevance as a matter of epistemology, while the latter view sees it as a matter of ontology. On the ontological view, what counts as the relevant parts of a system is independent of how we view the system. This does not hold for the latter view, where relevance is in the eye of the beholder. I find myself more convinced by the epistemological view of relevance, for the reason that relevance can only be assessed in relation to a particular purpose: it is tied to the wants, wishes and aims of the observers involved. Therefore, relevance can never be totally divorced from epistemological issues.

With these issues in mind, then, we have come to our first feature that is linked to the idea of *abstraction*. Let's call it *abstraction as omission* and define it as follows:

> <u>Abstraction as Omission</u>
>
> Abstraction can involve the omission of features of a system for a particular purpose of study. This **should not involve an explicit**

statement of falsehood, and the omitted features **should be irrelevant for the task at hand.**

With this first aspect in mind, we will now turn to a discussion of *abstraction* in relation to another term that is closely associated with it: *generalization*.

5.3.2. Abstraction and Generalization

Another common idea associated with abstraction, is that it involves a comparison of objects and the retention of common features. Turning back to Aristotle, another factor he attributed to abstraction was that it was process that lead to the discovery of a universal, which he defined as "that which is **common to many individuals**" (Coniglione, 2004: 60, emphasis added). A similar focus upon common properties played a major role in the empirical philosophy of Locke (1690/1975), who argued that abstraction, **understood as the identification of similarities between sensible objects** formed the basis of all general ideas. Similarly, Russell (1912) states that to understand a general property like 'whiteness' involved **abstracting the common property** from individual occurrences of objects that possess this property (c.f. Laurence and Marglois, 2012, for a critical discussion of this view). From a more modern perspective, Lewis (1986) defines his *way of abstraction* as forming categories by means of **comparing a number of experienced objects and only focussing on similarities.**

We have the idea, therefore, that in some cases the idea of *abstraction* is linked to the process of generalizing over a given number of cases and noting the similarities between such cases. This idea is the process of generalization. In this section we will look at some discussions on this topic: the first of these comes from mathematics and mathematical pedagogy, while the second comes from computer science, which makes heavy use of this notion. We will conclude by making explicit the differences between *abstraction as generalization* and *abstraction as omission*, as I believe that although they closely linked, they are not identical. This difference will prove to be important when the discussion moves to SFL, as I will use these two as independent criteria for determining the meaning of *abstraction* in the cases previously identified.

Within mathematics, the idea of *generalization* is taking to be important in understanding the nature of mathematical objects. This importance is discussed by Ferrari (2003) in a discussion of the formation of mathematical concepts. In discussing the ways in which abstraction can be characterized in mathematics, he says the following (*ibid*: 1226).:

> *A good share of mathematicians regard abstraction mainly as generalization. The following definition is taken from Well's (2002: 17) online glossary of mathematics:*
>
> > *An abstraction of a concept C is a concept C' that includes all instances of C and that is constructed by taking as axioms certain statements that are true of all instances of C.*

Here we have the explicit link between *abstraction* and *generalization* in mathematics, and also a definition of *generalization* as a theoretical object defined as what is true for a number of other concepts or objects. In particular, he discusses the abstract concept of QUADRILATERAL. On his account, this concept is formed on the basis of the comparison of a number of encountered geometric shapes, for example squares, rectangles, rhombuses, parallelograms, etc., all of which can be described as possessing a number of properties. The kinds of properties that can describe such geometric shapes include: length of sides, angles of corners, colour, number of sides, etc. For instance, consider the following toy picture:

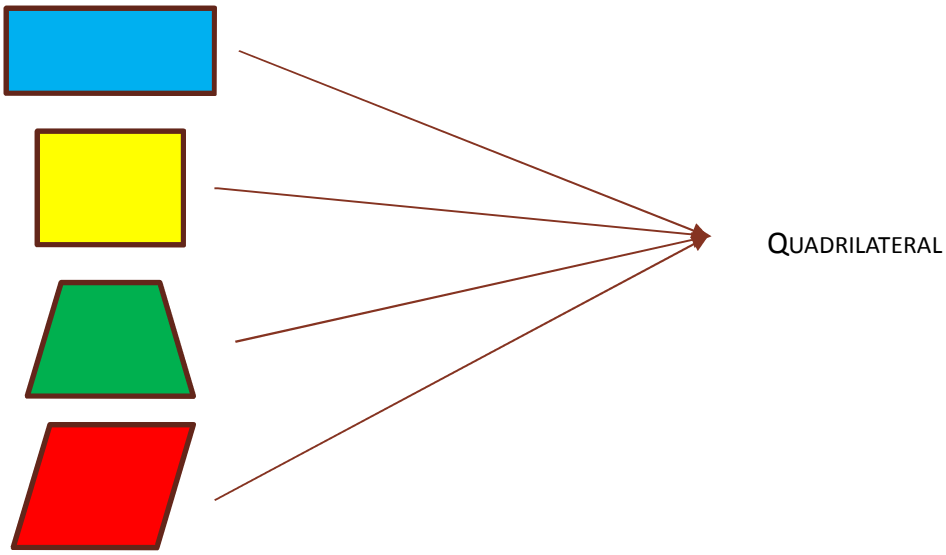


Figure 5.2: The abstract concept of QUADRILATERAL formed by generalizing over experienced objects

In figure 5.2. we have a number of different objects, which can all be described as consisting of a set of properties.[32] Looking at the various objects in the figure, it is clear that there are many differences between them. For instance, they are of different colours, the size of the angles differs in some of the cases (but not in all), and the side lengths differ in each case. What is the same, however, is that they all have four sides. On the basis of this similarity, one is able to form the concept of quadrilateral, for which the statement 'is a geometric shape and has four sides' is true for all objects that would qualify for set membership. This same example is discussed by Posnan et al. (2010), who discuss abstraction and generalization in much the same way as is discussed by Ferrari as above.

A further example from mathematics, but this time from mathematical pedagogy, comes from Michelmore and White (2004). In their proposal, knowledge of mathematical terms is based upon the learner noting similarities between some concrete examples, a process they dub as *empirical abstraction*. For instance, a child attempting to understand the process of addition will be presented with a certain mock situation, and by using the already known

[32] I will refer here to objects and concepts as able to be defined as a collection of properties. This is simply for expository purposes only. I do not intend to convey any further metaphysical claims about this, nor claim that in all cases it will be easy to define any object as a well-defined set of properties, or that such a set will exist in all cases.

principle of counting, will work out the answer in this specific, contrived example. Education then proceeds by presenting more and more situations to the learner, and over time the student comes to realise that the particular objects are irrelevant, and that the same principle is being applied in every case. Upon this realization, the student is now able to form the abstract concept of ADDITION that they have gleamed from the observed examples, and apply it to any new relevant situation that they encounter. An abstract understanding of this concept, then, is formed on the basis of noticing and keeping what is the same over various concrete instances, i.e. it is formed by generalizing over experience.

The idea that generalization as the comparison between objects also finds advocates within the realm of computer science. For instance, Posnan et al. (2010: 6) define a generalization operation as "[defining] similarities between objects." They also note, however, that such a strict definition in terms of necessary and/or sufficient conditions may not always hold. In light of this concern, they posit "a weaker definition of generalization [that] states we have good evidence that all [objects] behave in a similar way." This weaker version of generalization will prove useful in some cases when we turn our attention to theoretical dimensions in SFL.

We have seen, then, that abstraction has been linked with the closely associated notion of generalization. Furthermore, generalization is defined as a process that maintains the common properties between concrete examples, and this set of common properties defines the abstract concept. We can diagram this relationship as follows:

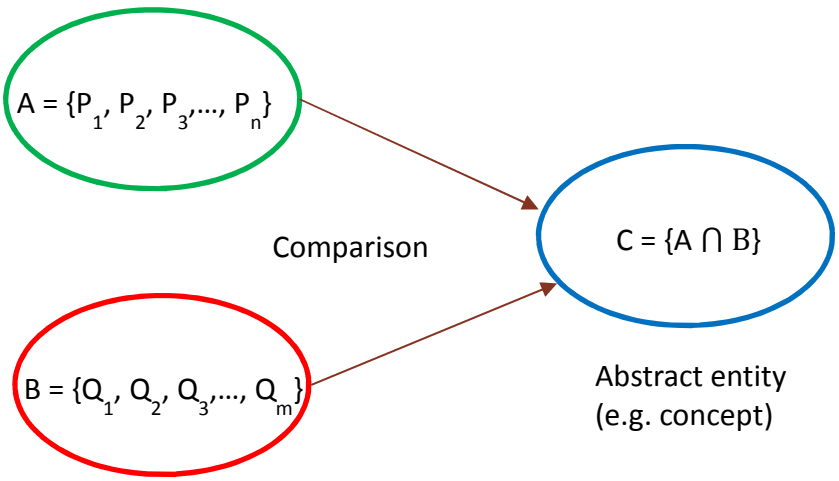


Figure 5.3: Abstraction as generalization: a blueprint

In the figure above, we have a number of objects described as a set of objects, and the abstract object is then formed as a set conjunction of these objects, so that only the properties common to all of the sets become properties of the abstract object after the process of abstraction.

We can then provide our second facet of abstraction that may prove useful in our discussion of the notion in SFL theory. It can be stated as follows:

> Abstraction as Generalization
>
> A move in abstraction involves **generalizing over more than one more concrete counterpart**, and **retaining what is common** between them.

At this point, it is worth demonstrating that the above definition is different to the one presented in the previous section. It might be thought that they are simply two sides of the same coin: in the process of retaining common information, it is necessary to omit other information. In this regards, then, the two ideas of *abstraction* provided so far may simply be seen to be two sides of the same coin, and therefore presenting them as two different criteria would be a mistake, as the presence of one would entail the presence of the other. To conclude this section, I intend to argue that this is not the case: although the ideas of *omission* and *generalization* are indeed similar and linked, they should be correctly thought to be distinct, thereby motivating their separate inclusion in this study.

Firstly, it will be recalled that only certain types of omission were seen to correctly be called *abstractions*, namely those that can appropriately be argued to be irrelevant facts and details. The purpose of this inclusion was to stop *any* omission without explicit misrepresentation being counted as an abstraction. In the case of generalization the question of relevance does not seem to enter the picture. In most cases, it will be hoped that these common properties will be the relevant to the task at hand, but that is in no way

assured by the generalization procedure. It is 'blind' to issues of relevancy. For instance, one may run a particular experiment a given number of times, and on the basis of common events between these experiments form an abstract explanation or description of what is taking place. It will be hoped that these common properties will indeed be relevant, but this would simply be a serendipitous consequence of the generalization process; it is not part of the process itself. These different conditions of inclusion, therefore, motivate including these two facts as separate conditions.

A further difference between these two is the number of concrete objects an abstract description by *omission* and *generalization* necessarily cover. By necessity, generalization involves the comparison of more than one object; it is impossible to form a generalization on only one concrete entity. On the description of omission given in this chapter, on the other hand, this possibility seems viable, as the process was described as simply ignoring some features of a given system. On the face of it, then, this seems to be a reason for further thinking that these two aspects should be kept separate, since the abstract entities that are formed from them seem to cover a potentially different number of concrete entities: omission covers at least one, while generalization must be linked to at least two concrete counterparts.

There is, however, a potential problem with this simple characterization: how do we know exactly what counts as relevant and irrelevant for the purposes of feature omission? Surely this cannot be gained from only one example of one concrete entity? If this is true, then the simple characterization just presented misses the mark, and this potential difference between the two procedures fails to hold. I think, however, that it would be hasty to jump to this conclusion. My reason for thinking this is as follows: although it is true that in most, if not all, cases the means of determining the irrelevant features will indeed be made on the basis of comparison with (at least one) other object(s), but **the abstract entity created on the basis of this omission procedure does not have to hold for all the concrete entities involved**. That is, although a comparison with something else may be made, the abstract entity may only count as covering the original entity, and not necessarily the one used for comparative purposes. This is not the case for generalization, whereby the abstract entity formed **must cover all entities that are involved in the generalization procedure**.

For these reasons, I believe this chapter is justified for maintaining that these first two criteria should be kept as separate criteria for the meaning of *abstraction*.

5.3.3. Abstraction and Decontextualization

The process of *abstraction* is thought to involve detaching specific observations and events from their immediate context. This process of *decontextualization* is thought to be the hallmark of many scientific constructs. For instance:

> *Science is a successful endeavour because its concepts are applicable to a wider range of phenomena the* ***less they are tied to the specifics of any context****, and representations* ***become more inclusive and abstract*** *the less contextual parameters they include.* (Roth, 2012: 95, emphasis added).

In this paper, Roth is discussing the misconceptions and problems that arise when scientists come to interpret the findings presented by others in scholarly published work. He claims, on the basis of experiments, that the data and graphs are often not understood well by scientists who were not directly involved in their creation. The reason he gives is that work in science is often highly decontextualized, and although the original researchers themselves are able to re-contextualize these observations, this is not as easy for others. This helps to reinforce the idea that science deals with highly decontextualized entities, and also the direct link between this factor and *abstraction* can be seen in the above quote. This view is also presented by Noss et al. (2002: 206, emphasis added):

> *A traditional view regards abstraction as an essential property* ***extracted from a situation but not contained in it****. From this perspective, an abstraction is deemed* ***"apart" from, even above, the situation*** *of its genesis.*

Again here we have the ascription of the view that *abstraction* involve removing certain things from any particular situation, in this sense they become detached and 'rise above' their original situation. This intuitively makes sense. An abstract description of many scientific processes, for instance the mathematical laws of physics, or chemical equation

denoting the results of certain chemical reactions, are not meant to only hold for the particular situation, but also in similar situations that may arise in the future. The idea, it seems, is that concrete entities are closely tied or linked to a particular time and a particular place, whereas after a process of *abstraction* this no longer holds.

As an example, let's consider one of the equations that are used to model action potentials. Proposed by Hodgkin and Huxley (1952), it models the current flows within excitable cells and how they come to create the basis for neuronal communication. It is stated below:

$$I = C_m \frac{dV}{dT} + \overline{g_k} n^4 (V - V_k) + \overline{g_{Na}} m^3 h (V - V_{Na}) + \overline{g_l} (V - V_l)$$

The above equation provides a model of how the current across the excitable cell depends upon the voltages provided by two electrically charged ions, potassium and sodium, along with some "leakage". These details do not really matter here, but are included for interpretative purposes. More relevant for the discussion here is to appreciate the situation in which the above abstract object was created. As discussed in a recent summary of the discovery by Schwiening (2012), this equation, as well as the others that form the famous Hodgkin-Huxley equations, was formed on the basis of experiments conducted on the squid giant axon, as it had the necessary dimensions to make electrical recordings practical. However, the model is currently not solely taken to refer to the ion flows within the neuronal cells of squids, but is no taken to provide a model of said processes in all types of neurons, and generally such a view has been taken to be successful. In many cases, then, abstract entities are not necessarily tied to any one particular situation.

As was the case in the last section, it can be difficult to ascertain on face value whether this criterion really deserves to stand alone, or whether in fact it is simply one of the previous criteria we have already discussed under a different guise. It is worth showing that it deserves this status in this chapter and in the later discussion.

Regarding the links between decontextualization and generalization, it may be thought that when you generalize, you are also necessarily producing an act of decontextualization. Again, let's consider the case of the abstract concept of QUADRILATERAL discussed above. It was seen previously that (at least) part of why this concept is abstract is because it is formed on the basis on comparison of a number encountered objects, with the properties of the

abstract object being defined as common features. Now, one can ask what context or situation is the concept of QUADRILATERAL linked to? The answer seems to be similar to the case of one of the Hodkin-Huxley equation just discussed: it doesn't seem to be linked to any particular situation or set of situations, as the concept convers all actual and future possible four sided geometric shapes. If this fact holds true in all cases, the motivation for keeping generalization and decontextualization as separate criteria disappears, as the former immediately entails the presence of the other. However, the view taken here is that generalization, taken as the retention of common properties, does not always have to involve detachment from particular situations.

To show this, consider the case presented in figure 5.2. again, and this time imagine that they were all presented to an individual on one particular occasion who has no notion of what a QUADRILATERAL is, but he does possess knowledge of basic numeracy. On this occasion, after being asked to reflecting on the shapes in front of him, he says the following: "all these shapes here have four sides". The important aspects to notice in this simple exclamation is that he has successfully performed a *generalization* operation, he has retained in his response all the common information, but crucially he has also indicated that this procedure is linked to only this situation, as indicated by the spatial deictic element. Here we have generalization without removal form a context. Of course, in many cases both generalization and decontextualization will be in evidence in the status of abstract objects. This fact need not worry us, since, as was indicated above, these criteria should not be considered to be mutually exclusive views, but simply different aspects of what the meaning of *abstraction* may be. Indeed, in the second half of this chapter, we will encounter abstraction processes that involve both *generalization* and *decontextualization*, but similarly we will encounter others that utilize one but not the other.

When comparing *decontextualization* and *omission*, it is harder to see a difference. In this case, I will proceed with the idea that they are both omission process, but operating on separate aspects of the real world situation. Let's reconsider the cannon ball example from §5.3.1. Part of the features of the situation that were omitted without misrepresentation in this example were aspects such as the time of day, and the particular place, i.e. some of the omission procedures seem to inevitably lead to decontextualization. But not all of them do, as should be clear from examples such as ball colour. I will then make the following

distinction, in the hope that it will be useful but does not necessarily allay fears that decontextualization is simply a certain kind of omission: *abstraction as omission* will refer to omission of features of the system itself, while *abstraction as decontextualization* will refer to omission of features of the environment that the system is embedded in.

We have, then, our third criterion for diagnosing the different types of *abstraction* processes that may be in evidence in SFL. We can sum it up as follows:

> Abstraction as decontextualization
>
> A move in abstraction can involve a certain degree of detachment from situations associated with its more concrete counterparts.

With these three aspects of *abstraction* outlined, explained and shown to be different from each other, this chapter will now turn to the same notion in SFL. The aim is to determine whether the features identified here can be used to determine if the same notion means the same as applied to the dimensions of *stratification*, *delicacy* and *instantiation*.

5.4. Abstraction in Systemic Functional Linguistics

5.4.1. Delicacy

The first dimension examined in this chapter is that of thesis is *delicacy*. To recap previous discussions, delicacy can be described as a hierarchy containing a number of different *levels of specificity*, ranging from the most general linguistic types, to the most specific. Also recall that in theory, there should be a dimension of delicacy operating at each level of interpretation (stratum) in the framework, as each of these is conceived as a system of potential options. In order to frame this discussion, I will be assuming that *abstraction* can be mapped onto the scale of delicacy in the following way, built upon the discussion in §5.1:

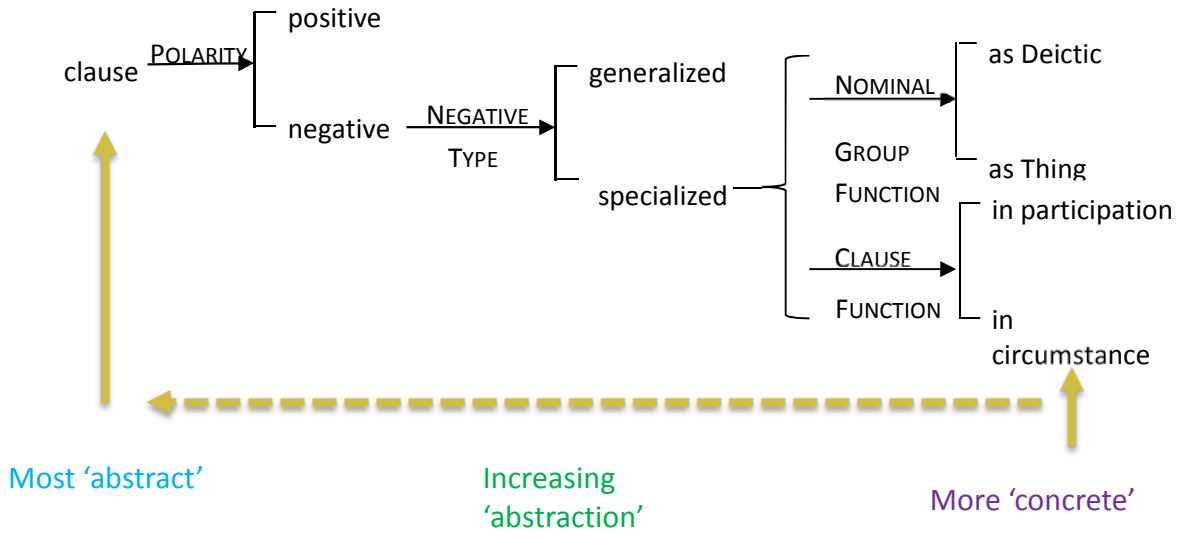


Figure 5.4: Abstraction and the scale of delicacy

It should be noted that the choice of system here, namely the polarity system at the lexicogrammatical level, is purely arbitrary, and the same should hold for any system that could be presented. The general picture is that those system features further down the cline are considered to be *more abstract*, than those that precede them.

We can now begin to interrogate this notion of *abstraction* with regards to the criteria that have been presented so far: namely, does the move from a more specific to a more general system feature involve *omission* of features without explicit misrepresentation? Does it involve *generalization*, and does is involve *decontextualization*?

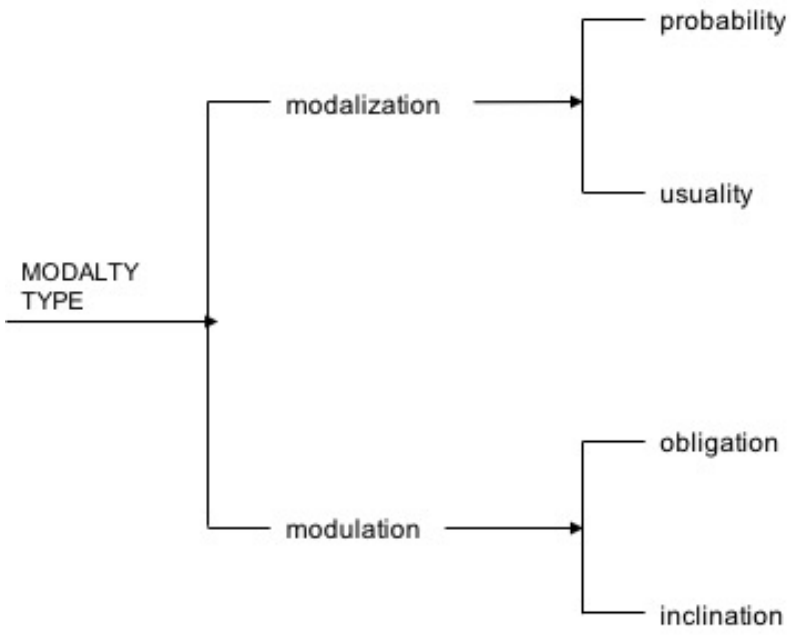
5.4.1.1. Omission?

The first question that needs to be asked is whether more general system features within a system network omit information that is found in the other system features that enter into this hyponymic scale. Furthermore, if this is in fact the case, then such an omission of detail should not amount to an explicit statement to the effect that the more delicate information is not present in a particular text.

How are we best to fathom the answer to the above question? In the first instance, it is necessary to discover whether more "abstract" points in delicacy really do omit information that is contained in the more specific counterparts. In order to do this, a few different

system networks will be presented, and a description of sentences at various points on these delicacy clines will be given. This will enable us to determine whether the criteria discussed in this section have been met. Once these descriptions have been given, if the answer to the first part of the question is positive, the more difficult matter of explicit misrepresentation will be tackled by looking at the nature of linguistic description and links to the text in SFL.

As an example, below is the lexicogrammatical system of MODALITY TYPE, a subsystem within the interpersonal system of MODALITY that distinguishes between different "flavours" of modality (adapted from Halliday and Matthiessen, 2014:182):



Network 5.1: System of MODALITY TYPE

The choice of this particular system network is, and indeed should be, arbitrary. In this discussion I have opted for system networks that are well founded in the literature, because as much as possible it is advantageous to distance the discussion away from more controversial system networks. In any case, the choice of system network should not make a difference to the arguments presented here; if they all do indeed operate over the same theoretical dimension, then the detail of the features and systems themselves should be of secondary importance to the theoretical dimension itself.

We can now consider some clauses and see where they fit into the above system network:

(5) John's car keys are probably in his coat pocket

(6) The lab members are usually in the pub on Friday evenings

(7) I can drive you to the airport tomorrow afternoon

(8) I have to meet my boss for a meeting later today

If one makes the most general distinction between the different types of modality in the system diagram above, then the following classifications can be made:

Modulation = (7) and (8)

Modalizaton = (5) and (6)

This is because the criterium for *modulation* is whether the grammatical markings are intended to carry the meaning of "prescribing or proscribing" (Halliday and Matthiessen, 2014: 177). However, being a system of modality, the assertion of prescribing or proscribing a certain action can be stated with varying degrees of certainty. In (7), the deontic modal verb <can> expresses a certain degree of inclination to bring a certain action about, whereas the verbal element <have to> indicates a strong sense of obligation that some event must take place. *Modalization*, on the other hand, is associated with propositions, and specifically by specifying the degree of commitment the speaker has to the assertion or denial she makes. In (5), the modal adjunct *probably* indicates a relatively high commitment to the assertion that John's car keys really are in his pocket, whereas the assertion in (6) is a strong commitment, but not a full commitment, to the early weekend whereabouts of this particular lab group based upon their past behaviour.

If one makes more delicate distinctions between these four clauses, however, these groupings no longer hold. Moving one step in delicacy the following categorizations now hold:

Modulation - obligation = (8)

Modulation - inclination = (7)

Modalization - probability = (5)

Modalization - usuality = (6)

These groupings now hold because, in the case of *modulation* SFL distinguishes between those prescriptions that are linked to commands (obligation), and those that are based upon requests (inclination). In the case of *modalization*, the distinction is made between those which are "equivalent to 'either yes or no' […] with different degrees of likelihood attached" (*ibid*), and those which are "equivalent to 'both yes and no', i.e. sometimes yes, sometimes no, with different degrees of oftenness attached" (*ibid*)[33]. These finer-grained criteria allow linguistic descriptions to further distinguish between different types of modality that is not available at the more abstract level of *delicacy*.

It is the case, therefore, that more general system features omit detail that is potentially present in a more specific description. This was shown by the fact that in the system networks more general divisions group certain texts together that can be treated as different if considered in more specific terms. In this way, the general features omit potential differences that may exist between texts. The remaining matter for our purposes is whether such an omission would qualify as an intentional misrepresentation or not. If the answer to this secondary question is negative, then the *abstraction* process that is involved in the scale of delicacy can correctly be classed as *abstraction as omission*.

To aid the discussion, it will prove useful to look at examples whereby grammatical analysis has been employed in order to draw conclusions about the nature of language use as pertaining to a particular theoretical question within SFL. By looking at the actual practice of functional linguists, it may be possible to investigate the claims they make as pertaining to the issue of potential explicit misrepresentation. The first example comes from Rafiei and Modirkhamene (2012), where the authors undertook a textual analysis of Iranian TEFL master's level thesis written in English. The focus of their study was to determine whether the textual elements of both the method and results sections of said theses correlated with whether the work resulted in either articles in scholarly journals, articles presented at academic conferences, or to no such further stage. The result of their study was that there were no significant differences between the three groups of theses stated above, but that is

[33] On a semantic note, the idea that the truth of a proposition can be "both yes and no" seems odd, assuming the law of the excluded middle. The alternative description of "sometimes yes, sometimes no" in the above quote seems more appropriate. We will gloss over this point here, as it does not directly relate to our discussion.

not important for the purposes here. What is more important is where on the scale of delicacy the authors situated their analysis.

In their analysis, the authors decided to analyse the thematic choices in the relevant sections of the master's work based upon the distinctions made by Halliday (1994). The reasons for choosing this version of Halliday's *Introduction to Functional Grammar*, instead of the later editions, are not given. Regardless, drawing from this work, the types of thematicity used in the study were as follows:

- Textual theme
- Interpersonal theme
- Marked theme
- Unmarked theme
- Simple theme
- Multiple themes

It should be noted that these choices are not necessarily as list of exclusive choices, a single clause may include more than one of the above choices (in these cases we would have *multiple themes*, if only one theme is present then this is a case of a *simple* theme). What should be noted in this above list, however, is that the authors have chosen not to distinguish between the different types of themes; they have set their analysis at quite a general level of delicacy. For instance, distinctions between more delicatetypes of interpersonal theme have not been made. Examples of such more delicate choices would be between modal themes and other types of interpersonal theme. In the case of textual themes, a distinction could be made between structural themes and other types of textual themes (c.f. Halliday and Matthiessen, 2014: 107).

However, although such distinctions have not been made in the above this does not amount to a statement that such distinctions do not exist. The reason stems from Halliday and Matthiessen's (1999: 325-326) comments about the links between linguistic analysis and pragmatic considerations about the nature of such descriptions: on their view pitching an analysis at a certain level of delicacy is a conscious choice on the part of the linguist that takes into account the use that can be made out of said description. There are two dangers

associated with analysing a group of texts at a certain point in delicacy: an analysis that is too general risks glossing over potential distinctions that may illuminate a certain theoretical question. A description which is too specific, on the other hand, is in danger of including too much detail; it risks missing out important regularities and tendencies because it is taking into account distinctions that are irrelevant for the study at hand.

A description at a particular delicate point amounts, I would argue, to an assertion along the lines of the following: "These are the paradigmatic options in language that are relevant to answering the theoretical problem at hand. Possible more general distinctions are not fine-grained enough to successfully resolve this issue, whereas the more specific options are irrelevant for the purposes at hand." What we have here, therefore, is not a misrepresentation that the more specific options do not exist in language, but instead that they are not important *for this particular description*, and as a result the linguist is claiming that they are justified in omitting this level of detail from their work. One could go further and say that positioning oneself in terms of delicacy could be interpreted as an implicit recognition that more specific options exist (in most cases, unless one is pitching themselves at the extreme end of the scale), but that they are not relevant and are so to be ignored.

With this is mind, I believe that part of the meaning of *abstraction* with respect to the scale of delicacy falls under the rubric of *abstraction as omission*. When a move is made from a more specific, or more concrete, point in delicacy to a more abstract one, some information is omitted. However, this happens without explicit misrepresentation about the existence of such omitted information.

5.4.1.2. Generalization?

The next criterion is whether a move in decreasing delicacy also involves a move in *generalization*. To recall, *generalization* has so far been described as the retention of common features between several more concrete alternatives, while discarding the features of those more concrete entities that differ.

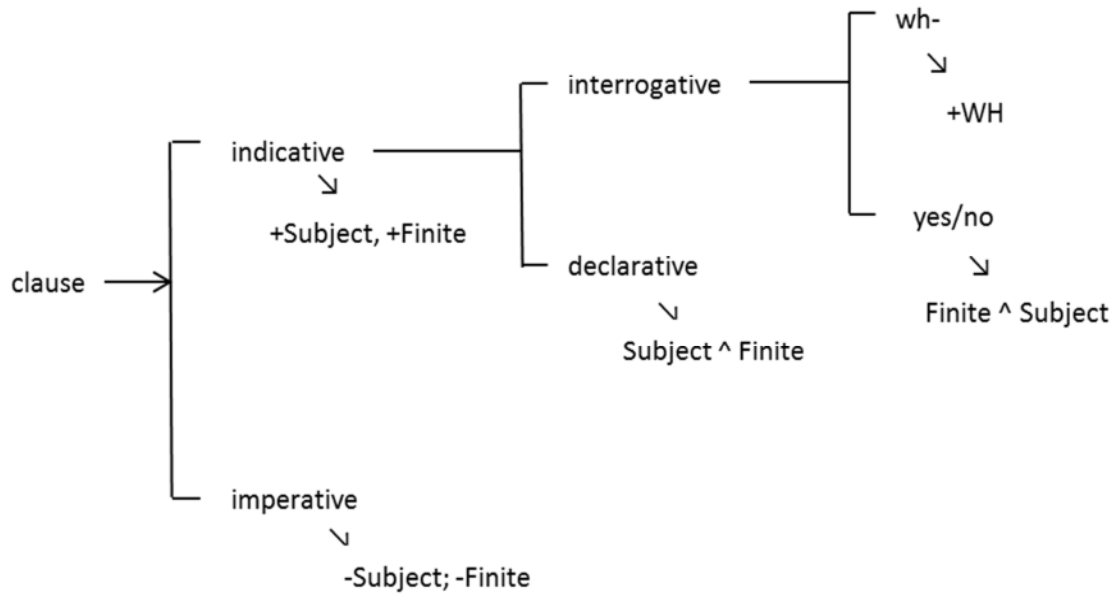
On the face of it, the fact that delicacy should involve *generalization* seems like a given; after all, the dimension itself is defined as operating from the *general* to the *specific*. However, this fact alone does not provide enough support for considering *delicacy* in terms of *abstraction as generalization*, since one of the main points of this chapter is that the use of the same theoretical term is no guarantee of identity in ideas. Therefore, it is still necessary to enquire about the links between *delicacy* and *generalization* in order to determine whether it carries that meaning relevant for our purposes here.

It seems as though less delicate system features do generalize over more specific ones by means of maintaining common properties. This view is backed-up by Halliday and Matthiessen's (1999: 326) thoughts on the nature of delicacy:

> *Any properties associated with less delicate options are inherited by those that are more delicate. For instance, 'simmering' is a kind of 'cooking', both of which are culinary phenomena*.

In this instance, the authors are discussing delicacy as it applies to the contextual notion of *field*, that is, the nature of the social action taking place in which language plays a part. But this over-arching principle should hold for the notion of delicacy at all points on the stratificational framework. The above view has also been formed by considering the delicacy relations as they run from the least delicate options to the more specific ones, but it is relatively easy to reverse this perspective and consider the same trend but as it operates in the other direction. Instead of thinking of properties being "inherited" by more delicate options, one can recast this same sentiment by stating that common properties amongst a certain number of more delicate options are retained in the more general one.

As an example, re-consider the following system network for the lexicogrammatical system of MOOD (repeated from §3.1



Network 5.2. System network of MOOD

Furthermore, consider the following clauses:

(9) Don't go into the basement!
(10) The right wing media tries to control the population.
(11) Are you going to the art exhibition tomorrow?
(12) Why are your keys in the fridge?

At the first level of delicacy, we can distinguish between the clause in (9) and those in (10)-(12); the former of these being an imperative clause, and the other three are all grouped under the heading of indicative. This is due to the realization statements given in the above system diagram: (9) does not have a *Subject* or *Finite* element, and both of these are found in the examples in (10)-(12). Moving a further step in delicacy, it is now possible to distinguish between the declarative sentence in (10), and the two interrogative sentences in (11) and (12). This distinction is based upon the relative ordering of the Subject and Finite elements: in declarative clauses the Subject precedes the Finite element, while the reverse is the case in interrogative clauses.

The important aspect of this example to note is that each of the above more general system features generalizes over those towards the more delicate end by means of retaining

common elements, even though those more specific options will differ from one another in other ways. All indicative clauses have a Subject and Finite element, although as is clear by comparing the examples in (10)-(12) there are other differences between them. The same is the case when one compares (11) and (12). Both of these clauses share the property of having the Finite element preceding the Subject, but they differ in the presence of the *Wh*-element. As a result only the former property is present in the system feature *interrogative*.

It seems then, that part of the meaning of *abstraction* as it relates to the scale of delicacy is that the theoretical dimension involves *generalization*, by means of focussing on common properties of the more delicate options a particular system feature is linked to.

5.4.1.3. Decontextualization?

The question of whether a move in decreasing delicacy involves decontextualization is a subtle one, but I believe the correct answer to this is in the negative.

The complicated factor for this issue is that on the face of it more delicate options in a system network initially do seem to be more closely tied to particular situations by means of representing a subset of cases that are covered by a more general counterpart. This can be seen by considering some of the system networks presented in the previous two sections. For instance, a system feature such as *modulation* would seem to cover more instances than the more delicate option of obligation. Therefore, one might be tempted to think that because the potential number of *modulation* clauses is greater than the more delicate *obligation* clauses, then the former of these is more decontextualized than the latter.

This, however, would be a mistake. It this discussion, it is critical to remember both the nature of the *potential* itself, as well as the means by which entities in paradigmatic relationships acquire their theoretical meaning. When introducing the basic notions of semiotics in §3.1, it was seen that the very notion of paradigmatic relations mean that signs only become meaningful in comparison with other signs that can legitimately occur in the same position. Reminding ourselves of this fact enables us to recognize the fact that the comparison above is not the correct one, we should not be comparing the potential of *modulation* clauses with the potential to utter *obligation* clauses; instead the comparison should be between the former and the potential for either *obligation* or *inclination* clauses,

since both of these only get their theoretical significance by means of comparison with the other.[34] Put in this way, it is clear from this that the potential of both of these is identical. We can sum up the picture as follows: no matter in how much detail you divide up the meaning making potential at a given level for interpretation, **you are still dealing, overall, with the same potential**. As a result, if the potential is the same, then it must cover and be associated with the same number of situations, even if, as is the case here, it is not possible or immediately obvious how many cases that may be.

From a more general picture, there is another reason to think the view noted at the beginning of this section is mistaken. This applies to the "grammarian's dream" discussed in §3.3, c.f. Halliday (1961) and Hasan (1987): the idea that grammar and lexis are not different systems, but instead the same system viewed at different points of delicacy. If a process of deceontextualization was in operation over the scale of delicacy, this would result in the system of lexical potential being more contextualized than the system of grammatical potential. This would bring into question the very notion of the two systems being identical, or in other words question the very notion of the grammarian's dream and the view in SFL that there are not separate systems of lexis and grammar.

Overall, I believe these arguments show that the notion of *abstraction* in the scale of delicacy does not carry the meaning of decontextualization.

In sum, then, I believe that the meaning of the term *abstraction* as applied to the scale of *delicacy* involves the notions of *omission* and *generalization*, but does not include the notion of *decontextualization*.

5.4.2. Instantiation

The next dimension to be discussed is the *cline of instantiation*. As has been previously discussed, this is a scale that operates over a number of different *levels of regularity*, running from the *system* to the *instance*. In between, we have a number of midpoints that describe regularities as they relate to different situations. As before, it is useful to make explicit how the notion of *abstraction* will be interpreted with regards to this dimension. Based upon the discussion given in §5.1., it will be stated in the following way:

[34] Technically, the comparison should be between the potential for *modulation* and *modalization* clauses with the subtypes of both of them. This has been omitted here for ease of explanation.

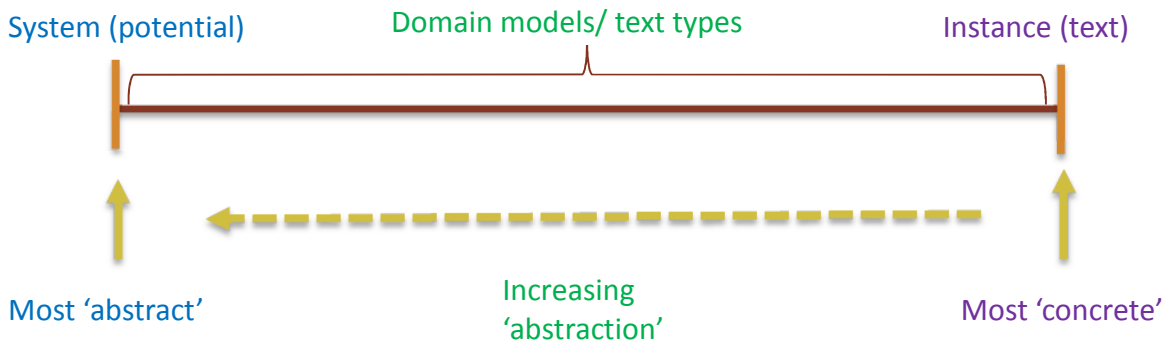


Figure 5.5: Abstraction and the cline of instantiation

To frame the discussion in the following sections, this work will refer to three points situated at various points on the cline of instantiation. These points all come from the example of cooking recipes, discussed in §3.5, and outlined in detail in Halliday and Matthiessen (1999: §8.3). The first of these will be a particular text, or occasionally a set of texts, in this case all being instances of cooking recipes. The second of these is a certain mid-point on the cline, which represents at the same time both the text type of cooking recipes and the variant of the overall ideational semantic potential for the production of new cooking recipes. The final point will be the system as a whole, in this case the overall ideational semantic potential. Many of the details can be found in the above cited section of this thesis, but I will outline some of the major characteristics below for ease of explanation.

The text itself is relatively easy to conceptualize, but it is important to remember that normally when one is considering the text end of the cline of instantiation, it is usually only at one stratum. For instance, at the level of lexicogrammar, the text end will consider a particular piece of language as being an instance of *wording*, and as a result only the information relevant to the study of lexicogrammar will be included in its description (see §5.4.3. for detailed discussion about this point).

The midpoint that describes the variant on the overall meaning-making potential as is relevant to a particular group of situations, in this case to the situations of cooking recipes. The main conceptual point about mid-points on the cline of instantiation is that it can be thought of as a particular version of the overall potential of language. As a result, the same system features that describe the overall meaning-making potential of language will be

present in the mind-point. The major difference is that the distribution of occurrence of those features will differ in a comparison of a domain model with the overall picture of language. For instance, in general there is a 0.9-0.1 split between positive and negative clauses in language in general (Halliday and James, 1993), but this distribution may change in a particular domain model at the lexicogrammatical stratum. In our discussion of cooking recipes in chapter 3, the major differences between the overall description of language and that of cooking recipes were as follows:

a) Participant roles will be exclusively filled by concrete "things"; there is an absence of "metathings" in the ideational semantics of cooking recipes. So, for instance, there is a lack of facts, ideas and locutions.
b) Even within the category of concrete "things", only a few sub-types are readily found in certain participant roles: Agents are predominately persons, while objects, animals, etc. occupy the Goal/ Patient role.
c) Certain types of figure are preferred: the majority of figures are of the "doing" type, and specifically "doing to/with" at a further level of delicacy. Less frequent are figures of sensing and figures of being, while figures of saying are extremely rare.
d) In terms of sequences: generally there exist more coarse-grained distinctions between projecting sequences and expanding sequences. In the sub-potential/ text type of recipes, only expanding types are attested.

Definitions and examples of the above semantic features can be found in the original discussion in §3.5. The important aspect to note is that the *same* features are in both the overall system and the sub-system, but the relative likelihood of a certain feature being present will be different in the two cases.

Finally, the third point to be utilized in our discussion will be the system as a whole, defined as it has been throughout this thesis as the various options that make up the potential or meaning-making linguistic activity, at a certain degree of delicacy.

Considering these, we can now turn to look at our criteria for the meaning of *abstraction* for this theoretical dimension within SFL.

5.4.2.1. Generalization?

Departing from the order of the discussion in the last section, it is necessary to begin our discussion of the term abstraction as it relates to the idea of generalization. Although the cline of instantiation has been linked to abstraction, it is more often linked to the term *generalization* (e.g. Martin 2009: 558; Matthiessen et al. 2010: 121-122). As we have seen, however, generalization forms one of the main meanings of abstraction, so making the distinction between abstraction and generalization does not necessarily mean that two distinct processes are being referred to.

To recap, §5.3.2 drew upon the definition of generalization as the creation of an abstract object by focussing on the shared properties of various observed entities. The example given above concerned the abstract category QUADRILATERAL, but other examples exist, for instance the taxonomic relations found in biology. As noted, abstraction as relevant for the cline of instantiation is often described in terms of generalization. Halliday and Matthiessen (1999: 323) describe a domain model as a "generalization across particular instances of text" when "seen from the instance end". Therefore, if the meaning of the term in SFL coincides with the way it is used in other disciplines, especially mathematics and philosophy, a domain model (when viewed from the instance end) should contain only those properties which are shared by all instances of text claimed to be a member of that domain. Focussing on the domain model of cooking recipes, it seems as though the use of the term generalization as intended by SFL practitioners is different from the process of generalization as just described. In order to show this, let's turn to the second and third characteristics identified above. I take this characteristic to mean, and this interpretation seems reasonable, that when considering all instances of this particular domainall instances will contain figures of doing, some instances will contain figures of sensing and/ or being, but some instances will contain neither, and only a few will contain figures of saying. Therefore, the only property which will be common to all these texts would be the presence of figures of doing, and if a domain model is considered to be a generalization over texts in the sense described above, then only this system feature should be present.

This, however, is not what SFL practitioners intend when defining a mid-point on the cline of instantiation. The term is this case is to be defined as "an accumulation of frequencies in text" (Wu 2000: 246). In other words, generalization as it is used here is encroached in

probabilistic terms, and is considered to be a form of statistical induction. This possibly could still be considered a form of generalization, it differs from the kind of generalization which is usually associated with abstraction. Also, it differs from the generalization process which operates over the dimension of delicacy, which was seen to be more similar to its use in mathematics and philosophy. Instead of making general statements which hold for all more concrete examples, they make statements concerning the features which they are more or less likely to exhibit.

5.4.2.2. Omission?

Unlike the dimension of *delicacy*, abstraction along the cline of instantiation does not include the omission of any factors found at the concrete end of the cline. As was discussed in the last section, omission was evident along the scale of delicacy because more abstract features in system networks omit features which are evident in their more concrete counterparts. It was linked to generalization, because the same process can be described as the focus on the similarities between two or more concrete features.

As just discussed, generalization along the cline of instantiation operates more in terms of a statistical induction; on the basis of past experience, the linguist builds up a statistical predication of what features are likely to be exhibited by a particular text. This is most clearly seen in domain models. The question now remains: do the more abstract points on the cline omit any (irrelevant) information that is found in their concrete counterparts, i.e. the set of texts that they are statistically generalizing over? The answer seems to be no. Let's say that an example was found in the creation of the cooking recipe domain model of a text containing a figure of saying (which was not attested in the sample given, but was thought to be a possibility). In this case, the variant of the ideational potential which defines that domain model must include the possibility of figures of saying, but it notes its relative rarity by attributing to it an extremely low probability of occurring. Similar processes would apply to other rare and atypical text features at all the strata which make up the SFL model. As a result, abstraction as applied to the cline of instantiation does not carry the meaning of omission.

5.4.2.3. Decontextualization?

The ways in which the cline of instantiation correlates with the notion of decontextualization mainly stem from consideration of points upon this dimension as potentials, as well as statistical regularities across texts. The notion of potentials is fundamental in systemic theorizing, because it allows the linguistic model to apply not just to previously encountered linguistic utterances, but to additionally try and explain how speakers produce and interpret novel instances of language. Without this conceptualization, there is a danger of language being considered as purely a generalization over previously encountered texts, a view deemed inadequate for linguistic theory ever since Chomsky's (1959) criticism of behaviourist psychology. To see how decontextualization links to abstraction along the cline of instantiation, let's take three points along this cline: firstly, a particular text found at the instance pole, secondly, the domain model of cooking recipes already introduced, and lastly, the overall description of language as system. If the notion of decontextualization applies here, then the concrete example (instance/ text) should be more closely linked to a particular context than either the domain model or the overall systemic potential, and the overall system model should be more decontextualized than the domain model. A text is produced, and often considered, within a particular context or situation. Indeed its very character is shaped, in part, from the situation the text is produced in and the relevant registerial configuration. As such, a text is fully contextualized. This is exemplified by the fact that texts are considered to be linguistic tokens (Halliday and Matthiessen 1999: 15). The system, on the other hand, represents the entire meaning making potential of languages. As such, it is not tied to any particular setting, but intended, theoretically, to cover all encountered situations and provide the means to successfully function linguistically in any new situation and relate it to the system as already conceived. Although the notion of system of a particular language itself cannot ever be considered to be fully decontextualized, since it has been shaped by the culture of the language users, it involves a higher degree of decontextualization than a particular instance of language. Intermediate points, lying between the two, are more removed from specific contexts than a text, but less so than the overall system. In the example already mentioned, the domain model of cooking recipes can be viewed as a variant of the overall ideational semantic potential of language. It is a characterization of the potential of creating meaning within a certain specified setting; in other words its characterization is field motivated (as the

ideational metafunction correlates with the contextual metafunction of field, see Hasan 1995; Thompson 1999). It is the potential of creating meaningful text in a restricted set of circumstances; it is not tied to any specified on particular context, but instead a characterization of the relevant aspects of the ideational potential which are activated in a set of circumstances of a particular type. Therefore, it is more contextualized than the overall system, which must cover all potential instances of language use, but less contextualized than a text, which is intimately tied to the particular setting it was created in. As a result, instantiational abstraction involves a degree of decontextualization as part of its meaning.

5.4.3. Stratification

The final theoretical dimension to be discussed in this chapter is that of stratification. In §3 this dimension was seen to operate over a number of *levels of interpretation*, and runs from *phonetics* as its lowest level to *context* as its highest, with *phonology*, *lexico-grammar* and *semantics* linking the two. In §5.1, the links between *stratification* and *abstraction* were noted, as they have been often linked in the existing literature. In order to discuss *abstraction* as it relates to this final dimension it is imperative that the various orders of *abstraction* in the stratificational model are made explicit. Based upon the previous characterization, this chapter views this term to apply to this dimension in the following way:

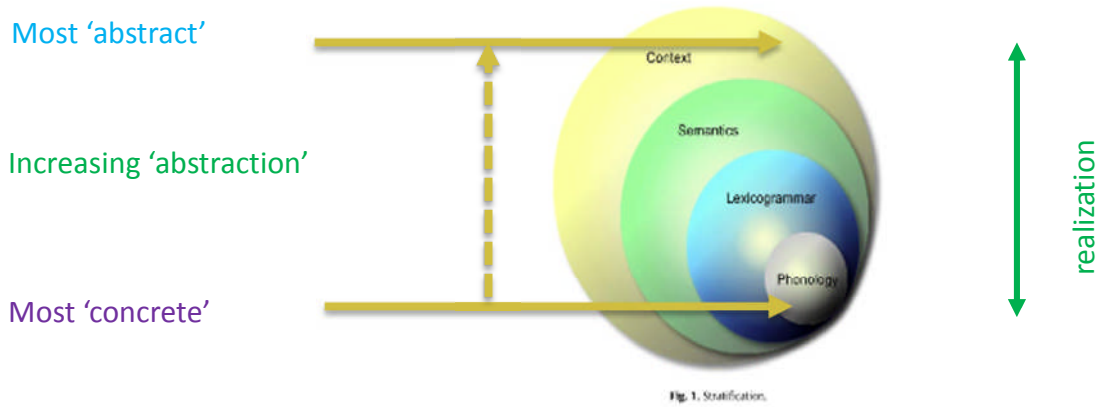


Figure 5.6: Abstraction and the Stratificational framework.

With this in mind, it is now possible to follow the same link of enquiry as taken with regards to the two previous theoretical dimensions.

5.4.3.1. Omission?

The relevant question in the first case is: "do the higher strata in the stratificational model merely omit information that is contained in the stratum immediately below them?" The answer to this question I hope to motivate in this section is in the negative. My reasons for thinking so is based upon two facts: (1) stratification is considered to be the process of "talking about language under different headings' (Halliday and Matthiessen, 1999: 24), and (2) the process of *abstraction as omission* can never add new information at the more abstract level of description, simply omit information that is found in the more concrete counterpart.

Regarding the first of these two points, Halliday and Matthiessen (*ibid*) state that the stratificational framework enables linguistics to think and talk about language in different ways. It can be described as an instance of "sounding" (phonology), a system of "wording" (lexicogrammar), or a system of "meaning" (semantics). The final stratum considers the semiotic environmental system that language is coupled to (context). In other words, each of these strata represents the different ways that it is possible to view the complex phenomenon of language; they are different perspectives that are useful to the tasks of the linguist.

With this in mind, one might be tempted to think that the notion of omission is well applied here: for surely, if each stratum only focusses on some aspects of the system of language, then they, by definition, they must be omitting others. Furthermore, the choice of description at a particular linguistic stratum is usually linked to a particular purpose in the mind of the researcher. Also, the nature of the holistic nature of the stratificational framework seems to counter any claims to explicit misrepresentation of omitted facts. As has been pointed out in both §3 and §4, one of the major ideas in SFL is that the information provided at any stratum in SFL can only be fully appreciated with reference to the features that exist at the strata both above and below the one under consideration, as well as in the context of the other features that are placed at the same level of interpretation. Therefore,

the information not contained at a particular level of representation is not idealized away, but is explicitly found in the other strata which features need to reference in order to gain their full worth in the theory.

Notice though, that this is not quite what the term *abstraction* is referring to in this case. Remember, the term as applied to this dimension moves *along* the stratificational framework. In the sense that it is being discussed here, higher strata are linked to the lower ones by means of *abstraction*, and given the definition at the end of §5.3.2., this means that the omission process should be on the basis of the immediately lower stratum, not from the text itself. In order to make this point clearer, repeated below isfigure 3.4 from §3.2:

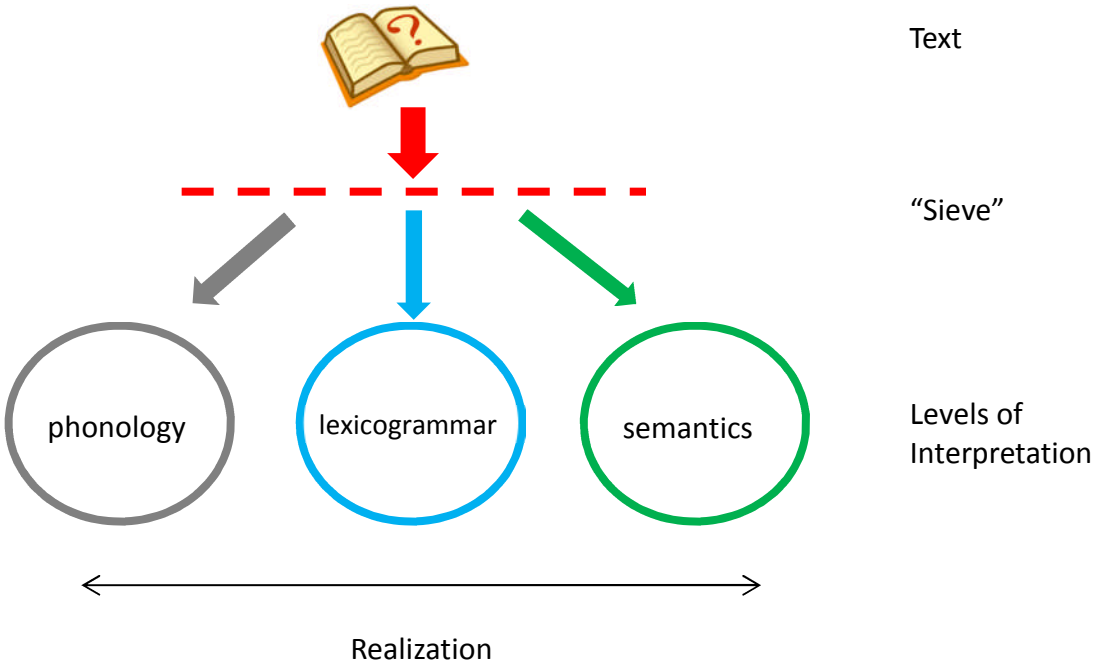


Figure 5.7: From text to stratification

The aim of the above figure to emphasizing that in a certain way, there is a process of *abstraction as omission* taking place in the stratificational framework, but that it is not involved between the strata themselves. Instead, it enters the figure in giving a description at a certain level of interpretation, that is, by choosing to focus only on certain features of the text out of all the possible features that it would be possible to include. However, this is

different to saying that the **move from (say) lexicogrammar to semantics involves omission of detail.**

My argument here is that no such process takes place when one comes to look at such inter-stratal relations. The reason for this is based upon the fact that a process of omission cannot add or re-introduce any information that is not found within the level of description that the operation is performed on. During the "filtering" process outlined above, a certain amount of information has been omitted at every stratum within the overall framework. On the view that each stratum represents a different 'perspective' one can take with regards to the system of language, at the stratum of, say, phonology the only information included are those that pertain to the view of "language as sounding". The other possible information that one could give for the linguistic system is not included. For instance, the information relevant for viewing language as a system of "sounding", or "meaning" does not enter into our description of the phonological system. This is precisely the reason why linguists set up the different levels of interpretation initially, so that the task of investigating such a complex system can simplified by certain aspects of it being dealt with at different places within the overall theory.

However, this view of piecemealing language into different levels of interpretation is problematic for the view that the link between the strata themselves involves omission of information. We can think of each stratum as consisting of a set of features, ordered in a certain way due to the principle of delicacy: as a paradigmatic system. But crucially, the various sets that make up the entire stratificational framework are non-overlapping, i.e. they each have no feature in common. A feature such as *declarative mood* only has meaning at the lexicogrammatical stratum, to place it elsewhere would amount to a level mismatch or category mistake. The only way, therefore, that a move upwards in terms of stratification could involve omission of detail would be if all the information at the higher level of interpretation is also contained at the level immediately below it. To be specific about the strata themselves, this would be tantamount to the proposal that taking the interpretative stance of phonology would involve not only viewing language as system of "sounding", **but also of "wording" and "meaning"**. In other words, the lowest stratum in the model would

involve describing language in all its detail, with information gradually being discarded as one moves further up the stratificational hierarchy.

This is, I believe, an untenable position, or at least a viewpoint that is contrary to the way the strata within SFL have been designed and the role they are meant to play within linguistic theorizing and description. For this reason, it seems as though the meaning of the term *abstraction* as applied to the stratificational framework does not involve the omission of information.

5.4.3.2. Generalization?

The next criteria to be discussed is whether higher strata, by virtue of being more abstract, generalize over those below them by means of either retaining common properties that are evident in a number of linguistic types at that lower level of interpretation. Again, I think the correct answer to this question is in the negative, and the reasons for this are similar to those just presented in §5.4.4.1., but before we turn to these arguments, it is worth discussing an occasion in the literature where stratification has been directly linked with the idea of generalization.

Martin and Matthiessen (1991: §2.5) indeed do consider the relations between different strata to involve generalization. For example, in their work they say the following (*ibid*: 353, emphasis in the original):

> *The most powerful generalization is, arguably, that there is a tendency towards **generalization** as we move down […] the stratal organization […]. That is, to put it crudely, a given number of semantic systems are realized by a smaller number of lexicogrammatical ones; similarly, a given number of lexicogrammatical systems is realized by a smaller number of phonological ones.*

It is important to note here, however, that regardless of whether the authors are right in their assertion, their use of *generalization* here is markedly different to how it has been used

so far in this work. In this case, they are not commenting on retention of common properties, but instead their proposal refers to the nature of the relations between systems at different points on the stratificational hierarchy. Their proposal is that multiple higher-level system networks stand in a realization relationship with a fewer number of lower-level systems. For this reason, this ascription alone is not evidence for the view that generalization is involved in the move from lower to higher levels of interpretation.

This, of course, does not count against the view that *generalization* as it is being used in this thesis is involved. However, there are reasons to doubt whether this interpretation of *generalization* is involved. The argument against generalization in this case is largely similar to that given above. The process of *generalization*, much like that of *omission* can not introduce information that is not contained in the base description. Following similar arguments to those just presented, if the generalization procedure cannot introduce or re-introduce new or previously discarded information then the lower strata within the framework would have to include all the information also found at the higher levels, so that they could be retained after the generalizing process. Similar to the last section, I take this to be antithetical to the nature, purpose and aims of the stratificational framework. For this reason, there is reason to believe that a move in *abstraction* upwards on the stratificational framework does not involve generalization, at least not in terms it has been discussed in this chapter.

5.4.3.3. Decontextualization?

The final issue to be discussed with regards to stratification is whether lower linguistic strata are more closely linked to particular situations than those higher up. As has been argued to be the case with the other two criteria dealt with so far, I believe the answer to this is again in the negative. This section will provide the arguments why I think this is so.

The system of language, as is clear from the discussion in chapter 3, is taken to model not just all previous language use, but also all possible language use. If the abstraction process between strata involved decontextualization, then one would expect the lower strata to be more closely tied to particular situations than the higher ones. In other words, the system of

sounding would have to be considered to describe a lower number of situations than the systems of both wording and (linguistic) meaning. This, however, would seem like an odd proposal to put forward. Both the descriptions of the meaning potential of language at the phonological and lexicogrammatical strata are meant to account for all the possible meanings that language users can make using linguistic resources. If decontextualization was in operation in this case, then we would expect to find (considering spoken language) actual instances of language use which exhibit grammatical properties but not phonological ones. The same would also hold between the lexicogrammatical and semantic strata, whereby there would be features in the semantic stratum that have no realization at the lower level of interpretation. Although it is a difficult question to measure the notion of *potentiality* in these cases, and no attempt will be made here, this does seem like an intuitively odd idea.

The situation seems even worse when one comes to consider the analysis of just a single text: by looking at a move in stratification at the instance end of the cline of instantiation. Given a certain (spoken) text, it is possible to provide a description in terms of the area of potential phonological space it instantiates. The same is possible at the other levels of interpretation, given the fact that a linguistic text simultaneously instantiates all of the linguistic sub-systems. If we were dealing with a move decontextualization then this would not be possible, because one would be forced to detach oneself from the immediate context of the text itself. In other words, it would be impossible to give a lexicogrammatical description of a single text if a phonological description is also available. I take this position to be absurd. Given a certain text, it is indeed possible to give a description at all levels of interpretation, and neither of these descriptions is more or less tied to particular situations, **because they are all descriptions of the same text**.

One further factor about the links between stratification and decontextualization needs to be briefly discussed, in order to clarify any potential misunderstandings. The highest level of interpretation within the stratificational framework is named *context*. One might be tempted to think, therefore, that this level would be highly contextualized, since it provides a description of contexts themselves. It is important to remember, however, that within SFL the contextual level is a genuine stratum; that is, it is also a system, as was discussed in the previous chapter in some detail. In the same way that the system of language is intended to describe the complete meaning-making potential that language possesses, the semiotic

contextual stratum aims to map the whole range of possible social situations that are relevant for language use. Although not completely detached form situations themselves, just like the linguistic system, it is highly decontextualized, since it provides the map of not all past experienced situations but also all future ones. Even if this was not the case, if one considered the highest point on the stratificational hierarchy to involve a higher degree of contextualization, then the picture that emerges is that the more *abstract* point on the framework would be more contextualized than the more *concrete* ones. Hopefully, these points can clarify the potentially problematic nature of the stratificational framework containing an explicitly contextual level of interpretation.

There is little more that can be said about this issue. It seems that *abstraction* as applied to the links between strata cannot carry the meaning of decontextualization. The reasons given for this viewpoint is that, if accepted, it would deductively lead to two untenable proposals. Firstly, the linguistic potential for semantics would cover a wider range of situations than either lexicogrammar or phonology. Secondly, simultaneous instantiation would be impossible on this view, as by necessity higher levels of interpretation would have to be less associated to any one situation than those further down. For these reasons, I do not believe that the meaning of *abstraction* in this case involves decontextualization.

5.4.4. A proposed typology of abstraction types in SFL

In sum, the previous section has argued that although each of the theoretical dimensions in SFL have been nominally associated with the notion of *abstraction*, there is really more than one 'type' of *abstraction* in operation. The results of this discussion can be summarized as follows:

	Omission	**Generalization**	**Decontextualization**
Delicacy	✓	✓	×
Instantiation	×	✓	✓
Stratification	×	×	×

Table 5.1: Abstraction Types in SFL

On the face of it, this raises a surprising situation. As was discussed in §5.1., out of all the theoretical dimensions discussed in this chapter, *stratification* is the one that is most often associated with abstraction. However, if the arguments presented here are correct, this is the one dimension that is the least similar to *abstraction* as it has been discussed in the previous literature from philosophical viewpoints on scientific practice. It is worth, then, returning to this final dimension. As will be seen, we may have already stumbled upon the relevant meaning of abstraction as it applies in this case.

5.4.5. A return to stratification

If the arguments presented in §5.4.3 hold, then we are yet to successfully identify the core meaning of *abstraction* in the case of stratification. The situation as it has been presented so far is as follows: each stratum involves the omission of some detail of the text, choosing only to focus upon certain aspects. This fact has led to the conclusion that the explicit ordering of the strata in SFL cannot be due to either *generalization* or *omission*, since neither of these processes can re-introduce the omitted information that is "lost" during the "filtering" procedure that leads us from the text to its simultaneous instantiation at each of the levels of interpretation. However, this by itself does not provide us with the motivation for the ordering of the strata themselves. Why is lexicogrammar 'lower' than semantics, and 'higher' than phonology? Furthermore, how does this link to notion of *abstraction* given that in this case it seems not to fit any of the criteria discussed previously? This may seem like an obvious question to answer; there is a long tradition of thinking of semantics as "higher" than grammar within the study of language. However, at risk of labouring the point, unless explicit criteria can be given, accepted thought must be questioned, regardless of how intuitive answers to foundational questions may be.[35]

Fortunately, it is my view that there is a means of thinking about *abstraction* in the case of the stratificational framework that can provide the ordering we require. One of the main proposals in §3 was the *spatio-temporal hypothesis of stratification*, and I will argue in this section that this forms a good candidate for understanding the nature of *abstraction* in the case of stratification.

[35] It is important to note that I am only raising this question with regards to the stratificational framework as found in SFL. Reasons why other linguistic theories consider the different linguistic levels to be ranked in a particular order may well differ, and indeed may map on to the criteria already presented. However, I do not have the space to discuss this issue in other theories here.

But why would one think that appeals to temporal and spatial localization could form the basis for *abstraction*? There are hints that an increase in abstraction involves a move in spatial and temporal extent when you look at some accounts in the scientific literature, and also from other aspects of linguistics. For example, Levins (2006: 744) provides a discussion on the various perspectives one can take on how fruit flies react to changes in temperature. The following is an adaption of the relevant abstractions of perspective:

Perspective	**Horizontal Scale**	**Temporal Scale**
Temperature tolerance	Individual Fly	Minutes to hours
Adaption to temperature	Individual Fly	Days to a week
Behaviour in relation to temperature	Population of flies of one species	Minutes
Demography	Population of flies of one species	Seasonal
Community	Ecosystem of interacting species	Months to years
Micro-evolutionary	Single species	Years

Levins (2006: 747) describes such differing perspectives as "nested sets of abstraction", implying a specific ordering between them. Focusing on the horizontal scale and temporal scale, there is a general trend for the particular perspectives to increase their scope in time or space (or both). The question is: which ones should be considered the "higher" levels? I take it that notions such as demography and community are more abstract than temperature tolerance. So, in this example at least, it seems as though a move in increasing abstraction correlates with a wider perspective in both space and time.

Further evidence for this approach stems from consideration of the difference between abstract and concrete concepts. For instance, consider the following hierarchy of nominal semantic types (Hinzen and Uriagereka 2006: 79; Hinzen 2012: 126):

abstract < mass < objectual/ count < animate

In their discussion, the authors argue that the above hierarchy is one of semantic complexity. The authors construct this hierarchy on formal simplicity: "in a deliberately naïve sense, the mental space denoted by an abstract noun like *beauty* is intuitively simpler than that denoted by a concrete mass noun like *beer*, which is in turn formally simpler than an objectual space denoted by a count noun like *mug*. (Hinzen and Uriagereka 2006: 79). Their appeal to formal complexity need not concern us here, but one way in which they justify the hierarchy is in terms of relative boundedness of concepts associated with abstract, mass and count nouns. Concepts associated with count nouns, such as *mug*, are bounded in the sense that there exists a relatively coherent boundary to the object the concept is grounded in. In the case of abstract concepts, however, the situation seems to be different, with no obvious coherent single object in which to ground concepts associated with abstract nouns. Instead, it has been proposed that abstract concepts are grounded in situations or events with a wider focus. For example, Barsalou and Wiemer-Hastings (2005:136) hypothesise that "for concrete objects, the focus is on a relatively local, spatially circumscribed region of a situation […] for abstract concepts, however, the focus is on multiple components that are not localized but distributed widely." They then go on provide empirical evidence to back up this hypothesis. Granito et al. (2015: 1-2) also comment that the acquisition of abstract concepts may be more difficult for children because whereas some concepts refer to "concrete single elements", more abstract ones pose a problem because "their referents are not perceptually similar and can be given by complex relations between items." The notion of complex relations, or relations in general, operate between singular objects and occur over time, and therefore involve an increase in spatial and temporal distribution.

We have evidence, therefore, that there is a link between *abstraction*, and spatial and temporal distribution. If the arguments given previously for the applicability of these notions to the stratificational framework, then I believe that this fact can account for at least part of the meaning of *abstraction* in this case.

5.5. Conclusion

The purpose of this chapter was to investigate the notion of *abstraction* in SFL to see if it has the same meaning in all its uses. Specifically, this relates to the theoretical dimensions of

stratification, *delicacy* and *instantiation*. It has been shown in this chapter, that this term has a different meaning in all of these cases. To conclude, a few comments will be made on how this fits into the general work presented in this thesis.

The reason for undertaking this work was to investigate whether being more specific about the nature of the theoretical dimensions in SFL can aid the development of the general cognitive framework under development. I believe that this discussion was worthwhile, since the notions of *generalization*, *omission*, and *decontextualization* are important parts of the theory building processes, as indicated in the range of studies overviewed in this chapter. Since the discussion of *stratification* is along the lines previously discussed in §3, this commentary will only involve *delicacy* and *instantiation*.

Delicacy operates over the dimensions of *omission* and *generalization*. It involves grouping together systemic types based upon their similarities, and thereby omitting the differences between them. From a practical perspective, if, for example, neurophysiological types begin to be identified, then this generalization process will prove useful for finding the right level of specificity for framing discussions of a certain cognitive phenomenon. Furthermore, the appeal to *omission* makes explicit the idea that it is optimal if the cognitive potential at any level of representation is stated at an appropriate level. Seemingly irrelevant details, or more specific options, should be omitted from the description, but enough detail should remain for theoretically meaningful patterns to emerge.

Instantiation, on the other hand, involves the processes of (statistical) generalization and decontextualization. This is the driving force behind making wider conclusions from the data. Its use of statistical concepts results in the levels of regularity discussed previously are couched in what is likely to occur. That is, given a certain contextually-motivated act, what options within the **psychological** and **neurophysiological potential** are likely to occur. Appeals to the environment can then be made to enquire about why this might be so.

Chapter 6

Conclusions and Looking Forward

"It is because you are a very small animal that you will be Useful in the adventure before us."

A.A. Milne, Winnie-the-Pooh

The purpose of this thesis has been to develop and defend a conceptual framework for cognitive science that is based upon treating such systems as *dynamic, open* and *semiotic*. The major claim in this thesis is that couching the cognitive theory in the terms of SFL allows the conceptual positives of the *4E approach* to cognition to be maintained, while providing a statement of the kinds of *levels of analysis* that are necessary for studying cognition from such a viewpoint. This specific statement was missing from previous approaches, since either the issue of *levels of analysis* had been either missing from some proposals, or had not been expanded upon sufficiently.

In arguing for this position, this thesis has concerned itself with a number of different disciplines, and has attempted to answer the research questions by examining a number of different avenues. The discussions in this thesis have encroached upon the disciplines of **linguistics**, **psychology**, **neuroscience** and **philosophy**. Additionally, the arguments made have concerned brain-behaviour relations (§3), *environment-organism* interactions (§4) and the notion of *abstraction* in building frameworks (§5). As such, it is worth ending this work with a statement of the main ideas and proposals made, and also to provide a means of bringing all these different elements together by means of an example.

In order to conclude this work in this way, this final chapter is organized as follows. It begins with a statement of the proposals in this thesis. This is provided so that the novel aspects of this work can be easily outlined. After this, the classical neuropsychological experiment of the *Morris Water Maze* is used as an example to show how the ideas can be used. It ends with a discussion of future work.

6.1. Proposals in this thesis

At this concluding point, it is worth stating explicitly the proposals in this thesis believed to contain some degree of novelty, thereby satisfying the requirements for a doctoral thesis. This work has worked within two disciplines, both SFL and cognitive science more generally, and as such different novel aspects of this thesis pertain to either one of these two fields.

6.1.1. Proposals in Theoretical SFL

In our discussion of SFL theory, I feel that two proposals have been made with regards to this subject. They are:

I. The Spatial-Temporal Hypothesis of Stratification
II. A Typology of Abstraction types in SFL

The first of these proposals pertains to a certain way in which the specific ordering of strata within the stratificational framework can be considered on a more objective basis; it is a justification of the ordering posited between strata, that is, not based upon any intuitive notions of what should count as *phonological*, *lexicogrammatical* or *semantic* facts. The basis of the proposal is that with each increase in stratification we are dealing with units that are defined in terms of increase spatial and temporal distribution. At the lowest level of phonology, the units are defined in terms of an *intrinsic* basis, i.e. on notions that are properties of the units themselves. Evidence for this came from the nature of *phonemes* and *tone types*. At the stratum of *lexicogrammar* the units are defined in terms of relations between the unit and others within a relatively restricted spatial extent, usually within the clause boundary itself. Furthermore, whereas features at the lower level mapped onto to units with a more easily defined boundary, this was seen to not necessarily hold at this higher *level of interpretation*. Evidence for these reasons came from the system networks of MOOD and THEME. Finally, the semantic stratum was seen to be based upon factors that extend into the wider semiotic context. This stems from the notion of semantics as the construal of the speaker's interaction with the semiotic context.

The second of these proposals concerned the nature of *abstraction* as relevant to the theoretical dimensions of *stratification*, *instantiation* and *delicacy*. It was argued that, although the same term is used in each case, on closer examination differences in the

specifics can be seen. The notions of *omission*, *generalization*, and *decontextualization* were used for this purpose. The three dimensions just mentioned were argued to differ in the combinations of these three factors they typified, or in the case of *stratification*, none of them. This is a useful finding, not just is the context of the work in these pages, but for the field more generally. This has aimed to bring some extra clarity to these notions, all of which were previously linked to the same notion. Hopefully, this can be of practical use to scholars using these terms in the future.

6.1.2. Proposals in Cognitive Science

The second group of proposals make up much of the content of this thesis, as they relate to the development of the particular conceptual framework under discussion:

i. **Stratification** in cognitive systems and **bi-directional realization**.
ii. The inclusion of novel dimensions of *delicacy* and *instantiation* in cognitive frameworks.
iii. Considering the cognitive environment as a **system**
    a. The environment is **amenable to description using the same terms** as other strata in the framework.
    b. The **environmental parameters** of *objective*, *relational* and *goal.*

Initially, the proposal was given that it is possible to consider cognitive systems as **stratified** systems, consisting of a number of *levels of interpretation*. Each of these was considered as a potential, in line with SFL's perspective on language. In this it shares certain conceptual ground with **dynamic proposals**. At the lowest stratum, this involves considering the **neurophysiological system**: the set of **possible brain states** the organism can be in. At a higher stratum the **psychological system** is stated, providing an account of the set of **possible cognitive acts** the same individual can exhibit.

Considering both these points as **systems**, they can then be discussed in terms of delicacy, i.e. at a number of different *levels of specificity*. This dimension divides up the vast potential in each case in more or less specific ways: it might state different types of oscillating behaviour, different firing codes, different ways of moving, or remembering. The **system**

perspective, however, should never be forgotten. In any account, the relevant question is why did such and such activity take place, and not the others.

Finally, cognitive system can be described in terms of the **cline of instantiation**. This captures the notion that cognitive organisms are not static entities, but are constantly being shaped by the output of their own cogntive system. The constant interplay between system and instance results in a constantly adapting system. It also allows for further theoretical notions such as *neurophysiological type*. This amounts to grouping individual instances together based upon common properties, and treating them all at as a whole. It is then possible to look towards the higher strata of the **psychological system**, and the **environmental system** to investigate whether there are motivating factors in terms of explanation. This does not only work from a "bottom-up" perspective. For instance, a particular **psychological type** can only be fully explained by considering the higher stratum of **environment,** and the lower stratum of **neurophysiology**. In this way, one part of the system can only be considered in terms the system as a whole, in line with the systemic perspective on language.

Another major proposal in this thesis is that the environment can be described in the same way as the other cognitive strata. This is because this highest stratum is intrinsically involved in the meaning-making process of cognition. In similar way to above, this means that we can think about the environment as a **system** of possible combinations of the environmental parameters of *objective*, *relational* and *goal*. These refer to, repsectively, the objects in the semiotic environment, and the realtions between them, the relations between the individual and those objects, and the situation the individual is trying to bring about. Using these as a starting point, it is possible to consider the environmental system as also operating over the dimensions of *stratification*, *instantiation* and *delicacy*, although only preliminary ideas could be put forward in this work.

6.2. Putting it all together – Morris Water Maze Experiments

In order to bring together the discussions in the previous chapters, it will be beneficial to give an example of an oft-performed psychological experiment that looks into spatial memory in rats: the Morris Water Maze (Morris, 1984). The reasons for choosing this particular example is because it is an oft-used paradigm that has been used in a number of

different cases, including the development of neural systems (Garthe and Kempermann, 2013), and neural degeneration (Ma et al., 2014), including studies of particular neurodegenerative models, such as Alzheimer's disease (Bromley-Brits et al., 2011). As such, there has always been a link between behavioural behaviour and the study of the underlying neurological correlates for such behaviour (c.f. D'Hooge and De Deyn, 2001 for a review).

The purpose of this summary section is to lay-out step by step how conceptualizing this particular experiment along the lines advocated in this thesis might initially take shape. As discussed in §6.3, however, this remains at the initial stage.

A water maze task involves placing an organism (usually a rat) inside a circular bath of opaque liquid. At a certain point within the bath is a submerged platform, and it is the task of the animal to find said platform by swimming and exploring the water bath. In order to hide the platform from direct visual identification, the water within the maze is made opaque by the addition of milk or another substance. Around the environmental setting are a number of visual cues that can be used at reference points for future location of the hidden platform in later trials.

An experiment generally unfolds as follows. On the first instance, an animal is lowered into the bath. The temperature of the water is set so as to be uncomfortable in order to encourage the animal to find the platform, but not so extreme as to cause undue anxiety. Upon exploration, the animal will eventually find the submerged platform. In the training procedures the animal is lowered additional times into the bath. The hope is that it can learn the position of the submerged platform using the visual cues as reference points. Experiments can then measure certain neurological signals in normal cases, and thereby identify neurological correlates to spatial memory and navigation, or test certain types of manipulation to see if they disrupt the animal's ability to re-find the goal area.

On the cognitive framework outlined in these pages, interpreting and studying cognition in this particular task should proceed as follows:

I. Start with the nature of the semiotic environment

As in the case of systemic linguistics, the first port of call in any interpretation should be the environmental setting in which a particular behavioural act is taking place. Instead of treating the environment as some unorganized, heterogeneous entity, the perspective taken in this work is that the environment can be structured along the *environmental parameters* of *objective*, *relational* and *goal*. The factors that make up these various notions have been given in the previous section. A particular *environmental instance* can then be considered to be a collection of values for these three notions. This would then form an *instance* of the environmental system for that particular organism. The caveat "for that particular organism" is important, because the nature of what can potentially enter the semiotic environmental system will depend upon physiological properties of said individual. This *environmental instance* can then form the basis of any instantiational analysis, as will be discussed shortly.

For our particular example here, we can consider these three environmental groupings in the context of the Morris Water Maze. It should be re-iterated from the discussion in §4.2.1 that the elements that count for inclusion in these sets are those that are potentially *meaningful* for the organism, and that are relevant for the task at hand. In this way, the semiotic environment can be considered an *abstraction* from the overall environment, in a certain sense of the term as made clear in §5. With these two factors in mind, for most iterations of a *Morris Water Maze* task, the environmental groupings can be given as follows:

Objective: Visual cues, their characteristics in terms of colour, shape, etc. How similar they are to one another.

The submerged platform, and its location.

Relational: On initial trials, all visual cues will be novel for the organism. They will also stand in a certain spatial configuration from the perspective of the animal.

Goal: To find relief from the water maze

Of course, the maze is not only attempted once. On future training and experimental trials, these values will differ:

Objective: Visual cues, their characteristics in terms of colour, etc.
How similar they are to one another.

The submerged platform

Relational: Visual cues will now be experienced. They will stand in the same spatial relations from the perspective of the animal.

Goal: To find the platform.

This characterization is essential for the framework defended in this thesis, as it takes to view that cognitive systems can only be fully understood in their environmental setting. This will always form the start point for any description or explanation. This change of focus from 'brain-focused' approached is stark: every explanation begins in such environmental terms.

II. Define the instance at both the psychological and neurophysiological levels

With the semiotic environment described along the above notions, the next step is to describe the activity of the animal at the other *levels of interpretation*. That is, interpreting the event as an instance within the overall behavioural potential, and within the overall neurological potential. During such experiments, electrical activity will be recorded from certain structures in the rat's brain, typically the hippocampus, as this has been associated with spatial memory storage (Neves et al. 2008). Such a recording is thus considered to be an **instance** of the underlying **neurophysiological system**. The system perspective forces the relevant question to be in terms of what other instances could possibly have been witnessed, instead of purely looking at any structural properties of the signal.

The same takes place at the higher level of interpretation, but in this case we interpret the task in terms of **cognitive acts**. The relevant measure in this case will be the movement of the rat during the experiment, the paths it takes. Again, this is an **instance** of the underlying **psychological system**. Similar questions as above then arise.

III. Comparison with future cases: deciding upon rank, delicacy and instantiation

One of the issues with the above framework is that there is little that can be gathered from only one instance. In order to make any reasonable conclusions, comparisons must be made between many different **instances** of the cognitive system in similar situations. There are a number of different ways in which this comparison can take place.

Firstly, the issue of the relevant unit of rank must be decided upon. Since the only rank scale considered in this work was at the lowest stratum, solely this case will be discussed here.

In terms of rank, it is necessary that the relevant meaningful unit of neuronal structure is identified. This may be the neuronal ensemble, or higher constituents such as macro-brain structures. This needs to be explicitly stated, because the options in the neurophysiological potential will be made with reference to this structural unit. Such facts will only be gathered from comparison of many cases.

Next the appropriate level of delicacy needs to be highlighted. In the neurobiological case, this means specifying the specificity in the neuronal responses that are deemed to be important. This might be in terms of oscillations, as outlined in §3, or may be more specific in terms of individual firing rates of the units within the structure. This cannot be known *a priori*, but must be at a level where patterns can be seen without being drowned out by noise.

Once a system network, at a particular level of rank, has been identified, it is then possible to compare individual cases based upon shared similarities, i.e. start to consider a move up the instantiational cline. As just outlined in §6.1, this allows a collection of similar instances to be explained in terms of the system as a whole; by comparing them to the other co-instantiating systems.

iv. Alter the environmental parameters to see if they have any effect

Since this is a contextually-driven framework, the final stage I would propose would be to alter the values for some of environmental parameters of *objective*, *relational* and *goal*. Not only will this allow preliminary accounts of these notions to be tested, but it would also allow for the creation of *environmental types*: different types of non-identical environments that nevertheless share common features. In this way, a picture of how an explicit embedded theory of cognition may look will begin to become evident.

6.3. Future Work

Given the amount of work that can legitimately be completed in the course of a doctoral work, and the effort which is necessary to fully develop and detail any new conceptual

framework, the work contained in this thesis can only be considered to be the beginning, not the end. Therefore, to conclude it is important to stake stock of the stated aims of this work, how far these goals have been realized, and what future work can be performed if the arguments contained here are found convincing.

To begin, the research questions that began this work are as follows:

- Can theoretical notions from SFL, especially as they pertain to the kinds of *levels of analysis* important for studying semiotic systems, be of use to cognitive science more generally?
- More specifically, is it (a.) possible, and (b.) useful to describe cognitive systems along the dimensions of *stratification*, *delicacy*, *instantiation* and *rank*.
- What are the features of such a framework? What view of cognitive systems does it provide us with? How does this picture differ from current proposals in cognitive science?
- Can such a framework overcome some of the potential problems with current approaches?

All of these research aims will be revisited in turn, with a commentary given to how successful each of them has been answered in this thesis.

The first two research questions can be answered together, as essentially the second is a more specific rendition of the first. It was seen that there were good reasons for including all of the theoretical dimensions in SFL to cognitive science more generally, either because they are already thought to be important, or because there is some motivating neurological or psychological reasons for thinking they may be useful. I cannot say that they *will* be useful at this stage, so on point (b.) on the second goal, I feel that this thesis has fallen short in that regard. However, future research will help to determine whether this is true or not.

The general picture of cognition presented here, is in line with many of the proposals that are characterized under the **4E approaches**. The focus has not been on computational architectures, or mental representations, but instead on looking for patterns in activity at various points in time and space. The foregrounding of the environment makes it *embedded*, the idea of **potentials** make it *embodied* and the interaction between **instance**

and **system** have the potential to make it *enacted*. The idea of **potential** has links to **dynamic** modelling in cognitive science, although how much remains to be seen (c.f. §6.3.1). The **bi-directional** nature of the links between the various levels of interpretation in the model force us to consider cognition in terms of the substrate being used; it cannot be considered in some amodal means, as was seen in some proposals in the **classical** approach.

As far as the last research question is concerned, it is too early to assess whether the ideas put forward in this work can overcome the potential problems in terms of having an explicit conceptual framework in cognitive science for studying *embedded*, *enacted* and *embodied* cognition. This will only be ascertained once the ideas are put to use, have some value in the interpretation of experimental work, and if found to have some solid basis, altered accordingly. As a working hypothesis, however, I believe it has some merit. The increased structural mappings of the environment should help those interested in *embedded* notions, the idea of **neurophysiological** and **psychological potentials** only being considered in terms of general physiological properties should find favour with *embodied* researchers; the interplay between system and instance may go some way to being more explicit about the nature of *enacted* cognition. However, at this early stage, the jury is still out.

However, this thesis should be seen as the beginning of an alternative means of describing cognitive processes. As such, I hope that there are many different avenues of future work that will arise out of this work. In this section, however, I will restrict myself to future work that concerns building upon this initial proposal; the avenues listed here can be thought as the 'next steps' in the journey to producing a full explicit framework.

6.3.1. Mathematical Formalization

At the beginning of §2, a distinction was made between *conceptual frameworks* and *theoretical models*. As used in this work, the former referred to the abstract, informal ideas that guide the study of certain phenomena, and are usually stated in terms of *levels of analysis*. It is this aspect that this thesis has concentrated on. However, part of the usefulness of any framework is the predictions that can be made from it.

One of the most common means of forming explicit predictions is by means of certain mathematical formalizations. Throughout this thesis, the idea of *probability* has featured

heavily in the discussions of both SFL and its application to cognitive science, although no specific statistical terminology has been employed. The appeal to **probability** has been merely conceptual, in order to talk about the fluid nature of links between strata. Within SFL, although probability features heavily in discussions, it is rarely presented in terms of highly formalized mathematical or statistical models. Exceptions include Halliday's (1992a.;b.) work into the relative distributions of competing system features, and Munro's (2004) more recent proposal to couch the link between system networks and the rank scale in terms of Baysian probability.

However, beyond these steps there is still much work to do in this regard. For instance, Halliday and Matthiessen (1999: 15) talk about the ideational semantic base in terms of a "semantic space". It is unclear whether the term is meant in the same way to how it is used in dynamic systems theory, which also discusses potentials in terms of state spaces (c.f. §2.5.1).

I leave this issue of mathematical formalization open at this stage. It is my belief that such work is needed within SFL, especially concerning the idea of inter-stratal realizations. Ideally, some of the links identified in this work could lead to collaborative research between dynamic cognitive modellers and systemic lingusits.

6.3.2. Further steps towards the 'dynamic' nature of cognition

Throughout this work, the idea of *dynamic, open systems* has been central to the way in which this proposal aims to conceptualize cognitive systems. In the work as presented, the majority of this dynamic nature arises through the *cline of instantiation*, where certain regularities at different distributions in time and space become theoretically important. The means by which the cognitive *system* is shaped by the nature of the instances that are directly observable, amounts to one way in which this framework can be considered to be *dynamic*.

There is a potential other way in which added dynamicism can enter into this framework. This comes from discussions in SFL about the nature of *dynamic* context. As overviewed thoroughly by O'Donnell (1999), there have been various attempts to introduce *dynamic models of context* into SFL work. The crux behind these proposals is that a semiotic context

may not be fixed over the whole course of a linguistic interaction, but instead the various parameters of *field, tenor* and *mode* may change over the course of a linguistic event. There are two different types of such *dynamic* accounts in SFL. In the first group, the context changes over time, but independently of the linguistic utterances made (Halliday, 1984). The second group of proposals looks at the ways in which the text and context co-operatively evolve over the course of the linguistic act (Hasan, 1981:118).

Adaption of these ideas has not been carried out because although initial ideas have been given in SFL, I believe that this is a situation that needs further research in the theory. The problems with many of these proposals are given by O'Donnell (1999). These relate to a perceived *prescriptive* trend in these works. O'Donnell (*ibid*) has sketched an alternative proposal that aims to overcome this issue, but to my knowledge this facet of systemic theory has not received much attention since, as efforts still continue to map the contextual level from a systemic perspective (Bowcher, 2014; Hasan, 2014).

6.3.3. Case-studies of particular phenomena

As can be seen from many of the avenues just discussed, one of the biggest shortcomings in this thesis is the lack of detail in many of the ideas advocated in this work. I do not believe that this is fatally detrimental to the success of this thesis, given the research questions stated at the beginning.

However, I will concede that there is a lack of detail in many of the proposals, with the arguments presented here only providing a hint towards a way to proceed. There is only so much that can be argued for from a **conceptual perspective,** without further experimental work directly performed within the conceptual framework outlined in these pages. From the perspective, and within the timeframe, of this thesis, this gap is an unavoidable shortcoming. The amount of further research to begin to fill in the empirical 'flesh' to the conceptual 'skeleton' provided in these pages would be vast. One of the main reasons for this is that, similarly to the way in which SFL views the study of language, the perspective put forward in this work is it puts **action** first. By action, I mean instantiation of both the neurological and psychological strata that make up the stratificational picture in actual experimental and real-life instances. This differs, especially, from the classical framework, that made certain assumptions concerning the universality of various higher-level functions

in a number of physical realizers, enabling neurological facts to be back-grounded in their proposals. This is not possible on this view, since any cognitive or behavioural act can only be fully understood by reference to the lower-level neurological interpretation, and the higher-level of the environmental parameters that the cognitive system is realizing. The price to pay in this view is that a large amount of data is needed to get the framework up and running.

Additionally, the *environmental parameters* proposed in this work remain a working hypothesis. For the reasons given in §4.3.2, I think the characterization given is a sensible initial step, but whether they are correct would depend on description of particular experimental or 'natural' cases, and manipulation of the proposed values to see if they have any effect.

In sum, in all of the ideas presented here the work has only just begun. Only by attempting to use the interpretations and proposals in this work can they be shown to be of any use to cognitive scientists or not. I feel that there are good **conceptual** reasons to think that this is a real possibility, but at this stage I am unable to provide anything more than a conceptual justification for this work. However, this is enough at this stage and in the context in which this work was performed. Beginning to test these ideas will be a collaborative effort from a number of different disciplines. As the proposal currently stands, there are a few areas that I feel need particular further elucidation. These will be outlined briefly below.

Firstly, although the lowest stratum of neurophysiological signals is relatively well understood in terms of the kinds of features that would constitute such a system, this situation is less clear as far as the higher stratum is concerned. One relevant debate to the nature of this higher level stratum concerns the usefulness of so-called *folk psychological models* (Ravenscroft, 2010). This debate concerns whether our everyday vocabulary can act as an indicator of the types of categories that pertain to psychological theory. Proponents of this idea include Lewis (1975) and Fodor (1975; 1990). Prominent amongst the critics of *folk psychology* are P.S. Churchland (e.g. 1986) and P.M. Churchland (e.g. 1981), amongst others. In §3.8.3, some possible applications of this work in the context of this issue were discussed.

I'd like to end this work with a short statement from Brenner's (1999) commentary on Wittgenstein's *Philosophical Investigations*. In setting out the business of doing philosophy

he opines that “writing philosophy is like presenting sketches of a landscape made in the course of long and involved journeys” (*ibid*: 9). Although this work does not claim to have achieved the lofty standards of academic philosophy, this perspective is one that rings true here. If this work is considered to only be the start of a journey of discovery into the nature of cognition and how it can be studied **systematically** in an *embedded*, *embodied* and *enacted* means, then I am content. The purpose of this work has to be to attempt to show that such a path is possible, plausible and to justify its inclusion as an option for those investigating the issues of the *mind* and its links to the *brain*. There is much still to do, and certain avenues on this path may yet be illuminated, or shown to be dead-ends. However, the journey remains an exciting one, and its completion will take the combined efforts of a number of researchers in a variety of disciplines. The hope, then, is that this work can prove useful as part of a road-map to the final goal.

Bibliography

Adams, Fred. 1979. 'Properties, Functionalism, and the Identity Theory.' *Eidos*, 1(2), pp. 153-179.

Adams, Fred and Kenneth Aizawa. 2001. 'The bounds of cognition.' *Philosophical Psychology*, 14(1), pp. 43-64.

Adams, Fred and Kenneth Aizawa. 2008. 'Why the Mind is Still in the Head.' In: Phillip Robbins and Murat Aydede (eds.), *The Cambridge Handbook of Situated Cognition*. Cambridge: Cambridge University Press, pp. 78-95.

Adams, Fred. 2010. 'Why we still need a mark of the cognitive.' *Cognitive Systems Research*, 11(4), pp. 324-331.

Allen, Timothy and Joshua May. 2014. 'Does Opacity Undermine Privileged Access?' *International Journal of Philosophical Studies*, 22(4), pp. 617-629.

Bache, Carl. 2010. 'Hjelmslev's Glossematics: A source of inspiration to Systemic Functional Linguistics?' *Journal of Pragmatics*, 42, pp. 2562-2578.

Bäck, Allen. 2014. *Aristotle's Theory of Abstraction.* Dordecht: Springer.

Baker, Lynne. 1999. 'Folk Psychology.' In: Rob Wilson and Frank Keil (eds.), *MIT Encyclopaedia of Cognitive Science*. Cambridge, MA: MIT Press, pp. 319-320.

Bakkum, Douglas J., Zenas C. Chao, and Steve M. Potter. 2008. 'Spatio-temporal electrical stimuli shape behaviour of an embodied cortical network in a goal-directed learning task.' *Journal of Neural Engineering*, 5(3), pp. 310-323.

Barsalou, Lawrence W. 2003. 'Abstraction in perceptual symbol systems.' *Philosophical Transactions of the Royal Society, B (Biological Sciences),* 358, pp. 1177-1187.

Barsalou, Lawrence W. 2005. 'Abstraction as Dynamic Interpretation in Perceptual Symbol Systems.' In: L. Gershkopf-Stowe and D. Rakison (eds.), *Building Abstract Categories.* Carnegie Symposium Series. Majwah, NJ: Erlbaum, pp. 389-431.

Bartlett, Tom. 2013. '"I'll manage the context": context, environment and the potential for institutional change.' In: Lise Fontaine, Tom Bartlett and Gerard O'Grady (eds.), *Systemic Functional Linguistics: Exploring Choice*. Cambridge: Cambridge University Press, pp. 342-364.

Bechtel, William. 1994. 'Explanation in Cognitive Science.' *Minds and Machines,* 4, pp. 1-25

Bechtel, William. 2002. 'Representations: From neural systems to cognitive systems.' In: William Bechtel, Pete Mandik, Jennifer Mundale and Robert S. Stufflebeam (eds.), *Philosophy and the Neurosciences: a reader*. Massachusetts, MA/ Oxford: Blackwell, pp.332-348.

Bechtel, William and Adele Abrahamsen. 2002. *Connectionism and the Mind*, 2nd Edition. Malden, MA: Blackwell.

Bechtel, William and Jennifer Mundale. 1999. 'Multiple Realizability Revisted: Linking Cognitive and Neural States.' *Philosophy of Science,* 66(2), pp. 175-207.

Beckermann, Ansgar. 1992. 'Supervenience, Emergence and Reduction.' In: Ansgar Beckermann, Hans Flohr and Jaegwon Kim (eds.), *Emergence or Reduction: essays on the prospects of non-reductive physicalism.* Berlin: Walter de Gruyter, pp. 94-118.

Beer, Randall D. 2000. 'Dynamical approaches to cognitive science.' *TRENDS in Cognitive Science,* 4(3), pp. 91-99.

Benson, James, and William Greaves (eds.). 2009. *Functional Dimensions of Ape-Human Discourse*. London: Equinox.

Bentler, P.M. and G. Speckart. 1979. 'Models of attitude-behaviour relations.' *Psychological Review,* 86, pp. 452-464.

Bi, Guo-qiang and Mu-ming Poo. 1998. 'Synaptic Modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength and postsynaptic cell type.' *Journal of Neuroscience*, 18(24), pp. 10464-10472.

Blackburn, Simon. 1984. *Spreading the Word.* Oxford: Oxford University Press.

Block, Ned and Robert Stalnacker. 1999. 'Conceptual analysis, dualism and the explanatory gap.' *Philosophical Review,* 108(1), pp. 1-46

Bowcher, Wendy L. 2007. 'Field and multimodal texts.' In: Ruqaiya Hasan, Christian M.I.M. Matthiessen and Jonathan J. Webster (eds.), *Continuing Discourse on Language: A functional perspective, Vol. 2*. London: Equinox, pp. 619-646.

Bowcher, Wendy L. 2014. 'Issues in developing unified systems for contextual field and mode.' *Functions of Language*, 12(2), pp. 176-209.

Brass, Marcel and D. Yves von Cramon. 2004. 'Selection for Cognitive Control: A Functional Magnetic Resonance Imaging Study on the Selection of Task-Relevant Information.' *Journal of Neuroscience*, 24(40), pp. 8847-8852.

Brenner, William H. 1999. *Wittgenstein's Philosophical Investigations*. New York, NY: SUNY Press.

Buccino, G., F. Binofski, G.R. Fink, L. Fadiga, L.Fogassi, V. Gallese, R.J. Seitz, K. Zilles, G. Rizzolatti and H.-J. Freund. 2001. 'Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study.' *European Journal of Neuroscience*, 13, pp. 400-404.

Burton, D. 1980. *Dialogie and Discourse: A Sociolinguistics Approach to Modern Drama Dialogue and Naturally Occurring Conversation.* London: Routledge.

Butler, Christopher. 1985. *Systemic linguistics: theory and applications*. London: Batsford.

Butler, Christopher S. 2003. *Structure and Function: A guide to three major structural-functional theories. Vol 1. Approaches to the simplex clause.* Studies in Language Companion Series, 63. Amsterdam/ Philadelphia, PA: John Benjamins.

Butt, David. 2004. *Parameters of Context*. Mimeo. Department of Linguistics, Macquarie University, Sydney.

Butt, David, Rebekah Wegener and Jörg Cassens. 2013. 'Modelling Behaviour Semantically.' In: P. Brézillon, P. Blackburn and R. Dapoigny (eds.) *CONTEXT 2013*. Berlin/ Heidelberg: Springer-Verlag, pp. 343-349.

Buzáki, G. and A. Draguhn. 2004. 'Neuronal Oscillations in cortical networks.' *Science,* 304(5679), pp. 1926-1929.

Caffarel, Alice, James R. Martin and Christian M.I.M. Matthiessen (eds.). 2004. *Language Typology: a functional perspective*. Amsterdam/ Philadelphia, PA: John Benjamins.

Caroni, Pico, Flavio Donato and Dominique Muller. 2012. 'Structural plasticity upon learning: regulation and functions.' *Nature Reviews Neuroscience*, 13, pp. 478-490.

Carter, Ellen C. and Robert C. Carter. 1981. 'Color and conspicuousness.' *Journal of the Optical Society of America*, 71(6), pp. 723-729.

Cartwright, Nancy. 1983. *How the Laws of Nature Lie.* Oxford: Oxford University Press

Cartwright, Nancy. 1989. *Nature's Capacities and their Measurement.* Oxford: Oxford University Press.

Chomsky, Noam. 1955. *The Logical Structure of Linguistic Theory.* Cambridge, MA: MIT Humanities Library, Microfilm. Reprinted. 1975. New York and London: Plenum Press.

Chomsky, Noam. 1959. 'A review of B.F. Skinner's Verbal Behaviour.' *Language*, **35(1)**, pp. 26-58.

Chomsky, Noam. 1965. *Aspects of a Theory of Syntax*. Cambridge, MA: MIT Press.

Chomsky, Noam. 1980. *Rules and Representations*. London: Fontana.

Chomsky, Noam. 1986. *Knowledge of Language*. New York, NY: Praeger.

Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.

Christie, Frances. 2004. 'Systemic Functional Linguistics and a theory of language education.' *Ilha do Desterro*, 46, pp. 13-40

Churchland, Paul. 1981. 'Eliminative Materialism and the Propositional Attitudes.' *Journal of Philosophy,* 78(2), pp. 67-90.

Churchland, Patricia S. 1986. *Neurophilosophy: towards a unified science of the mind-brain*. Cambridge, MA: MIT Press

Churchland, Patricia S., and Terrance J. Sejnowski. 1992. *The Computational Brain*. Cambridge, MA: MIT Press

Clark. Andy. 2008. *Supersizing the Mind: Embodiment, Action and Cognitive Extension.* Oxford: Oxford University Press.

Clark, Andy. 2010a. '*Memento's* Revenge: The Extended Mind, Extended.' In: Richard Menary (ed.), *The Extended Mind*. Cambridge, MA: MIT Press, pp. 43-66.

Clark, Andy and David J. Chalmers. 1998. 'The Extended Mind'. *Analysis,* 58(1), pp. 7-19.

Clarke, Ben. 2013. *Do patterns of ellipsis in text support systemic functional linguistics' 'context-metafunction hook-up' hypothesis? A corpus based approach.* PhD Thesis, Cardiff University.

Cleirigh, Chris. 1998. *A Selectionist Model of the Genesis of Phonic Texture: Systemic Phonology and Universal Darwinism*. Ms. PhD thesis, University of Sydney.

Cloran, Carmel. 1992. *Rhetorical Units and Decontextulization: an enquiry into some relations of context, meaning and grammar*. Monographs in Systemic Linguistics, 6. Nottingham: School of English, University of Nottingham.

Cloran, Carmel. 1999. 'Context, Material Situation and Text.' In: Mohsen Gahdessy (ed.), *Text and Context in Functional Lingusitics*. Amsterdam/ Philadelphia, PA: John Benjamins, pp. 177-218.

Collobert, Ronan and Jason Weston. 2008. 'A unified architecture for natural language processing: deep neural networks with multi-task learning.' *Proceedings of the 25th international conference on machine learning*, pp. 160-167.

Coniglione, Francesco. 2004. 'Between Abstraction and Idealization: Scientific Practice and Philosophical Awareness'. *Poznań Studies in the Philosophy of the Sciences and the Humanities,* 82, pp. 59-110.

Cooper, N.R., R.J. Croft, S.J.J. Dominey and A.P. Burgess. 2003. 'Paradox lost? Exploring the role of alpha oscillations during externally vs. internally driven directed attention and the implication for idling and inhibition hypotheses.' *International Journal of Psychopsychology*. 47(1), pp. 65-74.

Cowart, Monica. 2005. 'Embodied Cognition.' *Internet Encyclopaedia of Philosophy*. URL = http://www.iep.utm.edu/embodcog/. Accessed: 10/04/2015.

Cranston-Paull, R. and Theodore Sider. 1992. 'In Defense of Global Supervenience.' *Philosophy and Phenomenological Research*, 52, pp. 833-854.

Craver, Carl. 2007. *Explaining the Brain: Mechanisms and the monasic unity of neuroscience.* Oxford: Oxford University Press

Daffner, K.R., L.F.M. Scinto, A.M. Weitzman, R. Faust, D.M. Rentz, A.E. Budson and P.J. Holcomb. 2003. 'Frontal and Parietal Components of a Cerebral Network Mediating Voluntary Attention to Novel Events.' *Journal of Cognitive Neuroscience*, 15(2), pp. 294-313.

Danner, Unna N., Henk Aarts and Nanne K. de Vries. 2008. 'Habit vs. intention in the prediction of future behaviour: the role of frequency, context stability and mental accessibility of past behaviour.' *British Journal of Social Psychology*, 47, pp. 245-265.

Davidson, Donald. 1970. 'Mental Events'. In: Lawrence Foster and J. W. Swanson (eds.), *Experience and Theory*. Amherst, MA: University of Massachusetts Press, pp. 79-101.

Dayan, Peter and L. F. Abbott. 2001. *Theoretical Neuroscience: Computational and Mathematical Modelling of Neural Systems*. Cambridge, MA: MIT Press.

Deacon, Terrance. 1997. *The Symbolic Species: The co-evolution of language and brain.* New York, NY: W. W. Norton & Co.

Del Vecchio, Domitilla, Richard Murray and Pietro Perona. 2003. 'Decomposition of Human Motion into Dynamics based Primitives with applications to drawing tasks.' *Automatica*, 39(12), pp. 2085-2098.

DeMers, D. and G. Cottrell. 1993. 'Non-linear dimensionality reduction.' In: Morgan Kaufmann (ed.), *Advances in Neural Information Processing Systems*, pp. 580-587.

Demeter, Elise, Luis Hernandez-Garcia, Martin Sarter and Cindy Lustig. 2010. 'Challenges to attention: a continuous arterial spin labelling (ASL) study of the effects of distraction on sustained attention.' *Neuroimage*, 54, pp. 1518-1529.

Dennett, Daniel. 1991. *Consciousness Explained*. London: Little, Brown & Co.

Dennett, Daniel. 1996. *Kinds of Minds*. New York, NY: Basic Books.

Descartes, René. 1641/1999. *Meditations on the First Philosophy.* Reprinted as: *Meditations and Other Metaphysical Writings,* Desmond M. Clarke (trans.). London: Penguin

Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press

Dretske, Fred. 1988*. Explaining Behaviour: Reasons in a World of Causes.* Cambridge, MA: MIT Press

Dretske, Fred. 1996. 'Phenomena Externalism or If meanings ain't in the head, where are qualia?' *Philosophical Issues*, 7, pp. 143-158.

Dretske, Fred. 2003. 'Experience as Representation.' *Philosophical Issues,* 13, pp. 67-82.

Eagleman, David. 2001. 'Visual illusions and neurobiology.' *Nature Reviews Neuroscience*, 2, pp. 920-926.

Edelman, Gerald. 1993. *Brilliant air, Brilliant Fire: On the matter of the mind*. New York, NY: Basic Books.

Edelman, Shimon. 2008. *Computing the Mind: how the mind really works*. Oxford: Oxford University Press.

Eliasmith, Chris. 2003. 'Moving beyond metaphors: Understanding the mind for what it is.' *Journal of Philosophy,* 100(10), pp. 493-520.

Eliasmith, Chris. 2013. *How to build a brain: a neural architecture for biological cognition.* Oxford: Oxford University Press.

Eliasmith, Chris. 2015. 'Marr's attacks: On reductionism and vagueness.' *Topics in Cognitive Science*. 7(2), pp. 323-335.

Eggins, Suzanne. 2004. *An Introduction to Systemic Functional Linguistics*. London: Continuum

von der Emde, Gerhard. 1999. 'Active Electrolocation of objects in weakly electric fish.' *Journal of Experimental Biology*, 202: 1205-1215.

Fantz, Robert L. 1964. 'Visual Experience in Infants: Decreased Attention to Familiar Patterns Relative to Novel Ones.' *Science*, 146, pp. 668-670.

Fawcett, Robin. 2008.*Invitation to Systemic Functional Linguistics through the Cardiff Grammar: an extension and simplification of Halliday's systemic functional grammar.* 3rd Edition. London: Equinox.

Fawcett, Robin and Gordon Tucker. 1990. 'Demonstration of Genesys: a very large, semantically based systemic functional grammar. *13th International Conference on Computational Linguistics (COLING-90),* Vol 1, pp. 47-49. Helsinki, Finland.

Firth, John R. 1957. *Papers in Linguistics 1934-1951*. Oxford: Oxford University Press.

Fish, William. 2010. *Philosophy of Perception: A contemporary introduction*. London: Routledge.

Flash, Tamar and Binyamin Hochner. 2005. 'Motor Primitives in vertebrates and invertebrates.' *Current Opinion in Neurobiology*, 15(6), pp. 660-666.

Fodor, Jerry A. 1974. 'Special Sciences (or: The disunity of science as a working hypothesis). *Synthese*, 28(2), pp. 97-115.

Fodor, Jerry A. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.

Fodor, Jerry A. 1987. *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press.

Fodor, Jerry A. 1990. *A Theory of Content and other Essays*. Cambridge, MA: MIT Press.

Fodor, Jerry A. 1997. 'Special Sciences: Still Autonomous After All These Years.' *Noûs,* 31 (Supplement: Philosophical Perspectives, 11, Mind, Causation, and World), pp. 149-163.

Fodor, Jerry A. 2008. *LOT 2: The Language of Thought Revisited.* Oxford: Oxford University Press.

Fodor, Jerry A. and Zenon Pylyshyn. 1988. Connectionism and cognitive architecture: a critical analysis. *Cognition*, 28(1-2), pp. 3-71.

Fontaine, Lise, Tom Bartlett and Gerard O'Grady (eds.). 2013. *Systemic Functional Linguistics: Exploring Choice*. Cambridge: Cambridge University Press.

Frankish, Keith and William M. Ramsey. 2012. 'Introduction.' In: Keith Frankish and William M. Ramsey (eds.), *The Cambridge Handbook of Cognitive Science.* Cambridge: Cambridge University Press, pp. 1-6.

Frege, Gottlob. 1884. *Die Grundlage der Arithmetik.* Translated by J.L. Austin as *The Foundations of Mathematics*. Oxford; Blackwell, 1959.

Friedenberg, Jay and Gordon Silverman. 2012. *Cognitive Science: An introduction to the Study of Mind.* Los Angeles, CA: Sage Publications.

Frith, Uta. 1995. 'Dyslexia: can we have a shared theoretical framework?' In: N. Frederickson and R. Reason (eds.), *Phonological assessment of specific learning difficulties, Educational Child Psychology,* 12, pp. 6-17.

Frith, Uta and Francesca Happé. 2001. 'Why specific developmental disorders are not specific: On-line and developmental effects in autism and dyslexia. *Developmental Science,* 1(2), pp. 267-272.

Fusella, Paul V. 2013. 'Dynamic Systemc Theory in Cognitive Science: Major Elements, Applications and Debates Surrounding a Revolutionary Metatheory.' *Dynamical Psychology*. URL: http://dynapsyc.org/2013/Fusella.pdf. Accessed: 17/04/2015.

Gazdar, Gerald. 1979. *Pragmatics: Implicature, presupposition and logical form.* New York: Academic Press.

van Gelder, Timothy. 1995. 'What Might Cognition Be, if not Computation?' *Journal of Philosophy*, 92(7), pp. 345-381.

van Gelder, Timothy. 1998a. 'The Dynamical Hypothesis in Cognitive Science.' *Behavioural and Brain Sciences,* 21, pp. 615-628.

van Gelder, Timothy. 1998b. 'Disentangling Dynamics, Computation and Cognition.' *Behavioural and Brain Sciences,* 21, pp.654-665.

van Gelder, Timothy and Robert F. Port. 1995. 'It's About Time: An Overview of the Dynamical Approach to Cognition.' In: Timothy van Gelder and Robert F. Port (eds.), *Mind as Motion; Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press, pp. 1-43.

Gigerenzer, Gerd and Daniel G. Goldstein. 1996. 'Mind as Computer: Birth of a Metaphor.' *Creativity Research Journal,* 9(2+3), pp. 131-144.

Glenberg, Arthur M. and Michael P. Kaschak. 2002. 'Grounding Language in Action.' *Psychonomic Bulletin and Review*, 9(3), pp. 558-565.

Godfrey-Smith, Peter. 2005. 'Folk Psychology as a Model.' *Philosopher's Imprint*, 5(6). www.philosophersimprint.org/005006/ . Accessed: 08/06/2015.

Godfrey-Smith, Peter. 2009. 'Abstractions, Idealization and Evolutionary Biology.' In: A. Baraberousse, M. Morange and T. Pradeu (eds.), *Mapping the Future of BiologyL Evolving Concepts and Theories*. Dordecht: Springer, pp. 46-56.

Gozzano, Simone and Christopher S. Hill (eds.). 2012. *New Perspectives of Type Identity: The Mental and the Physical*. Cambridge: Cambridge University Press.

Graf, Peter and Daniel Schacter. 1985. 'Implicit and explicit memory for new associations in normal and amnesic subjects.' *Journal of Experimental Psychology*, 11(3), pp. 501-518.

Greaves, William. 2007. 'Intonation in systemic functional linguistics.' In: Ruqaiya Hasan, Christian Matthiessen and Jonathan J. Webster (eds.), *Continuing Discourse of Language: a functional perspective, volume 2.* London: Equinox, pp. 978-1025.

Green Christopher D. 1998. 'The Degrees of Freedom Would be Tolerable if Nodes were Neural: Reply to Lamm on Connectionist-Explanation.' *Psycoloquy,* 9(26), p. 23.

Green, Christopher D. 1999. 'Are Connectionist Models Theories of Cognition?' In: Wolfgang Maiers, Betty Bayer, Barbara Duarte Esgalhado, René Jorna and Ernst Schraube (eds.), *Challenges to Theoretical Psychology*. Ontario: Captus Press, pp. 139-146.

Gregory, Michael and Susanne Carroll. 1978. *Language varieties and their social contexts*. London: Routledge.

Grice, Paul. 1989. *Studies in the ways of words.* Cambridge, MA: Harvard University Press.

Grush, Rick. 2002. 'The Architecture of Representation.' In: William Bechtel, Pete Mandik, Jennifer Mundale and Robert S. Stufflebeam (eds.), *Philosophy and the Neurosciences*. Massachusetts, MA/ Oxford: Blackwell, pp. 349-368.

Gurney, Kevin. 2003. *An introduction to Neural Networks*. Boca Raton: CRC Press.

Hainline, Louise. 1998. 'Summary and Commentary: Eye Movements, Attention and Development.' In: John E. Richards (eds.), *Cognitive Neuroscience of Attention: A developmental perspective*. Mawah, NJ: Lawrence Erlbaum, pp. 163-178.

Halliday, Michael A.K. 1961. 'Categories of the Theory of Grammar'. *Word,* 17(3), pp. 241-92.

Halliday, Michael A.K. 1966. 'Some notes on 'deep' grammar'. *Journal of Linguistics,* 2(1), pp. 57-67.

Halliday, Michael A.K. 1972 [2003]. 'Towards a Social Semiotics.' Reprinted in Jonathan J. Webster (ed.), *The Collected Works of M.A.K. Halliday, Vol 3. On Language and Linguistics*. London: Continuum, pp. 323-354.

Halliday, Michael A.K. 1975. *Learning how to mean.* London: Arnold.

Halliday, Michael A.K. 1977a. [2003]. 'Ideas on Language.' *Occasional Papers I*, pp. 32-55. ' Reprinted in Jonathan J. Webster (ed.), *The Collected Works of M.A.K. Halliday, Vol 3. On Language and Linguistics*. London: Continuum, pp.92-115.

Halliday, Michael A.K. 1977b. 'Text as Semantic Choice in Social Contexts.' In: Teun A. van Dijk and János S. Petofi (eds.), *Grammars and Description*. Berlin: Walter de Gruyter, pp. 176-226.

Halliday, Michael A.K. 1979. *Language as social semiotic: The social interpretation of language and meaning.* London: Edward Arnould.

Halliday, Michael A.K. 1981 [2005]. 'Text Semantics and Clause Grammar: how is a text like a clause?' In: Jonathan J. Webster (ed.), *The Collected Works of M.A.K. Halliday, Vol 1. On Grammar.* London: Continuum, pp. 219-261.

Halliday, Michael A.K. 1984. 'Language as code and language as behaviour: a systemic-functional interpretation of the nature and ontogenesis of dialogue.' In: Robin P. Fawcett et al. (eds.), *The semiotics of culture and language. Vol.1.* London: Francis Pinter, pp. 3-35.

Halliday, Michael A.K. 1985. 'Systemic Background.' In: James D. Benson and William S. Greves (eds.), *Systemic Perspectives on Discourse, Volume 1. Selected Theoretical Papers from the 9th International Systemic Workshop.* Norwood, NJ: Ablex Publishing Corporation, pp. 1-15.

Halliday, Michael A.K. 1991 [2005]. 'Towards Probabilistic Interpretations.' Reprinted in Jonathan J. Webster (ed.), *The Collected Works of M.A.K. Halliday. Vol. 6: Computational and Quantitative Studies*. London: Continuum, pp. 42-62.

Halliday, Michael A. K. 1992a. [2005]. 'The act of meaning.' In: James E. Alatis (ed.), *Georgetown University Round Table on Languages and Lingusitics:Language, Communication and Social Meaning*. Washington, DC: Georgetown University Press, pp. 7-21. Reprinted in Jonathan J. Webster (ed.), *The Collected Works of MAK Halliday, Vol. 3. On Language and Linguistics*. London: Continuum, pp. 372-389.

Halliday, Michael A.K. 1992b. 'Language as system and Language as instance: the corpus as a theoretical concept.' In Jan Svartvik (ed.), *Directions in Corpus Linguistics: Proceedings of Nobel Symposium 82*, Berlin: Mouton de Gruyter, 1991.

Halliday, Michael A.K. 1992c. 'Systemic Grammar and the Concept of a "Science of Language".' *Waiguoyu (Journal of Foreign Languages)*. Reprinted in: Jonathan J. Webster (ed.), *The Collected Works of M.A.K. Halliday, Vol. 3: On Language and Linguistics*

Halliday, Michael A.K. 1993. 'Towards a Language-Based Theory of Learning'. *Linguistics and Education* 5, pp. 93-116.

Halliday, Michael A.K. 1994. *An Introduction to Functional Grammar, 2nd Edition*.' London: Edward Arnold.

Halliday, Michael A.K. 1995a. [2005]. 'Fuzzy Grammatics: a systemic functional approach to fuzziness in natural language'. *Proceedings of FUZZ-IEEE/IFES '95*. Yokohama: Yokohama University. Reprinted in: Jonathan J. Webster (ed.) *The Collected Works of MAK Halliday, Volume 6: Computational and Quantitative Studies*. London: Continuum, pp. 213-238.

Halliday, Michael A.K. 1995b. 'On Language in Relation to Fuzzy Logic and intelligent Computing.'

Halliday, Michael A. K. 1996. 'On grammar and grammatics.' In: Ruqaiya Hasan, Carmel Cloran and David Butt (eds.), *Functional Descriptions: theory into practice*. Ansterdam/ Philadelphia, PA: John Benjamins, pp. 1-38.

Halliday, Michael A.K. 1997. 'Linguistics as Metaphor.' In: A-M. Simon-Vandenbergen, K. Davidse and D.Noel (eds.), *Reconnecting Language: Morphology and Syntax in Functional Perspectives,* pp. 3-27. Reprinted in: Jonathan J. Webster (ed.) *The Collected Works of M.A.K. Halliday, Volume 3: On Language and Linguistics*. London: Continuum: pp. 248-270

Halliday, Michael A.K. 1999. 'The Notion of "Context" in Language Education.' In: Mohsen Ghadessy (ed.), *Text and Context in Functional Linguistics.* Amsterdam/ Philadelphia, PA: John Benjamins, pp. 1-24.

Halliday, Michael A.K. 2002. 'Computing meaning: some reflections on past experience and present prospects.' In: G.W. Huang, Z.Y. Wang (eds.), *Discourse and Language Functions.* Shanghai: Foreign Language Teaching and Research Press, pp. 3-25.

Halliday, Michael A.K. 2003. 'Introduction: On the 'architecture' of human language.' In Jonathan J. Webster (ed.), *The Collected Works of M.A.K. Halliday, Volume 3: On Language and Linguistics*. London: Continuum, pp. 1-32.

Halliday, M.A.K. 2009. 'Methods-Techniques-Problems.' In: M.A.K. Halliday and Jonathan J. Webster (eds.), *Continuum Companion to Systemic Functional Linguistics*. London/ New York: Continuum, pp. 59-86.

Halliday, Michael A. K., and Ruqaiya Hasan. 1985. *Language, Context and Text: Aspects of Language in a Social-semiotic perspective*. Geelong: Deakin University Press.

Halliday, M.A.K. and Christian M.I.M Matthiessen. 1999. *Construing Experience through Meaning: a language-based view of cognition.* New York/ London: Continuum.

Halliday, M.A.K. and Christian M.I.M. Matthiessen. 2014. *Halliday's Introduction to Functional Grammar*, 4th Edition. Abingdon: Routledge.

Hansen, Stephen José and David J. Burr. 1991. 'What connectionist models learn: learning and representation in connectionist networks.' *Behavioural and Brain Sciences*, 13, pp. 471-518.

Harnard, S. 1990. 'The Symbol Grounding Problem.' *Physica D.*, 42, pp. 335-346.

Harrison, David and Marc de Kamps. 2011. 'A dynamical neural simulation mapping feature-based attention to location in non-linear circuits.' *BMC Neuroscience*, 12(1), p.55.

Hasan, Ruqaiya. 1981. 'What's going on: A dynamic view of context.' *Seventh LACUS forum*. Columbia, SC: Hornbeam Press.

Hasan, Ruqaiya. 1987. 'The grammarian's dream: lexis as most delicate grammar'. In: M.A.K. Halliday and Robin Fawcett (eds.), *New Developments in Systemic Linguistics. Vol 1, Theory and Description.* London/ New York: Pinter, pp. 184-211.

Hasan, Ruqaiya. 1995. 'The conception of context in text.' In: Peter Howard Fries, Michael Gregory and Michael A.K. Halliday (eds.), *Discourse in Society: Systemic Functional Perspectives*. *Meaning and Choice in Language: Studies for Michael Halliday.* Westport, CT: Ablex Publishing, pp. 183-284.

Hasan, Ruqaiya. 1999. 'Speaking with reference to context.' In: Mohsen Ghadessy (ed.), *Text and Context in Functional Linguistics*. Amsterdam/ Philadelphia, PA: John Benjamins, pp. 219-328.

Hasan, Ruqaiya. 2009. 'The place of context in a systemic functional model.' In: Michael A. K. Halliday and Jonathan J. Webster (eds.), *The Continuum Companion to Systemic Functional Linguistics*. London, Continuum, pp. 166-189.

Hasan, Ruqaiya. 2014. 'Towards a paradigmatic description of context: systems, metafunctions, and semantics.' *Functional Lingusitics,* 1(9), doi:10.1186/s40554-014-0009-y. Accessed: 12/05/2015.

Haugeland, John. 1978. 'The Nature and Plausibility of Cognition.' *Behavioural and Brain Sciences*, 1, pp. 215-226.

Haugeland, John. 1991. 'Representational Genera.' In: W. Ramsey, S.P. Stich and D.E. Rumelhart (eds.), *Philosophy and Connectionist Theory*. Hillside, NJ: Erlbaum, pp. 61-89.

Hauk, O., I. Johnsrude and F. Pulvermueller. 2004. 'Somatotopic representation of action words in human motor and premotor cortex.' *Neuron*, 41(2), pp. 301-307.

Haykin , Simon. 2008. *Neural Networks and Learning Machines*. 3rd Edition. New York: Prentice Hall.

Heil, John. 2004. *Philosophy of Mind: a contemporary introduction.* Second Edition. New York, NY/ Oxon: Routledge.

Heim, Irene. 1983. 'On the projection problem of presupposition.' In: M. Barlow, D. Flickinger and M Westcoat (eds.), *Second Annual West Coast Conference on Formal Linguistics,* Stanford University, pp. 114-126.

Heim, Irene and Angelika Kratzer. 1997. *Semantics in Generative Grammar*. London: Blackwell.

Hempel, Carl G. 1966. *Philosophy of natural science*. Englewood Cliffs, NJ: Prentice Hall.

Hill, Christopher and Brian McLaughlin. 'There are fewer things in reality than are dreamt of in Chalmers's philosophy.' *Philosophy and Phenomenological Research*, 59(2), pp. 445-454.

Hjelmslev, Louis. 1947/ 1961. *Prolegomena to a theory of language*. University of Wisconsin Press.

Hoffmann, Vera and Albert Newen. 2007. 'Supervenience of Extrinsic Properties.' *Erkenntnis,* 67, pp. 305-319.

Horgan, Terence and James Woodward. 1985. 'Folk Psychology is Here to Stay.' *The Philosophical Review*, 94(2), pp. 197-226.

Horgan, Terence and John Tienson. 1993. 'Levels of Description in Nonclassical Cognitive Science.' *Royal Institute of Philosophy Supplement*, **34**, pp. 159-188.

Houghton, George. 2005. 'Introduction: Structures, Processes and Algorithms.' In: George Houghton (ed.), *Connectionist Models in Cognitive Psychology.* Hove/ New York, NY: Taylor and Francis, pp. 1-41.

Hsieh, L.T, and C. Ranganath. 2014. 'Frontal midline theta oscillations during working memory maintenance and episodic encoding and retrieval.' *Neuroimage*, 15(85), pp. 721-729.

Huang, Liqiang and Harold Pashler. 2005. 'Attention capacity and task difficulty in visual search.' *Cognition*, 94, pp. B101-B111.

Hubel, D.H. and T.N. Wiesel. 1962. 'Receptive fields, binocular interaction and functional architecture in the cat's visual cortex.' *Journal of Physiology*, 160, pp. 106-154.

Hubel, D.H. and T.N. Wiesel. 1964. 'Effects of monocular deprivation in kittens.' *Naunyn-Schmiedebergs archives of experimental pathology and pharmacology*, 248, pp. 492-497.

Hubel, D.H. and T. N. Wiesel. 1965. 'Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat.' *Journal of Neurophysiology*, 29, pp. 229-289.

Hummel, John E. and Keith J. Holyoak. 2003. 'A Symbolic-Connectionist Theory of Relational Inference and Generalization.' *Psychological Review*, 110(2), pp. 220-264.

Hurley, Susan. 1998a. *Consciousness*. Cambridge: Cambridge University Press.

Hurley, Susan. 1998b. *Consciousness in Action.* Cambridge, MA: Harvard University Press.

Hurley, Susan. 2010. 'Varieties of Externalism.' In: Richard Menary (ed.), *The Extended Mind*. Cambridge, MA: MIT Press, pp. 101-153.

Hutto, Daniel D. and Erik Myin. 2013. *Radicalizing Enactivism: Basic Minds without Content*. Cambridge, MA: MIT Press.

Jackson, Frank. 2000. *From Metaphysics to Ethics: a defence of conceptual analysis.* Oxford: Oxford University Press.

Jain, Anil K., Jianchang Mao and K.M. Mohiuddin. 1996. 'Artificial Neural Networks: A Tutorial.' *Computer,* 29(3), pp. 31-44.

Joanisse, Mark F. and James L. McCelland. 2015. 'Connectionist perspectives on language learning, representation and processing.' *Wiley introductory reviews: cognitive science*, 6(3), pp. 235-247.

Johnson, M. and J. Morton. 1991. *Biology and Cognitive Development: The case of face recognition*. Oxford: Blackwell.

Jones, Martin R. 2005. 'Idealization and Abstraction: A Framework.' In: Martin R. Jones and Nancy Cartwright (eds.), *Idealization XII: Correcting the Model; Idealization and Abstraction in the Sciences* (Poznań Studies in the Philosophy of the Sciences and Humanities, 86). Amsterdam/ New York: Rodopi, pp. 173-217.

Kahana, Michael J., D. Seelig and Joseph R. Madsen. 2001. 'Theta returns.' *Current Opinion in Neuroscience*, 11(6), pp. 739-744.

Kahana, Michael, Robert Sekuler, Jeremy Caplan, Matthew Kirschen and Joseph Madsen. 1999. 'Human Theta oscillaions exhibit during task dependence during virtual water maze navigation.' *Nature*, 399(24), pp. 781-784.

Kaplan, R. M. and J. Bresnan. 1982. ' Lexical Functional Grammar: A formal system for grammatical representation.' In: J. Bresnan (ed.), *The Mental Representation of Grammatical Relations*, pp. 173-281.

Karttunen, Lauri. 1974. 'Presuppositions and Linguistic Context.' *Theoretical Linguistics 1*, pp. 181-194.

Kempson, Ruth, Tim Fernando and Nicholas Asher (eds.), *Philosophy of Linguistics*. Oxford/ Amsterdam: Elesvier.

Kersten, Alan W. and Linda B. Smith. 2002. 'Attention to Novel Objects during verb learning.' *Child Development*, 73(1), pp. 93-109.

Kim, Jaegwon. 1984/1993. 'Concepts of Supervenience.' *Philosophy and Phenomenological Research,* 45, pp. 153-176. Reprinted in: Jaegwon Kim. *Supervenience and Mind: Selected Philosophical Papers.* Cambridge: Cambridge University Press, pp. 53-78.

Kim, Jaegwon. 1987/ 1993. '"Strong" and "Global" Supervenience Revisited.' *Philosophy and Phenomenological Research,* 48, pp. 315-326. Reprinted in: Jaegwon Kim. *Supervenience and Mind: Selected Philosophical Papers.* Cambridge: Cambridge University Press, pp. 79-91.

Kim, Jaegwon. 2005. *Physicalism or something near enough.* Princeton, NJ: Princeton University Press.

Kim, Jaegwon. 1996. *Philosophy of Mind*. Boulder: Westview Press.

Koch, Christoff. 2004. *The Quest for Consciousness.* Englewood, CO: Roberts and Company.

Kress, Gunther and Theo van Leeuwen. 1996. *Reading Images: The Grammar of Visual Design.* London: Routledge.

Lakoff, George and Mark Johnson. 1980. *Metaphors we live by*. Chicago: Chicago University Press.

Lakoff, George and Mark Johnson. 1999. *Philosophy in the Flesh: The Embodied mind and its Challenge to Western Thought.* New York, NY: Basic Books.

Lamb, Sydney. 1966. *Outline of stratificational grammar*. Washington, DC: Georgetown University Press.

Lasnik, Howard. 2005. 'Grammar, levels, and biology.' In: James McGilvrey (ed.), *The Cambridge Companion to Chomsky*. Cambridge: Cambridge University Press, pp. 60-83.

van Leeuwen, Theo. 1998. 'Music and ideology: Notes towards a sociosemiotics of mass media music.' *Popular Music and Society,* 22(4), pp. 25-54.

Leitan, Nuwan D. and Lucian Chaffey. 2014. 'Embodied Cognition and its applications: a brief review.' *Sensoria: a journal of mind, brain & culture*. DOI: http://dx.doi.org/10.7790/sa.v10i1.384. Accessed: 22/04/15.

Lemke, Jay. 1990. *Talking Science: Language, Learning, and Values*. Norwood, NJ: Ablex.

Lemke, Jay. 1995. *Textual Politics: Discourse and Social Dynamics*. London: Taylor and Francis.

Lewis, David. 1970. 'How to define theoretical terms.' *Journal of Philosophy*, 67, pp. 427-446.

Lewis, David. 1986. *On the plurality of worlds*. Oxford: Basil Blackwell.

Lowe, E.J. 2000. *An introduction to the philosophy of mind.* Cambridge: Cambridge University Press.

Ludlow, Peter. 2011. *The Philosophy of Generative Linguistics.* Oxford: Oxford University Press.

Lynott, Dermot, Louise Connell and Judith Holler. 2013. 'The Role of Body and Environment in Cognition'. *Frontiers in Psychology*. URL: http://journal.frontiersin.org/researchtopic/823, DOI: 10.3389/fpsyg.2013.00465. Accessed: 09/04/2015

MacDonald, Graham and David Papineau (eds.) 2006. *Teleosemantics: New Philosophical Essays*. Oxford: Clarendon Press.

MacKay, D. 1967. 'Ways of looking at perception.' In: W. Watthen-Dunn (ed.), *Models for the Perception of Speech and Visual Form.* Cambridge, MA: MIT Press.

Mahboob, Ahmar and Naomi Knight. 2010. *Appliable Linguistics*. London: Continuum.

Malenka, Robert C. 'Synaptic Plasticity'. In: Kenneth L. Davis, Dennis Charney, Joseph Coyle and Charles Nemeroff (eds.), *Neuropsychopharmacology: 5th Generation of Progress*. Philadelphia, PA: Lippincott, Williams and Wilkins, pp. 147-157.

Mantik, Pete. 2011. 'Supervenience and Neuroscience'. *Synthese*, 180, pp. 443-463.

Marr, David. 1982. *Vision: A computational investigation into the human representation and processing of visual information.* New York: W.H. Freeman and Company.

Marr, David and Thomas Poggio. 1977. 'From understanding computation to understanding neural circuits.' *Neurosciences Research Progress Bulletin,* 15, pp. 470-488.

Marraffa, Massimo and Alfredo Paternoster. 2012. 'Functions, levels, and mechanisms: Explanation in cognitive science and its problems.' *Theory and Psychology*, **23(1)**, pp. 22-45.

Martin, James R. 1985. 'Process and Text: two aspects of human semiosis.' In: James Benson and William Greaves (eds.), *Systemic Perspectives on Discourse, Vol. 1*. Norwood, NJ: Ablex Publishing.

Martin, James R. 1992. *English Text: system and structure*. Amsterdam/ Philadelphia, PA: John Benjamins.

Martin, James R. 1997. 'Analysing genre: functional perspectives.' In: Francis Christie and James R. Martin (eds.), *Genre and institutions: Social processes in the workplace and school.* London/ New York, NY: Continuum.

Martin, James R. 1999. 'Mentoring semogenesis: 'Genre-based' literacy pedagogy.' In: F. Christie (ed.), *Pedagogy and the Shaping in Consciousness: Linguistic and Social Processes.* London: Cassell, pp. 123-155.

Martin, James R. 2014. 'Evolving systemic functional linguistics: beyond the clause.' *Functional Linguistics*, 1(3). URL: http://www.functionallinguistics/com/content/1/1/3. Accessed: 27/05/2015.

Martin, James R. and Christian M.I.M. Matthiessen. 1991. 'Systemic Typology and Topology.' In: Francis Christie (ed.), *Social Processes in Education: Proceedings of the first Australian systemic network conference.* Darwin: Centre for Studies of Language in Education, Northern Territory University.

Martin, James R. and David Rose. 2007. *Working with Discourse: Meaning beyond the clause*. London: Continuum.

Martin, James R. and David Rose. 2008. *Genre relations: mapping culture*. London/ Oakville, CA: Equinox.

Martin, James R. and Joan Rothery. 1980. 'Writing Project Report, Number 1.' *Working Papers in Linguistics, No. 1*. Department of Linguistics, University of Sydney.

Martin, James R. and Joan Rothery. 1981. 'Writing Project Report, Number 2.' *Working Papers in Linguistics, No. 1*. Department of Linguistics, University of Sydney.

Martinez, Sergio F. and Xiang Huang. 2011. 'Epistemic Groundings of Abstraction and their Cognitive Domain.' *Philosophy of Science*, 78(3), pp. 490-511.

Matten, Mohan. 2014. 'Debunking enactivism; a critical notice of Hutto and Myin's Radicalizing Enactivism.' *Canadian Journal of Philosophy*, 44(1), pp. 118-128.

Matthiessen, Christian M.I.M. 1987. 'Representational Issues in Systemic Functional Grammar.' *Proceedings of the 12th Annual Systemic Workshop in Ann Arbor, MI, August 21-24 1985.*

Matthiessen, Christian M.I.M. 2007a. 'The "architecture" of language according to systemic functional theory: developments since the 1970s.' In: Ruqaiya Hasan, Christian Matthiessen and Jonathan J. Webster (eds.), *Continuing Discourse of Language: a functional perspective, volume 2.* London: Equinox, pp. 505-561.

Matthiessen, Christian M.I.M. 2007b. 'Lexicogrammar in systemic functional linguistics: descriptive and theoretical developments since the 1970s.' In: Ruqaiya Hasan, Christian Matthiessen and Jonathan J. Webster (eds.), *Continuing Discourse of Language: a functional perspective, volume 2.* London: Equinox, pp. 765- 858.

Matthiessen, Christian M.I.M. 2012. 'Systemic Functional Linguistics as Appliable Linguistics: Social Accountability and Critical Approaches. *D.E.L.T.A,* 28, pp.435-471.

Matthiessen, Christian M.I.M. and John A. Bateman. 1991. *Text Generation and Systemic Functional Linguistics: Experiences from English and Japanese.* London: Pinter Publishing.

Matthiessen, Christian M.I.M. 2004. 'The evolution of language: a systemic functional exploration of phylogenetic phases.' In: G. Williams and A. Lukin (eds.), *Language Development: functional perspectievs on evolution and ontogenesis.* London: Continuum, pp. 45-90.

Matthiessen, Christian M.I.M., Kazuhiro Teruya and Marvin Lam. 2010. *Key Terms in Systemic Functional Linguistics.* London/ New York: Continuum.

McCelland, James, David Rumelhart and the PDP research group. 1986. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1.* Cambridge, MA: MIT Press.

McCelland, James, David Rumelhart and Geoffrey E. Hinton. 1986. 'The Appeal of Parallel Distributed Processing.' In James McCelland, David Rumelhart and the PDP research group. 1986. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1.* Cambridge, MA: MIT Press, pp. 3-44.

McClamrock, Ron. 1991. 'Marr's Three Levels: a re-evaluation.' *Minds and Machines,* **1**, pp. 185-196.

McGinn, Colin. 1991. *The Problem of Consciousness*. London: Blackwell.

McGurk, Harry and John MacDonald. 1976. 'Hearing lips, seeing voices.' *Nature*, 264, pp. 746-748.

McKinsey, Michael. 1991. 'Anti-individualism and privileged access.' *Analysis,* 51(1), pp. 9-16.

McLaughlin, Brian and Karen Bennett. 2014. "Supervenience". In: Edward N. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy* (Spring 2014 Edition). URL = http://plato.stanford.edu/archives/spr2014/entries/supervenience/. Accessed: 24/02/15.

McMullin, Ernan. 1985. 'Galilean Idealization.' *Studies in History and Philosophy of Science, Part A,* 16(3), pp. 247-273.

McKinsey, M. 1991. 'Anti-individualism and privileged access.' *Analysis*, 51, pp. 9-16.

Melrose, Robin. 2003. 'Having things both ways'. In: A.M. Simon-Vandenbergen, Miriam Taverniers, and Louise Ravelli (eds.). *Grammatical Metaphor: Views from Systemic Functional Linguistics*. Amsterdam: John Benjamins, pp. 417-442.

Menary, Richard. 2010. 'Introduction: The Extended Mind in Focus.' In: Richard Menary (ed.), *The Extended Mind*. Cambridge, MA: MIT Press, pp. 1-26.

Merabet, Lotfi, B., and Alvaro Pascual-Leone. 2010. 'Neural Reorganization Following Sensory Loss: The Opportunity of Change.' *Nature Reviews Neuroscience*, 11(1), pp. 44-52.

Mitchelmore, Michael and Paul White. 2004. 'Abstraction in Mathematics and Mathematical Learning.' *Proceedings of the 28th Conference of the International Group for the Psychology of Mathematics Education,* 3, pp. 329-336.

Miller, George A. 2003. 'The cognitive revolution: a historical perspective.' *TRENDS in Cognitive Sciences*, 7(3), pp. 141-144.

Millikan, Ruth G. 1984. *Language, Thought and other Biological Categories*. Cambridge, MA: MIT Press.

Morris, Richard. 1984. 'Developments of a water-maze procedure for studying spatial learning in the rat.' *Journal of Neuroscience Methods*, 11, pp. 47-60.

Morton, John and Uta Frith. 1995. 'Causal Modelling: A Structural Approach to Developmental Psychology.' In: D. Cicchetti and D. H. Cohen (eds.), *Manual of developmental Psychopathology*. New York: Wiley, pp. 357-390

Mumford, Stephen. 2003. *Dispositions*. Oxford: Oxford University Press.

Mundale, Jennifer and William Bechtel. 1996. 'Integrating neuroscience, psychology and evolutionary biology through a teleological conception of function.' *Minds and Machines,* 6(4), pp. 481-505.

Munro, Robert. 2004. 'A probabilistic Representation of Systemic Functional Grammar.' *31st International Systemic Functional Congress (ISFC31)*, Kyoto.

Munro, Paul W. 1986. 'State-Dependent Factors Influencing Neural Plasticity: A Partial Account of the Critical Period.' In: David Rumelhart, James McClelland and the PDP research group. 1986. *Parallel distributed processing: explorations in the microstructure of cognition. Volume 2, Psychological and Biological Models.* Cambridge, MA: MIT Press, pp. 471-502.

Nagel, Ernst. 1961. *The structure of science: problems in the logic of scientific explanation*. London: Routledge and Kegan Paul.

Nagel, Thomas. 1974. 'What is it Like to be a Bat?' *Philosophical Review*, 88(4), pp. 435-450.

Nalbantian, Suzanne. 2011. 'Introduction.' In: Suzanne Nalbantian, Paul M. Matthews and James L. McCelland (eds.), *The Memory Process: Neuroscientific and Humanistic Perspectives.* Cambridge, MA: MIT Press, pp. 1-26.

Neeleman, Ad. 2013. 'Comments on Pullum.' *Mind and Language,* 28(4), pp. 533-531.

Neisser, U. and R. Becklan. 1975. 'Selective looking: Attending to visually specified events.' In: A.D. Pick (ed.), *Perception and its Development: A Tribute to Eleanor J. Gibson*. Hillsade, NJ: Lawrence Erlbaum, pp. 201-219.

Neves, Guilherme, Sam F. Cooke and Tim Bliss. 2008. 'Synaptic Plasticity, memory and the hippocampus: a network approach to causality.' *Nature Reviews Neuroscience*, 9, pp. 65-75.

Newell, Allen. 1990. *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Newell, Allen and Herbert A. Simon. 1976. 'Computer Science as an Empirical Enquiry: Symbols and Search.' *Communications of the ACM,* 19(3), pp. 113-126.

Newmeyer, Frederick. 1998. *Language Form and Language Function*. Cambridge, MA: MIT Press.

Noë, Alva. 2004. *Action in Perception*. Cambridge, MA: MIT Press

Noë, Alva. 2009. *Out of our Heads*. New York, NY: Hill and Wang.

Northoff, Georg. 2014. *Minding the Brain: a guide to philosophy and neuroscience.* New York, NY: Palgrave Macmillan.

Opie, Jon. 2010. 'Levels and Explanations.' *ASCS09: Proceedings of the 9th Conference of the Australasian Society for Cognitive Science*, pp. 270-276.

Oppenheim, Paul and Hilary Putnam. 1958. 'Unity of Science as a Working Hypothesis.' *Minnesota Studies in the Philosophy of Science*, 2, pp. 3-36.

O'Donnell, Mick. 1999.'Context in Dynamic Modelling.' In: Mohsen Ghadessy (ed.), *Text and Context in Functional Linguistics*. Amsterdam/ Phiadelphia, PA: John Benjamins, pp. 63-100.

O'Keefe, John and Michael Reece. 1993. 'Phase relationship between hippocampal place units and EEG theta rhythm.' *Hippocampus,* 3(3), pp. 317-330.

O'Regan, J. Kevin and Alva Noë. 2001. 'A sensorimotor account of vision and visual consciousness.' *Behavioural and Brain Sciences*, 24, pp. 939-1031.

O'Regan, J. Kevin and Alva Noë. 2002. 'What it is like to see: A sensorimotor theory of perceptual experience.' *Synthese*, 79, pp. 79-103.

O'Toole, Michael. 1990. 'A systemic-functional semiotics of art.' *Semiotica,* 82(3/4), pp. 185-209.

Ouellette, J.A. and W. Wood. 1998. 'Habit and intention in everyday life: The multiple processes by which last behaviour predicts future behaviour.' *Psychological Bulletin,* 124, pp. 54-74.

Papineau, David. 2013. 'Causation is macroscopic but not reducible.' In: Sophie Gibb and Roegnvaldur Ingthorsson (eds.), *Mental Causation and Ontology*. Oxford: Oxford University Press.

Pastra, K. and Y. Aloimonos. 2012. 'A Minialist Grammar of Action.' *Philosophical Transactions of the Royal Society of London, B: Biological Sciences*, 367(1585), pp. 103-117.

Platchias, Dimitris. 2010. *Phenomenal Consciousness*. Durham: Acumen.

Poggio, Tomaso. 2012. 'The Levels of Understanding framework, revisted.' *Computer Science and Artificial Intelligence Laboratory Technical Report*. Massachusetts Institute of Technology

Polger, Thomas W. 2002. 'Putnam's Intuition.' *Philosophical Studies: an international journal for philosophy in the analytic tradition*, 109(2), pp. 143-170.

Portner, Paul. 2009. *Modality*. Oxford: Oxford University Press.

Potter, Steve M., Ahmed El Hady and Eberhard E. Fetz (eds.). 2014a. 'Closing the Loop around Neural Systems.' *Frontiers in Neural Circuits, Research Topic*. URL: http://journal.frontiersin.org/researchtopic/785/closing-the-loop-around-neural-systems. Accessed: 14/07/15.

Potter, Steve M., Ahmad El Hady and Eberhard E. Fetz. 2014b. 'Closed-loop neuroscience and neuroengineering.' In: Steve M. Potter, Ahmed El Hady and Eberhard E. Fetz. (eds.). URL: <http://dx.doi.org/10.3389/fncir.2014.00115>. Accessed: 14/07/15

Priebe, Nicholas J. and David Ferster. 2012. 'Mechanisms of Neuronal Computation in Mammalian Visual Cortex.' *Neuron*, 75, pp. 194-208.

Putnam, Hilary. 1975. 'The nature of mental states.' In: Hilary Putnam. *Mind, Language and Reality: Philosophical Papers, Volume 2*. Cambridge: Cambridge University Press, pp. 429-440.

Pylyshyn, Zenon. 1984. *Computation and Cognition.* Cambridge, MA: MIT Press.

Rafiei, Khadijeh and Sima Modirkhamene. 2012. 'Thematicity in Published vs. Unpublished Iranian EFL Theses.' *Theory and Practice in Language Studies,* 2(6), pp. 1206-1213.

Rieke, Fred, David Warland, Rob de Ruyter can Steveninck and William Bialek. 1997. *Spikes: exploring the neural code.* Cambridge, MA: MIT Press.

Riesen, A.H. and D.E. Zilbert. 1975. 'Behavioural consequences of variations in early sensory environments.' In: A.H. Riesen (ed.), *The Developmental Neuropsychology of Senosry Deprivation*. Academic Press, New York, pp. 211-252.

Ripley, Brian D. 2008. *Pattern Recognition and Neural Networks*. Cambridge: Cambridge University Press.

Rizzolatti, G. and M.A. Arbib. 1998. 'Language within our grasp.' *Trends in Neuroscience*, 21(5), pp. 188-194.

Robinson, Howard. 2012. "Dualism". In: Edward N. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy* (Winter 2012 Edition). URL = http://plato.stanford/edu/archives/win2012/entries/dualism/ . Accessed: 16/02/15.

Robinson, Howard. 2014. "Substance". In: Edward N. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy* (Spring 2014 Edition). URL = http://plato.stanford.edu/archives/spr2014/entries/substance. Accessed: 11/10/2015.

Roediger, H. L., F. M. Zaromb, and M. K. Goode. 2008. 'A Typology of Memory Terms'. In: R. Menzel (ed.), *Learning Theory and Behaviour*. Volume 1 of J Byrne (ed.), *Learning and Memory: A Comprehensive Reference*. Oxford: Elsevier, pp. 11-24.

Rose, David. 2006. *Learning to Read, Reading to Learn: Scaffolding the English curriculum for Aboriginal secondary schools*. Ms. University of Sydney.

Roth, Wolff-Michael. 2012. 'Undoing Decontextualization or How Scientists Come to Understand Their Own Data/Graphs.' *Science Education*, 97(1), pp. 80-112

Roth, Martin. 2013. 'Folk Psychology as Science.' *Synthese*, 190, pp. 3971-3982.

Rowlands, Mark. 2009. 'Extended Cognition and the mark of the cognitive.' *Philosophical Psychology*, 22(1), pp. 1-19.

Rowlands, Mark. 2010. *The New Science of the Mind: From extended mind to embodied phenomenology.* Cambridge, MA: MIT Press.

Royce, Terry. 1999. *Visual-Verbal Intersemiotic Complementarity in the Economist Magazine*. PhD Thesis, University of Reading.

Rumelhart, David, Geoffrey E. Hinton and James McCelland. 1986. 'A General Framework for Parallel Distributed Processing.' In: James McCelland, David Rumelhart and the PDP research group. 1986. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1.* Cambridge, MA: MIT Press, pp. 45-76.

Rumelhart, David, Geoffrey E. Hinton and Ronald J. Williams. 1986. 'Learning Internal Representations by Error Propagation.' In: James McCelland, David Rumelhart and the PDP research group. 1986. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1.* Cambridge, MA: MIT Press, pp. 318-362.

Rumelhart, David, James McCelland and the PDP research group. 1986. *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and Biological Models.* Cambridge, MA: MIT Press

Rupert, Robert D. 2004. 'Challenges to the Hypothesis of Extended Cognition.' *Journal of Philosophy*, 101(8), pp. 389-428.

Rupert, Robert D. 2008. 'Causal Theories of Mental Content.' *Philosophy Compass,* 3(2), pp. 353-380.

Saitta, Lorenze and Jean-Daniel Zucker. 2013. *Abstraction in Artificial Intelligence and Complex Systems.* New York

Samuel, Arthur. 1996. 'Phoneme Restoration.' *Language and Cognitive Processes*, 11(6), pp. 647-654.

Sarter, Martin, Ben Givens and John P. Bruno. 2001. 'The cognitive neuroscience of sustained attention: where top-down meets bottom-up.' *Brain Research Reviews*, 35, pp. 146-160.

Sarter, Martin, William J. Gehring and Rouba Kozak. 2006. 'More attention must be paid: The neurobiology of attentional effort.' *Brain Research Reviews,* 51, pp. 145-160/

de Saussure, Ferdinand. 1916 [1983]. *A Course in General Linguistics*. Charles Bally and Albert Sechehaye (eds.). Trans. Roy Harris. La Salle, IL: Open Court.

Scheich, Hennig, André Brechmann, Michael Brosch, Eike Budinger, and Frank W. Ohl. 2007. 'The cognitive auditory cortex: Task-specificity of stimulus representations.' *Hearing Research*, 299, pp. 213-224.

Schneider, Susan. 2013. 'Non-Reductive Physicalism and the Mind Problem.' *Noûs*, 47(1), pp. 135-153.

Scholtz, Barbara C., Francis Jeffrey and Geoffrey K. Pullum. 2014. 'Philosophy of Linguistics.' *The Stanford Encyclopaedia of Philosophy* (Summer 2014 edition), Edward N. Zalta (ed.), URL = http://plato.stanford.edu/archives/sum2014/entries/linguistics/. Accessed: 10/09/14.

Schwiening, Christof J. 2012. 'A brief historical perspective: Hodgkin and Huxley.' *Journal of Physiology*, 590(11), pp. 2571-2575.

Searle, John. 1980. 'Mind, Brains and Programs.' *The Behavioural and Brain Sciences*, 3, pp. 417-457.

Searle, John. 1999. *Mind, Language and Society: Philosophy in the Real World*. London: Weidenfield and Nicolson

Seale, John. 2002. 'Why I am not a property dualist.' *Journal of consciousness studies*, 9(12), pp. 57-64.

Sejnowski, Terrence J. and Charles R. Rosenberg. 1987. 'Parallel Network that Learn to Pronounce English Text.' *Complex Systems,* 1, 145-168.

Shapiro, Lawrence A. 2000. 'Multiple Realizations.' *Journal of Philosophy*, 97(12), pp. 635-654.

Shaprio, Lawrence A. 2004. *The Mind Incarnate*. Cambridge, MA: MIT Press.

Shapiro, Lawrence A. 2007. 'The Embodied Cognition Research Program.' *Philosophy Compass*, 2(2), pp. 338-346

Shapiro, Lawrence A. 2009. 'Review of "The bounds of cognition."' *Phenomenology and the Cognitive Sciences*, 8, pp. 267-273.

Shapiro, Larry. 2011. *Embodied Cognition*. Oxon, UK: Routledge.

Shore, Susanna. 1992. *Aspects of a Systemic Functional Grammar of Finnish.* PhD Thesis. Sydney: Macquarie University.

Simons, Daniel J. and Christopher F. Chabris. 1999. 'Gorillas in our midst: sustained inattentional blindness for dynamic events.' *Perception*, 1999, pp. 1059-1074.

Skaggs, William, Bruce McNaughton, Matthew Wilson and Carol Barnes. 1996. 'Theta Phase progression in hippocampal neuronal populations and the compression of temporal sequences.' *Hippocampus,* 6, pp. 149-172.

Skinner, B.F. 1957. *Verbal Behaviour.* Actom, MA: Copley Publishing Group.

Smolensky, Paul. 1986. 'Neural and Conceptual Interpretation of PDP Models.' In: David Rumelhart, James McClelland and the PDP research group. 1986. *Parallel distributed processing: explorations in the microstructure of cognition. Volume 2, Psychological and Biological Models.* Cambridge, MA: MIT Press, pp. 390-431.

Smolensky, Paul. 1988. 'On the Proper Treatment of Connectionism.' *Behavioural and Brain Sciences*, **11**, pp. 1-23.

Sobert, Elliott. 1999. 'The Multiple Realization Argument against Reductionism'. *Philosophy of Science*, 66(4), pp. 542-564.

Spivey, Mark. 2007. *The Continuity of Mind*. Oxford: Oxford University Press.

Stalnacker, Robert. 1973. 'Presuppositions.' *The Journal of Philosophical Logic,* 2: pp. 447-457.

Stalnacker, Robert. 1996. 'Varieties of Supervenience.' In: James E. Tomberlin (ed.), *Metaphysics, 1996*. Cambridge, MA: Blackwell, pp. 221-241.

Stalnacker, Robert. 1998. 'On the representation of context.' *Journal of Logic, Language and Information*, 7, pp. 3-19.

Stalnacker, Robert. 2002. 'Common Ground.' *Linguistics and Philosophy*, 25, pp. 701-721.

Steels, Luc. 2008. 'The symbol grounding problem has been solved, so what's next?' In: m. de Vega (ed.), *Symbols and Embodiment: Debates on meaning and cognition*. Oxford: Oxford University Press, pp. 223-244.

Stein, J.F. and C.J. Stoodley. 2006. *Neuroscience: an introduction*. Chichester: John Wiley

Sutton, John, Celia B. Harris, Paul, G. Keil and Amanda J. Barnier. 2010. 'The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences,* 9(4), pp. 521-560.

Taverniers, Miriam. 2008. 'Hjelmslev's semiotic model of language: an exegesis.' *Semiotica,* 171(1-4), pp. 367-394.

Taverniers, Miriam. 2011. 'The syntax-semantics interface in systemic functional grammar: Halliday's interpretation of the Hjelmslevian model of stratification.' *Journal of Pragmatics*, 43, pp. 1100-1126.

Teich, Elke. 2003. *Cross-linguistic variation in system and text: a methodology for the investigation of translations and comparable texts*. Berlin: Mouton de Gryter.

Thagard, Paul. 2014. "Cognitive Science". *The Stanford Encyclopaedia of Philosophy*, Edward N. Zalta (ed.), URL = http://plato.stanford.edu/archives/fall2014/entries/cognitive-science/. Accessed = 27/10/15.

Thelen, Esther, Gregor Schoener, Christian Scheier and Linda Smith. 2001. 'The dynamics of embodiment: A field theory of infant perseverative reaching.' *Behavioural and Brain Sciences*, 24, pp. 1-86.

Thibault, Paul J. 2004. *Brian, Mind and Signifying Body: An Eco-Social Semiotic Theory.* London: Continuum

Thompson, Evan. 2007. *Mind in Life: biology, phenomenology and the sciences of mind.* Cambridge, MA: Harvard University Press.

Thompson, Evan and Mog Stapleton. 2008. 'Making Sense of Sense-Making: Reflections on Enactive and Extended Mind Theories.' *Topoi*, 28, pp. 23-30.

Triesman, A. and G. Gelade. 1980. 'A feature integration theory of attention.' *Cognitive Psychology*, 12, pp. 97-136.

Tucker, Gordon H. 1999. *The Lexicogrammar of Adjectives: a systemic functional approach to lexis.* London: Continuum.

Tucker, Gordon H. 2007. 'Between lexis and grammar: towards a systemic functional approach of phraseology.' In: Ruqaiya Hasan, Christian Matthiessen and Jonathan J. Webster (eds.), *Continuing Discourse of Language: a functional perspective, volume 2.* London: Equinox, pp. 953-977.

Tulving, Endel. 2007. 'Are there 256 kinds of memory?' In: J.S.Nairne (ed.), *The Foundations of Remembering: Essays in honour of Henry L. Roediger, III*. New York, NY: Psychology Press, pp. 39-52.

VanRullen, R. and Simon Thorpe. 2001. 'Rate Coding versus temporal coding: what the retinal ganglion cells tell the visual cortex.' *Neural Computation*, 13(6), pp. 1255-1283.

VanValin, Robert D. and Randy J. LaPolla. 1997. *Syntax: structure, meaning and function.* Cambridge: Cambridge University Press

Varela, F., E. Thompson and E. Rosch. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.

Vogelstein, Joshua T., R. Jacob Vogelstein, and Carey E. Priebe. 2011. 'Are mental properties supervenient on brain properties?' *Scientific Reports*, 1(100). DOI: 10.1038/srep00100.

Vogt, Brent. 2005. 'Pain and Emotion Interactions in Subregions of the Cingulate Gyrus.' *Nature Reviews Neuroscience*, 6(7), pp. 533-544.

Walter, Sven. 2010. 'Cognitive extension: the parity argument. Functionalism and the mark of the cognitive.' *Synthese,* 177, pp. 285-300.

Ward, L.M. 2003. 'Synchronous neural oscillations and cognitive processes.' *Trends in Cognitive Science*, 7(12), pp. 553-559.

Warren, Richard M. 1970. 'Perceptual Restoration of Missing Speech Sounds.' *Science*, 167(3917), pp. 392-393.

Washburn, David A. and R. Thompson-Putney. 2001. 'Attention and Task Difficulty: When is Performance Facilitated?' *Learning and Motivation*, 32, pp. 36-47.

Wells, J.C. 2006. *English Intonation: an introduction*. Cambridge: Cambridge University Press.

Whitehead, Albert . 1929. *Process and Reality: an Essay in Cosmology*. New York: McMillan.

Wilson, Andrew D. and Sabrina Golonka. 2013. 'Embodied Cognition is not what you think it is.' *Frontiers in Psychology*. URL: http://journal.frontiersin.org/article/10.3389/fpsyg.2013.00058/abstract. DOI: 10.3389/fpsyg.2013.00058. Accessed: 10/04/2015.

Wilson. Margaret. 2002. 'Six views of embodied cognition'. *Psychonomic Bulletin and Review*, 9(4), pp. 625-636.

Wilson, Robert A. and Andy Clark. 2008. 'How to Situate Cognition: Letting Nature Take Its Course.' In: Phillip Robbins and Murat Aydede (eds.), *The Cambridge Handbook of Situated Cognition*. Cambridge: Cambridge University Press, pp. 55-77.

Wimsatt, William C. 1976. 'Reductionism, Levels of Organization and the Mind-Body Problem.' In: Gordon G. Globus, Crover Maxwell and Irwin Savodnik (eds.), *Consciousness and the Brain: A Scientific and Philosophical Inquiry.* New York, NY/ London: Plenum Press, pp. 205-268.

Wu, Canzhong. 2000. *Modelling Linguistic Resources: A systemic functional approach.* Ms, PhD Thesis: Macquarie University.

Zhu, Yan. 2013. 'The *Drosophila* visual system: from neural circuits to behaviour.' *Cell Adhesion and Migration*, 7(4), pp. 333-344.

Zipser, David. 1986. 'Biologically Plausible Models of Place Recognition and Goal Location.' In: David Rumelhart, James McClelland and the PDP research group. 1986. *Parallel*

distributed processing: explorations in the microstructure of cognition. Volume 2, Psychological and Biological Models. Cambridge, MA: MIT Press, pp. 432-470.