Towards the detection of tonotopic reorganisation in the human auditory cortex



Benjamin Jason Gurer

MRC Insitute of Hearing Research

School of Psychology

University of Nottingham

A thesis presented for the degree of

Doctor of Philosophy

September 2019

Towards the detection of tonotopic reorganisation in the human auditory cortex

Abstract

Peripheral hearing loss not only causes a reduction in hearing sensitivity (increase in hearing threshold) but also leads to problems in suprathreshold hearing – such as difficulty understanding speech in noise – and is often associated with tinnitus. Animal research has suggested that such impairments may be related to changes in central auditory processing arising as a consequence of a reduction in peripheral auditory output (Noreña and Farley 2013). The auditory cortex is tonotopically organised, with neighbouring parts of the cortex responding to neighbouring frequencies. In animals, peripheral hearing loss has been shown to cause large-scale tonotopic reorganisation in auditory cortex, whereby neurons responsive to frequencies affected by the hearing loss (typically higher frequencies), shift their tuning towards less affected (lower) frequencies.

Tonotopic reorganisation in auditory cortex has been investigated in several studies using functional magnetic resonance imaging (fMRI) to non-invasively assess cortical frequency tuning properties in humans. However, it is unclear whether these measurements are sufficiently sensitive

to detect hearing loss induced reorganisation. This thesis presents three complementary studies concerned with developing fMRI methods for effective measurement of tonotopic properties of the human auditory cortex. Voxel-wise measures of cortical frequency preference and selectivity, and summary measures of organisation are evaluated in normal-hearing subjects. The impact of hearing loss on these measurements is assessed.

To estimate cortical frequency tuning properties, blood oxygen level-dependent (BOLD) fMRI responses to sounds with a wide range of frequencies were recorded from normal-hearing participants in three studies to derive the preferred frequency and selectivity of discrete cortical areas. The three studies presented serve to: 1) optimise the fMRI tonotopic mapping procedure, 2) measure tonotopic cortical magnification, and 3) develop methods for measuring tonotopic reorganisation in subjects with hearing loss.

The first study investigated sources of error from experimental factors and suggests how procedures can be optimised for tonotopic mapping fMRI experiments to maximise the validity of tonotopic estimates. The impact of acoustic scanner noise on estimates of tonotopic properties was assessed by comparing tonotopic property estimates derived from data collected using either a sparse acquisition protocol, in which acoustic stimuli are presented in silence but fewer data points are acquired, or a continuous acquisition protocol, where stimulus are presented during data acquisition with typical acoustic scanner noise and more data points are acquired. Frequency preference and selectivity were also estimated using different statistical methods of fitting models to BOLD activity, measured in response to narrowband noise (NBN) stimuli from 8 normal-hearing subjects at 7T; population receptive field modelling (pRF) or calculating the centroid and spread of frequency response profiles estimated

using general linear modelling (GLM). Acoustic scanner noise was found to systematically bias preferred frequency estimates, presumably because it masked some frequencies of the acoustic stimuli. The distribution and range of preferred frequencies derived from the GLM analysis were more biased than those derived from the pRF analysis but the spatial location of high- and low- frequency regions, relative to the frequency range afforded by the analysis, appeared to be robust to these biases. The results suggest that using a sparse acquisition protocol and pRF analysis achieves preferred frequency estimates that best reflect the true neuronal frequency tuning properties.

It has been shown in other sensory systems, that the cortical surface allocated to a given region in sensory space is proportional to its behavioural importance. For example, visual studies have demonstrated that proportionally more cortical area dedicated to the centre of the visual field than to the periphery (Wandell, Dumoulin, and Brewer 2007). This cortical magnification function has been well characterised in visual cortex using fMRI (Dumoulin and Wandell 2008). Tonotopic magnification in the human auditory cortex (the cortical distance spanned by a given frequency) is still unknown but we expect it to follow behavioural relevance as in the visual system.

The second study of this thesis aimed to measure and describe tonotopic cortical magnification in normal-hearing subjects. The auditory cortical magnification function was calculated in 8 normal-hearing subjects from voxel-wise estimates of preferred frequency derived using a pRF analysis from fMRI data acquired using a sparse acquisition protocol at 7T. Cortical magnification was estimated for two mirror-symmetric tonotopic gradients located on Heschl's gyrus in each hemisphere by characterising the relationship between cortical distance and preferred frequency with

the best fitting of four functions: linear, logarithmic, cochlear frequency tuning and behavioural frequency discrimination. Preferred frequency transformed by a behavioural frequency discrimination scale provided the best fit, suggesting that the cortical tonotopic magnification function directly reflects frequency discrimination. Estimates of the cortical magnification function provide a global measure and characterisation of cortical organisation of the auditory cortex which may be able to detect and summarise tonotopic reorganization in hearing loss subjects. Understanding the retinotopic organisation and cortical magnification of the human visual cortex was critical to investigate retinotopic reorganisation following localised periphery deprivation (Wandell and Smirnakis 2009; Baseler et al. 2011). Likewise, for auditory research, understanding tonotopic magnification will benefit studies of tonotopic reorganisation and its possibility as a generation mechanism for chronic tinnitus.

The third study of this thesis develops a better control for future studies of tonotopic reorganisation in subjects with hearing loss, and its potential as a mechanism for supra-threshold hearing deficits and chronic tinnitus, which is not addressed directly in the current literature. In participants with hearing loss, BOLD responses to frequencies affected by the hearing loss are expected to be reduced, potentially skewing the shape of the measured frequency response functions and biasing cortical preferred frequency estimates. Two methodological procedures were proposed to address these issues. The first is to simulate hearing loss in normal-hearing controls in order to replicate the biases expected for hearing-impaired subjects. The second, tested here in normal-hearing controls with simulated hearing loss, is a novel data analysis method that accounts for variations in sensation level due to hearing loss to alleviate the biases of hearing loss on estimates of preferred frequency. Tonotopic proper-

ties of auditory cortex were estimated in seven normal-hearing subjects using a sparse acquisition sequence at 3T and a pRF analysis, either in the presence or absence of a continuous threshold-elevating noise simulating a steeply sloping high-frequency hearing loss. Steeply sloping high-frequency hearing loss has been suggested to lead to tonotopic reorganisation (Sereda, Hall, et al. 2011), providing a realistic case for assessment of the proposed methods. In addition, a modified population receptive field modelling method was implemented that attempted to minimize the effect of hearing loss on tonotopic property estimates by taking into account the hearing thresholds in its fitting procedure. The results show that using the unmodified pRF method for simulated hearing loss data results in expected biases in preferred frequency and tuning width estimates, and that the modified method reduces these biases and results in tonotopic maps closer to ones estimated from normal-hearing data. These results represent a key step towards studying the effects of hearing loss on the tonotopic organisation in human auditory cortex using BOLD fMRI, and understanding their potential perceptual consequences.

The work presented in this thesis provides actionable methodologies that could potentially improve the study of hearing-loss-induced tonotopic reorganisation in human auditory cortex.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification.



Benjamin Jason Gurer

Acknowledgements

First, I would like to thank my supervisors: Katrin, for giving me the opportunity, Denis, for his invaluable guidance, and Julien, for his ceaseless support.

If Katrin hadn't taken a chance on me, I would have never have had the privilege of researching questions I find truly interesting. If Denis hadn't taken me on, I would have quit long ago. If Julien hadn't continued to support me from Beirut, I'd have no idea what to do.

I want to thank and congratulate the friends I made along the way. We made it. Special thanks to: Martin, the best office mate and a truly great friend. Alex, three degrees together must be a record. Michael, for his help scanning and shared love of Matlab.

I also need to thank my friends and family for not forgetting me when I disappeared into a PhD shaped hole.

Most of all I would like to thank my beautiful wife, Bec. She's listened to me vent and encouraged me when I felt like an imposter. Her support has kept me sane and most importantly, I want to thank her for reminding me that there are more important things in life than work and a PhD.

Contents

1 General Introduction **1**

1.1 Hearing 1

1.2 Peripheral and central auditory processing 2

1.3 The Human Auditory Cortex 5

1.3.1 Anatomical landmarks of the human auditory cortex 5

1.3.2 Tonotopic organisation of the human auditory cortex 5

1.4 Clinical impact of tonotopic research 11

1.5 Magnetic Resonance Imaging 13

1.5.1 Nuclear Magnetic Resonance (NMR) 13

1.5.2 Relaxation - equations 18

1.5.3 Magnetic Resonance Imaging 21

1.5.4 functional MRI 26

1.5.5 Why go to higher field? 29

1.5.6 fMRI data analysis 30

1.5.7 fMRI of the auditory system 31

1.6 Overview of studies 32

2 General Methods **35**

2.1 General outline 35

2.1.1 Study 1 - Tonotopic mapping of the human auditory cortex: influence of acquisition protocol and data analysis 36

2.1.2 Study 2 - Cortical magnification of the human auditory cortex 37

2.1.3 Study 3 - How can we reliably detect hearing loss induced cortical tonotopic reorganisation in humans? 38

2.2 Data acquisition 39

2.2.1 Structural data acquisition 39

2.2.2 Functional data acquisition 40

2.2.3 Auxiliary data acquisition 41

2.2.4 Optimising data acquisition 41

2.2.5 Auditory stimulus 42

2.2.6 Paradigm 45

2.2.7 Accurate acoustic stimulus presentation 46

2.3 Data Analysis 47

2.3.1 Initial processing of functional data 47

2.3.2 Estimating voxel frequency preference and selectivity 50

2.3.3 Calculating gradient reversals 58

2.3.4 Estimation of the hemodynamic response function . . 58

2.3.5 Defining Regions of Interest (ROI) 64

2.3.6 ROI analysis 67

3 Tonotopic mapping of the human auditory cortex: influence of acquisition protocol and data analysis 70

3.1 Introduction 70

3.2 Methods 75

3.2.1 Data collection 76

3.2.2 Data analysis 76

3.3 Results 78

3.3.1 Impact of analysis methods on estimates of cortical frequency tuning properties 78

3.3.2 Comparison of Acquisition Protocols 87

3.4 Discussion 99

3.4.1 Impact of analysis methods on estimates of cortical frequency tuning properties 99

3.4.2 Impact of scanner noise on cortical estimates of frequency preferred and selectivity 105

4 Cortical magnification of the human auditory cortex 111

4.1 Introduction 111

4.2 Methods 117

4.2.1 Dataset 117

4.2.2 Data Analysis 118

4.3 Results 123

4.3.1 Change in cortical distance as a function of frequency 123

4.3.2 Cortical magnification factor of the auditory cortex . 124

4.3.3 Point image in the auditory cortex 127

4.3.4 Differences between anterior and posterior gradients . 127

4.4 Discussion 127

5 How can we reliably detect hearing loss induced cortical tonotopic reorganisation in humans? 133

5.1 Introduction 133

5.2 Methods 140

5.2.1 Data collection 140

5.2.2 Hearing Loss Simulation 141

5.2.3 Data Analysis 143

5.2.4 Modified population receptive field analysis 143

5.2.5 Data Visualisation 144

5.3 Results 145

5.3.1 Impact of hearing loss simulation on preferred frequency estimates 145

5.3.2 Impact of hearing loss simulation on average voxel frequency response profiles 146

5.3.3 Impact of hearing loss simulation on average voxel frequency response profiles and tuning width 149

5.4 Discussion 154

6 General Discussion **157**

6.1 Accurate measurement of tonotopic properties 157

6.2 Cortical tonotopic magnfiction 159

6.3 Measuring hearing loss induced tonotopic reorganisation . . 160

6.4 Future Directions 161

List of Figures

1.1 Models of frequency organisation in the monkey and human auditory cortex. 7

2.1 Volumetric to cortical surface to flat map 49

2.2 Population receptive field (pRF) modelling procedure 57

2.3 Average hemodynamic response function 60

2.4 HRF to different centre frequencies as a function of distance from preferred frequency 61

2.5 Group-average fitted HRF functions 63

3.1 Comparison of the spatial distribution of voxel-wise preferred frequency estimates using the four different estimation methods 80

3.2 Distribution of preferred frequency estimates across AC voxels for different estimation methods 82

3.3 Comparison of spatial distributions of voxelwise frequency selectivity computed using three different estimation methods 84

3.4 Impact of analysis methods on the spatial distribution of voxelwise goodness of fit 86

3.5 Impact of acquisition protocol on the spatial distribution of voxelwise preferred frequency 89

3.6 Impact of acquisition protocol on average BOLD response in primary auditory cortex as a function of stimulus frequency 91

3.7 Comparison between the difference in average BOLD response between acquisition protocols 92

3.8 Impact of acquisition protocol on the distribution of preferred frequencies 93

3.9 Frequency response profiles averaged according to voxel's preferred frequency for continuous and sparse acquisition data 95

3.10 Effect of acquisition protocol on the spatial distribution of voxelwise goodness of fit of the pRF models to the fMRI time courses 97

3.11 Effect of acquisition protocol on the spatial distribution of voxelwise frequency selectivity 98

4.1 ERB and DLF as a proportion of frequency 114

4.2 Mapping of frequency to cortical distance 121

4.3 Change in average cortical distance as a function of frequency 125

4.4 Average cortical distance, tuning width, cortical magnification factor and point image as a function of frequency 126

4.5 Cortical distance as a function of preferred frequency in anterior and posterior AC gradients 128

4.6 Distribution of preferred frequency in anterior and posterior AC gradients 129

4.7 Average voxel frequency response profiles of the anterior and posterior AC gradients 130

5.1 Impact of high-frequency hearing loss on estimates of pRF properties 136

5.2 A flow chart describing the modified population receptive field (pRF) modelling procedure 144

5.3 Spatial distribution of voxel-wise preferred frequency estimates from normal-hearing and hearing loss simulation conditions 147

5.4 Distribution of preferred frequency estimates from normal-hearing and hearing loss simulation conditions 148

5.5 Average percentage signal change in response to stimuli at different centre frequencies in the presence and absence of hearing loss simulation 150

5.6 Difference between average beta weights collected in the presence and absence of hearing loss 151

5.7 Impact of hearing loss simulation on average voxel frequency response profiles averaged 152

5.8 Spatial distribution of voxel-wise frequency selectivity estimates from normal-hearing and hearing loss simulation conditions 153

Abbreviations

AC ROI	Primary Auditory Cortex Gradients ROI
AP	Anterior-Posterior
B_0	Static Magnetic Field
B_1	Rotating Magnetic Field
BOLD	Blood-Oxygenation-Level-Dependent
CR	Critical Ratio
dB	Decibel
dB HL	Db Hearing Level
dB SPL	Db Sound Pressure Level
dHb	Deoxygenated Haemoglobin
EMF	Electromagnetic Field
EPI	Echo-Planar Imaging
ERB	Equivalent Rectangular Bandwidth
f	Frequency
FDR	False Discovery Rate
FFT	Fast Fourier Transform
FH	Foot-Head
FID	Free Induction Decay
fMRI	Functional Mri
FOV	Field Of View
FWHM	Full-Width-Half-Maximum

GE	Single-Shot Gradient Echo
GLM	General Linear Modelling
Hb	Oxygenated Haemoglobin
HG	Heschl'S Gyrus
HRF	Hemodynamic Response Function
Hz	Hertz
J	Spin Angular Momentum
kHz	Kilohertz
M_{xy}	Rotating Magnetisation
M_0	Net Magnetisation
MPRGE	3D Magnetization-Prepared Rapid Gradient Echo
MRI	Magnetic Resonance Imaging
μ	Magnetic Moment
NBN	Narrow-Band Noise
nERB	Number Of Cochlear Filter Bandwidths
NMR	Nuclear Magnetic Resonance
PAC	Primary Auditory Cortex
PD	Proton Density -Weighted)
PP	Planum Polare
pRF	Population Receptive Field Modelling
PSIR	3D Phase-Sensitive Inversion Recovery
PT	Planum Temporale
RF	Radio-Frequency
RL	Right-Left
ROI	Region Of Interest
ROI	Regions Of Interest
RSS	Sum Of Squared Residuals
SE	Spin Echo Sequence
SL	Sensation Level

SNR	Signal-To-Noise
STG	Superior Temporal Gyrus
STP	Superior Temporal Plane
T	Tesla
TE	Echo Time
TFE	Turbo-Field Echo
TI	Inversion Times
TR	Repetition Time

Chapter 1

General Introduction

This introductory chapter will detail the basic principles underlying the studies presented in this thesis. First, I will present a brief overview of the physics of sound and the neuroscience of the mammalian auditory system. In particular, I will describe how sounds are thought to be represented in the human cortex, focusing on the spatial organisation of sound frequency in the brain. I will then discuss the physics underlying the imaging method used for the experimental data acquisition.

1.1 Hearing

The sense of hearing is the ability to perceive sound. Hearing is vital to many experiences that are crucial to the quality of life, for example, communication with others, enjoying music and reacting to sudden changes in our environment. Sound is produced by the vibration of objects, which in turn causes motion of the molecules in the surrounding air. This motion compresses and rarefies the air molecules, creating sound pressure waves that radiate from the object (Everest and Pohlmann 2009a; Ever-

est and Pohlmann 2009c). Humans perceive sound vibrations with frequencies between 20 to 20,000 cycles per second (Hertz, Hz), the audible range.

Changes in the frequency of these sound pressure waves – through a set of transformations at different stages in the auditory system – ultimately lead to *perceived* changes in pitch. The intensity or amplitude of the sound pressure waves, the amount of difference between the condensations and rarefactions of the air, is perceived as the auditory quality of *loudness*. These two perceptual qualities – pitch and loudness – are the main characteristics that define our perception of sound. There are other qualities of sound such as *timbre*, which relates to the complexity and number of frequency components of the sound source. The relationship between the physical source and its perception, however, is even more complex than the mapping between the frequency and perceived pitch, and intensity and perceived loudness (Moore 2012; Everest and Pohlmann 2009d).

1.2 Peripheral and central auditory processing

At the coarsest scale, the auditory system can be thought of as having two parts, peripheral and central. The peripheral auditory system can further be divided into three parts: the outer ear, middle ear and inner ear. The outer ear or the external ear is called the pinna. The pinna is a spatial dependent frequency filter and an important component of spatial hearing. The basic principle of this process is that the shape of the pinna will alter the frequency profile of a sound depending on its relative spatial location. In addition, pinna focuses sound into the ear canal, which

contains the tympanic membrane, the boundary to the middle ear (Pickles 2012; Moore 2012).

Sound entering the ear will cause the tympanic membrane to vibrate. Behind the tympanic membrane is the middle ear, which contains three bones, the ossicles. The first bone, the malleus, connects to the tympanic membrane and through the next two bones of the ossicles, the incus and stapes, transmits the vibrations of the tympanic membrane to the inner ear. The middle ear converts sound from acoustic pressure to mechanical movement that is optimised for transmission to the inner ear, overcoming the impedance mismatch of transferring energy travelling in the medium of air to the liquid-filled inner ear. The stapes of the middle ear connects to the oval window of the cochlea, a labyrinth-like bone structure that contains the organ of Corti (Pickles 2012; Moore 2012).

The inner ear converts the mechanical energy from the middle ear into electrical energy in the form of neuronal activity. The neuronal output from the organ of Corti is the brain's interface between physical acoustic sound and the sense of hearing, forming the input to the central auditory system, the auditory-responsive and -processing parts of the brain. The organ of Corti contains the basilar membrane, a membrane that varies in thickness and rigidity as it progresses along the cochlea (Pickles 2012; Moore 2012).

When an external oscillating, mechanical force is applied to a physical object, the relationship between the frequency of the force and the material properties of the object will determine the efficiency of energy transfer. And so, from the varying mechanical properties arises a relationship between frequency and position on the membrane, with the locations of membrane displacement depending on the frequency components of the

sound (Everest and Pohlmann 2009b).

Specialised hair cells, called *inner hair cells*, are placed along the length of the membrane that converts the mechanical movement of the basilar membrane to neuronal activity. As a result, the inner hair cells are tuned to different frequencies, progressing in a continuous gradient from higher frequencies at the base of the cochlea, where the basilar membrane is rigid and thin, to lower frequencies as the apex, where the basilar membrane is supple and thick. This frequency and location relationship is referred to as *tonotopic*, tono = frequency, topic = place. The cochlear also features *outer hair cells* that increase tonotopic resolution through active mechanisms. These outer hair cells are often the first to be damaged in hearing loss (Pickles 2012; Moore 2012).

The inner ear can be thought of as a sound frequency analyser, converting sound into numerous frequency components and transmitting this information to the central auditory system in parallel channels. This tonotopic organisation is set up in the peripheral auditory system and maintained throughout the auditory system. It is a fundamental organisational principle of the auditory system.

One can think of the central auditory system as processing the electrical output from the inner ear to determine the acoustical properties of the sound, which ultimately leads to the perception of the sound and the sense of hearing. From the cochlea, the auditory pathway passes through several structures of the midbrain, each performing specific transformations on the original cochlear input signal before finally projecting from the thalamus of the midbrain to the auditory cortex (Pickles 2012; Moore 2012). The work presented in the thesis will focus on the function of the auditory system at the cortical level. The following section will cover the

findings of the functional properties of the human auditory cortex, focusing on fMRI studies of cortical tonotopic organization.

1.3 The Human Auditory Cortex

1.3.1 Anatomical landmarks of the human auditory cortex

The human auditory cortex is located on the superior temporal plane (STP) within the lateral (Sylvian) fissure, as well as laterally on the superior temporal gyrus (STG). The medial two-thirds of Heschl's gyrus (HG), the main macro-anatomical feature of the STP, has been identified to be the approximate location of the primary auditory cortex (PAC) through anatomical measured features associated with primary sensory cortices; dense myelination, properties of cortical layers and thalamic connectivity (Brodmann and Gary 1909; Wallace, Johnston, and Palmer 2002). The auditory cortex extends from HG anteriorly to planum polare (PP), separated by the first transverse sulcus, and posteriorly to planum temporale (PT), separated by Heschl's sulcus. The anatomical structure of HG varies between subjects, and within-subjects between hemispheres, with partial and complete duplications of the gyrus also being observed, as well as the more common single HG (Tzourio-Mazoyer and Mazoyer 2017).

1.3.2 Tonotopic organisation of the human auditory cortex

Neuroimaging techniques allow for non-invasive, in-vivo measurement of human brain structure and activity, opening insights into the anatomi-

cal and functional properties. The development of neuroimaging methods has allowed for increasingly accurate measurements of the brains structure and functional activity, allowing for greater insights and understanding. The tonotopic organisation of the human auditory cortex has been investigated using functional magnetic resonance imaging (fMRI) since the early days of its inception as a research tool. The organisation of tonotopic gradients in the human auditory cortex has required decades of research to determine. The difficulties stem not only from experimental sources, such as the resolution and sensitivity of fMRI, the relatively small size of the auditory cortex in relation imaging resolution, and the confounding effects of acoustic scanner noise, generated by the gradient coils used in the data acquisition process, but also the difference between tonotopic organisation in non-human primates and humans.

Much of what is known about the human auditory cortex has been extrapolated from animal studies. Tonotopic organisation has been observed in numerous species, throughout the auditory pathway, and reliably measured in the auditory cortex by electrophysiological studies. In primates, invasive studies have established a model of the rhesus macaque auditory cortex (Kaas and Hackett 2000). Early fMRI studies investigated tonotopy under pretences extrapolated from this model. However, this assumption is incorrect; it has taken modern advances in MRI and decades of research to determine that the organisation of the human auditory cortex is fundamentally different from non-human primate models. With hindsight, the results from auditory fMRI tonotopic measurements present a consistent pattern of tonotopic organisation. However, the interpretation of findings has been difficult and inconsistent, resulting in two possible configurations being proposed. Parallel, with tonotopic gradients running parallel to HG, and orthogonal, with tonotopic gradients

running orthogonally to HG. Only recently has the organisation of the human AC been determined.

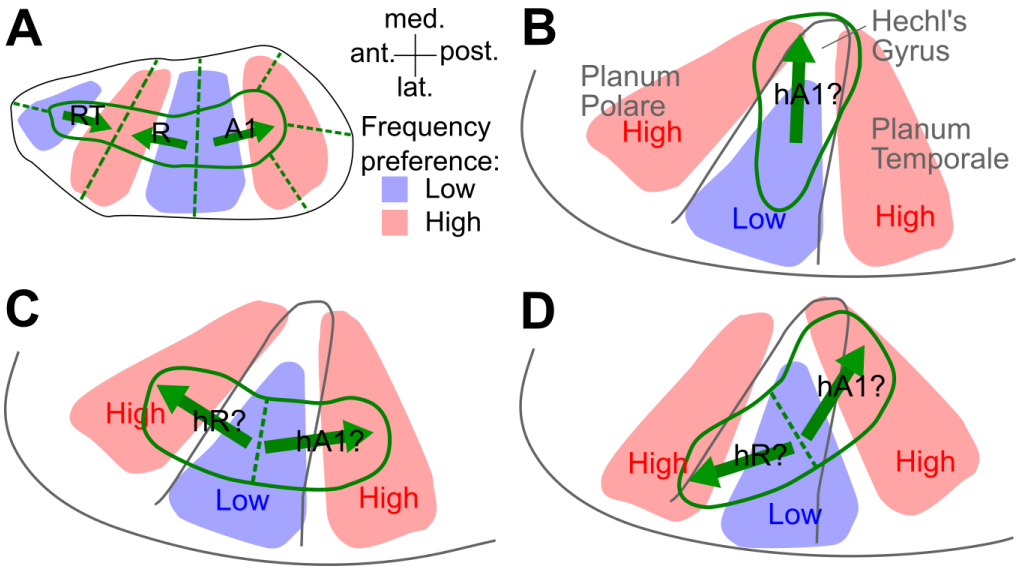


Figure 1.1: Models of frequency organisation in the monkey and human auditory cortex.

A) The macaque auditory cortex proposed by Kaas and Hackett (2000) containing three core areas (A1, R, RT). B) The "parallel" model based on anatomical findings and the assumption the organisation of human auditory cortex would follow the macaques with the human A1 (hA1) orientated parallel to HG. C) The "orthogonal" model proposed by Humphries, Liebenthal, and Binder (2010) and contemporary human fMRI studies with hA1 and the human R area (hR) extending orthogonally across HG and tonotopic gradients. D) The "oblique" model proposed by Baumann et al. 2013 in an attempt to reconcile the two conflicting models, "parallel" and "orthogonal". Here hA1 and hR extend diagonally across HG and tonotopic gradients. Figure used with permission from Besle et al. (2019).

Tonotopic organisation has been observed throughout the auditory pathway in non-human primates by numerous electrophysiological studies. In the rhesus macaque, invasive studies have established a model of the auditory cortex, which has traditionally be assumed to be a good basis for a model of the human auditory cortex. The macaque auditory cortex is lo-

cated on the superior temporal plane (STP) and comprises three primary areas, A1, R, RT, referred to as the "core", which are surrounded by secondary regions, collectively referred to as the "belt". The primary areas of the core are adjacently stacked, running anteroposteriorly from A1, located on the posterior STP, to R and RT, ending at a slight protrusion on the edge of the circular sulcus near the middle of STP, elongating the auditory cortex. Tonotopic gradients run parallel to the core, along an anterior-posterior axis and are shared between core and belt areas. Reversals in the tonotopic gradient mark the boundaries between the three primary areas, as well as some belt areas. Gradient reversals, however, do not mark the boundaries between core and belt areas and require further measures to distinguish between them (Kaas and Hackett 2000; Morel, Garraghty, and Kaas 1993).

The major anatomical difference in the human and rhesus macaque STP is HG, which is absent in the rhesus macaque. Prior to the advances in non-invasive neuroimaging, post-mortem anatomical human studies were the primary means of investigating the organization of the human auditory cortex. Anatomical studies have consistently reported a primary-like core region located on the medial two-thirds of HG, surrounded by belt-like regions, through measures of cytoarchitecture (Wallace, Johnston, and Palmer 2002). These studies were taken to suggest that the organisation of the human auditory cortex was similar to that of non-human primates, with core and belt regions, but rotated to be elongated lateral-medially along HG.

The rise of non-invasive neuroimaging methods allowed the tonotopic gradients of the human AC to be investigated for the first time. Under the assumption that the human AC would follow that of the non-human primates, the axis of tonotopic gradients would be expected to follow the

direction the anatomically identified elongate auditory core like regions of HG, parallel to HG. Early studies supported this hypothesis. Talavage et al. (2000) found a low-frequency responsive node laterally located on HG and a high-frequency node located medially on HG. This was interpreted to be in support of the hypothesis that tonotopic gradients ran parallel with HG, akin to the gradients of non-human primates progressing in parallel with the core, but with the entire AC rotated medial-laterally (Talavage et al. 2000). The study by Talavage et al. (2000) was limited not only by early MRI technology, in terms of the resolution and sensitivity of measurements, but also experimentally. The study presented stimuli of two different frequencies, which didn't allow for the assessment of tonotopic gradients, only the location of regions responding significantly to the stimulus frequencies. The organisation proposed in this study was supported by other early studies (Engelien 2002; Formisano et al. 2003).

This proposed organisation was soon brought into question by findings from more advanced fMRI studies. Schönwiesner, Yves Von Cramon, and Rü (2002) questioned these findings and presented a study using a range of frequencies, rather than just high- and low-frequency stimuli, to assess the tonotopic gradient of the human AC. However, the results from this study did not produce clear tonotopic gradients that could be simply assessed to determine if the organisation of the human AC followed that of non-human primates. Formisano et al. (2003) presented the first auditory fMRI study performed at 7 Tesla (T) and took advantage of the improved resolution afforded by the high magnetic field to produce detailed maps of tonotopic gradients in individual subjects. Results from this study appeared to support the non-human primate model, with mirror gradients axis-aligned roughly with HG. Later studies restricted the analysis of tonotopic gradients to predefined core regions, based on struc-

tural or independent functional measurements, corroborated these findings and supported this interpretation (Dick et al. 2012; Moerel, Martino, and Formisano 2012; Schönwiesner, Dechent, et al. 2015).

Humphries, Liebenthal, and Binder (2010) were the first to suggest an alternate interpretation, with the axis of the tonotopic gradient being orthogonal to HG running perpendicularly across HG, as opposed to parallel as previously suggested (Humphries, Liebenthal, and Binder 2010). For this arrangement to corroborate with the non-human primate model, it would require the auditory cortex to be elongated along the same axis, anterior-posteriorly, across HG. This, however, is in disagreement with findings from post-mortem anatomical human studies.

Since the findings by Humphries, Liebenthal, and Binder (2010), a consistent pattern of tonotopic gradients has been found by recent fMRI studies (Humphries, Liebenthal, and Binder 2010; Da Costa et al. 2011; Langers 2014), from both group-averaged and individual subject level data. The auditory cortex has a low-frequency wedge, located on the centre of the lateral part of HG with two mirror-reversed tonotopic gradients, one running medioanteriorly and the other posteriorly towards different high-frequency regions, forming a V-shaped gradient centred on HG, the angle of which varies between subjects. It is possible that earlier studies may have lacked the spatial resolution to differentiate between the two gradients. Baumann et al. (2013) proposed an oblique orientation of the auditory cortex in an attempt to reconcile the conflicting models with opposing orientations of tonotopic gradients, with the human auditory cortex and the axis of the tonotopic gradient oriented at an angle to HG (Baumann et al. 2013).

Only recently have these three models been assessed against each other

using the same data set. Besle et al. (2019) tested all three proposed organisations of tonotopic gradients in the human AC, "parallel", "orthogonal" and "oblique". Based on their interpretation of structural MRI measures of myelin and functional MRI measures of frequency selectivity and tonotopic gradients the authors suggest that none of the previously proposed arrangements agrees with the data. They observe a contradiction between the direction of the tonotopic gradient and the orientation of the core, and therefore claim that the organisation of tonotopic gradients in the human AC does not follow that of non-human primates (Besle et al. 2019). The exact model of the tonotopic organisation in the human auditory cortex may slight not be agreed upon, but the consistent pattern of tonotopic gradients robustly found in numerous studies suggests it is appropriate to study changes in this pattern.

1.4 Clinical impact of tonotopic research

Despite the shortcomings in our knowledge and the unknown tonotopic organisation of the human auditory cortex, tonotopic mapping experiments have found clinical research application in the investigation of chronic tinnitus, the perception of phantom sound, aiming to determine the possible neuronal mechanisms underlying it. One model of tinnitus generation suggests that peripheral hearing loss, the reduced sensitivity to sounds due to damage to hair cells within the inner ear, has been suggested to lead to reorganisation of tonotopic maps in the auditory system through plasticity and homeostatic mechanisms (Pienkowski and Eggermont 2011). Animal models have suggested that through these mechanisms, and the resulting reorganisation, over-representation of frequencies at the boundary of hearing loss can occur, which can lead to the perception of phantom sounds (Eggermont and Roberts 2004; Eggermont and

Roberts 2014). Tonotopic mapping experiments provide a powerful tool to investigate this hypothesis but as of yet, results remain inconclusive (Langers, Kleine, and Dijk 2012; Papanikolaou et al. 2015). This is partly due to the lack of knowledge about the human auditory cortex and the disagreement in models of tonotopic organisation, but also due to limitations of experimental methods employed. The work in this thesis aims to address both of these issues. The first study presented systematically investigates how to accurately measure the frequency properties of the human auditory cortex using fMRI. Frequency is the principle organisation domain of the auditory cortex and therefore accurate measurements of frequency properties are vital for further investigation. The second study measures the cortical magnification function of the human auditory cortex, which relates, different sound frequencies with the corresponding cortical surface area, a property that could potentially be used to measure cortical reorganisation in clinical populations. The final study presents a holistic experimental procedure for measuring tonotopic map reorganisation in subjects with hearing loss, with the aim to minimise the confounding effects of variation of hearing sensitivity as a function of frequency on tonotopic mapping. The studies presented in this thesis add to the auditory fMRI literature by suggesting actionable improvements to experimental methods, both in the collection and analysis of functional data, presenting measurements of cortical magnification in a normal-hearing population, a measure that may change with hearing loss, and demonstrating how hearing loss can be accounted for in the design and analysis of auditory fMRI studies.

1.5 Magnetic Resonance Imaging

The research questions I address in my thesis require determining the spatial organisation of frequency selectivity in auditory cortex. This, in turn, requires measurements of how brain activity changes over time – in response to a stimulus – and across the cortical surface – depending on variations in the properties of the stimuli. To achieve this, we need a measurement technique that meets the following criteria. It needs to:

- be non-invasive: we wish to measure brain activity in healthy human participants, repeatedly and with no side effects.
- provide high spatial resolution data: we need to distinguish how the functional and structural properties of the brain change across the cortical surface.

Magnetic resonance imaging (MRI) meets these criteria, with a superior spatial resolution to other non-invasive brain imaging techniques. The following chapter will present the basic physical principles of MRI and functional MRI (fMRI). I will then discuss the unique challenges faced by fMRI of the auditory system and some experimental design considerations.

1.5.1 Nuclear Magnetic Resonance (NMR)

Nuclear magnetic resonance (NMR) is a phenomenon that is related to an intrinsic property of atomic nuclei, spin, which is the source of the signal measured by MRI. All matter is made of atoms, which each comprise of one nucleus and a number of electrons. The nucleus itself contains

nucleons – protons and neutrons. Nuclei have four fundamental properties: Charge, Spin, Mass and Magnetic moment. The spin of a nucleus depends on its number of nucleons: only nuclei with an odd number of either protons or neutrons possess spin. Nuclei with spin have a spin angular momentum (J) and a magnetic moment (μ).

$$\mu = \gamma J \tag{1.1}$$

Where γ is the gyromagnetic ratio, a property that is specific to each type of nucleus.
NMR experiments measure the aggregate magnetic moment (μ) of many, many nuclei. Any nuclei with spin can be used for NMR experiments. However, for imaging biological sources, protons (^1H) are most commonly used due to their abundance in the human body and their large gyromagnetic ratio, and therefore generate the greatest NMR signal possible from a natural source.

The basic idea of the nuclear magnetic resonance phenomenon can be explained with the following classical picture: When a nucleus with a magnetic moment is placed in a strong magnetic field, it experiences torque, a twisting force that causes rotation of the nuclei and will align the nucleus with the applied field. In combination with the nucleus' spin, torque will cause the nucleus to precess like a gyroscope, rather than the nuclei simply rotating and aligning. The earth's motion is also a precession, spinning around its own axis (daily) and precessing (seasons) while travelling around the sun. A precessing nucleus releases energy that can be measured as an electromagnetic wave at the Larmor frequency. The frequency of this precession is given by the Larmor equation:

$$\omega = \gamma B_0 \tag{1.2}$$

Where ω is the frequency of precession, also known as the Larmor frequency, and B_0 is the magnetic field. For ^1H, the gyromagnetic ratio is $42.58\,\mathrm{MHz}/T$, and so, at a magnetic field of $7T$ the Larmor frequency is $298.02\,\mathrm{MHz}$.

One key idea for understanding the NMR phenomenon is the so-called Zeeman splitting, the quantum view of a population of nuclei with spin. This is the fact that in the presence of a magnetic field, the spin angular momentum (J) of protons is quantised to two possible Zeeman energy levels, a more stable low energy state, spin up, and a less stable high energy state, spin down. A useful analogy to this is a gymnast on a high-bar aligned with the gravitational field: there are two relatively stable states: in a hand-stand (high-energy) or simply hanging (low energy).

In the absence of a magnetic field, the spin states for a population of nuclei are uniformly distributed and the magnetic moments cancel each other out and so no signal can be measured. However, in a strong magnetic field, B_0, the distribution of spins into the two states is appreciably biased towards the low-energy state. The signal in an NMR experiment relies on this unequal distribution of spins states: many nuclei with magnetic moments in a high spin states are aligned parallel to the large static magnetic field B_0, but even larger numbers of nuclei in a low spin state are aligned antiparallel to B_0, which leads to a net magnetisation (M_0). The proportions of low and high spin states are determined by the so-called Boltzmann distribution.

$$\frac{\eta\downarrow}{\eta\uparrow} = e^{\frac{\gamma h B_0}{k_b T}} \tag{1.3}$$

Where

h: planck's constant divided by 2π

k_b: boltzmann's constant

B_0: magnetic field strength

γ: the gyromagnetic ratio (value dependent on the nuclei)

$\eta \downarrow$: number of spins in low energy state

$\eta \uparrow$: number of spins in high energy state

As the equation illustrates, the distribution of spin energy levels between high and low states is determined by both the magnetic field strength and the temperature of the system: The larger the magnetic field strength and the lower the temperature, the larger the ratio of low to high spin states and the greater the signal. To create larger signals, NMR therefore uses a large, static magnetic field (B_0) as it is not viable to lower the temperature to obtain a meaningful effect on M_0 when imaging a biological subject.

The net magnetisation of the nuclei (M_0) is many times smaller than the static field (B_0). At first sight, it might seem impossible to measure the tiny net magnetisation of the nuclei (M_0) in the presence of the large, static field (B_0). Magnetic resonance imaging overcomes this, by making M_0 measurements in the transverse plane (orthogonal to the direction of B_0). By convention, MRI uses a coordinate system in which the direction of B_0 defines the z-axis, and the x and y axes correspond to the transverse plane in which measurements are taken.

In steady-state, M_0 is aligned to B_0. Applying energy to the population of nuclei, using a radio-frequency (RF) pulse, excites transitions to occur between the two Zeeman energy states. This tilts M_0 out of alignment with the z-axis (i.e. in line with B_0), towards the transverse plane. In order to excite a particular population of nuclei, the RF pulse must provide energy at the resonance (Larmor) frequency of these nuclei. For ^1H, this is done using a rotating magnetic field (B_1) whose rotation frequency corresponds to the Larmor frequency of protons. The duration and strength

of B_1 determine the angle by which M_0 will tip – this is known as the flip angle. A useful analogy is the idea of pushing a child on a swing. If energy is provided at the correct resonance frequency (of the swing), then it will result in the child swinging higher and higher.

Applying this RF pulse (B_1) has an additional effect on top of altering spin energy states: it brings the proton spins into phase (all spins rotate at the same rate in harmony), resulting in a net magnetization in the transverse plane. The end result of applying the right amount of RF energy at the right frequency is to tilt the net magnetisation vector from the longitudinal (B_0) axis into the transverse plane.

Signal detection

Immediately following an RF pulse that flips M_0 by 90°, all proton spins are in phase and equally distributed between Zeeman energy levels, resulting in a net magnetisation in the transverse (x-y) plane which precesses around B_0 at the Larmor frequency. Importantly, this rotating magnetisation M_{xy}, produces an oscillating electromagnetic field (EMF) that can induce a current in a radio frequency pickup coil surrounding the sample (head). This signal generation forms the basis of MRI.

This oscillating EMF is not a stationary signal however: as soon as the RF pulse ceases, the net magnetisation starts "relaxing" back into alignment with B_0. This relaxation has two components, corresponding to two different processes. One is related to alignment of the net magnetisation vector back into the main magnetic field direction (B_0 or z), in other words, the return to an unequal distribution of spin states. This process is called T_1 relaxation and happens on the order of hundreds of milliseconds to seconds. The other process, called T_2 relaxation, is related

to the loss of coherence between individual proton spins resulting in a decrease in net transverse magnetization. Signal decay due to T_2 relaxation is much quicker – on the order of tens of milliseconds. Both these effects are well described by empirically derived functions with particular time-constants. The overall signal relaxation (decay) is the combination of those two effects – a double-exponential decay.

Free Induction Decay (FID)

The simplest form of MR signal detection is called free induction decay (FID). This is the detection of oscillating EMF in an RF pick-up coil, following the application of a simple RF pulse (usually 90°). The measured signal will oscillate with a frequency close to the Larmor frequency and will decay at a rate determined (dominated) by time constant T_2^* (described below).

1.5.2 Relaxation - equations

The quantitative picture of signal relaxation is as follows: A 90° RF pulse will cause nuclei to transition between energy levels and change spin states, which causes a reduction in the net magnetisation in the longitudinal plane, M_z. Nuclei experiencing an RF pulse will come into phase with each other, which causes a net rotating magnetisation in the transverse plane, M_{xy}. Once the RF pulse ceases, the system will return to an equilibrium state in which net magnetization is aligned with the longitudinal axis.

The recovery of longitudinal magnetisation ("spin-lattice") depends on the time constant, T_1. The decay of the transverse magnetisation ("spin-spin") depends on the time constant T_2 or T_2^*. The rate of change of the

magnetisation of each component (M_x, M_y, M_z) is described by the Bloch equations, which were empirical derived by Felix Bloch (Bloch 1946).

$$\frac{\delta M_z}{\delta t} = \frac{M_z - M_0}{T_1} \tag{1.4}$$

$$\frac{\delta M_x}{\delta t} = \frac{-M_x}{T_2} + \gamma M_y B_z \tag{1.5}$$

$$\frac{\delta M_y}{\delta t} = \frac{-M_y}{T_2} - \gamma M_x B_z \tag{1.6}$$

Longitudinal relaxation

The time constant of M_z recovery is called T_1. Once RF pulse excitation ceases, longitudinal (M_z) and transverse (M_{xy}) will return to Boltzmann distribution. The rates at which M_z and M_{xy} return to equilibrium are dependent on different factors. Recovery of longitudinal net magnetisation (M_z) is due to the Brownian motion of molecules and the exchange of spin energy between their nuclei. The rate of this recovery depends on the molecular environment the nuclei are in, with reduced freedom of molecular movement reducing recovery time.

Measuring the T_1 time constant can be done "indirectly" using a 180° RF pulse that inverts the longitudinal magnetisation. After this initial "inversion pulse" longitudinal magnetisation will recover, with a time constant T_1. After a particular (chosen) time, called TI (inversion time), a 90° RF pulse is applied, flipping the longitudinal magnetisation at that point in time into the xy-plane where the net rotating magnetisation can be detected by the RF coil. The signal measured is therefore proportional to the longitudinal magnetisation at the time of the 90° RF pulse application and hence measures T_1.relaxation. By taking measurements at different values of TI (at a repetition time, TR, much greater than T_1 ($5T_1$)), the time course of T_1 relaxation can be measured.

Many anatomical imaging methods rely on signals related to T_1. Quantitative T_1 measurements have been used to determine the relative myelin content of cortical areas to determine human auditory core regions versus belt homologues (Dick et al. 2012).

Transverse relaxation

Applying an RF pulse brings all nuclei spins into phase, resulting in a net magnetization M_{xy} rotating at the Larmor frequency. The decay of transverse net magnetisation (M_{xy}) occurs due to a loss of coherence (dephasing) of spins until their phases are uniformly distributed in the xy-plane. This occurs with a time constant, T_2 or T_2^*. Dephasing of the nuclei spins arises through several mechanisms:

1. Natural longitudinal recovery causes spins to flip, which in turn causes phase incoherence. This means that processes that determine T_1 also contribute to T_2 processes and T_2 decay can never be longer than T_1 decay.

2. T_2 - Time-varying processes that change the local magnetic fields of spins. This can be caused by slow molecular motion or magnetic shield of protons from electrons within molecules. The degree of shielding depends on the molecular environment, resulting in different T_2 time constants for protons in different materials and tissues.

3. T_2^1 - Static processes that cause differences in local magnetic fields of spins but do not vary in time. These can arise due to inhomogeneities in the static magnetic field B_0 across the sample. This is termed imperfect shimming. Changes in the magnetic susceptibility of a sample, for instance between water and air at tissue boundaries, will also cause a local magnetic field gradient.

The actual transverse decay is a combination of T_2 and T_2^1 decay processes, referred to as T_2^*. Individual nuclei experience different magnetic fields so the Larmor frequency will vary across the sample.

The time course of T_2 decay can be measured using a spin echo (SE) sequence. A spin echo sequence first uses a 90° RF pulse to flip the net magnetisation M_0 into the transverse plane. This transverse magnetisation is coherent (the spins of nuclei in the population are in phase), but then decays with time constant T_2^* as soon as the RF pulse ceases. Net magnetization is then reversed, at time TE/2 after the 90° RF pulse has ceased by applying a 180° RF pulse, called a refocusing pulse. Because T_2^1 processes are static, spins precessing at different speeds due to these processes will come back into phase with each other at time TE, i.e. TE/2 after the refocusing pulse. Variation in spin precession rate due to T_2 processes, on the other hand, is not static and therefore will continue to change with time despite the refocusing pulse. This means that the signal measured at time TE will exclusively reflect the decay due to T_2 processes, and not T_2^1. The time course of T_2 can be measured by collecting data at various echo times (TE).

1.5.3 Magnetic Resonance Imaging

MR Imaging

The previous section explained the basis of the "bulk" signal from a sample in an NMR measurement. To obtain actual images of the sample, the spatial location of contributions to this bulk NMR signal must be resolved. In this section, I will explain how *images* are formed by tagging different parts of the measured signal with different electromagnetic frequencies according to their location in space.

Resolving space is normally achieved through a combination of a) slice selection and b) phase and frequency encoding. Slice selection works by restricting the locations of where the RF pulse is effectively applied to a 2D plane (a slice). Phase and frequency encoding works by tagging parts of the measured signal according to where in that plane they came from. In both cases, magnetic resonance imaging (MRI) encodes the spatial location of nuclei within the bore by varying the magnetic field strength spatially, and therefore the Larmor frequency.

Slice Selection

A "slice" of the sample can be selected by varying the magnetic field strength experienced by the sample along a given spatial direction (usually the longitudinal z-axis) during the excitation pulse. By adding a gradient of magnetic field along the z-axis, B_0 can be changed to vary smoothly from $B_0 - \delta B_0$ to $B_0 + \delta B_0$. This, in turn, makes the Larmor frequency spatially dependent, such that protons at one end of this smoothly varying magnetic field will precess at $f - \delta f$ at one end and $f + \delta f$ at the other. Nuclei can then be excited at a given location in the sample by applying an RF pulse with a narrow band of frequencies δw corresponding to the selected location. Only nuclei with Larmor frequencies within δw, and therefore in that spatial location will be excited and any subsequently measured signal will, therefore, be generated by nuclei within that spatial location of the sample.

Frequency and phase encoding

Selecting a slice of the sample and limiting spin excitation to it is only the first part in generating an image in MRI. To spatially resolve the signals within the slice, frequency and phase encoding are normally used.

The general approach here is often referred to as spin-warp imaging.

A spin-warp imaging sequence uses the frequency and phase of spins to encode spatial location. The MR signal measured from the selected slice is generated from all spins excited by the RF pulse in that slice. To encode the location of nuclei within the slice the precession frequency (and phase) can be varied over the slice. To achieve this, magnetic field gradients are applied across the slice during the acquisition of the MR signal.

Following the selective excitation of the slice, a different – frequency encoding – gradient is applied to vary the magnetic field in a given direction (e.g. along the x-axis) across the slice. In the presence of this gradient, the Larmor frequency will vary across the slice, making spin precession frequency-dependent (and therefore location dependent) along this axis of the slice. Note that frequency encoding only allows for resolving one dimension across the slice.

To also determine the location in the second dimension across the slice (e.g. along the y-axis), phase encoding is used: an additional phase encoding gradient is applied orthogonally to the frequency encoding gradient, but only for short durations between applications of the longer frequency encoding gradient. The brief application of the phase encoding gradient causes the phases of spin precession within the slice to vary as a function of location along that dimension: When this phase-encoding gradient is then turned off and the frequency encoding gradient is applied, spins along the phase-encoding direction will be at different phases, due to previously having different precession frequencies, while spins along the frequency-encoding direction will precess at different frequencies. Therefore, spins at any x,y coordinate within the slice will vary in both frequency and phase and the frequency and phase components of the

measured MR signal will correspond to different x,y locations of nuclei spins. Images are constructed from the measured signal using 2D Fourier transform to take the signal from the time domain to the frequency and phase domain.

MRI contrasts and Pulse Sequences

The contrast of an MR image is determined by the contribution of T_1 and T_2 processes to the measured signal. The contributions of T_1 and T_2 processes, in turn, depend on the TR and TE of the pulse sequence. T_1 and T_2 time constants depend on the molecular environment of nuclei and therefore are different in different tissues of the human body. This can be utilised to create images of different parts of the human body. The difference in T_1 and T_2 time constants between tissues allows sequences to optimised to have the largest differences in the signal generated by different tissues.

Proton-density contrast

An image generated with a proton-density contrast reflects the density of protons at each voxel, MR images are comprised of Voxels, which can be thought of as three-dimensional pixels. The voxel is an individual data point within the image. In such an image, the intensity at each voxel is determined by the number of protons within that voxel. To obtain proton-density contrasts, the pulse sequence must minimise T_1 and T_2 contrasts between tissues. T_1 contrast can be minimised by using either a very long or very short TR, such that longitudinal magnetization will have recovered equally for any type of tissue. T_2 contrast can be minimised by using either a very long or very short TE such that protons spins in all tissue are either equally in phase or equally out of phase. In

practice, a long TR and a short TE are used to maximise the MR signal recorded.

Proton density contributes to the intensity of all images generated in MRI because protons are the source of the MR signal.

T_1 contrast

Images are referred to a T_1 weighted when the intensity of a voxel mainly depends on the T_1 value of the tissue. The TR of the sequence is optimised to have the greatest difference in signal between tissues due to T_1 processes. This requires an intermediate TR, approximately equal to the T_1 of the tissue for the largest contrast is desired. To minimise the effect of T_2 decay, a very short TE is used and so the image is acquired when very little T_2 decay has happened (protons spins are still equally in phase in all types of tissues.

T_2 contrast

Images are referred to a T_2 weighted when the intensity of a voxel depends mainly on the T_2 value of the tissue. An intermediate TE is used so that some signal is lost due to T_2 decay processes, but not all of it. The amount of signal loss due to T_2 processes depends on the type of tissue and therefore the intensity of the signal will reflect the susceptibility of different tissues to these processes. A long TR is used so that the longitudinal magnetisation is almost completely recovered and the effects of T_1 processes are minimal.

T_2^* contrast

Blood-oxygenation-level-dependent (BOLD) fMRI imaging relies on T_2^* contrast, which is closely related to T_2. As discussed above, transverse relaxation depends on T_2 process, which reflects interactions between spins, and T_2^1 processes that are due to inhomogeneities the magnetic field. The combination of these processes results in a signal decay characterized by a time constant T_2^* that is shorter than T_2. T_2^*-weighted images are the foundation of functional MRI, the details of which are discussed in the next section. A pulse sequence optimised for T_2^* contrast will have a long TR and medium TE, similarly to T_2-weighted images. The difference is in how the signal echo is generated. Instead of using a refocusing RF pulse that rephases spins and removes T_2^1 effects. Instead, magnetic field gradients are used to generate the signal echo. This is known as a gradient echo. Following the application of frequency encoding gradient for a time, T, the inverse of this gradient is applied for twice the duration. The gradient rephases the spins and generates an echo signal.

An echo-planar imaging (EPI) sequence samples the entire slice in one FID. This dramatically reduces the time taken to collect data from a given area of the body. EPI sequences allow dynamic processes to be measured, for example, BOLD responses to stimuli as indirect measures of brain activity.

1.5.4 functional MRI

Functional magnetic resonance imaging allows for the indirect imaging of brain activity and the measurement neuronal responses to experimental stimuli.

Physiology of blood-oxygenation-level-dependent (BOLD) mechanism

Brain activity requires energy. Because no energy is stored within the brain, neurons are permanently supplied with glucose and oxygen via the vascular system to meet the metabolic demands of their activity. Neuronal activity is thought to be localised within the brain based on function, with different areas of the brain specialising in different tasks. Neuronal activity requires oxygen metabolism and therefore reduces in the amount of oxygen in the blood supplying the active region of the brain. In response, the blood flow to this area is increased to deliver oxygen rapidly. This delivery of oxygen is actually far greater than the amount of oxygen metabolised, resulting in a local increase of oxygen in active areas of the brain.

Physics of BOLD detection

Oxygen is transported and delivered through blood by haemoglobin. Oxygenated haemoglobin (Hb) has no unpaired electrons and zero magnetic moment, and is, therefore, diamagnetic, i.e. weakly repulsed from a magnetic field. Deoxygenated haemoglobin (dHb) has unpaired electrons and a magnetic moment, therefore it is paramagnetic, i.e. attracted to a magnetic field (although more weakly than ferromagnetic substances). As a result, deoxygenated blood is about 20% more magnetically susceptible, i.e. more magnetised when in a magnetic field, than oxygenated blood. Different brain regions, depending on whether they are active or not, will therefore show different levels of magnetisation, depending on the oxygen levels within the local blood supply. When placed in the static magnetic field of the MRI scanner, they will perturb it differently: more active (i.e. magnetized) regions will introduce more inhomogeneities in the mag-

netic field. T_2^* decay of a signal measured from the brain is sensitive to these inhomogeneities in the magnetic field because inhomogeneities increase T_2^1 processes and therefore shorten the T_2^* time constant. T_2^* contrast will therefore be sensitive to the level of blood oxygenation (more blood oxygenation results in less magnetic inhomogeneities and therefore more MR signal when using T_2^*-weighted sequences). This is called blood-oxygenation-level-dependent (BOLD) contrast and is the basis of functional MRI. BOLD contrast gives researchers a powerful endogenous contrast with which they can measure the amount of brain activity in localised regions of the brain. BOLD contrast is an indirect measure of neuronal activity, that is, it measures the changes in blood/ oxygen supply in response to brain activity (relative to baseline brain activity). How good an estimate of neuronal activity this is, is still open to debate, the details of which are beyond the scope of this thesis.

When a region of the brain responds to stimulation or is involved in any task, the BOLD contrast changes in time in a stereotypical way. The function that describes the time course of this BOLD response is called the hemodynamic response function (HRF).

Hemodynamic response function (HRF)

The HRF comprises three key components. The 'deoxygenation dip', a transient decrease in the signal from the baseline. This results from an increase of deoxygenated haemoglobin following oxygen consumption in neurons, which increases T_2^* decay effects due to magnetic field inhomogeneities. This dip is small, short-lived and difficult to observe. A positive peak occurring with a latency referred to as the hemodynamic delay. This positive peak results from the excess of oxygen brought by oxygenated blood in response to neural activity. This increases oxygenated

haemoglobin, pushing away de-oxygenated haemoglobin, which decreases magnetic inhomogeneities, reduces T_2^* decay effects and increases the MR signal. A prolonged post-stimulation undershoot. The physiological mechanism underlying this phenomenon is still debated and beyond the scope of this thesis.

1.5.5 Why go to higher field?

As discussed earlier, the net magnetisation of a population of nuclei with spin is proportional to the strength of the static magnetic field, B_0 due to increased energy difference between Zeeman levels. This net magnetisation is the source of the raw signal from an MRI scanner, and therefore by increasing the main magnetic field strength, the raw signal-to-noise (SNR) ratio increases. Turner et al. (1993) demonstrated, using 1.5T and 4T MRI systems, that this increase in raw SNR at higher field strength translates to an increase in percent BOLD signal change relative to baseline, increasing the functional SNR. This improved functional SNR allows smaller voxel sizes to be used, improving spatial resolution and, as a result, reducing partial volume effects (multiple tissue types present in one voxel). The improved SNR also increases the number of voxels found to be significantly active in response to stimulation, for the same statistical threshold and voxel size.

Increasing field strength also improves the spatial specificity of BOLD activity. To improve the localisation of brain activity, the signal measured wants to be dominated by BOLD activity from sources as close the neurons as possible, and with minimal contributions from the vascular sources. BOLD activity measured at a given voxel has intravascular and extravascular, components, both form large and small vessels. How much each of these components contributes to the measured BOLD activity

changes with field strength. At higher field strengths, a greater proportion of the signal comes from the extravascular component, and the contribution from small vessels also increases. This means that more of the signal measured in a given voxel comes from extravascular sources and small vessels within grey matter, and less from larger vessels.

Imaging at higher field strengths is not without difficulties. Signal drop out from areas of the brain near air pockets can occur do to changes in magnetic susceptibility.

1.5.6 fMRI data analysis

The goal of neuroimaging is to determine the functional properties of localised parts of the brain. Presenting sensory stimuli and measuring brain activity using fMRI is a common method employed when investigating sensory cortices. The time course at each sampled location is the measured hemodynamic activity corresponding to the aggregate activity from a population of neurons, plus unknown sources of signal. We construct and test statistical models to estimate the underlying properties of the brain from this indirect measure.

fMRI studies often employ general linear modelling (GLM) to estimate the magnitude of brain activity in response to different experimental stimuli. GLM modelling assumes a measured BOLD fMRI time course can be predicted from the time of each stimulus presentation. Stimulus presentations are treated as events that occur independently of each other, for which the time of is known but the amount brain responds, and therefore its contribution to the measured time course, is unknown. The analysis determines the extent that each stimuli contributes to measured fMRI time course by predicting a time course that minimises the

unexplained signal, residual error. fMRI statistical analysis techniques and GLMs are discussed in further detail in Chapter 2 - General Methods.

1.5.7 fMRI of the auditory system

Investigating the auditory system using fMRI poses unique challenges. Applying current to the gradient coils to generate the magnetic field gradients produces extremely loud acoustic noise (+110 dB SPL). MRI scanners use gradient coils to linearly vary the magnetic field strength across the sample in the scanner. The amplitude of the currents applied to the gradient currents determines the gradients of the magnetic field within the scanner. The electrical currents cause movement of these gradient coil wires in the magnetic field. This is the source of the loud acoustic noise when imaging with MRI, a factor which must be considered in auditory fMRI.

This noise must be considered in the design of an auditory fMRI experiment due to its interaction with the auditory stimuli presented for the experiment. Sparse imaging is a popular experimental method designed to minimise the interaction of acoustic scanner noise and presented auditory stimuli. In sparse imaging, data acquisition is done with a TR that is longer than the time necessary to acquire the full image. Between image acquisition, no data is acquired and therefore gradient coils are not used. Acoustic stimuli are presented between image acquisition without the presence of scanner noise, taking advantage of the hemodynamic delay to image the activity in response to an auditory stimulus, and allowing for the disentanglement of the BOLD response to the scanner noise and the stimuli of interest. The drawback of sparse acquisition is that a longer TR is needed to present the stimuli between data acquisition, re-

ducing the amount of data that can be collected within a given time and the sample resolution of the HRF. The effects of this scanner noise and sparse sampling are investigated in chapter 3 of this thesis.

1.6 Overview of studies

Frequency is a fundamental property of physical sounds and also a key organising factor of the auditory system that arises from neuronal encoding at the periphery of the auditory system. Within the inner ear, the cochlea performs a time-frequency decomposition of the acoustic stimuli, converting mechanical movement into electrical signals. Afferent auditory nerve fibres are distributed along the cochlear and tuned to different characteristic frequencies (CF), the frequency of which varies as a function of distance along the cochlear spiral (Békésy 1949). This cochleotopic organisation is maintained as tonotopic maps throughout the central auditory system up to the level of the auditory cortex. Such a topographic organisation of sensory characteristics into *maps* is common for sensory cortical areas. In the visual modality, **retino**topic maps have been widely reported and extensively studied (for a review see (Wandell, Dumoulin, and Brewer 2007)) and retinotopic mapping with functional magnetic resonance imaging (fMRI) has lead to a greater understanding of the organisation of the visual cortex, and the perceptual consequences of visual impairment and deficits (Baseler et al. 2011). It will be of great – ultimately clinical – benefit for the same to be done in auditory fMRI research. For example, the study of tonotopic reorganisation of the auditory cortex induced by peripheral hearing loss has only been demonstrated in animal models (Eggermont and Roberts 2004), but evidence from neuroimaging studies of humans remains inconclusive (Langers, Kleine, and Dijk 2012). Therefore, there is only indirect evidence of such

reorganisation through psychoacoustic studies (Thai-Van et al. 2003).

This thesis presents three studies that use fMRI to investigate the human auditory cortex. Their principal aim was to develop and evaluate experimental methods for the measurement of hearing loss induced cortical changes, such as tonotopic reorganisation.

The *General Introduction* chapter first describes the current state of auditory fMRI research and the challenges associated with accurately estimating tonotopic properties of the human auditory cortex. Next, that chapter explains the physics underlying magnetic resonance imaging (MRI) and limitations of auditory fMRI due to the loud acoustic noise generated by image acquisition.

The first set of experiments (chapter 3) aims to improve the estimation of *preferred* frequency, the frequency that elicits the greatest response, and frequency *selectivity*, the range of frequencies that elicit a response above baseline from normal hearing subjects. First, the impact of fMRI data acquisition protocols is assessed, focusing on the trade-off between presenting auditory stimuli in silence and the required reduction in data sampling points needed to achieve this. Second, the impact of statistical methods of estimating preferred frequency and selectivity from BOLD fMRI cortical responses to acoustic auditory stimuli is assessed and compared. From these results, recommendations are made for making the most accurate estimates of tonotopic properties of the human auditory cortex.

The second set of experiments (chapter 4) implements the methods evaluated in the first experiment chapter to estimate the cortical tonotopic magnification in the human auditory cortex. *Tonotopic magnification* refers to the cortical distance spanned by a given frequency, and the rela-

tive over-representation of some frequencies compared to others. Changes in cortical tonotopic magnification in patients with hearing loss may provide a simple measure for detecting reorganisation in the human auditory cortex.

The third set of experiments (chapter 5) presents a method to improve the estimation of tonotopic properties in subjects with hearing loss. This is addressed by two methods. The first method implements hearing loss simulation and uses it to modify the stimuli presented to healthy participants during the acquisition of auditory fMRI data. The second method aims to correct for hearing loss induced biases during data analysis, at the point of estimation of tonotopic properties.

The following chapter (2) provides a detailed description of the methodology shared by each study. Each of the data chapters is presented with an introduction to background literature, an overview of specific methods implemented for that study, followed by the results and a critical discussion of the findings.

A final discussion chapter relates the findings of each of the experiments and suggests avenues for follow-up studies and future research.

Chapter 2

General Methods

2.1 General outline

Normal-hearing participants took part in the studies presented in the following chapters. Normal-hearing was self-reported, as assessed by a questionnaire that assesses hearing ability in the context of real-life situations. All of the volunteer participants gave written consent before undertaking the study and reported no history of tinnitus, neurological disorders or hearing difficulties. The study was approved by the Ethics Committee of the University Of Nottingham Medical School. The studies presented in this thesis shared the following common methods and follow a group-of-case studies approach.

For all studies, the *preferred frequency*, the frequency that elicits the greatest response, and frequency *selectivity*, the range of frequencies that elicit a response above baseline was estimated at each voxel. Both were estimated using two different data analysis strategies: using the framework of the generalized linear model and another using population receptive field modelling.

2.1.1 Study 1 - Tonotopic mapping of the human auditory cortex: influence of acquisition protocol and data analysis

Summary:

- Measured BOLD responses to tonotopic stimuli using Sparse and Continuous acquisition protocols
- Compared statistical methods of estimating frequency preference and selectivity

The first study presented in this thesis aims to determine the most reliable way of measuring preferred frequency and frequency selectivity in the human auditory cortex using fMRI. The effect of the acquisition protocol and the analysis method on the measurement of these properties is assessed. Preferred frequency and frequency selectivity was estimated using general linear modelling (GLM) and population receptive field modelling (pRF). These estimates were compared and assessed. The effect of scanner noise and the number of data points was assessed by comparing the estimates from data collected using each acquisition protocol.

The reliability of each method is estimated by the repeatability of estimates between scans. Possible biases in the estimates of frequency preference are assessed through the distribution of voxels estimated frequency preference. All datasets were collected in a single session, apart from one subject, where additional functional and auxiliary datasets were collected in a second session. BOLD activity was measured in response to NBN stimuli in 8 normal-hearing subjects (4 males; mean age = 29.38 SD = 7.34) using sparse (TR = 7.5) and continuous (TR = 1.5) acquisition

protocols on the Philips Achieva 7T MRI system using a head volume transmit coil and a 32-channel SENSE receiver coil.

2.1.2 Study 2 - Cortical magnification of the human auditory cortex

Summary:

- Measured BOLD responses tonotopic stimuli
- Estimated voxel preferred frequency
- Estimated gradient reversals around Heschl's gyrus - the primary auditory cortex
- Calculated cortical distance from each reversal at each point of estimated of preferred frequency
- Characterised relationship between the cortical distance between reversals and preferred frequency

The second study aims to measure and describe tonotopic cortical magnification in normal-hearing subjects. We estimate the cortical magnification function of 16 hemispheres in 8 normal-hearing subjects (Estimated from the data collected in study 1 using the sparse acquisition protocol). Preferred frequency was estimated voxel-wise using a pRF method. Cortical magnification was estimated in the anterior and posterior tonotopic gradients on Heschl's gyrus in each hemisphere by characterizing the relationship between preferred frequency at each point of the cortical surface and the shortest cortical distance between this point and the low-frequency reversal.

2.1.3 Study 3 - How can we reliably detect hearing loss induced cortical tonotopic reorganisation in humans?

Summary:

- Measured BOLD responses tonotopic stimuli in the presence and absence of hearing loss simulation
- Estimated tonotopic properties of the cortex for both conditions
- Quantified the bias resulting from hearing loss simulation on tonotopic estimates
- Implemented data analysis methods to reduce this bias

The final study presented in this thesis aims to develop methods for measuring tonotopic reorganisation in subjects with hearing loss. BOLD activity was measured in 7 subjects (5 male; Mean age = 28.71, STE = 7.95) in response to NBN stimuli in the presence and absence of high-frequency hearing loss simulation, implemented using threshold elevating noise, using a sparse (TR = 7.5) acquisition protocol. The effects of hearing loss simulation were quantified by comparing estimates of preferred frequency and frequency selectivity with and without hearing loss simulation. In an attempt to minimise these effects by implementing a modified pRF modelling method that factors hearing threshold in its fitting procedure. Scanning was performed on a Philips Ingenia 3T MRI system with 32-channel SENSE receive coil. All datasets were collected within a single session.

2.2 Data acquisition

2.2.1 Structural data acquisition

A whole-head T1-weighted volume was collected for each subject at a resolution of 1 mm isotropic.

Philips Achieva 7T Parameters (Studies 1 and 2): A 3D phase-sensitive inversion recovery (PSIR) sequence (Gowland and Leach 1991) (TE = 6.1 ms, TR = 15 ms, FA = 8°, shot-to-shot interval= $5000ms$, Field of View (FOV) = 200x180.95x139.8 mm (Anterior-Posterior (AP), Right-Left (RL), foot-head (FH)) covering the whole head was used for structural data acquisition (acquisition time = 9 min). A PSIR sequence consists of two turbo-field echo (TFE) trains acquired at different inversion times (TI) after a single inversion pulse: one earlier (T_1-weighted) and one later (proton density (PD)-weighted). The first TFE readout is acquired at TI =742 ms, at the approximate time point when grey and white matter will generate signals of equal but opposite amplitude, maximising T_1weighting. The second readout is acquired at TI = 2685 ms when the contrast between tissues is minimal and the signal is mainly weighted by proton density. The PSIR volume is obtained by voxelwise division of the T_1-weighted volume by the PD-weighted volume, which both reduces PD-weighting and corrects inhomogeneities in the inversion pulse field strength, B_1, because the $T_1$and PD-weighted volumes are affected identically by B_1 inhomogeneities. The magnitude of the resulting PSIR volume is proportional to the T_1 of the underlying tissues.

Philips Ingenia 3T Parameters (study 3): A 3D magnetization-prepared rapid gradient echo (MPRAGE) sequence (TE = 3.8 ms, TR = 8.2 ms, FA = 8°, Field of View= 200 × 180.95 × 139.8 mm (AP × RL × FH))

covering the whole head) was used for structural data acquisition (acquisition time = 9 min).

2.2.2 Functional data acquisition

Functional data was obtained using a 2D T_2^*-weighted, single-shot gradient echo (GE), echo-planar imaging (EPI) sequence (1.5 mm isotropic resolution; 24 slices, phase encoding direction: AP) with a field-of-view covering, and oriented parallel to the Sylvian fissure and an outer-volume-suppression slab placed anterior to the volume. Two dummy scans, where no data was acquired, were performed at the start of each run to minimise the effects of start-up magnetization transients.

4 runs of functional data were collected, each run taking around 8 minutes to complete. Two runs of 73 volumes using a sparse acquisition protocol (TE = 25 ms, TR = 7.5 s, FA = 90°) and 2 runs of 361 volumes using a continuous acquisition protocol ($TE = 25ms, TR = 1.5s, FA = 65^{\circ}$). For the continuous protocol, slices were acquired at equal intervals uniformly filling the entire TR to reduce the salience of the scanner noise. For the sparse protocol, slice acquisition was clustered at the start of the TR, with each volume taking 1.25 seconds to acquire, leaving a 6.25 seconds silent period between acquisitions, in which acoustic stimuli were presented. Parameters were optimised following a pilot study (not presented in this thesis).

Philips Ingenia 3T Parameters: Parameters at 3T as the sparse protocol at 7T except for TE = 40 ms, FA = 90°, TR = 7.5s.

2.2.3 Auxiliary data acquisition

A high resolution structural T2*-weighted volume (3D FLASH sequence, $0.5x0.5x1.5mm$ resolution), with the same FOV as the function data was acquired to aid spatial alignment between the functional and structural data (7T Parameters: TE = 9.9 ms, TR = 370 ms, FA = 32°; 3T Parameters: TE = 20 ms, TR= 30 ms, FA = 12°).

Time-varying inhomogeneities in the main magnetic field, B_0, arise due to the motion and respiration of the subject and result in distortion of the functional EPI images. To correct for this, we acquired addition functional volumes with no stimulus presentation, one with the same echo time as the functional data sets (TE = 25 ms), and one with a slightly different echo time (TE = 28 ms). The phase and magnitude information from these scans, as well as the phase information of the functional data, was used to correct the magnitude image of the functional data sets (Besle et al. 2019). Distortion correction was only applied to data collected using the Philips Achieva 7T.

2.2.4 Optimising data acquisition

We optimised functional data acquisition protocols by iteratively changing scanning parameters to maximise the temporal signal-to-noise ratio of the functional data and minimise the acoustic scanner noise of the scanner. Acoustic scanner noise was measured using a Bruel & Kjaer (Naerum, Denmark) Type 4134 Pressure field microphone placed in the scanner on foam, held by the subject to their chest. The output from the microphone was fed into a G.R.A.S. Sound and Vibration (Holte, Denmark.) Type 26AM Preamplifier, then into a Sinus Apollo (Sinus Messtechnik GmbH, Leipzig, Germany) interface, and the signal was

recorded using the AcSoft SAMURAI software package at a sampling rate of 50 kHz. The temporal signal-to-noise ratio (standard deviation of ROI average time series divided by ROI average BOLD amplitude) of the functional data was calculated within a region-of-interest encompassing Heschl's gyrus, using a custom Matlab script created at the Sir Peter Mansfield Imaging Centre (SPMIC), University of Nottingham.

2.2.5 Auditory stimulus

Acoustic stimuli comprised narrow-band noises (NBNs) at 32 different centre frequencies, spaced between 0.1 and 8 kHz. NBN stimuli elicit greater neuronal activity compared to pure tones with minimal loss in frequency selectivity due to more energy being presented to the auditory system. Stimuli were spaced equally on a cochlear frequency scale. This ensured uniform excitation of the cochlea. The bandwidth of each NBN was equal to one normal cochlear filter bandwidth at its centre frequency, expressed as equivalent rectangular bandwidth (ERB). The cochlear ERB (in kHz) increases with frequency according to the following equation (calculated from behavioural measurements of cochlear filter resolution; Glasberg and Moore, 1990):

$$f_{ERB} = A(Bf_{kHz} + 1) \tag{2.1}$$

Where $A = 24.7/1000$ and $B = 4.37$.

NBN centre frequencies were spaced equally on this cochlear frequency scale, such that successive centre frequencies were separated by a fixed number of cochlear ERBs. To achieve this, the limits of the frequency range were converted from kHz to the number of cochlear filter bandwidths (nERB) by taking the integral of the ERB function from 0 to the

desired frequency in kHz [REF]:

$$f_{nERB} = \frac{1}{(AB)log(Bf_{kHz} + 1)} \tag{2.2}$$

Where $A = 24.7/1000$ and $B = 4.37$.

Stimulus centre frequencies were then calculated by linearly spacing thirty-two frequencies between the stimulus limit values in nERB. These nERB values were then converted to kHz, a more familiar unit, using the following formula:

$$f_{kHz} = \frac{1}{B(exp(ABf_{nERB}) - 1)} \tag{2.3}$$

Where $A = 24.7/1000$ and $B = 4.37$.

This achieves a high sampling resolution of the auditory system's frequency properties (within the stimulus limits: 0.1 kHz or 3.3589 nERB to 8 kHz or 33.1892 nERB, frequency range: roughly 30 nERB) by using 32 stimuli with tuning widths of one ERB, resulting in minimal overlap but effective coverage of the sampled frequency range. The frequency range was limited by the transducers used for the presentation of stimuli.

Stimuli were random sequences of 50-ms and 200-ms NBN bursts at a constant centre frequency with an inter-burst interval of 50 ms, resulting in a "Morse Code" like pattern of long and short narrowband noises (Da Costa et al. 2011; Thomas et al. 2015). Such random amplitude modulated stimuli increase the salience of the stimuli in comparison to the consistent background noise of the scanner and promote larger responses from non-primary regions in the auditory cortex (Overath et al. 2012; Alho et al. 2014).

The number of NBN bursts in a stimulus depended on the type of acquisition protocol (Sparse: TR = 7.5 s, stimulus duration 4.25s, number of 200-ms bursts = 16, number of 50 ms bursts = 10; Continuous: TR = 1.5 s, stimulus duration 1.25s, number of 200-ms bursts = 4, number of 50 ms bursts = 2). Individual bursts were gated with 10-ms quarter-cosine ramps to avoid acoustic artefacts due to transients. In the sparse protocol, stimuli were presented in the silent period between each functional volume acquisition (2s after the start of each acquisition). In the continuous protocol, stimuli were presented during the acquisition of a function volume (0.25s after the start of each acquisition). For each functional data acquisition run, thirty-two stimuli were presented in a pseudorandomized order and interleaved with silent trials (Sparse = 9 out 73 were silent trials, Continuous = 45 out 361 were silent trials), needed to estimate baseline BOLD activity. Maximum efficiency of the stimulus sequence, the ability to statistically detect a response to an individual stimulus, was achieved by generating 1000 random sequences and selecting the most efficient sequence using the following equation, as described by Buračas and Boynton (2002) (Buračas and Boynton 2002; Dale 1999):

$$E = \frac{1}{trace((X^T C_n^{-1} X)^{-1})} \tag{2.4}$$

Where X is the design matrix and C_n the covariance matrix.

To remain alert, subjects watched a silent, subtitled movie of their choice. To equalise the sensation and loudness level of the stimuli across centre frequencies and subjects, a continuous background broadband noise, filtered to contain equal energy per ERB, was presented throughout the acquisition period. The background noise level depended on the study.

Acoustic stimuli were created using Matlab (The MathWorks, Natick,

MA, USA) and presented binaurally using MRI compatible insert earphones (S14, Sensimetrics) via a TDT System 3 (Tucker Davis Technologies, Alachua, FL, USA) at a sample rate of 25 kHz. Digital signal processing and digital to analogue conversion was performed using a TDT RP2.1 real-time processor, controlled using Matlab via ActiveX and RPvdsEx. An inverse of the earphones transfer function (see the section: *Accurate acoustic stimulus presentation* below for how this was measured) was applied digitally to account for the non-linear frequency response of the transducers and to maximise the accuracy of stimulus presentation. Amplification and impedance matching between the digital to analogue converter and headphones was performed using a TDT HB7 headphone amplifier with stepped attenuation.

2.2.6 Paradigm

Study 1 and 2

Data sets were collected in four runs, with two runs for each acquisition protocol (sparse or continuous). Runs for each acquisition protocol were interleaved, with a sparse run always acquired first, and the same NBN stimuli were presented for all data acquisition. The duration of the Morse code burst depended on the acquisition protocol and the order was randomised for each run. Stimuli were presented at 80 dB SPL and a constant sensation level of 45 dB SL was achieved by presenting threshold elevating broadband noise at 35 dB SPL/ERB throughout the acquisition period. The second study presented in this thesis used the data collected from study one using the sparse acquisition protocol.

Study 3

Four runs of functional data were acquired using a sparse protocol (73 volume acquisitions per run) in two different conditions (two runs per condition). The same NBN stimuli were presented in both conditions but the order was randomised for each run. In the "normal hearing" condition, stimuli were presented at a constant sensation level (SL) of 50 dB, achieved by presenting stimuli at 75 dB SPL and background broadband noise at 25 dB SPL/ERB. This condition served as a within-subject control to the hearing loss simulation condition. In the "simulated hearing loss" condition, the background noise was designed to alter the sensation level of the stimuli as a function of frequency (See the method section of the third study chapter for more details). The final spectrum of the masking noise was determined by subtracting the critical ratio (CR) from the desired masking level. The CR represents the minimum signal-to-noise ratio at which a stimulus can be detected and varies as a function of frequency (Desloge et al. 2012; Hawkins and Stevens 1950). Runs in the "normal hearing" and "simulated hearing loss" conditions were alternated and their order was counterbalanced across subjects.

2.2.7 Accurate acoustic stimulus presentation

To account for the non-linear frequency response of the transducers and improve the accuracy of sound generation, auditory stimuli were multiplied in the frequency domain with the inversion of a measured acoustic transducer transfer function.

The frequency transfer function of the acoustic transducers, the filtering of the signal resulting from the non-linear conversion of the electric to acoustic signal by the acoustic transducer, was measured by presenting a

broadband white noise at 70 dB SPL to the headphones for 120 seconds using the same setup as the experimental procedure. The acoustic signal from the transducer was measured using a Bruel & Kjaer (Naerum, Denmark) Ear Simulator Type 4157 (complies to ISO 389-2:1997) for insert earphones. The output from the microphone fed into a G.R.A.S. Sound and Vibration (Holte, Denmark.) Type 26AM Preamplifier and a Sinus Apollo (Sinus Messtechnik GmbH, Leipzig, Germany) interface. The signal was recorded using the AcSoft SAMURAI software package at a sampling rate of 50 kHz. The frequency transfer function of the acoustic transducer was obtained by averaging fast Fourier transform (FFT) spectrums (computed using AcSoft SAMURAI) across all time samples.

2.3 Data Analysis

Data analysis was performed using Matlab (MATLAB 2015a, The MathWorks, Natick, 2015), and mrTools, a Matlab toolbox for the processing of fMRI data (http://gru.stanford.edu/mrTools; (Gardner et al. 2018a)). Pre-processing of data used additional tools from FSL (https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FS

2.3.1 Initial processing of functional data

Alignment, distortion correction and spatial smoothing

Prior to analysis, functional volumes were aligned to the structural whole-head T1-weighted volume. First, the high-resolution T_2^* FLASH volume was linearly aligned to the whole-head using mrTools' mrAlign (Nestares and Heeger 2000). The functional EPI volume acquired closest in time to the high-resolution T_2^* FLASH volume (reference EPI) was then aligned to the high-resolution T_2^* FLASH volume using a linear registration algorithm (FSL FLIRT). Motion correction (alignment of all functional

EPI volumes to the reference EPI volume) was performed using mrTools (Gardner et al. 2018a). Functional data collected using the Philips Ingenia 3T MRI scanner was then smoothed with a 3 mm FWHM 3D Gaussian using FSL's fslmaths. The voxel time series of each run were high-pass filtered to eliminate low-frequency signals (such as scanner drift) and converted to percent signal change by dividing all data points by the mean of the time series. Functional runs were then concatenated by condition (sparse vs continuous, "normal hearing" vs "simulated hearing loss") before statistical analysis.

Projection onto cortical surfaces and flat patches

Volumetric data was projected onto a flattened representation of the cortex surrounding auditory cortex on the superior temporal plane in each hemisphere, centred on Heschl's gyrus. Freesurfer v6 ((Dale 1999); http://surfer.nmr.mgh.harvard.edu) was used to create cortical surface meshes of the grey/white matter boundary (inner surface), and outer, pial surface of each hemisphere from the participant's structural whole-head volume. These surfaces were then interpolated to create an additional nine surfaces equally spaced at intermediate cortical depths, resulting in eleven surfaces for each hemisphere of each participant. Flattened representations of the cortical patch around auditory cortex at 11 cortical depths were created from these surfaces using the mrFlatMesh algorithm (Vista software, http://white.stanford.edu/software/). The resolution of the flatmaps was upsampled to a target pixel size of 0.33×0.33 mm^2 (actual resolution depended on distortions due to flattening and varied between hemispheres and participants). Voxel estimates from functional data sets were then projected onto the flat cortical representations using nearest-neighbour interpolation.

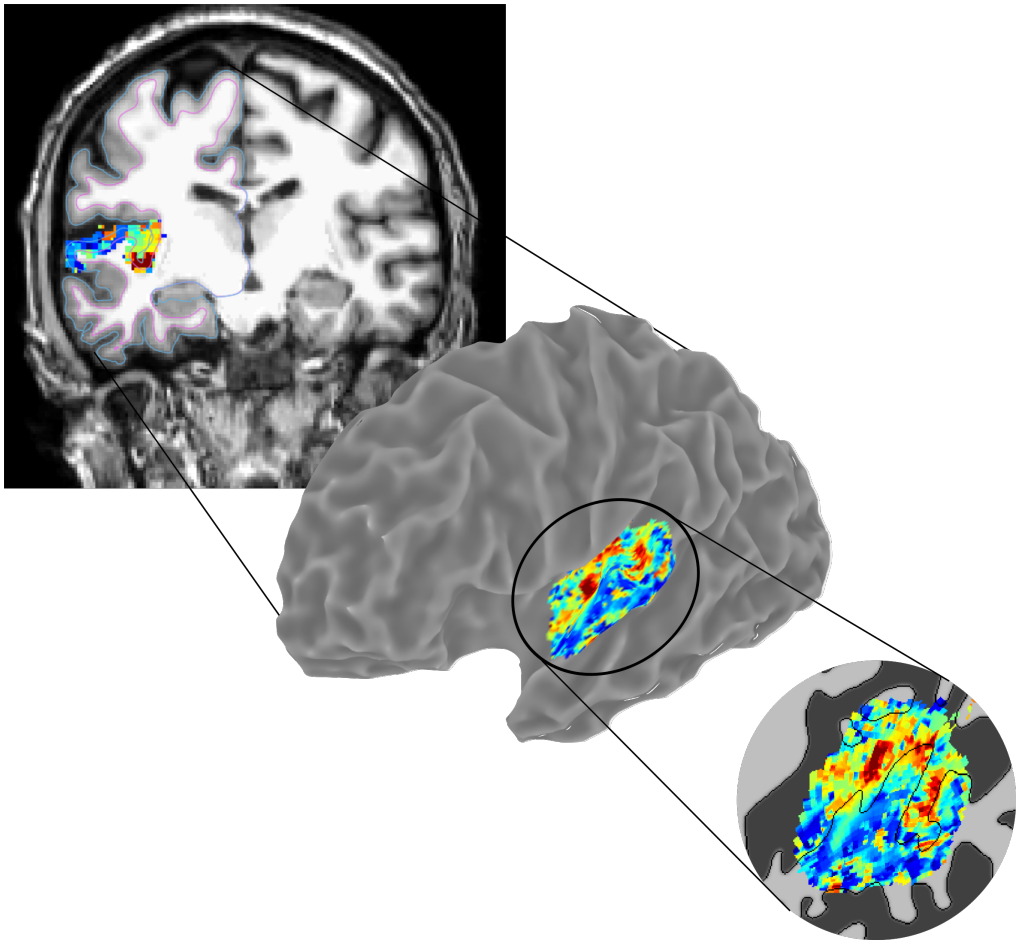


Figure 2.1: Volumetric to cortical surface to flat map

Functional data (multi-coloured pixels) shown in the original volumetric image, and on the derived cortical surface and cortical flat maps. Cortical surface meshes were created using Freesurfer from structural whole-head volumetric images. Flattened representations of the cortical patch around auditory cortex were created from these surfaces using the mrFlatMesh algorithm.

2.3.2 Estimating voxel frequency preference and selectivity

Estimating frequency preference and selectivity of individual voxels is the foundation of the functional data analysis used for the following studies. The preferred frequency of a voxel is defined as the frequency component eliciting the greatest level of brain activity. The frequency selectivity of a voxel determines the range of frequencies a voxel responds to.

Following preprocessing of the functional data, we estimated voxels' tuning properties from fMRI time series using two different models: general linear modelling and population receptive field modelling. Analysis of individual subject data was performed using custom-written Matlab scripts and mrTools (Gardner et al. 2018b). Data visualisation was performed using `R` (R Core Team 2013) and the tidyverse library (Wickham 2016).

General Linear Modelling

To estimate the brain activity in response to the acoustic stimuli, a separate general linear model (GLM) was fitted voxelwise to the functional data using mrTools GLM V2 plugin. Separate models were constructed and independently fitted both to individual functional runs and concatenated functional time series corresponding to each condition.

A GLM applied to fMRI data analysis assumes that the measured BOLD time course in a voxel can be predicted from the stimulus event matrix representing the time of presentation of different experimental stimuli. Known temporal sequences of different experimental stimuli, convolved with a model of the hemodynamic response function, are treated as independent variables whose respective contributions to the voxel's BOLD

time course are unknown. The GLM is fitted using a least-squares method in order to determine the weights (known as beta weights) of the linear combination of these variables (known as regressors) that best explains the observed time course.

$$y = \beta_0 + \beta_1 x_1 + \ldots + \beta_n x_n + \epsilon \tag{2.5}$$

where y is the measured BOLD time course, x_i are the regressor variables, β_i are the corresponding beta weights and ϵ is the residual error and signal that can't be explained that will be minimized in the least-squares fitting procedure.

It is commonly expressed in its matrix form:

$$Y = X\beta + \epsilon \tag{2.6}$$

where Y is a column vector of n recorded fMRI time-samples, β is a row vector of p estimates, X is the design matrix, i.e. a matrix of size n x p where each column represents a variable and ϵ is a column vector of n random values.

The coefficient of determination R^2 is a measure of the goodness of fit of the model to the data, indicating the proportion of variation in the measured time course that is explained by the GLM model.

$$R^2 = 1 - \frac{RSS}{TSS} \tag{2.7}$$

Where RRS is the residual sum of squares (sum of squared residuals) and TSS is the total sum of squares (sum of squared residuals to the averaged timecourse $\bar{y}$)

$$RSS = \sum \epsilon^2 = \sum (Y - X)^2 \tag{2.8}$$

Where Y is the measured time course and X is the modelled time course.

$$TSS = \sum (Y - \bar{Y})^2 \tag{2.9}$$

Where Y is the measured time course and $\bar{Y}$ is the mean of the time course.

Two design matrices were created for each functional dataset. The first assigned each stimulus to an individual regressor. This provided the highest resolution for estimates in terms of frequency. For the second design matrix, neighbouring stimulus frequencies were binned into groups of four adjacent frequencies, resulting in eight stimulus frequency bins, creating a design matrix with eight regressors corresponding to each of the possible stimulus frequency bins. This has the advantage of increasing statistical power, as more events occur within each regressor, but at the expense of frequency resolution.

A double gamma function was used to model the hemodynamic response function with parameters estimated from the data, as described below. The design matrix was normalized by setting the integral of each column in the design matrix, a column corresponding to a regressor, to one, and therefore estimated beta weights should be expressed in units of percent signal change. Setting the integral to one ensures that the scaling factor between the indicator function, following convolution with an HRF model, and the percent signal change time series doesn't change depending on the HRF model.

Estimation of preferred frequency and selectivity from voxels tuning curves

Voxelwise preferred frequency and frequency selectivity were estimated from individual voxels' frequency profiles, defined as the thirty-two regression beta weights calculated from the GLM as a function of frequency (expressed in ERBs). We used three different methods to estimate preferred frequency:

- Selecting the frequency of the beta weight which elicits the largest response
- Calculating the centroid of the frequency profiles
- Calculating a modified (bias-corrected) centroid to account for the limited stimulus frequency range, as described below.

Two methods were used to estimate frequency selectivity:

- Calculating the spread of the frequency profiles
- Calculating a modified (bias-corrected) spread to account for the limited stimulus frequency range, as described below.

The centroids of the voxel frequency profiles were computed as the mean of the 32 stimulation frequencies (expressed in nERBs), weighted by their corresponding beta estimates (in other words as the mean of the frequency profiles seen as a distribution histogram):

$$Centroid = \frac{\sum_i \beta_i f_i}{\sum_i \beta_i} \tag{2.10}$$

where f_i are the stimulus frequencies in nERBs and β_i are the corresponding response sizes. Negative β_i were replaced with 0. The spreads were computed as the standard deviation of the frequency profile seen as a distribution histogram

$$Spread = \sqrt{\frac{\sum_i (f_i - Centroid)^2}{\sum_i \beta_i}} \tag{2.11}$$

The centroid and spread of distribution within a finite range will be biased towards the centre of the range. The following procedure was implemented to correct for this bias (Besle et al. 2019). First, we defined a set of hypothetical Gaussian tuning curves with centre frequencies spanning the entire audible range of human hearing (0.01 - 15.05 kHz) (on the same frequency scale as the presented stimuli), and spread parameters ranging from 0.1 to 32 normal cochlear filter bandwidths (0.01 to 4.2 octaves at 1 kHz), for all combinations of centre frequencies and tuning widths. The hypothetical tuning curves were sampled at the thirty-two stimulus frequencies and the biased centroids and spreads of the sampled curve were calculated. We then determined, for each voxel, which of these hypothetical tuning curves (and associated biased centroid and spread) best matched the actually measured centroid and spread by minimizing the sum of squared differences between the hypothetical and measured biased centroids and spreads. The debiased centroid and spread were then taken to be the centre frequency and spread parameter of the best-matching Gaussian hypothetical curve.

Population Receptive Field Modelling

Voxelwise preferred frequencies (f_0) and tuning widths (σ) of voxels were obtained using population receptive field (pRF) modelling. Population receptive modelling models the frequency response profile of a voxel using

a Gaussian function. pRF models were fitted voxelwise to the functional data within a restricted ROI (described in a later section), implemented using mrTools and a modified version of the pRF plugin. A general linear model framework was used to fit the predicted time series to the measured BOLD time series, detailed below.

$$y(t) = p(t)\beta + e(t) \tag{2.12}$$

Where y is the recorded fMRI time series, p is the predicted time series, β is the scaling factor of the prediction and e is a noise term.

The predicted time series is constructed in the following way, illustrated in figure 2.2. First, we define the stimulus time series ($s(f,t)$). Here the simplification is made that the Morse code presentation of a stimulus within each TR is equivalent to a single event due to the low temporal resolution of BOLD fMRI. Events are indicated as binary values in time and ERB space in figure 2.2.

The population receptive field is modeled using a one-dimensional Gaussian function ($g(f)$) defined in nERB space, with the parameters of population centre frequency (f_0) and population tuning width (σ).

$$g(f, f_0, \sigma) = e^{-(\frac{f-f_0}{2\sigma^2})} \tag{2.13}$$

A predicted response is then created by calculating the linear summation of the stimulus time course, at each time point across the frequency dimension, and the Gaussian population receptive field model.

$$r(t, f_0, \sigma) = \sum_f s(f,t)g(f, f_0, \sigma) \tag{2.14}$$

This predicted response is then convolved with an HRF to give a predicted voxel time series. A gamma function was used to model the HRF using the fitted parameters derived from the GLM analysis.

$$p(t, f_0, \sigma) = r(t, f_0, \sigma) * h(t) \tag{2.15}$$

Preferred frequency and frequency selective of the population receptive field model are estimated using a non-linear fitting algorithm (Levenberg-marquardt, Matlab function lsqnonlin) minimizing the sum of squared residuals (RSS) between the prediction time course, $p(t)$, and the measured data, $y(t)$. A scaling factor, β, is used to account for the unknown amplitude of the MR signal.

$$RSS = \sum_t e(t)^2 = \sum_t (y(t) - p(t)\beta)^2 \tag{2.16}$$

The fitting algorithm then iteratively change the parameters of the pRF model until the best fitting model is found (until a minimum RSS value is found).

Initial starting parameters were determined by creating a range of predicted time series from possible pRF models, at 64 different centre frequencies and 16 different tuning widths between the fitting boundaries ($f = 0.026 - 16.934kHz$; $\sigma = 0.04 - 20ERB_n$), and then calculated which combination of parameters has the largest correlation (Pearson's linear correlation coefficient) with the measured fMRI time series.

The R^2 was calculated for each voxels fit as follows:

$$r^2 = 1 - \frac{\sum (e - \bar{e})^2}{\sum (y - \bar{y})^2} \tag{2.17}$$

Where $e = y(t) - p(t)\beta$.

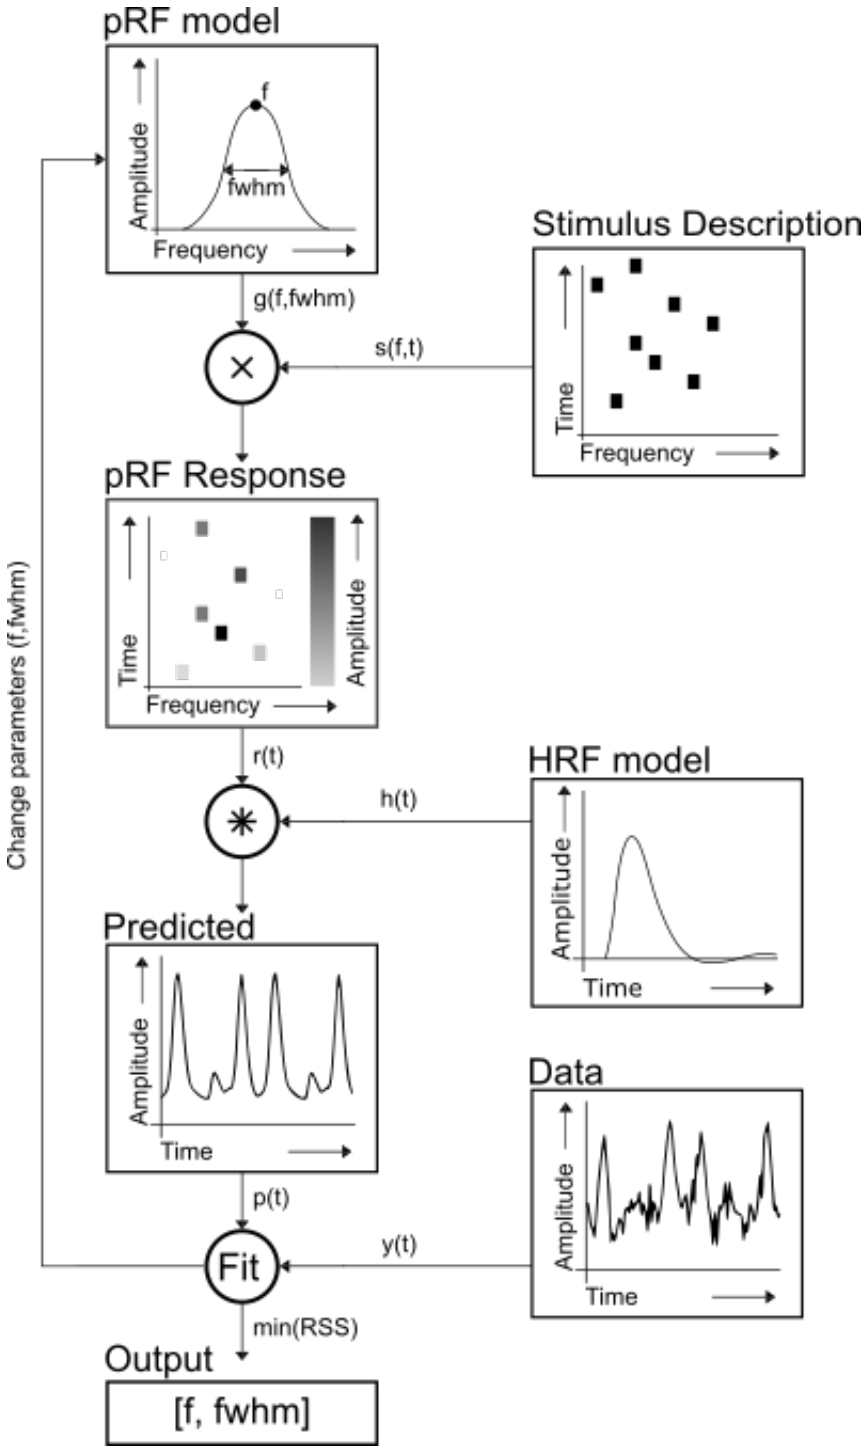


Figure 2.2: Population receptive field (pRF) modelling procedure

A flow chart describing the population receptive field (pRF) modelling procedure. A predicted BOLD response time course is fitted iteratively to the measured data by altering the parameters of a modelled auditory population receptive field.

2.3.3 Calculating gradient reversals

Low- and high-frequency gradient reversals were calculated from the GLM debiased centroid estimates of preferred frequency using an unbiased and automated procedure proposed by Schönwiesner, Dechent, et al. (2015). First, the local gradient of each voxel, between itself and its immediate neighbours along the flattened cortical depth, was calculated using a 2D Sobel operator (imgradient in Matlab). These gradient vectors were then averaged across all 11 cortical depths and smoothed with a 2D Gaussian kernel of 6mm (18-pixel) full-width-half-maximum (FWHM). Next, the relative angles between neighbouring local gradients were calculated by counting the number of gradient sign reversals detected along 180 orientations (equally spaced between 0 to 179°). Gradient reversals were identified by thresholding the resulting relative angle maps (10°; the exact value doesn't affect the location or shape of detected reversals).

2.3.4 Estimation of the hemodynamic response function

Both the conventional GLM and the pRF analyses require a design matrix to be convolved with a model of the HRF, the time course of BOLD signal in response to neuronal activity, assumed to be identical for all stimuli and at all voxels. How well this model described the true BOLD time course determines how well the model fits the data. To achieve the best fit possible, subject-specific HRF models were derived from the collected functional data.

HRF was estimated in individual subjects using the same data used to estimate tonotopic properties. Estimation of the HRF was done in the "Auditory Responsive region of interest (ROI)" (see the section "Defining

ROIs"). First, a general linear model was fitted voxelwise to the concatenated functional data of the continuous acquisition protocol in which different regressors represented different times after each stimulation. Continuous datasets were used for sufficient sampling resolution of the HRF in the time domain. Here, accuracy in estimating the frequency response properties of voxels isn't a concern and so bias from scanner noise is acceptable. Stimuli with similar centre frequencies were binned by four, as described previously, and the stimulus time sequences for each of the 8 bins were used to construct the design matrix. Ten regressors were assigned to each of the eight frequency bins, one regressor being an indicator function indicating at which volume acquisition a stimulus was presented, and the nine other being duplicates of this first regressor shifted ahead by 1 to 9 volume acquisitions. Fitting the GLM results, for each stimulus frequency bin, in an estimated HRF made of ten beta weights at different points in time starting from the presentation of a stimulus. Given that the TR was 1.5s, each estimated HRF was $1.5 \times 10 = 15$s long.

An estimate of the average HRF for all auditory responsive voxels was then derived by averaging HRF estimates across voxels of the "Auditory Responsive ROI" according to each voxel's frequency preference. Voxel beta weights were aligned by the frequency that elicited the largest beta weight at the time point 3.75 seconds, which had the largest average beta weight for all voxels. Following alignment of voxels, the mean across voxels for all time and frequency points was taken, see figure 2.4 for the group average HRF as a function of distance from preferred frequency. The average HRF in response to a voxels preferred frequency is assumed the average HRF in response to any auditory stimulus, see figure 2.3 for the group average of this function. We then fitted two different models of

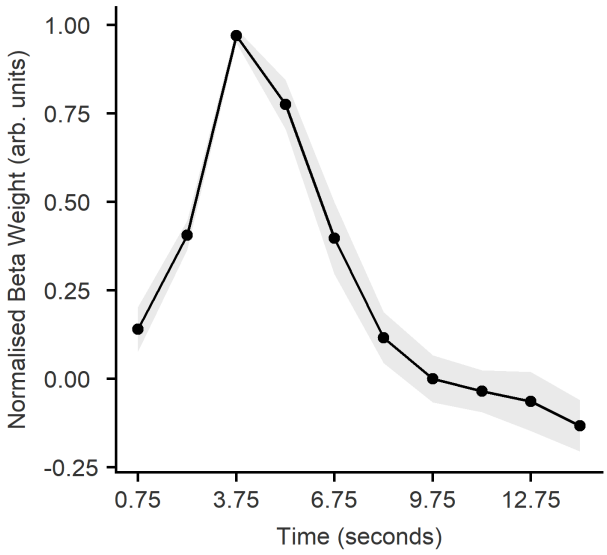


Figure 2.3: Average hemodynamic response function

The estimated HRF in response to voxels preferred frequency averaged across for all subjects. The average hemodynamic response function (HRF) of auditory responsive voxels from all subjects features a peak at a delay of around 4-5 seconds. The HRF was estimated using a voxel-wise GLM at 10 time points following stimulus presentation. The estimated voxelwise HRFs were averaged across all auditory responsive voxels following the alignment of frequency preference (Stimulus frequency that elicited the greatest response at time point 3.75 seconds, the time point with the largest average percentage signal change). The solid line is the mean and the ribbon is the standard error.

the hemodynamic response function to average HRF, the parameters of which were used in subsequent conventional GLM and pRF analyses.

Fitting procedure

Two functions that are commonly used to model the HRF, a gamma and double gamma, were fitted to the estimated HRF using a non-linear fitting procedure (Matlab function: fminsearch) (see figure 2.5). Each derived HRF was normalised by its maximum value across time points prior to fitting to avoid an additional scaling fitting parameter.

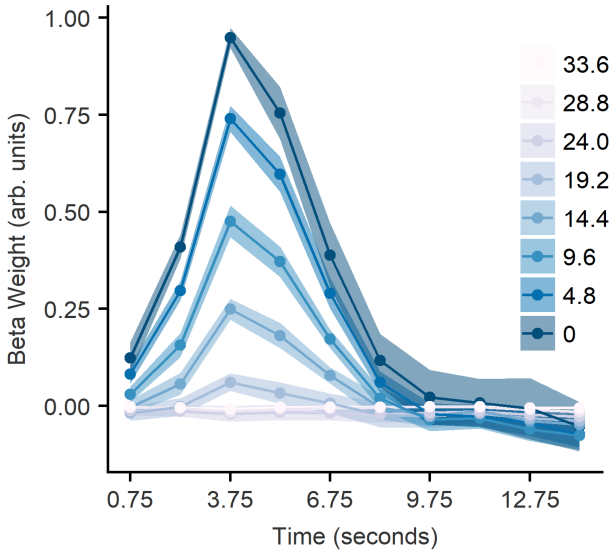


Figure 2.4: HRF to different centre frequencies as a function of distance from preferred frequency

The relationship between the HRF of a voxel and its preferred frequency. The maximum response is attenuated as a function of stimulus frequency distance from the preferred frequency of a voxel, but the general gamma function shape of the HRF is maintained until the population of neurons in the voxel no longer responsive to the presented stimulus frequency. Solid lines are the mean values, ribbons are the standard error, colour is the distance from the preferred frequency in nERB.

The gamma function models the initial peak and its delay but doesn't model the post-oxygenation dip.

$$HRF_g(x) = \Gamma(x, t, \tau, n) = (\frac{x-t}{\tau})^{n-1} e^{\frac{-\frac{x-t}{\tau}}{\tau(n!-1)}} \tag{2.18}$$

Where x = time points on which the function is defined, t = time lag of the gamma function peak, τ = the width at full half maximum (FWHM) of the HRF, and n = exponent parameter.

Initial fitting parameters were: time lag = 1, τ = 0.6, exponent = 4. Fitting boundaries were: time lag = [0 15], τ = [0 15], exponent = [1 15].

The double gamma function subtracts an additional gamma function from the simple gamma model to model the negative undershoot.

$$HRF_{dg} = \Gamma(x, t_p, \tau_p, n) - \frac{\Gamma(x, t_n, \tau_n, n)}{a} \tag{2.19}$$

Where t_p = positive peak time lag, τ_p = FWHM of positive peak, t_n = negative dip time lag, τ_n = FWHM of negative peak, and a represents the ratio of peak amplitudes of the two gamma functions.

Initial fitting parameters were: positive peak time lag = 6, FWHM of positive peak = 0.6, negative dip time lag = 8, FWHM of negative peak = 0.6, exponent = 1. Fitting boundaries were: positive peak time lag = [0 15], negative dip time lag= [0 15], ratio= [1 15],

The fitted parameters were used in subsequent HRF models for voxel-wise statistical analysis using both the GLM and pRF analyses. Parameters were fitted for individual subjects for data collected using the Philips Achieva 7T system. The group median of these fitted parameters (from the 7T datasets) were used for analysis of the data collected using the Philips Ingenia 3T system.

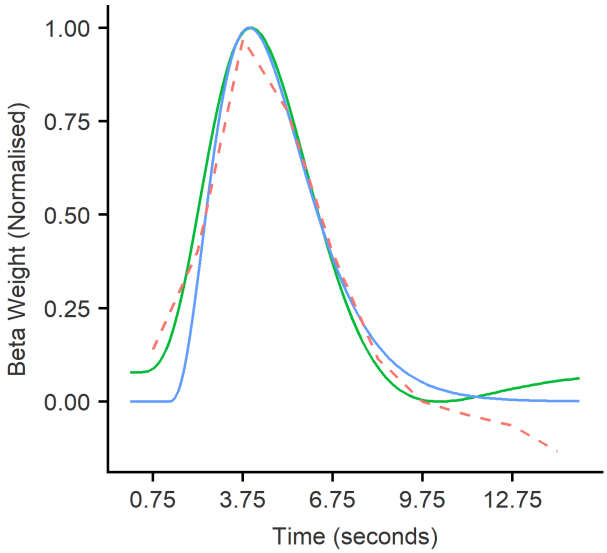


Figure 2.5: Group-average fitted HRF functions

Gamma (solid blue line) and double gamma (solid green line) functions were fitted to the estimated HRFs averaged across auditory responsive voxels in each subject and HRF models were calculated from the group median fitted parameters. The group average estimated HRF (presented in figure 2.4) is plotted in dashed red line for comparison.

Why model the HRF?

Following neuronal activity, the HRF takes at least 16 seconds to return to baseline. When stimuli are presented at repetition times of 1.5 and 7.5 seconds, HRFs to different stimuli overlap and the measured BOLD signal at each acquisition is the summation of responses to several stimuli. The HRF is used to deconvolve the neuronal response to each factor in the models created under the assumption that the activity measured at each time point is a linear summation of responses to the stimuli. The accuracy of the hemodynamic model function plays an important role in our ability to estimate the neuronal activity elicited by each stimulus frequency.

2.3.5 Defining Regions of Interest (ROI)

All ROIs were defined on flattened representations of the cortical sheet and restricted to the middle 5 cortical depths (0.3 to 0.7) to improve consistency of tonotopic maps (Ahveninen et al. 2016).

Auditory Responsive ROI

An "auditory responsive" region of interest (ROI) was created to include voxels in auditory cortex responding significantly to any of the NBN stimuli, presented at the same sensation level, for each subject. A GLM was fitted voxelwise to the concatenated functional data collected using a sparse acquisition protocol with the intention of minimising the effects of scanner noise. In this GLM, one regressor was assigned to each of the thirty-two stimulus frequencies, to generate a design matrix indicating whether a given frequency was presented at a given time point. The design matrix was then convolved with a boxcar HRF (a function that

takes value 1 a certain delay and for a certain duration after stimulation and 0 otherwise), simply delaying the onset of the function from the presentation of the stimulus and sustaining the duration of the response. Voxels which responded significantly (p-values < 0.05, adjusted for false discovery rate) to any of the thirty-two frequency stimuli were defined as auditory responsive. This was determined using an F-test on the set of the sound-verses-silence contrasts at each of thirty-two conditions. The "auditory responsive" ROI was defined as the set of auditory-responsive, cortical, voxels contiguous to auditory-responsive voxels located on Heschl's gyrus (HG).

Primary Auditory Cortex (AC) gradients ROI

In the human auditory cortex, the preferred frequency of a given location depends upon its location and changes in a continuous tonotopic gradient across the cortex. Two mirror-reversed tonotopic gradients have consistently been observed to run approximately perpendicular to HG and form an obtuse angle resulting in a V-shaped gradient reversal (Humphries, Liebenthal, and Binder 2010; Da Costa et al. 2011; Langers 2014). These gradients extend from the same lateral low-frequency region on HG towards two high-frequency regions, respectively located medio-anteriorly in Heschl's sulcus and posteriorly on the planum temporale. Reversals in tonotopic gradients, going from low-high to high-low, are thought to delimit functional areas of the auditory cortex based on non-human primate models (Kaas and Hackett 2000). Two mirror gradients are observed in each hemisphere measured in the presented studies. The gradient reversals of these tonotopic gradients are used to define the boundaries of the AC gradients ROI. According to the monkey model, frequency reversal only delimits a subset of the borders between areas. As a result, although

the two gradients probably correspond to different cortical areas, there may be more than one area within a given gradient, and only part of each gradient might belong to primary auditory cortex

Two observers independently defined the AC gradients ROIs on the flattened representation of the cortical sheet, taking into account gradient reversals, cortical morphology, unsmoothed tonotopic maps and statistical significance. Gradient reversals were calculated from estimates of voxel preferred frequency, as described in the "Estimating voxel frequency selectivity" section. The conventional GLM rather than the pRF analysis was used due to the computational time that would have been required to fit a pRF model to all cortical voxels.

The general procedure was to follow the two high-frequency gradient reversals, one anterior and one posterior to Heschl's gyrus, which runs medial-lateral along the gyrus, and the low-frequency gradient reversal that runs medial-lateral along the centre of Heschl's gyrus. These gradient reversals were then connected at medial and lateral ends using a straight line. ROIs were created by following continuous gradient reversal lines under the assumption of that gradient reversals follow cortical morphology and travel the shortest cortical distance possible. Continuous gradients which changed in direction were considered to have ended at the point of change following the assumption of one direction of reversal. Non-continuous or disconnected gradient reversals, which appeared to correspond to the same reversal, based on tonotopic maps and cortical morphology, were connected using straight lines. Gradient reversal maps were thresholded by F-test p-values (adjusted for false discovery rate (FDR)), which were cross-referenced to unsmoothed tonotopy maps. Three ROIs were defined per hemisphere. One containing both gradients, another containing the anterior gradient and one containing the posterior

gradient. In addition, the high and low gradient reversals were defined for estimates of cortical distance between low and high frequency reversals.

pRF Analysis Restriction ROI

Population receptive field analysis was limited to a region of interest (ROI) to save computational time. The auditory responsive ROI (described above) was convolved with a sphere with a diameter of 3 voxels in volumetric space. The ROI was then restricted to voxels within the cortical sheet, restricting the analysis to the grey matter of, and surrounding, Heschl's gyrus. For one subject of the 7T dataset, ROIs were further expanded to include HG.

2.3.6 ROI analysis

ROI analysis was restricted to voxels within gradient reversal ROIs from all subjects.

Averaging across cortical depths

ROI analysis was performed on voxel estimates averaged across the central 5 cortical depths. To average voxel estimates from the voxel-wise statistical analysis, voxel coordinates were converted from volumetric space to the space of the flattened cortical representation and then estimates were averaged over the middle 5 cortical depths.

Data Visualisation

Data from the ROIs of all subjects were included for data visualisation, which were created using the R(https://www.r-project.org/) package, ggplot2 (https://ggplot2.tidyverse.org/).

Distribution of preferred frequency

Presented preferred frequency distributions were calculated using a kernel density estimate (geom_density: default parameters), for each experimental condition. The area under each point of the curve represents the density of voxels with a given preferred frequency, creating a smoothed and continuous histogram of the voxel count as a function of frequency.

Average frequency response profile

The average frequency response profiles of all subjects were computed to assess differences between experimental conditions (effect of scanner noise in study one and simulated hearing loss in study three). To reduce noise, frequency response profiles were smoothed voxel-wise in the frequency dimension using a moving average window of 4 stimuli, resulting in 29 new data points from the 32 frequency regressors, for each voxel.

Average voxel frequency response profiles

Average voxel frequency response profiles were calculated from the beta weights estimated by the binned-frequency GLM analysis for each condition. A split-half averaging method was employed to minimise artificial sharpening due to noise. Two data sets were collected for each experimental condition, run A and run B from each subject. First, each voxel was categorised by the frequency that elicited the greatest response from run A, resulting in each voxel being labelled by one of a possible eight frequencies. The frequency profiles measured in run B were then averaged across voxels, according to voxels' frequency label. This was then repeated, with run B categorising the voxels frequency preference, and frequency profiles from run A were averaged. The mean of the resulting

two tuning curves was taken as frequency-specific average tuning curves.

Estimation reliability

Reliability was estimated by computing the Pearson's correlation coefficient of voxel preferred frequency estimates from all subjects between functional data acquisition runs.

$$r = \frac{\sum_i^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i^n (x_i - \bar{x})^2} \sqrt{\sum_i^n (y_i - \bar{y})^2}} \tag{2.20}$$

where: n is the sample size, x_i, y_i are the voxel estimates, $\bar{x} = \frac{1}{n} \sum_{i=1}^{n} x_i$ (the sample mean); and analogously for $\bar{y}$

Chapter 3

Tonotopic mapping of the human auditory cortex: influence of acquisition protocol and data analysis

3.1 Introduction

In order to understand the pathological changes in the auditory cortex as a result of peripheral hearing loss, it is crucial to have a comprehensive understanding of normal, non-pathological frequency representation in the auditory cortex. For this, effective methodologies need to be developed that facilitate accurate measurement of frequency properties of the human auditory cortex. This is not a trivial task and many factors must be considered. fMRI is an excellent brain imaging methodology due to its high spatial resolution and ability to non-invasively measure brain activity, allowing for the study of human populations. However, neuronal

brain activity can only be measured indirectly through blood-oxygen-dependent-level fMRI, and thus estimates of functional properties represent best estimates, rather than the true underlying properties of the brain. Potential methodological errors and measurement bias can arise at any point in the experimental procedure. The work in this chapter aimed to determine the most reliable way of measuring tonotopic properties in the human auditory cortex using fMRI.

There are two fundamental properties reflected in the responses measured from individual voxels that form the foundation of tonotopic and auditory fMRI research: (1) the preferred frequency, the frequency that elicits the greatest response, and (2) the frequency selectivity, the range of frequencies that elicit a response above baseline brain activity. Investigation of the auditory cortex and estimates of its properties are based on these parameters. Therefore, it is vital to understand possible shortcomings and limitations to the estimation method used to calculate them, and the validity of the estimated value.

There are two stages where the potential confounds in estimating preferred frequency and selectivity are greatest; data acquisition and the statistical method used to the estimate of frequency properties of a voxel. Estimation error can arise from each of these stages in the form of systematic bias and noise - the contribution of signals not related to the phenomenon that is intended to be measured - reducing the accuracy and reliability of estimates. Here, we compared estimates obtained using two popular schemes for auditory fMRI data acquisition and two analysis methods. The presented work highlights an approach for obtaining representative estimates of the true underlying neuronal properties of the human auditory cortex.

The choice of acquisition protocols should be carefully considered for tonotopic mapping experiments. Acoustic scanner noise is a potential confounding factor of auditory fMRI experiments (Hall et al. 1999; Scarff et al. 2004). The sound generated by the read-out gradients evokes activation in the auditory cortex during functional measurements. Sparse scanning protocols, where periods of scanner noise are minimised, are often employed to deal with the problem of acoustic scanner noise but continuous scanning protocols have also been used in tonotopic experiments (Langers, Sanchez-panchuelo, et al. 2014). Sparse protocols employ a short fMRI acquisition followed by prolonged periods of silence, allowing for stimuli to be presented in near silence (Hall et al. 1999) This results in greater sensitivity to auditory responsive voxels compared to continuous acquisitions, where acoustic stimuli are masked to some degree by the presence of scanner noise. In the study presented here, the effect of scanner noise on estimates of tonotopic properties has been assessed by comparing the estimates of tonotopic properties from data collected using each acquisition protocol.

The drawback of sparse protocols is a reduced number of images that can be acquired in a given time, compared to continuous scanning protocols, due to the time between data acquisitions needed for stimulus presentation in silence. Reducing the number of data points increased noise and reduces reliability. In the work presented here, this effect is assessed by measuring the consistency of data across acquisition protocols.

As well as fMRI acquisition, stimulus presentation can have a profound effect on the results of a tonotopic mapping experiment. Presenting acoustic stimuli as a sequential frequency progression has been associated with detrimental effects, such as habituation, expectation artefacts and spatiotemporal BOLD nonlinearities (Langers, Krumbholz, et al. 2014). To

overcome this, a pseudo-randomised event-related design has been employed for data collection, in which acoustic stimuli of numerous frequencies are presented in random, non-sequential order.

Beyond stimulus presentation and data acquisition, inferring functional neuronal properties from BOLD fMRI activity also requires statistical analysis of the data. The analysis methods employed can be another potential source of bias on estimates of the cortical properties. The second objective of the study presented here was, therefore, to assess and compare two statistical methods for estimating frequency preference and selectivity from BOLD fMRI activity measured in response to acoustic stimuli. The first, general linear modelling (GLM) is a popular method that has been used in numerous studies to estimate the weight of the contribution of each stimulus to the measured BOLD activity, and from the weights derive functional properties such as preferred frequency and frequency selectivity. We compare this method with a second estimation method based on non-linear least-squares fitting, so-called population receptive field (pRF) modelling method.

Briefly, a GLM models the BOLD fMRI time course of a voxel as a linear combination of effects from different sets of events, known as explanatory variables. The timing of the events are known but the amplitudes (or weights) of their contributions to the fMRI response are to be estimated. A fundamental assumption of this method is that the fMRI response time course in a voxel can be explained by a linear combination of fMRI responses to the different sets of events and that all events are independent of each other. The analysis results in a weight (beta weight) for each of the variables that give the linear combination that "best" explains the given data. "Best" in this context means the choice of weights that minimises the square root of the sum of the squared difference be-

tween the modelled and measured BOLD time course (L2 norm).

In tonotopic mapping experiments, these beta weights can then be interpreted as a voxel's frequency response profile, from which we can estimate the preferred frequency and selectivity. The method employed to do this needs to be carefully considered. The presented study compares three methods of estimating preferred frequency from voxel response profiles: selecting the frequency with the largest beta weight, calculating the centroid of the beta weights and calculating a modified, unbiased centroid estimate. Selecting the frequency with the largest beta weight is a winner-takes-all method that is limited by the discrete nature and range of the frequency stimuli, as voxels cannot be assigned a frequency that was not presented. It is particularly susceptible to noise due to using only one data point in its estimate. The centroid of a distribution of responses is expected to be biased towards the centre of the sampled frequency range, due to the inherent property of the calculation, but has the benefits of producing estimates on a continuous scale and being less susceptible to noise by using all of the data points available in the calculation. The final method, described in more detail in Chapter 2 - General Methods, aims to alleviate this bias. The three methods are compared and assessed in terms of systemic bias and reliability.

Thomas et al. 2015 demonstrated that population receptive field modelling (pRF), which was developed in the visual research field by Dumoulin and Wandell (2008), could be used to estimate tonotopic properties of the human auditory cortex from fMRI data. pRF modelling creates a model of a voxels population receptive field, the extent of BOLD activity elicited by a stimulus as a function of stimulus space, to predict an fMRI time course. This model is iteratively fitted to the measured fMRI time course using non-linear methods. In the auditory domain, the

pRF is the frequency response profile of a voxel, the extent of BOLD activity elicited by a stimulus as a function of frequency, which is modelled by a continuous function, such as a Gaussian with unknown mean, corresponding to frequency preference, and unknown spread parameter sigma, corresponding to frequency selectivity. Frequency preference and selectivity are therefore directly estimated by fitting the model to the measured fMRI time course (Dumoulin and Wandell 2008; Thomas et al. 2015), in contrast with the GLM-based estimates which require further analysis. It has the additional benefit of overcoming the discrete nature of a winner-take-all method, the ability to assign values beyond the presented stimulus range and is not biased towards the centre of a stimulus frequency range while using all data points available in its estimation.

This study compares estimating the tonotopic properties of the auditory cortex using a pRF modelling technique and with three methods of estimating frequency preference from GLM frequency response profiles. Systematic bias on estimates of frequency preference is assessed by comparing and discussing the preferred frequency distribution for estimates from each method. Correlation between frequency preference estimates from datasets split in half is used as a measure of reliability (See Chapter 2 - General Methods for details). The comparisons and assessments presented in the study will allow for informed suggestions and experimental design for fMRI tonotopic mapping of the human auditory cortex.

3.2 Methods

Eight normal-hearing subjects (4 males; mean age = 29.38 SD = 7.34) voluntarily took part in the presented study.

3.2.1 Data collection

For individual subjects, blood-oxygen-level-dependent (BOLD) activity was measured in response to amplitude-modulated narrowband noise (NBN) stimuli at thirty-two different centre frequencies equally distributed on a quasi-logarithmic ERB scale (range = 0.1 - 8 kHz) and presented in a pseudo-random order. In addition to the functional data, a whole-head structural dataset and auxiliary datasets to aid data preprocessing were obtained for each participant. Scanning was performed on a Philips Achieva 7T MRI system using a volume transmit coil and a 32-channel receiver coil. All datasets were collected within a single session, apart from subject 7, for whom functional data were collected in two separate sessions.

Two functional data sets were collected for each subject using a sparse acquisition design, with a long TR (7.5s) and stimulus presentation between data acquisition in silence, and another two datasets were acquired using a continuous acquisition design, with a shorter TR (1.5s), in which stimulus presentation and data acquisition was simultaneous and therefore in the presence of acoustic noise from the scanner. Acoustic scanner noise was measured, in a single scanning session, using a pressure field microphone placed in the scanner on foam, held by a subject to their chest. For further details, see Chapter 2 - General Methods.

3.2.2 Data analysis

This study is concerned with assessing the impact of the acquisition protocol and analysis method on the estimation of frequency properties of the auditory cortex. Data analysis followed the pipeline detailed in Chapter 2 General Methods.

To assess the acquisition protocols, the effect of scanner noise (i.e. sparse vs continuous acquisition) on estimates of frequency preference and selectivity for voxels was assessed within the AC gradients ROI, defined based on frequency gradient reversals. The reliability of estimates was compared between acquisition protocols to determine if reducing the number of observations, as a result of a sparse acquisition protocol, is detrimental to the consistency of estimates between data acquisitions.

Data analysis was restricted to voxels within the AC gradients ROI. We estimated the preferred frequency and frequency selectivity of each voxel using four different methods. The first three estimated the frequency preference and selectivity of voxels from their individual frequency response profiles to all stimuli (the result of the voxelwise GLM analysis). In order of increased computation complexity, these were: (1) selecting the frequency that elicited the greatest response from the voxel with no estimate of selectivity, (2) calculating the centroid and spread of the response profile, (3) calculating the centroid and spread using a method that aims to correct the bias due to the limited range of stimulation frequencies (see Chapter 2 - General Methods for more details). The final estimation method assess was taken the preferred frequency (f_0) and spread (σ) parameter estimates from fitting population receptive field model. The preferred frequency distribution within the AC gradient ROI of each estimation method was calculated and compared to investigate possible biases. The reliability of each method was measured by performing a Pearson correlation of voxel estimates of preferred frequency between functional data acquisition runs.

The effect of acquisition protocol was assessed by comparing the estimates of tonotopic properties from data collected using each acquisition protocol. The preferred frequency and frequency selectivity of vox-

els were estimated using population receptive field modelling (pRF). Voxel frequency response profiles, the strength of the response elected to each stimulus centre frequency, were estimated using a generalized linear model (GLM), with the estimated beta weights being a measure of the strength of response to each stimulus. Subsequent group-level data analysis was restricted to surface ROIs defined based on gradient reversals in the tonotopic maps of the auditory cortex estimated from data collected using the sparse acquisition protocol using an unbiased and automated procedure. The effects of scanner noise were assessed through data visualisation and statistical analysis. Possible biases in the estimates of frequency preference were assessed through the distribution of estimated frequency preference for voxels within the AC gradients ROI. The effects of scanner noise on the response of the auditory cortex were assessed by calculating the average frequency response profile and average tuning curves. The reliability of each method is estimated by the repeatability of estimates between scans.

3.3 Results

3.3.1 Impact of analysis methods on estimates of cortical frequency tuning properties

Impact of analysis methods on spatial distribution of preferred frequency

Four methods of estimating preferred frequency were compared. Taking the stimulus with the largest beta, *Max*, calculating the centroid of the voxel frequency profile, *centroid*, a method that attempts to correct the bias of calculating a centroid within boundaries, *bias-corrected centroid*,

and populational receptive field analysis, *pRF*. All preferred frequency estimates were derived from data acquired using the sparse acquisition protocol to avoid the masking effects of scanner noise.

Tonotopic maps derived from all estimation methods, presented in figure 3.1, show spatially consistent topographic organisations, with high and low-frequency foci consistently located anatomically. All methods yield the same tonotopic organisation: a low-frequency response focus located on the central-lateral part of HG with two mirror-reversed preferred-frequency gradients running towards two elongated high-frequency foci, one anterior-medial and the other posterior to HG, in a V like shape. The difference between the estimates generated by each method lies in the range of the preferred frequency estimates within these tonotopic gradients, and the function describing the change in frequency between the high and low-frequency regions. Visualising the data as tonotopic maps does not allow more detailed observations and inferences made can be influenced by the choice of colour scales (fooling the eye to see patterns not representative of the data), and therefore conclusions from these figures should be done in the context of other analysis' and visualisation of the data.

Impact of analysis methods on the distribution of preferred frequency

Estimates of preferred frequency were highly correlated (> 0.8) between all methods (See table 3.1) across voxels from all subjects. However, they showed obvious differences in distribution across voxels. The distribution of preferred frequency was determined for the four estimation methods discussed in the presented study. In figure 3.2 and the area under each point of the distribution curve represents the density of voxels with a given preferred frequency.

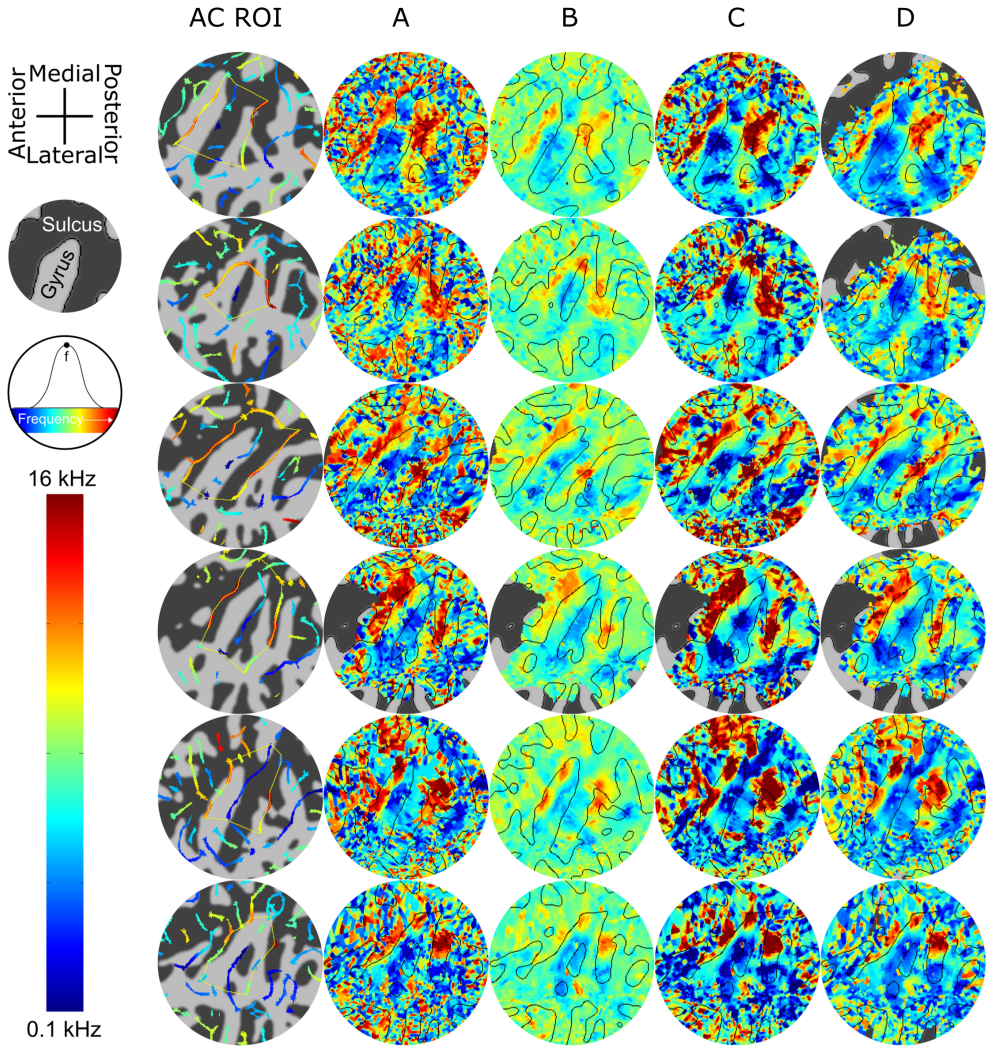


Figure 3.1: Comparison of the spatial distribution of voxel-wise preferred frequency estimates using the four different estimation methods

Maps of preferred frequency estimated using four methods. From left to right: sulci/gyri boundaries and frequency gradient reversals, maps of voxel-wise preferred frequencies estimated as (A) frequency with largest beta weight, (B) centroid of the frequency response profile, (C) de-biased centroid of the frequency response profile, and (D) centre frequency parameter of the population receptive field model. Each row is an individual hemispheres, with ones from the right flipped horizontally to align the orientation of HG across all hemispheres. The low-frequency gradient reversal on and high-frequency surrounding Heschl'S Gyrus were used to delineate an auditory core ROI (depicted by yellow line in the first column). The preferred frequencies estimates from each analysis method show a high degree of agreement in tontopic organisation.

The method used to estimate preferred frequency has a substantial impact on the distribution of estimates from voxels with the human PAC. The following interpretations are made under the assumption that preferred frequencies should be (roughly) uniformly distributed across frequency on a logarithmic scale and extend beyond the stimulus range to cover the entire human range of hearing. This assumption is based on cochlear frequency resolution varying in a quasi-logarithmic manner with frequency across the range of human hearing (Glasberg and Moore 1990).

Calculating the centroid of the frequency response profile biases estimated preferred frequencies towards the centre of the presented stimulus frequency range and limits them to values within this range. Selecting the stimulus which elicits the greatest response from a voxel reduces this systematic bias but values are still limited to the stimulus range. In addition, values are discrete, limited to the centre frequencies of the presented stimuli, and the distribution appears to be less smooth than the other analysis methods. The bias-corrected centroid method reduces the pronounced peak of the uncorrected centroid method and extends the distribution beyond the stimulus range, but representation of high and low frequencies beyond this range seems exaggerated, resulting in a multimodal distribution. The distribution of preferred frequencies from the pRF analysis follows the general profile of that of selecting the greatest response, a Gaussian-like shape that is skewed towards lower frequencies but extends beyond the stimulus range. In addition, estimates are on a continuous scale and appear smoother.

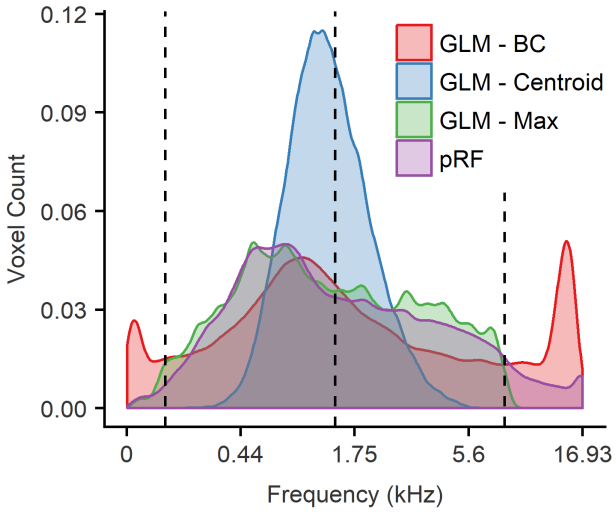


Figure 3.2: Distribution of preferred frequency estimates across AC voxels for different estimation methods

The distribution of preferred frequency estimates for voxels within the primary auditory cortex derived using different the four estimation methods were the maximum beta weight across EVs of the GLM model (GLM - Max), the centroid of the frequency response profile from the GLM analysis (GLM - Centroid), the bias-corrected centroid from the GLM analysis (GLM - BC), and the f_0 parameter estimate of the fitted population receptive field model (PRF). The preferred frequency is plotted in kHz on an nERB scale. All prepared frequency estimates were derived from concentrated runs of data acquired using a sparse acquisition protocol at the subject level. The preferred frequency is plotted in kHz on an axis space equally in terms of nERB. The height at a given point represents the density of estimates within a kernel (Gaussian function, sigma = 1 ERB). Dashed vertical black lines mark the limits and centre of the stimulus frequency range. Estimating preferred frequency from GLM frequency response profiles is biased by the method employed. Calculating the centroid is biased towards the centre of the stimulus frequency range. The method employed to correct this bias result is over representation of high and low frequency limits. Selecting the maximum beta weight is more evenly distributed across the frequency range but limited to stimulus frequencies. pRF estimates are similarly distributed but extends beyond the stimulus frequency range.

Table 3.1: Correlation between preferred frequency estimates from different estimation methods

	Centroid	bias-corrected	Max	pRF
Centroid	1.00	0.89	0.82	0.86
bias-corrected	0.89	1.00	0.81	0.89
Max	0.82	0.81	1.00	0.81
pRF	0.86	0.89	0.81	1.00

Impact of analysis methods on the spatial distribution of frequency selectivity

Frequency selectivity has been shown to be sharper, narrower tuning widths, for core areas in animal models of the auditory cortex (Kaas and Hackett 2000) and has been suggested as a potential marker of human auditory core(Moerel, Martino, and Formisano 2012). The frequency tuning width of voxels was estimated using three methods; calculating the sample spread of the frequency response profile estimated from the GLM analysis, the bias-corrected spread, and the spread parameter (sigma) of the fitted population receptive field model. Maps of frequency selectivity derived from GLM frequency response profiles (see figure 3.3) show that frequency selectivity is narrower for voxels located on HG, compared to the surrounding cortex and are highly correlated (0.84) with each other. Frequency tuning width estimates from the pRF analysis appear similar within the AC gradients ROI but are extremely narrow in areas of poor fit. The average spread of the AC gradients ROI voxels frequency response profiles was 9.81 nERBs (SD=1.28), which increase to 16.20 (SD=6.54) nERBs when the debiasing procedure was implemented. The Gaussian model fitted by the pRF analysis is narrower, with a sigma of 5.89 nERBs (SD=4.57).

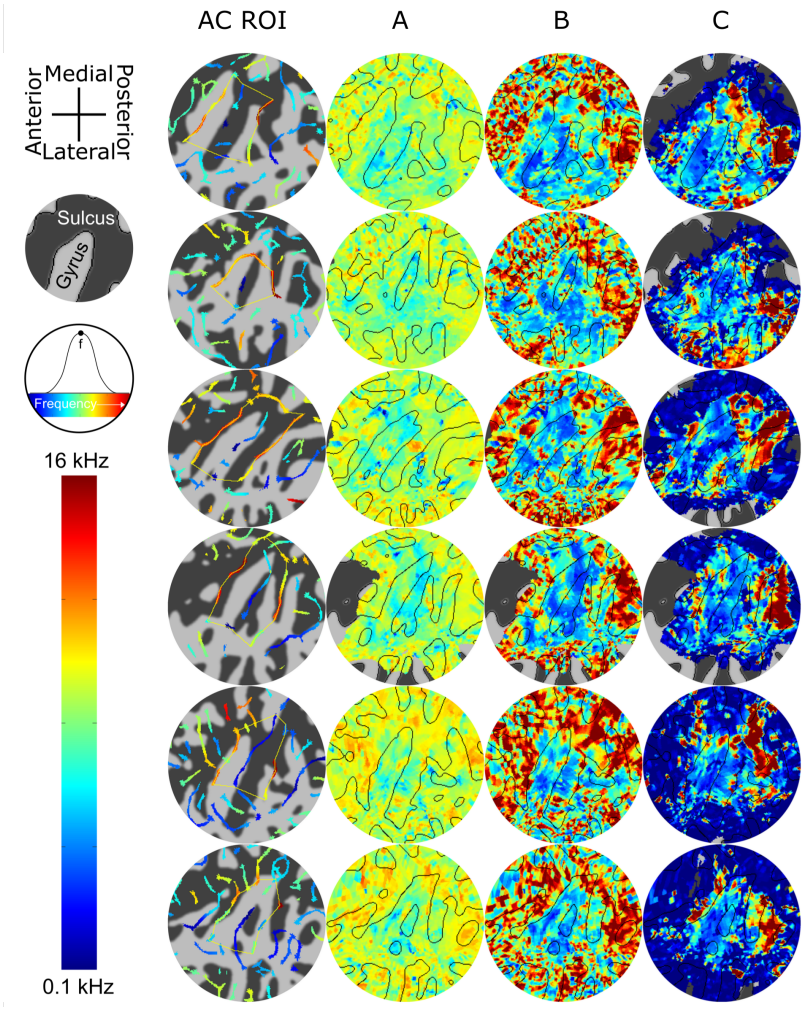


Figure 3.3: Comparison of spatial distributions of voxelwise frequency selectivity computed using three different estimation methods

From left to right: AC ROI) sulci/gyri boundaries and frequency gradient reversals, A) the spread of the GLM frequency response profile, B) the bias-corrected spread parameter, C) and the spread parameter of the population receptive field model fitted to the voxel. Both estimation methods from GLM frequency response profiles feature narrower tuning widths on HG. This trend is seen from the pRF estimates for but voxels with poor fits (areas further from HG) are extremely narrow, making this relationship hard to see. Each row is an individual hemispheres, with ones from the right flipped horizontally to align the orientation of HG across all hemispheres. The low-frequency gradient reversal on and high-frequency surrounding Heschl'S Gyrus were used to delineate an auditory core ROI (depicted by yellow line in the first column).

Impact of analysis methods on the spatial distribution of goodness of fit

Goodness of fit describes how much of the measured BOLD fMRI time-course a fitted model explains. Voxels with high R^2 values have most of their variance explained by the model and it can be assumed that the estimates of preferred frequency and selective are representative. Comparing the R^2 values between different analysis methods directly is not valid or informative. However, comparing the spatial distribution of the goodness of fit can help determine which voxels respond strongly to the stimulus (see figure 3.4). The best fitting voxels from GLM and pRF models are generally located in the same regions of the cortex. For all subjects, the best fitting voxels were located on HG.

Impact of analysis methods on the reliability of preferred frequency estimates

To assess the reliability of each preferred frequency estimation method, both GLM and pRF models were fitted to each of the two functional runs, and estimated run-specific voxelwise preferred frequency and frequency selectivity using each of the estimation methods presented above. The Pearson's correlation coefficient was computed between run-specific values across voxels for each estimation method. Estimating preferred frequency as the centroid from GLM beta weights resulted in the most highly correlated tonotopic maps (r = 0.64), followed by the bias-corrected centroid (r = 0.59), pRF estimates (r = 0.54) and finally selecting the frequency with the largest beta weight (r = 0.41). The reliability of frequency selectivity estimates was much lower for methods, with GLM-based estimates slightly less reliable, r = 0.227 and 0.279, than the pRF estimates, r = 0.283.

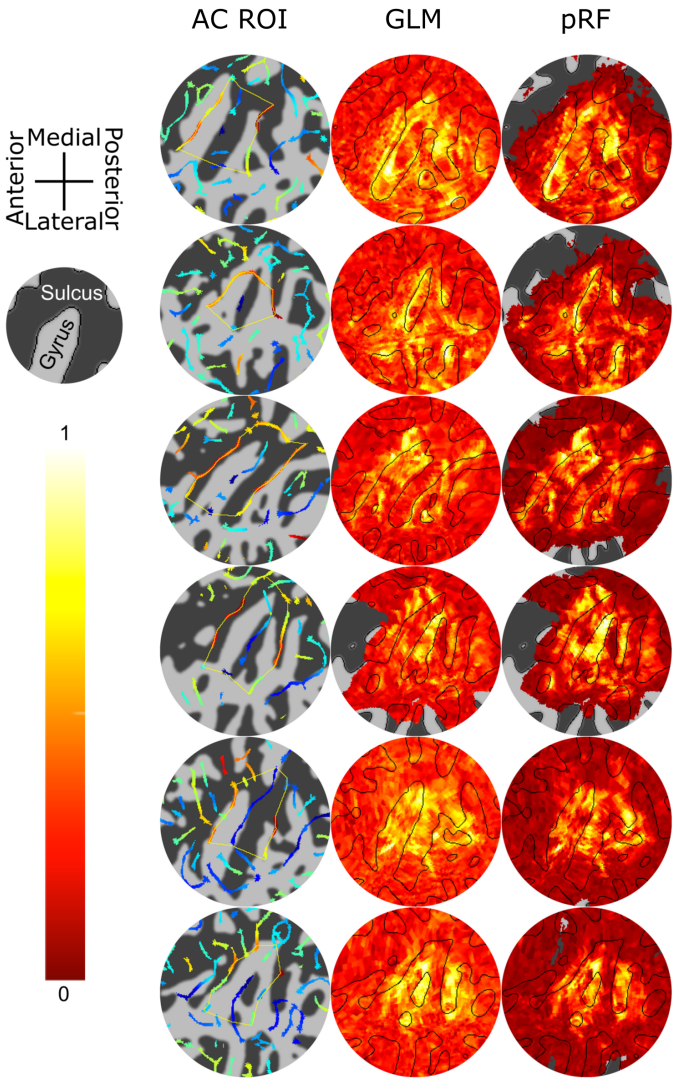


Figure 3.4: Impact of analysis methods on the spatial distribution of voxelwise goodness of fit

Spatial distribution of voxelwise goodness of fit for the GLM and pRF models, measured using R^2. Left, the R^2 values of GLM models fitted to the voxel time courses and right, the R^2 values from the voxelwise fitted pRF models. Model fits are best on HG for both analysis methods. The GLM model achieves higher R^2 values but as it has more degrees of freedom isn't a fair direct comparison. Each row is an individual hemispheres, with ones from the right flipped horizontally to align the orientation of HG across all hemispheres. The low-frequency gradient reversal on and high-frequency surrounding Heschl'S Gyrus were used to delineate an auditory core ROI (depicted by yellow line in the first column).

3.3.2 Comparison of Acquisition Protocols

Two acquisition protocols are compared, sparse and continuous, to determine which is the most appropriate for collecting data for tonotopic investigation. The impact of scanner noise on estimates of frequency properties of the human auditory cortex was assessed by comparing estimates derived from data collected in response to stimuli presented either in noise (continuous acquisition) or in silence (sparse acquisition). The impact of reducing the number of observations, necessary to achieve a sparse acquisition in the scanning time, on estimation reliability can not be disentangled from that of the acquisition protocol but the assumption is that will be a major factor. Here, reliability has been estimated as how consistent estimates are between function data acquisition runs. Both the GLM and pRF analyses were fitted to both datasets and estimates from either analysis pipeline were used depending on the particular comparison.

Impact of acquisition protocol on the spatial distribution of preferred frequency

Voxelwise preferred frequency was estimated by fitting population receptive field models, the least biased analysis method presented in this study, independently to each acquisition data set (sparse or continuous). Figure 3.5 shows the tonotopic organisation of preferred frequency in the auditory cortex of individual subjects. The spatial organisation of frequency is consistent between acquisition protocols, corresponding in the high- and low- frequency cortical areas on Heschl's gyrus. Tonotopic maps computed from sparse or continuous data showed a fairly high correlation (r=0.65) across voxels, however, there appears to be an increase in high-frequency representation in the maps derived from the continuous

data.

Impact of acquisition protocol on the average activation as a function of frequency

For each acquisition protocol, the average BOLD response amplitude was calculated as a function of stimulation frequency by averaging GLM frequency response profiles across all voxels of the AC gradients ROI of each subject. This quantifies the impact of the acquisition protocol, most likely due to acoustic scanner noise, on the average amount of BOLD activity elicited across the frequency range. ROI-averaged frequency response profiles significantly, based on non overlapping error bars, differed between the two datasets, see figure 3.6 and figure 3.8. The continuous response profile appears to be the product of the frequency profile obtained in silence (the sparse dataset) and an inverse of the scanner noise. The scanner noise has two characteristic peaks at 0.57 and 1.42 kHz which correspond with two pronounced dips in the continuous response profile, where the percentage signal change has been attenuated.

The maximum BOLD response amplitudes across stimulus frequencies obtained did not significantly differ between acquisition protocols (p-value = 0.76, t=0.32), but they corresponded to frequencies at opposite ends of the spectrum (low frequencies for sparse: 1.58% signal change at 0.572kHz, high frequencies for continuous: 1.71% signal change at 6.810 kHz). The average BOLD response amplitude across stimulus frequencies did not significantly differ between acquisition protocols (sparse:1.15, Continuous 1.11 p-value = 0.45).

Figure 3.8 depicts the difference between BOLD response amplitude between acquisition protocols as a function of stimulus frequency and the

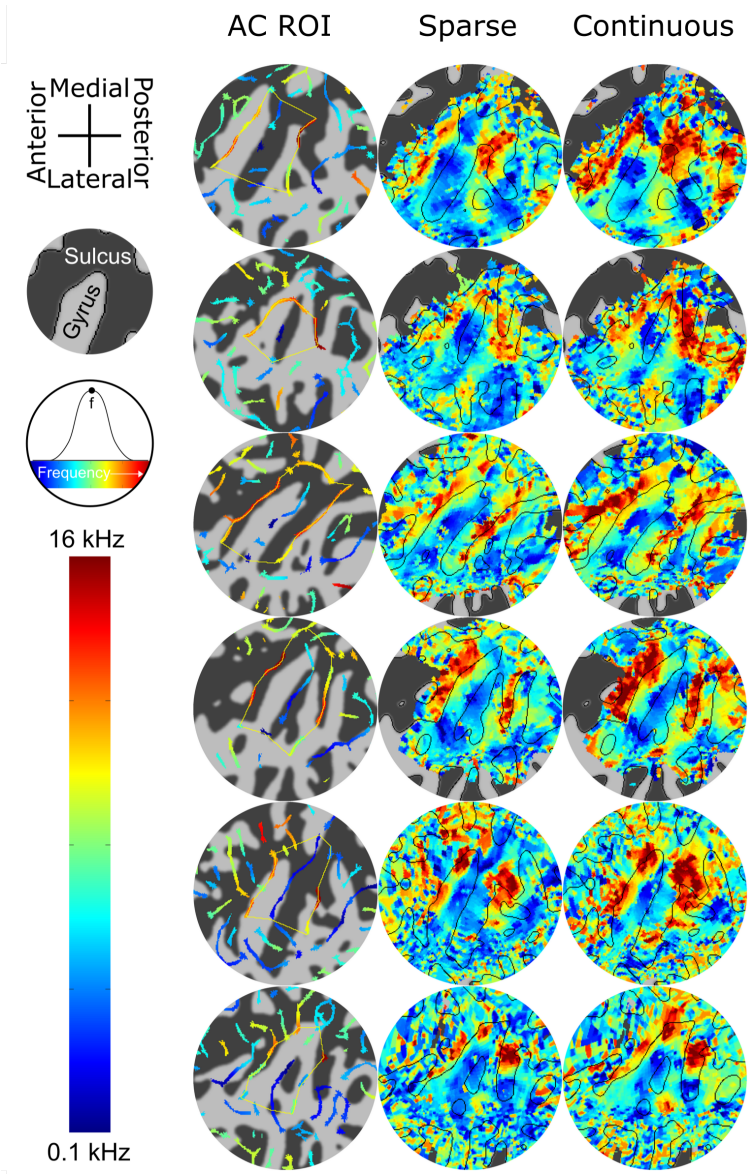


Figure 3.5: Impact of acquisition protocol on the spatial distribution of voxelwise preferred frequency

Spatial distribution of voxelwise preferred frequency estimated using a pRF analysis from data collected using either a continuous or sparse acquisition protocols. The estimated tonotopic organisation appears broadly the same from data collected using either acquisition protocol but continuous data appears to result in more voxels with a high frequency preference. Each row is an individual hemispheres, with ones from the right flipped horizontally to align the orientation of HG across all hemispheres. The low-frequency gradient reversal on and high-frequency surrounding Heschl'S Gyrus were used to delineate an auditory core ROI (depicted by yellow line in the first column).

frequency profile of the scanner noise. The largest difference is between 0.4 and 0.5 kHz, where there is a reduced response to stimuli in the presence of scanner noise, which corresponds to a large peak in the scanner noise. Above about 3 kHz, there is an increase in BOLD response for stimuli presented in the presence of scanner noise.

Impact of acquisition protocol on the distribution of preferred frequency

The distribution of preferred frequencies across voxels of the AC gradient ROI was calculated for each acquisition protocol from the pRF estimates (figure 3.8).

The distribution of preferred frequencies estimated from the sparse data is roughly Gaussian and skewed towards low-mid-frequencies. The distribution of preferred frequency estimated from the continuous condition was multimodal in shape, skewed towards higher frequencies than the sparse distribution, with a profile reflecting the inverse of the acoustic scanner noise frequency profile, with the two dips in the distribution possibly corresponding to the peaks in the acoustic noise of the scanner.

Impact of acquisition protocol on the average voxel frequency response profiles

To determine if scanner noise has the same effect on all voxels irrespective of their preferred frequency, average voxel frequency profiles were estimated for each dataset (see figure 3.9). The previous frequency response profiles give an idea of how acquisition protocol impacts the response of the auditory cortex as a whole to frequency but loses the frequency tuning of individual voxels. Here, voxels were first binned by pre-

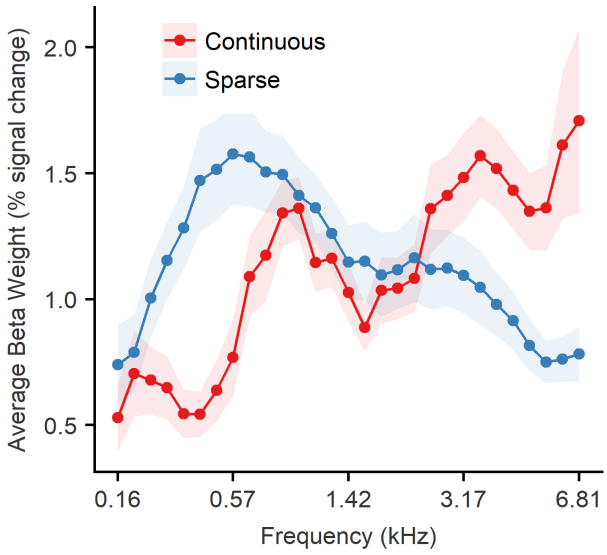


Figure 3.6: Impact of acquisition protocol on average BOLD response in primary auditory cortex as a function of stimulus frequency

Impact of acquisition protocol on average BOLD response in primary auditory cortex as a function of stimulus frequency. Average BOLD response as a function of frequency was computed by averaging voxelwise GLM response profiles (beta weights) across voxels with the AC gradient ROI. The average BOLD response level across frequency is different for each acquisition protocol. Sparse data results in a greater response from low-frequency stimuli, whereas continuous data elciates a greater response to high-frequency stimuli. Blue represents estimates from data acquired using a sparse acquisition protocol and red from data acquired using a continuous acquisition protocol. The solid line is the mean for all hemispheres (n =16) from all subjects AC gradients ROIs and the ribbons represent the standard error.

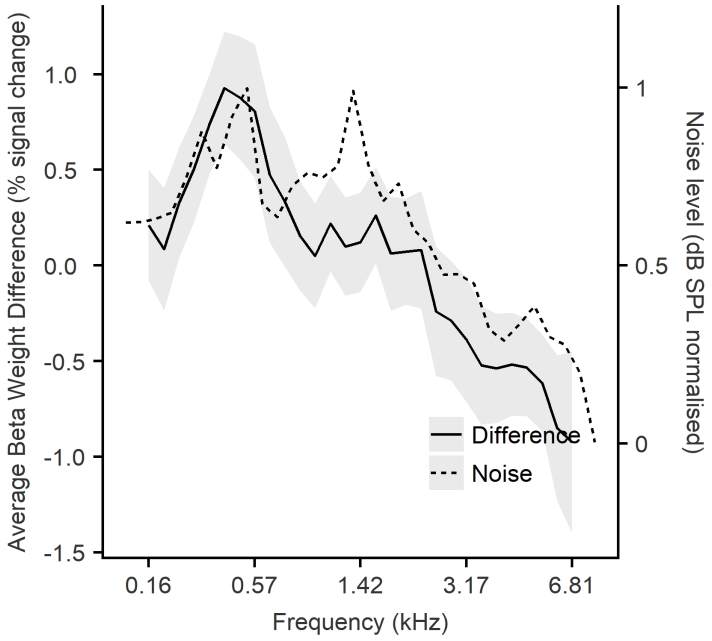


Figure 3.7: Comparison between the difference in average BOLD response between acquisition protocols

Comparison between the difference in average BOLD response between acquisition protocols in the AC gradient ROI and the frequency profile of the scanner noise. The black line is the difference between the average frequency response profiles of the two acquisition protocols (see figure 3.6), obtained by subtracting the continuous acquisition profile from the sparse acquisition profile. The dashed line is the frequency spectrum of the scanner noise measured in the scanner bore. The frequency profile of the difference between the BOLD response and that of the scanner noise is similar in shape.

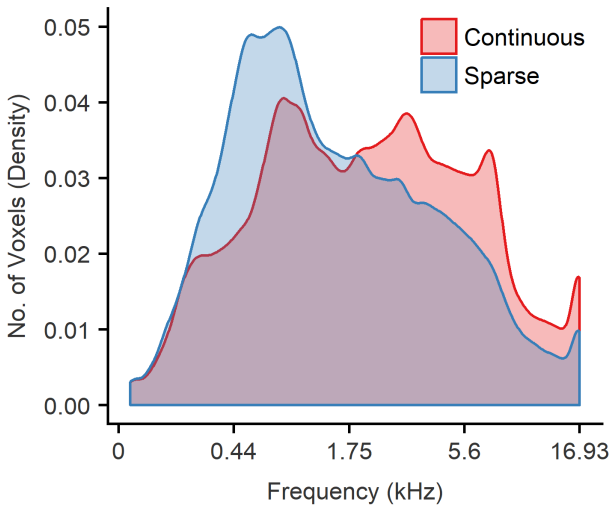


Figure 3.8: Impact of acquisition protocol on the distribution of preferred frequencies

Distribution of preferred frequencies across voxels of the AC gradient ROI for the sparse and continuous data. The greater response to high frequency stimuli for data from continuous acquisition protocols (see figure 3.6) appears to result in a greater proportion of voxel with high-frequency preference. Preferred frequency estimates were derived using the pRF analysis.

ferred frequency before averaging.

Frequency tuning was observed both in the presence and absence of scanner noise, but tuning curves differed between the sparse and continuous data in a way that probably reflected the effect of scanner noise on the continuously-acquired data. Frequency-specific tuning curves show that the amount of attenuation is dependent on the preferred frequency of the voxel. The presence of scanner noise reduces the BOLD response to stimuli with centre frequencies within scanner noise peaks, regardless of the voxels preferred frequency. The implications of this for estimates of frequency selectivity depend on the preferred frequency of the voxel. For voxels with preferred frequencies outside the peaks of the scanner noise (two highest preferred frequencies), the frequency response profiles are narrower in the continuous compared to sparse acquisition. Conversely, for voxels with preferred frequencies close to peaks of the scanner noise (lower preferred frequencies), frequency response profiles experience the maximum attenuation in response to low frequencies and increased responses to high-frequencies.

Impact of acquisition protocol on reliability of preferred frequency estimates

To assess the reliability of preferred frequency estimates for each acquisition protocol, the correlation was computed across voxels from all subjects between voxelwise preferred frequency estimates derived from each functional run of each acquisition dataset (splitting the collected data in half). The continuous acquisition protocol produces slightly more reliable estimates ($r = 0.61$) than the sparse acquisition protocol ($r = 0.54$).

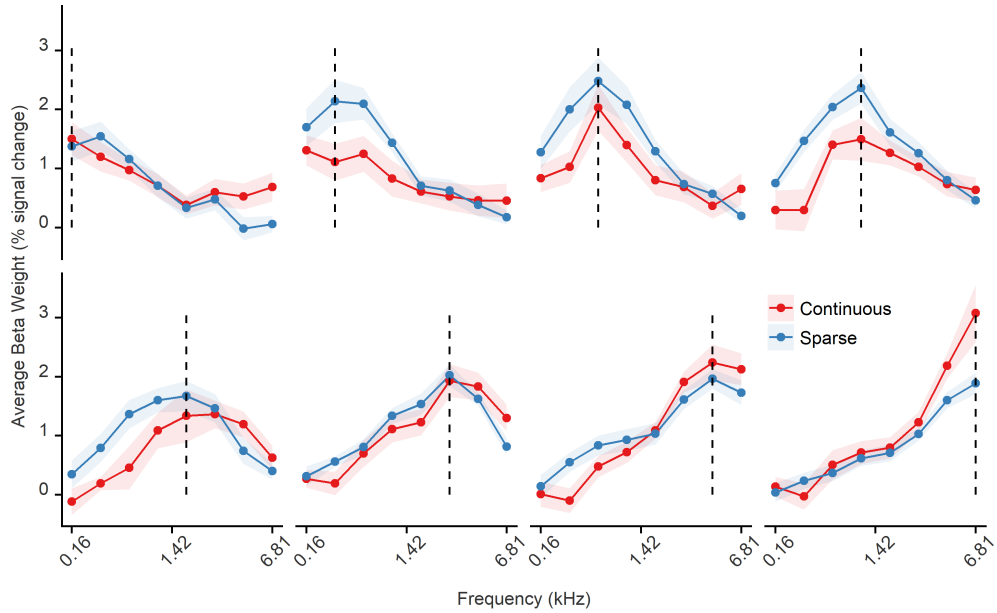


Figure 3.9: Frequency response profiles averaged according to voxel's preferred frequency for continuous and sparse acquisition data

Average frequency response profiles were derived by averaging voxelwise response profiles after grouping AC gradient ROI voxels in 8 preferred frequency bins. Voxelwise preferred frequencies and resulting average frequency response profiles were derived using a split-half method to avoid artificially narrow frequency response profiles (for details, see General Methods). The frequency profiles estimated from continuous data are less consistent in shape that the sparse data. Points connected by solid lines represent the average beta weight associated with each stimulus frequency bin across all hemispheres from all subjects (n =16) and the ribbon represent the corresponding standard error.

Impact of acquisition protocol on the spatial distribution of goodness of fit

Voxels with the best fitting pRF models (highest R^2 values) were consistently located medially on HG for all subjects, (as seen in the presented subset of subjects in figure 3.10), indicating an area of the cortex responsive to the presented auditory stimuli. The goodness of fit for models fitted to sparse data was better than data collected in the presence of scanner noise. This is despite the reduced number of time points sampled in the sparse protocol, 146 frames compared to 722 in the continuous protocol. The ROI-averaged model fit for data collected in silence ($R^2 = 0.21$) was significantly better than for data collected in the presence of scanner noise ($R^2 = 0.07$) (Welch two-sample t-test $p < 0.005$).

Impact of acquisition protocol on the spatial distribution of frequency selectivity

Voxel frequency selectivity was estimated as the spread parameter of the fitted pRF model (figure 3.11). A broadening of frequency selectivity was observed laterally on Heschl gyrus, corresponding to the area with a preference to low-frequencies, when estimated from data acquired in silence for all subjects. The frequency selectivity averaged across all voxels of the AC gradient ROI was significantly broader for estimates from data collected in silence (pRF sigma = 5.89 nERB) than measurements made in the presence of scanner noise (pRF sigma = 4.85 nERB) (p-value $< 2.2e - 16$). There was a slight correlation, 0.33, between the two conditions.

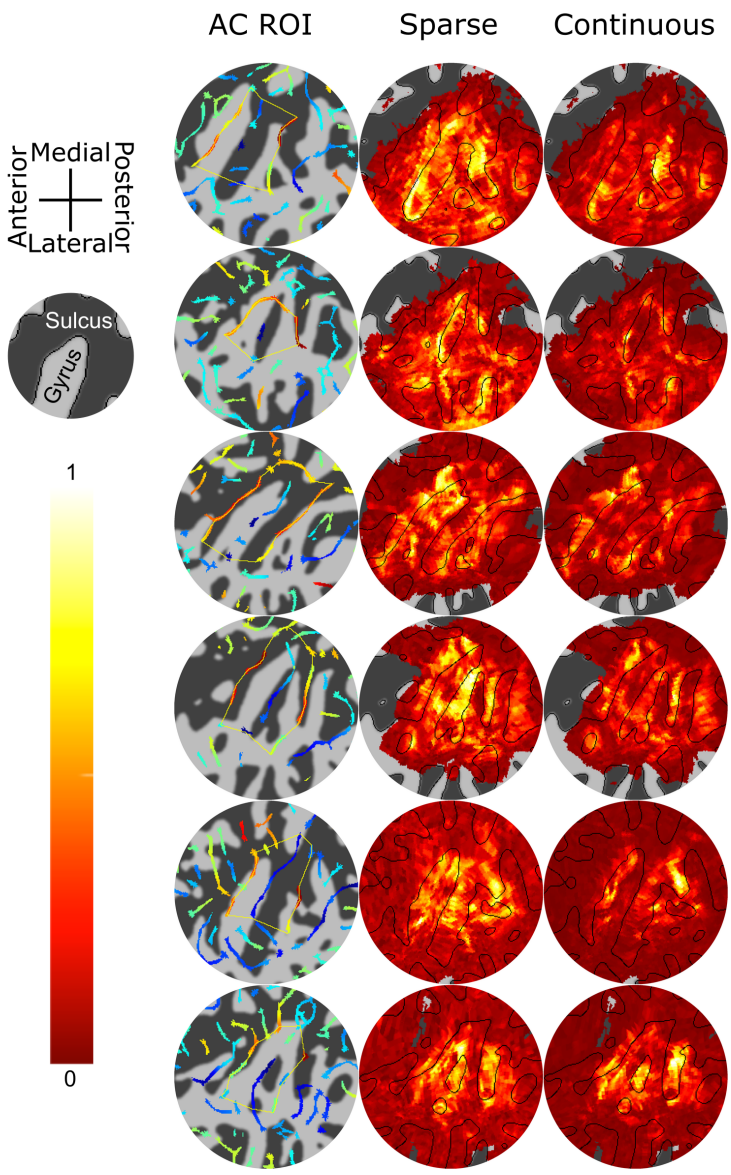


Figure 3.10: Effect of acquisition protocol on the spatial distribution of voxel-wise goodness of fit of the pRF models to the fMRI time courses

Effect of acquisition protocol on the spatial distribution of voxelwise goodness of fit of the pRF models to the fMRI time courses, measured using R^2. Left: pRF models fitted to sparse data; Right: pRF models fitted to continuous data. Sparse have a greater number of high R^2 values, located more centrally on HG. Each row is an individual hemispheres, with ones from the right flipped horizontally to align the orientation of HG across all hemispheres. The low-frequency gradient reversal on and high-frequency surrounding Heschl'S Gyrus were used to delineate an auditory core ROI (depicted by yellow line in the first column).

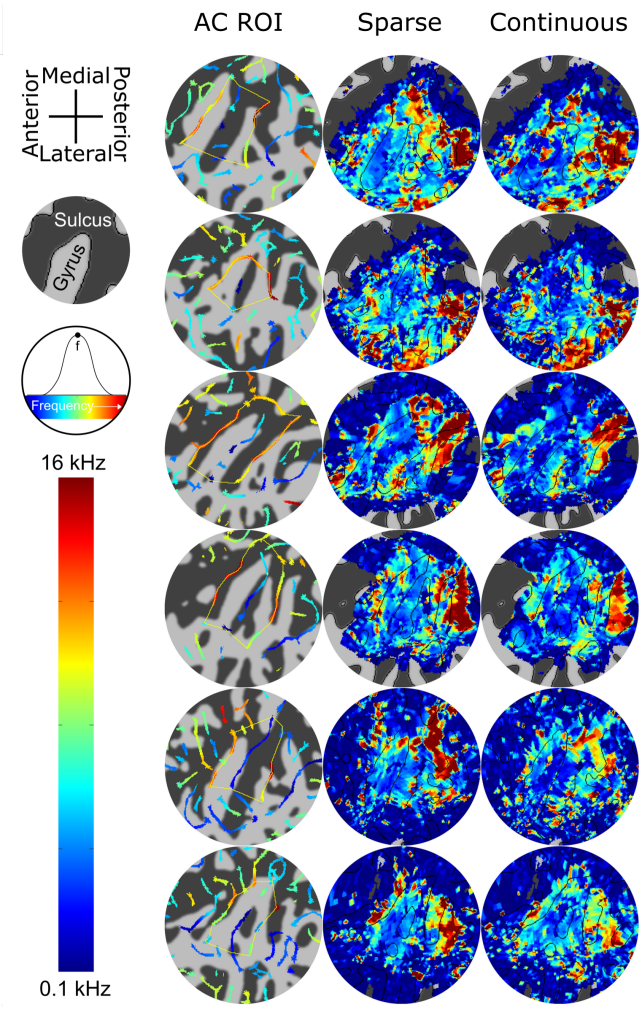


Figure 3.11: Effect of acquisition protocol on the spatial distribution of voxel-wise frequency selectivity

The spatial distribution of voxelwise frequency selectivity, estimated as the spread parameter of the pRF model fitted to the voxel timeseries for data collected using sparse and continuous acquisition protocols. Estimates on HG appear narrower than surrounding regions, for well fitted voxels. A broadly tuned region can consistently been seen posteriorly to HG. Poorly fitted voxels in surrounding areas are extremely narrow, which is unlikely to reflect the true properties of the cortex in that area. Each row is an individual hemispheres, with ones from the right flipped horizontally to align the orientation of HG across all hemispheres. The low-frequency gradient reversal on and high-frequency surrounding Heschl'S Gyrus were used to delineate an auditory core ROI (depicted by yellow line in the first column).

3.4 Discussion

The preferred frequency and frequency selectivity of an auditory responsive voxel is a fundamental property that forms the foundation of tonotopic research, and much of the wider auditory field. As further investigation and estimates are based on these parameters, it is vital to understand possible shortcomings and limitations to the estimation method used to calculate them, and the validity of estimated values. This study compared four estimation methods used to estimate functional tuning properties of voxels that have been applied in the auditory domain.

3.4.1 Impact of analysis methods on estimates of cortical frequency tuning properties

The preferred frequency and frequency selectivity of voxels sampling the human auditory cortex have been estimated. Estimations from four analysis methods have been compared and the impact of presenting stimuli in the presence of scanner noise has been assessed by comparing continuous and sparse acquisition protocols.

Estimates of preferred frequency are considerably impacted by the analysis method employed. The distribution of preferred frequency estimates is biased by the limitations of the method. Estimating the preferred frequency from voxel frequency response profiles (beta weights of GLM frequency regressors) is limited by the impact of the frequency boundaries. For preferred frequency estimates, using the maximum beta weight or the centroid restricts estimates to frequencies within the stimulus frequency range. Therefore cortical regions that prefer frequencies beyond this range will be assigned preferred frequencies at the boundary of the

presented stimulus frequency range. As it is not practically possible to present stimuli for the full frequency range of human hearing, this will result in frequencies outside the stimulus range, but within the human hearing range to be under-represented and stimulus boundary frequencies to over represented. Implementing a method to alleviate this bias from centroid estimation allows preferred frequencies beyond the stimulus range.

Despite the impact on the distribution of preferred frequency, the spatial cortical organisation remains consistent between estimation methods, suggesting that estimating tonotopic organisation is robust to estimation methods. In particular, the location of low- and high-frequency gradient reversals is consistent between analysis methods. This suggests that delineating cortical areas based on gradient reversals is robust to the procedure employed to estimate preferred frequency, making it a valid method of delineation. ROIs defined using gradient reversals provide consistent cortical areas to compare estimates between studies and appear to be robust to the preferred frequency estimation method employed.

Estimating preferred frequency to be the frequency of the stimulus that elicits the greatest response makes the least assumptions but is the most susceptible to noise due to being based on the least data, a single point from the frequency profile being larger than the rest. This noise can be seen visually in the unsmoothed distribution of preferred frequencies (see figure 3.3.1) but the discrete nature of this method also contributes to this. Split-data voxelwise within-subjects correlation coefficients show that this method is the least reliable of all methods tested. Calculating the centroid of voxels frequency response profiles appears to be the most biased, with estimates heavily biased towards the centre of the stimulus frequency range, but is also the most reliable method based on split-

data voxelwise within-subjects correlation coefficients. The centroid of a given distribution is heavily biased by its boundaries, as information is absent beyond. Therefore, the increased reliability of this method is to be expected as a consequence of this bias. First, a consistent bias all estimates experience will increase the correlation between data acquisition runs. Secondly, more information from all data points is used in the calculation of a centroid, reducing susceptibility to measurement noise. Attempting to correct for this bias, through modelling and fitting hypothetical frequency response profiles is only partially successful. Voxel preferred frequencies are assigned beyond the range of the frequency limits of the presented stimuli but the distribution of the estimates is still biased towards the centre of this range. The improvements afforded by this method appear worthwhile, with little loss of reliability, the correlation between data acquisition runs reduces from 0.64 to 0.59, and high correlation with the other estimation methods ($r > 0.81$). Estimates from a population receptive field analysis appear to be the least biased. The distribution reflects that of estimates from selecting the frequency of the stimulus that elicits the greatest response, the method that makes the least assumptions, and its distribution of preferred frequency doesn't appear to be biased by the stimulus frequency limits. In addition, estimates can be assigned continuous values that extend beyond stimulus boundaries, resulting in a smooth distribution. Estimates are highly correlated between analysis methods ($r > 0.81$) which suggests all methods are estimating the same functional property and physiological phenomena; this is further supported by the spatial consistency between the goodness of fit of the different models (see figure 3.4).

Frequency selectivity has been proposed as a functional marker of the primary, or core, auditory cortical region, which has a narrower tuning

than non-primary regions (Morel, Garraghty, and Kaas 1993). For the present study, three methods have been employed to estimate frequency selectivity. The first method calculates the spread of the frequency response profile estimated from the GLM analysis, which makes fewer assumptions about the shape the auditory responsive voxels response profile than fitting a Gaussian. This measure provides little information beyond the relative selectivity to other voxels and individual estimates have little meaning. However, changes across the cortex are informative and demonstrate how selectivity changes with cortical location. Using this method, a region of high selectivity is consistently seen centred on Heschl's gyrus for all hemispheres, corresponding well with what has been found by other fMRI tonotopy studies (Dick et al. 2012; Moerel, Martino, and Formisano 2012; Wasserthal et al. 2014; De Martino et al. 2015; Thomas et al. 2015; Besle et al. 2019) and human intracranial studies (Nourski et al. 2012; Gander et al. 2019). Heschl's gyrus is expected to approximately correspond to the location of the primary auditory cortex and a narrow tuning is suggestive of a primary sensory cortex, which has a finer sampling resolution of sensory space than non-primary areas (Harvey and Dumoulin 2011). This spatial pattern of narrow tuning on Heschl's gyrus and broader tuning in cortical areas surrounding it is clearest when frequency selectivity is estimated using the bias-corrected spread of the frequency response profiles. This makes the assumption that the shape of an auditory responsive voxel can be modelled as a Gaussian function on an nERB scale. The results of estimating average frequency response profiles, see figure 3.8, suggests that this a valid assumption, with each frequency bin appearing bell-like in shape on an nERB scale. Minimal assumptions have been made for in estimating average frequency response profiles. It is assumed that voxels with the same preferred frequency share similar frequency profile properties and that averaging their

responses is representative of individual voxel properties. The Gaussian shape of the frequency response profiles suggests that it is a suitable function to model of the receptive field of a population of neurons. The final method of estimating frequency selectivity using a pRF analysis, which fits a Gaussian function measured BOLD time course for each voxel. Despite fitting the same function to the same data, the results are not consistent on a macro scale. Narrow tuning widths dominate voxels surrounding Heschl's gyrus, with voxels being fitted model with the minimum parameter values allowed by the fitting procedure. This suggests these voxels have a poor fit, which is confirmed by the distribution of R^2 where the largest values and best fitting voxels are located on Heschl's gyrus and the surrounding areas have low R^2 values and poor fits. This spatial relationship that is consistent between general linear models and population receptive field models suggesting that the stimulus used elicit greater activity from primary regions than non-primary and aren't suitable for inferring the properties of cortical areas beyond the primary auditory cortex. However, at a micro-scale, within tonotopic gradients, there is variation in frequency selectivity across the cortex and as a function of preferred frequency that is clearer than from estimates derived from frequency response profiles. It is not clear if this variation is physiological in origin or is due to noise.

In summary, the analysis method employed does impact estimates of tonotopic properties but the importance of this will depend on a study's objectives and the required scale of measurements. Studies investigating properties on a micro-scale, within tonotopic gradients will see the greatest impact of analysis methods. The analysis method has the greatest impact on the distribution of preferred frequency and the possible frequency range of estimates. This distribution is heavily influenced by the

frequency range of stimuli, but also may be influenced by factors such as the distribution of stimuli within this range.

On a macro scale, looking at the overall properties of a cortical area will be more resilient to analysis methods. The spatial location of high- and low-frequency gradient reversals and regions of high-frequency selectivity on the cortex is consistent between methods and would have less impact on interpretations made from the delineation or summary statistics of cortical areas based on these functional properties.

This study also highlights the importance of stimulus presentation and the frequency properties of stimuli. The frequency range of stimuli will heavily impact estimates and the susceptibility to this will be dependent on the analysis employed. To test this, one could alter the frequency range of stimuli and compare the obtained estimates.

The presented methods appear to complement each other well, the shortcomings of one being the strength of the other. I would advise the use of both to take advantage of this for a holistic approach and view of the data. pRF estimates of preferred frequency appear to be the most representative of the true neurological properties of the cortex that a given voxel is sampling. However, voxelwise estimates of frequency selectivity appear noisy. ROI Average voxel frequency response profiles estimated from GLM beta weights provide insight into the selectivity of voxels with minimal assumptions.

It's not possible to determine the true functional properties of the cortex, due to the indirect nature of brain imaging methods and therefore it is hard to assess the accuracy of any estimate. The presented study, therefore, makes assumptions about the underlying functional properties of the auditory cortex. The first assumption is a near-uniform dis-

tribution of preferred frequency on an nERB scale, within the range of human hearing, within a cortical area. It is also assumed that the maximum amplitude in response to a preferred frequency is the same across the cortex and the reduction in activity elicited from frequencies away from this can be described by the same function. Therefore, frequencies presented at the same sensation level would create the same level of activity from a voxel with that preferred frequency to another frequency and a voxel with the same preferred frequency. To achieve this, all stimuli need to be presented at the same sensation level which has been attempted through threshold elevating noise, which allows stimuli to be presented at the same sound pressure level. This method was chosen over varying the sound pressure level to compensate for inter-individual or intervoxel differences in threshold to avoid the complications of the outer ears' nonlinear relationship with sound pressure. This has been implemented using broadband noise with a frequency profile of the average threshold of hearing. While on average this provides the flattest sensation level possible across frequencies, within individuals the sensation level of stimuli will vary. This could have been improved by measuring the sensation level of subjects and filtering the broadband threshold elevating to reflect this profile. This would create more accurate results but would have been very time-intensive for most likely a minor improvement with little impact on the results of the study.

3.4.2 Impact of scanner noise on cortical estimates of frequency preferred and selectivity

The impact of scanner noise has been assessed by comparing estimates preferred frequency and frequency selectivity from cortical responses to NBN noises, presented in the presence and absence of noise. The impact

of scanner noise can be seen in the average frequency response profile of voxels from the auditory cortex. The shape of the frequency response profile in the presence of scanner noise is considerably different from that of silence. In silence, the response profile is Gaussian-like in shape, with a peak shifted towards lower frequencies, 0.6 kHz, within the presented frequency range, a steeper slope towards lower frequencies and a gradual tail towards higher frequencies. In contrast, a bimodal distribution is observed in the presence of scanner noise, with a large high-frequency peak around 6kHz, and second mid-frequency peak around 1kHz.

The difference in the shape of the frequency response profile in the presence of noise and silence can be explained by the shape of the frequency profile of the scanner noise and auditory masking. The perception of a sound source in the presence of another sound source depends on the relationship between the frequency spectra of the two sound sources. At a single frequency, if one sound source is sufficiently louder than the other sound, only the louder sound will be heard and the quieter sound will be masked (Oxenham 2013). In addition, the more similar two sounds are in their frequency spectrum, the less energy is required for masking to occur. Masking can be partial, with a reduction in the perception of a target sound in the presence of masking noise, with the amount of attenuation depending on the relationship between acoustic properties of the sound sources and the attention of the listener (Alho et al. 2014; Pugh et al. 1996). Figure 3.8 depicts scanner noise, the dashed line, and the difference between the frequency response profile in silence and in the presence of the scanner noise, the solid line. The similarity between the peaks of the scanner noise and the frequency regions of greatest attenuation suggests that the scanning noise masks the auditory stimulus. However, the impact of masking isn't simply a larger reduction in cortical activity

to frequencies that experience the greatest masking.

It appears that frequencies experiencing less masking, where the scanner noise has less energy, elicit greater cortical BOLD response than when measured in silence and the average activation level within the auditory cortex is roughly equivalent in the presence of additional masking noise and in its absence. This could be due to cortical mechanisms, such as adaptation, compensating for the continuous masking noise and responding greater to the novel stimuli, which are modulated in time and frequency domain. However, it may not reflect neuronal properties and be due instead to the spread of BOLD activity across the cortex or even a change in the baseline activity of the cortex, which in turn will alter the baseline mean activity subtracted from the measured signal. It isn't possible to determine from this experiment but it suggests that the impact of scanner noise isn't simple and should be carefully considered in experimental design and analysis.

The effects of scanner noise are seen in estimates of preferred frequency. A shift in the distribution of preferred frequency estimates from Gaussian-like to bimodal-like can be observed, with a shift towards higher frequencies, reflecting what is seen in the frequency response profiles. The effect of acquisition protocols does not appear to be linear and the frequency-specific attenuation changes depending on voxels' preferred frequency. To further complications, the acoustic noise generated by a scanner depends on many factors, such as the parameters of the acquisition protocol and the model of scanner. The observed variation in estimates due to scanner noise would make comparing results between studies invalid or very difficult at best. Scanner noise can be minimised using hearing protection but this adds another level of complication, with variation in hearing protection complicating comparisons between studies. A greater issue is

comparisons between-subject. The actual attenuation will vary due to fit and this complicates comparison between subjects and group averages. It would be possible to measure the scanner noise in the ear canal of each subject in real-time which could be used to correct for biases in measurements. However, this would be costly in time, financially and computationally. In addition, the exact relationship between masking noise and cortical activity is unknown, limiting the ability to model and account for its impact. Furthermore, reducing the number of scans does not reduce reliability and pRF models fit the data better on average, which suggests the measurements of cortical responses using a sparse acquisition protocol are more consistent and less impacted by noise, compared to similar measurements using a continuous acquisition protocol. In addition, scanner noise will have impact on the perception of stimuli beyond simply changing the sensation level which cannot be accounted for.

The results are interpreted under the assumption that the frequency profile of the attenuation afforded by the hearing protection was flat. This will not be true but properly-fitted earplugs are likely to attenuate the range of frequencies of the presented stimuli by a uniform amount and I feel its a safe assumption to make that on average, across subjects, this would have minimum impact on our result. I also assume that the scanner noise was constant between sessions and the single measurement taken was representative. This is reasonable because the protocol was constant and any other variations would have been negligible. It would be a fairer test to keep the repetition time of data acquisition the same between conditions and present stimuli either at the same time as acquisition or in silence. However, this is not a logical experimental design and therefore wouldn't have been informative.

The results presented in this study demonstrates that the choice of ac-

quisition protocol can dramatically impact the estimation of tonotopic properties of the human auditory cortex and should be made with consideration. The observations made suggest that a sparse acquisition protocol and corresponding experimental design reduces the impact of scanner noise on measurements of cortical activity and estimates of tonotopic properties. The accuracy of these measurements is critical for the detection of hearing loss induced tonotopic reorganisation and therefore this should be considered in the design of such a study. Presenting stimuli in silence, using a sparse acquisition protocol, allows for lower stimulus sensation levels to be used, which may afford more accurate measurements of tonotopic properties (Langers and Dijk 2012). The study by Ghazaleh et al. (2017) which failed to detect tonotopic reorganisation in unilateral hearing loss subjects may have been confounded by the use of a continuous acquisition protocol and masking by, or other nonlinear effects from, scanner noise.

To achieve a more complete view of the human auditory cortex, multiple analysis methods should be employed to interpret the measured BOLD fMRI data. Each method has inherent strengths and weaknesses and can be used in conjunction to complement each other. Langers, Kleine, and Dijk (2012) employed a principal component analysis on the beta weights derived from a general linear model using stimulus frequencies as regressors. These estimates will most likely have been biased by the stimulus range as demonstrated by the results of the study presented here. This may explain why no tonotopic reorganisation was detected in subjects with tinnitus but audiometrically normal hearing, compared to normal hearing controls. Frequencies at the stimulus limits will be over-represented and the upper limit of the stimulus frequency range is most likely frequency region of hearing loss, and therefore reorganisa-

tion, to occur (Norena et al. 2002; Sereda, Edmondson-Jones, and Hall 2015). Therefore, any reorganisation would most likely be hard to distinguish due to the bias of the estimation method. In conclusion, to increase chances of detecting hearing loss induced tonotopic reorganisation, a sparse acquisition protocol and multiple, complementary analysis approaches should be implemented.

Chapter 4

Cortical magnification of the human auditory cortex

4.1 Introduction

The auditory cortex is tonotopically organised, with neighbouring parts of the cortex responding to neighbouring frequencies. A topographic, map-like organisation of sensory characteristics is common for sensory cortical areas. One particularly interesting aspect of this mapping, which has been shown in other sensory systems, is that the cortical surface allocated to a given division of sensory space is proportional to its behavioural importance. For example, visual studies have demonstrated that proportionally more cortical surface area is dedicated to the centre of the visual field, corresponding to the fovea in the retina, than to the periphery (Schira, Tyler, et al. 2009).

The mathematical description that relates sensory space to the area taken up on the cortical surface is often called the *cortical magnification func-*

tion. It has been well characterised in visual cortex, initially with methods such as 2-deoxy-d-glucose (2-DG) and post-mortem reconstructions in animals (Daniel and Whitteridge 1961; Tootell et al. 1988). In vivo measurements in animals and humans using fMRI are now routinely done for studies of visual cortex (Dumoulin and Wandell 2008). In contrast, the corresponding mapping for auditory frequencies on the cortical surface - and the *tonotopic magnification function* in the human auditory cortex is still unknown. One of the main hypotheses of the work presented here is that it should follow behavioural relevance, as is the case in the visual system.

In the visual modality, where retinotopic maps have been widely reported and extensively studied (and for a general review see (Wandell, Dumoulin, and Brewer 2007)), visual cortical magnification has been shown to correlate with behavioural measures of acuity (Duncan, Boynton, and Jolla 2003). Similarly, in the somatosensory domain, Duncan and Boynton (2007) found within-subject measures of tactile hyperacuity correlated to cortical magnification in primary sensory cortex. In line with this result, Tamè et al. (2019) found that the behaviorally determined shape of "skin" space matched its cortical representation in both primary somatosensory and motor cortices.

The mathematical framework for mapping sensory to cortical space (and vice versa) has been well developed for the visual domain. Schira, Wade, and Tyler (2007), for example, compared and contrast different models of the cortical magnification for the intact primary visual cortex, V1. Quantitative descriptions such as this, allow researchers to assess changes in pathological situations. For example, Baseler et al. (2011) investigated retinotopic reorganisation of the human visual cortex following localised deprivation of the peripheral visual/sensory apparatus. Clearly, under-

standing the retinotopic organisation and cortical magnification of the human visual cortex is critical for assessing changes in disease (Wandell and Smirnakis 2009). Likewise, for auditory research, understanding tonotopic magnification in healthy participants, will benefit studies that aim to quantify tonotopic reorganisation. In particular, estimates of the normal range of cortical magnification functions could provide a global measure and characterisation of cortical organisation of the auditory cortex that could then help quantify tonotopic reorganisation in hearing loss subjects.

Tonotopic magnification has not yet been systematically investigated, so providing a set of robust measurements is one of the key aims of the work in this thesis chapter. It is plausible that (but not yet known if) tonotopic magnification is based on extrapolation of cochlear mapping and is a direct consequence of the cochlear magnification function.

Cochlear frequency filters correspond to the range of frequencies that single inner hair cell in the cochlea responds to, which perceptually result in the range of frequencies within which two tones interfere with the perception of each other. The width of the cochlear frequency filters increases in approximate proportion with the filter frequency, see figure 4.1. This means that the relative filter width varies monotonically and is approximately constant above 1 kHz (Glasberg and Moore 1990), see figure 4.1 Panel B.

However, it has been shown in *visual* studies that cortical magnification correlates with Vernier acuity, the ability to detect a small offset between two parallel line segments (Westheimer 1975). Vernier acuity is a hyperacuity - the ability to detect changes as finer resolution than the sampling of sensory space by the sensory periphery (Geisler 1984). One can

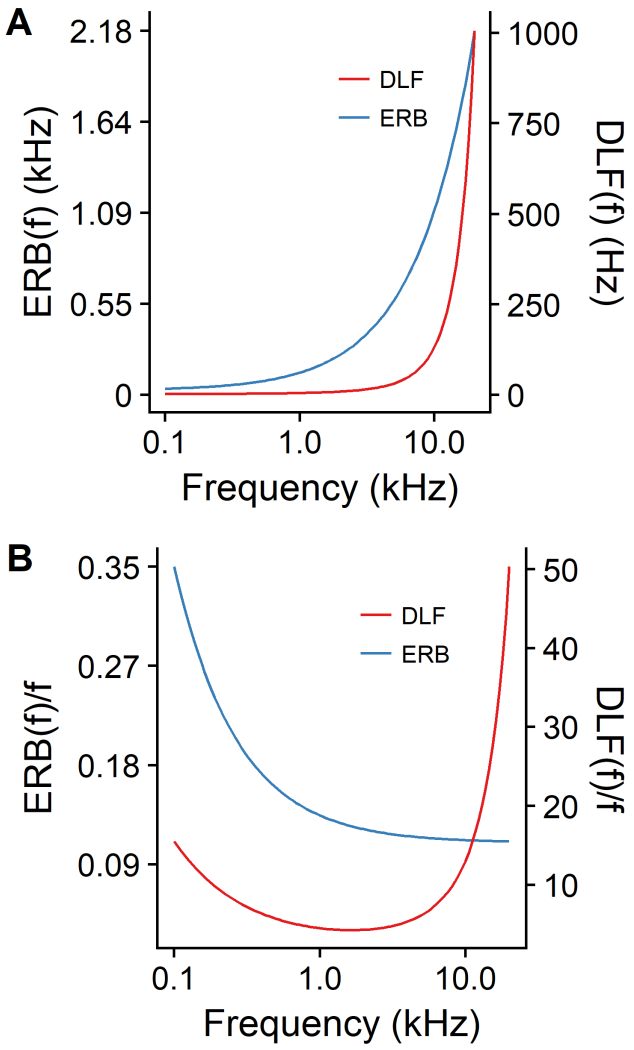


Figure 4.1: ERB and DLF as a proportion of frequency

A) Cochlear filter widths (ERB) and frequency discrimination thresholds (DLF) as a function of frequency (kHz) on a logarithmic frequency scale. B) Cochlear filter widths and frequency discrimination as a proportion of frequency (ERB/f and DLF/f respectively) as a function of frequency (kHz) on a logarithmic frequency scale. The width of the cochlear frequency filters increases in approximate proportion with the filter frequency, so that the relative filter width varies monotonically and is approximately constant above 1 kHz. In contrast to cochlear frequency resolution, the frequency discrimination threshold, as a proportion to frequency, varies strongly and non-monotonically with frequency, showing a broad minimum around 2 kHz and increasing sharply above 4 kHz.

think of *frequency discrimination* as the auditory equivalent of this visual task. The frequency discrimination threshold is defined as the minimal discriminable frequency difference (increment or decrement) as a proportion of the baseline frequency. In contrast to cochlear frequency resolution, the frequency discrimination threshold, as a proportion to frequency, varies strongly and non-monotonically with frequency, showing a broad minimum around 2 kHz and increasing sharply above 4 kHz (Wier, Jesteadt, and Green 1977; Moore and Glasberg 1987; Nelson, Stanton, and Freyman 1983), see figure 4.1 Panel B.

One particularly interesting neuroimaging result from the visual domain relates to the population receptive field size 2.3.2 and cortical magnification. The so-called *point image*, the cortical area corresponding to any given population receptive field, has been shown to be near-constant across visual space (Harvey and Dumoulin 2011). This arises from an inverse relationship between cortical magnification and population tuning width in the visual cortex, i.e. the greater the magnification, the narrower the tuning width. It is unknown if the same holds for auditory cortex: to what extent is the same amount of cortex dedicated to a particular range of frequencies across the range of human hearing.

In humans, cochlear frequency tuning and frequency discrimination are measured through psychophysical experiments. Cochlear frequency tuning is measured in equivalent rectangular bandwidths (ERB) and estimated through the detection of pure tones in *notched-noise* (Glasberg and Moore 1990). To create such *notched-noise* stimuli, a range - notch - of frequencies is removed from broadband noise. In a typical experiment, the threshold of the pure tone is measured as a function of the width of the notch (Patterson 1976).

Frequency discrimination, the ability to tell apart two frequencies, is usually measured as minimum frequency change detectable from a reference frequency. A typical experiment uses a two- or four- interval forced-choice adaptive procedure to estimate 71% correct performance (Levitt 1971; Wier, Jesteadt, and Green 1977; Nelson, Stanton, and Freyman 1983). In the literature, the thresholds are usually referred to as difference limens for frequency (DLF) (Wier, Jesteadt, and Green 1977; Moore and Glasberg 1987; Nelson, Stanton, and Freyman 1983). Together with the cochlear frequency tuning, these behaviourally-derived functions are used in the presented study to determine the relationship between cortical magnification and frequency acuity (Glasberg and Moore 1990; Nelson, Stanton, and Freyman 1983).

The present study aims to measure tonotopic cortical magnification in normal-hearing subjects and determine if auditory cortical magnification function reflects frequency resolution, ERB, or frequency discrimination, DLF. Here we assume that the extent of cortex devoted to a certain frequency is proportional to perceptual abilities around this frequency. This means that with greater behavioural performance at a given frequency (i.e. greater frequency resolution or lower frequency discrimination thresholds), the cortical area should increase and therefore cortical distance per unit frequency along the tonotopic gradient should increase too.

The auditory cortical magnification function was estimated by characterising the relationship between voxelwise preferred frequency in the two mirror-symmetric tonotopic gradients located on Heschl's gyrus in each hemisphere of 8 normal-hearing subjects and the voxelwise cortical distance along these gradients. Measured group cortical magnification functions were then fitted with models derived from two behaviorally-derived

functions, cochlear frequency tuning and frequency discrimination thresholds, to determine which of the two cortical magnification is most likely to reflect. The auditory point image was derived from estimates of voxel-wise cortical magnification and population tuning width.

4.2 Methods

We estimated the cortical magnification function of 16 hemispheres in 8 normal-hearing subjects (4 males; mean age = 29.38 SD = 7.34).

4.2.1 Dataset

The analysis described in this chapter extends that of the previous chapter and makes use of a subset of the same high-resolution 7T data set. To review briefly:

- data was acquired with the sparse acquisition protocol (TR=7.5s), allowing for stimulus presentation in the silence between data acquisitions.
- BOLD responses were measured while subjects were presented with narrowband noise (NBN) stimuli.
- stimulus details: amplitude modulated narrowband noise (NBN) stimuli —centred at thirty-two different frequencies (range = 0.1 - 8 kHz), binaural presentation, pseudo-random order.
- scanner: Philips Achieva 7T MRI system using a volume transmit coil and a 32-channel receiver coil. Two functional data sets were collected for each subject.

- in addition to the functional data, a whole-head structural dataset and auxiliary datasets to aid data preprocessing were obtained for each participant.
- all datasets were collected within a single session, apart from subject 7, for whom functional data were collected in two separate sessions.

4.2.2 Data Analysis

The preprocessing pipeline was the same as described in the previous chapter, Chapter 3, and the General Methods. The preferred frequency and the tuning of frequency selectivity of voxels were estimated using population receptive field modelling (pRF). Subsequent data analysis to produce across-subject statistics was restricted to regions of interest capturing the AC gradients as described previously, with the notable change that were split ROIs by the low-frequency gradient reversal to create two mirror-symmetric tonotopic gradients per hemisphere.

The cortical distance between gradient reversals was measured at each voxel within a tonotopic gradient ROI by measuring the shortest distance on the cortical surface using Dijkstra's algorithm (as implemented in mrTools (Gardner et al. 2018b)). Specifically, we measured the shortest distance from its vertex coordinate in 3D to the nearest vertex coordinate of the low and high-frequency gradient reversals defining the ROI. Each voxel (and associated measures of preferred frequency) in the dataset were assigned the vertex coordinate of the closest vertex on the mesh model of the cortical surface, which was derived with Freesurfer (http://surfer.nmr.mgh.harvard.edu; (Dale 1999)).

To combine across multiple hemispheres and subjects, we also normalised

the cortical distance measurements. Each path was normalised by the sum of the Dijkstra distances to both gradient reversals, i.e. the shortest cortical distance between reversals going this through this particular voxel. The normalised distance was then multiplied by the average greatest distance across subjects to give a subject average cortical distance in millimetres.

Prefered frequency transformation functions

We evaluated the effect of transforming the referred frequency estimates using four different functions:

1. linear (no transformation)
2. logarithmic (base 10)
3. cochlear filter resolution, ERB (calculated from behavioural measurements of cochlear filter resolution collected by Glasberg and Moore (1990)
4. frequency difference limen, DLF (calculated from behavioural measurements of frequency discrimination determined by Nelson, Stanton, and Freyman (1983)

To predict the cortical magnification function from the behavioural functions, the integral of the reciprocal of each function, see figure 4.2, from 0 to the desired frequency in kHz was calculated. The reciprocal of each function was taken to compute the predicted number of behavioural units per unit frequency (which would correspond to the local cortical surface dedicated to local frequencies). The integral of this was taken to predict the cortical distance relative to an arbitrary low-frequency starting point. See below for details.

Cochlear frequency tuning, ERB.

$$ERB(f_{kHz}) = A(Bf_{kHz} + 1) \tag{4.1}$$

The integral of the reciprocal of ERB (Glasberg and Moore 1990).

$$nERB(f_{kHz}) = \int_0^f \frac{1}{ERB(f)}\delta f = \frac{1}{(AB)\log(Bf_{kHz} + 1)} \tag{4.2}$$

Where $A = 24.7/1000$ and $B = 4.37$.

Difference limens for frequency, DLF.

$$DLF(f_{Hz})) = 10^{a\sqrt{(}f_{Hz})+k+m(\frac{1}{SL})} \tag{4.3}$$

The integral of the reciprocal of DLF.

$$nDLF(f_{Hz}) = \int_0^f \frac{1}{DLF(f)}\delta f \tag{4.4}$$

$$= -\frac{(a\sqrt{(}f_{Hz})log(10) + 1)2^{a(-\sqrt{(}f_{Hz}))-k-m(1/SL)+1}5^{a(-\sqrt{(}f_{Hz}))-k-m(1/SL)}}{(a^2(log(10))^2)} \tag{4.5}$$

Where $k = -0.15$, $m = 5.056$, $a = 0.0214$ and $SL = 45$ (SL of stimuli used in experiment).

Cortical distance as a function of preferred frequency

The cortical magnification function describes the change in cortical distance along a tonotopic gradient as a function of preferred frequency. This relationship was modelled as a linear relationship (using `lm`, in `R`

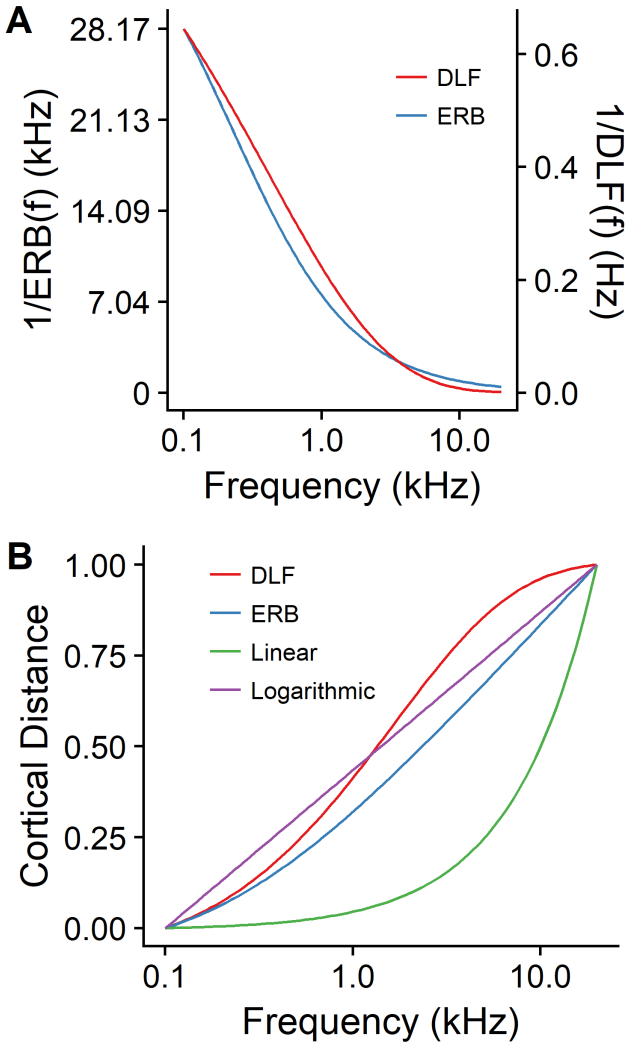


Figure 4.2: Mapping of frequency to cortical distance

A) The reciprocal of cochlear filter widths (ERB) and frequency discrimination (DLF). B) Linear, logarithmic, cochlear filter widths (ERB) and frequency discrimination (DLF) mapping of frequency to cortical distance. Preferred frequency estimates were transformed by each function before fitting a linear model. The transformation with the best fitting linear model is proposed to be the most likely candidate for the cortical magnification function of the human auditory cortex.

(R Core Team 2013)) - after first transforming preferred frequency estimates using each of the four different functions described above in turn. After transformation of the frequency axis, cortical distance values were binned into 32 equally spaced bins across the transformed frequency and then averaged (arithmetic mean). Finally, the results were averaged across participants and plotted (arithmetic mean +/- standard error of the mean). One of our working hypotheses was that the transformation which resulted in the best fitting linear model is the most likely candidate for the cortical magnification function of the human auditory cortex.

Tonotopic cortical magnification factor

The tonotopic cortical magnification factor, CMF, captures the change in cortical distance along a tonotopic gradient per a constant change in preferred frequency. Data was combined across all tonotopic gradients and estimated for each of thirty-two bins of preferred frequency spaced equally on the transformed frequency scale that produced the best linear fit with cortical distance, the nDLF scale. The cortical magnification factor was calculated by taking the derivative of cortical distance with respect to frequency in order to measure the change in distance per unit frequency (nDLF). A line was fitted to the data using a linear model in `R` (`lm` function).

Frequency selectivity as a function of preferred frequency

Frequency selectivity was derived from the voxel-wise estimates of the pRF analysis, using the standard deviation, σ of the Gaussian in the pRF model (see Section 2.3.2 in the General Methods for details). The average voxel tuning width as a function of frequency was calculated at 32 equally-spaced preferred frequencies on an nDLF scale. For each bin, the

(arithmetic) mean was calculated of the estimates. Data from all voxels from the AC gradients ROIs of all hemispheres were included. To provide a summary fit, a second-order polynomial was fitted to these data (`R` function, `lm`). A quadratic function was empirically found to give the best description.

Point image

The average point image was calculated as a function of preferred frequency (nDLF scale). Following Harvey and Dumoulin (2011), averaged estimates of tuning width were divided by the average cortical magnification factor of each preferred frequency bin. To provide a summary fit, a second-order polynomial was fitted to these data. Again, a quadratic function was empirically found to give the best description.

4.3 Results

4.3.1 Change in cortical distance as a function of frequency

The magnification of frequency representation in the human auditory cortex, tonotopic magnification, was estimated from eight normal-hearing subjects. The average change in cortical distance along anterior and posterior tonotopic gradients located on HG as a function of average preferred frequency was assessed.

Estimates of preferred frequency were transformed by four magnification functions; linear, logarithmic, cochlear filter widths and frequency discrimination. The resulting cortical magnification functions were fitted with a linear model. The goodness of fit value was compared in order

to compare the adequacy of these different cortical magnification functions to describe the cortical magnification function actually derived from the data. Preferred frequency transformed by a frequency discrimination function produced the most linear relationship between cortical distance and preferred frequency, see figure 4.3, with a root mean square error (RMSE) of 0.48. Frequency on log scale produced a relatively poor fit (RMSE = 1.03), and on a linear scale produced the poorest fit (RMSE = 1.47). Prefered frequency transformed by a cochlear filter function processes a better fit than a logarithmic transformation (RMSE = 0.74) but is outperformed by a frequency discrimination transformation for higher frequencies (> 1 kHz), see figure 4.3, where the two functions deviate, see figure 4.2. Variance in preferred frequency, plotted as the standard error in figure 4.3, was consistent across the frequency spectrum suggesting the average cortical distance of each bin was consistently representative of the voxels.

4.3.2 Cortical magnification factor of the auditory cortex

The average cortical tonotopic magnification factor of the human auditory cortex (CMF) was estimated as a function of preferred frequency. A second-order polynomial was fitted to the data, resulting in linear function with a negative slope. CMF is decreased with increasing preferred frequency from roughly 5 mm/kHz to 1 mm/kHz. Variability of estimated CMF, in terms of standard error, was consistent across preferred frequency.

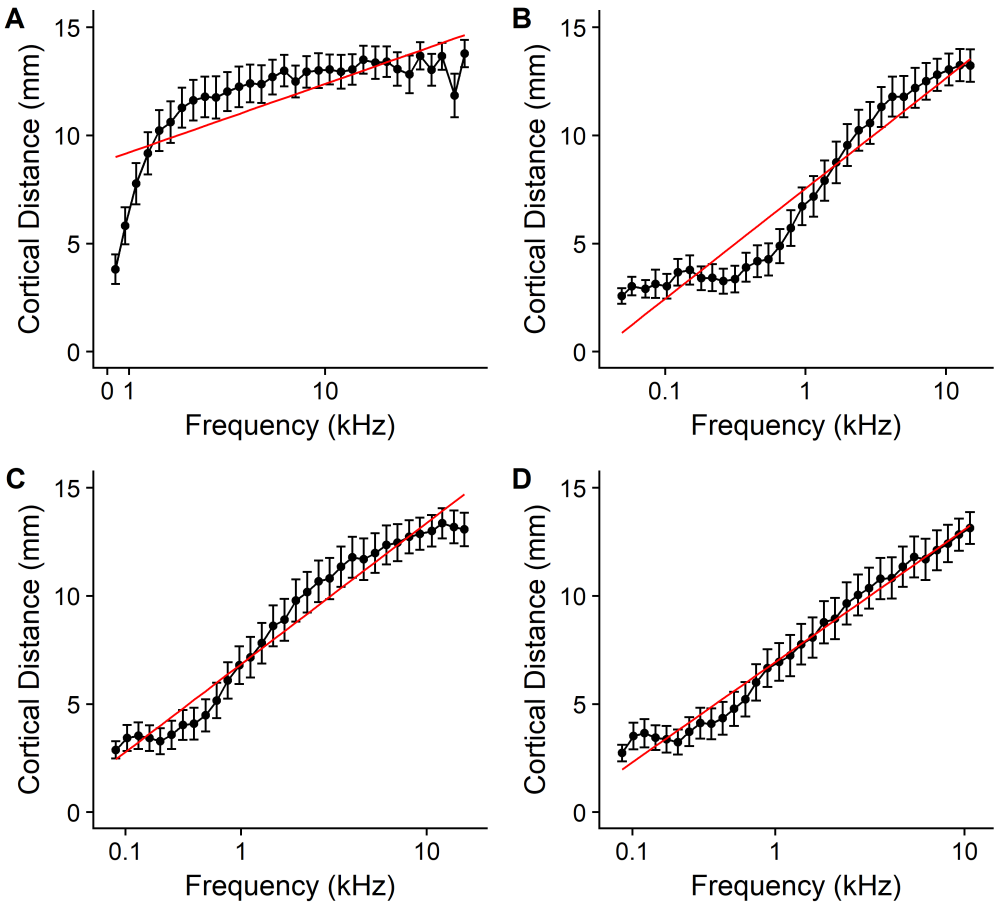


Figure 4.3: Change in average cortical distance as a function of frequency

Frequency bins were spaced equally on four different scales, A) linear, B) logarithmic, C) cochlear filter widths (ERB) and D) frequency discrimination limens (DLF). Connected black points represent the mean normalised cortical distance at a given preferred frequency bin (n=32) from all gradients of all hemispheres and vertical lines with bars represent the standard error. The red line represents a fitted line of best fit used to assess the best fitting spacing of frequency. Transforming frequency by the frequency discrimination limen scale produces the most linear relationship between cortical distance and frequency, and therefore is proposed to be the function that describes the cortical magnification of the human auditory cortex.

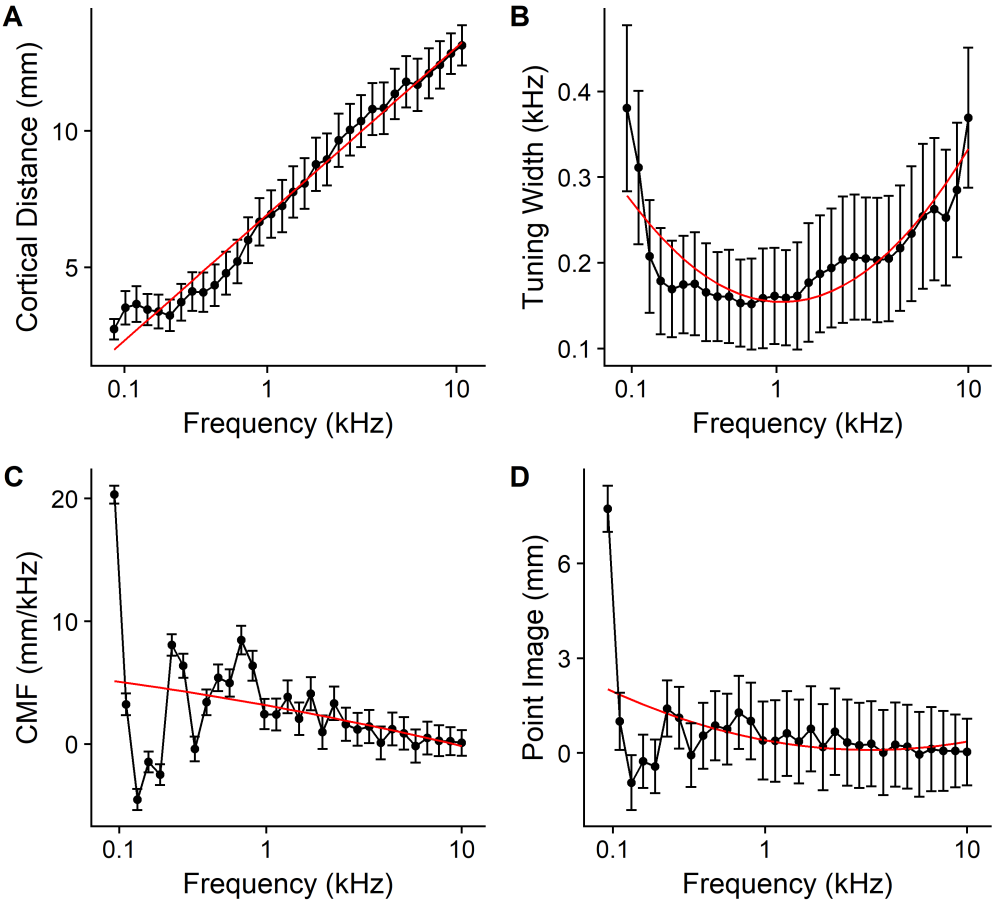


Figure 4.4: Average cortical distance, tuning width, cortical magnification factor and point image as a function of frequency

A) Average cortical distance as a function of frequency on an nDLF scale. B) Average pRF tuning width as a function of frequency. C) Average cortical magnification factor (the derivative of cortical distance with respect to frequency) as a function of frequency on an nDLF scale. D). Point image (mm) as a function of preferred frequency on an nDLF scale. Connected black points represent the mean value at a given preferred frequency bin (n=32) from all gradients of all subjects, vertical bars represent the standard error and red lines represent fitted functions to the data. Cortical magnification appears to decrease with frequency, see panel C. Tuning width appears to be narrowest for mid-frequencies (0.5-2kHz). The product of these two measures, the cortical area dedicated to a point in frequency space (referred to as point image in visual literature), appears to decrease with frequency to become roughly constant with frequency above 1kHz.

4.3.3 Point image in the auditory cortex

The average point image was calculated by dividing the estimates of average tuning width by the average cortical magnification factor of each preferred frequency bin. Tuning width is sharpest for preferred frequencies between 0.4 and 1kHz, see panel B in figure 4.4, and broadens towards the lower and upper limits of the human audible range. A second-order polynomial was fitted to the data, producing a concave curve with a minimum tuning width value at a preferred frequency around 1 kHz.

Point image appears to decrease as a function of preferred frequency but be roughly constant above 0.5 kHz. The second-order polynomial fitted to the data suggests point image decreases from roughly more than 2 mm to less than 1 mm.

4.3.4 Differences between anterior and posterior gradients

No meaningful differences could be ascertained between anterior and posterior gradients, from independently derived measures of cortical distance as a function of preferred frequency (see figure 4.5), distribution of preferred frequency (see figure 4.6), average voxel frequency response profiles (see figure 4.7) and tuning width, cortical magnification factor and point image as a function of preferred frequency.

4.4 Discussion

The presented study has derived the average cortical tonotopic magnification functions in two tonotopic gradients of human auditory cortex from estimates of preferred frequency and cortical distance, and in turn

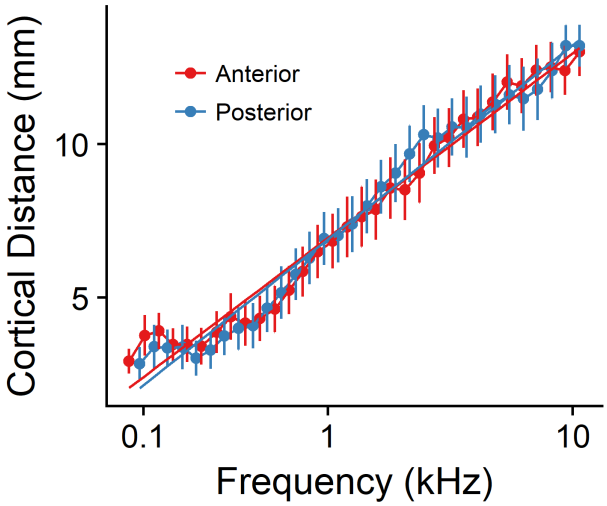


Figure 4.5: Cortical distance as a function of preferred frequency in anterior and posterior AC gradients

Cortical distance as a function of preferred frequency on a nDLF scale for anterior and posterior gradients averaged across all subjects hemispheres. Points are the mean cortical distances at a given preferred frequency are vertical lines are the standard error. Linear functions were fitted to estimates from each hemisphere, the solid lines. No significant differences between anterior and posterior gradients can be seen.

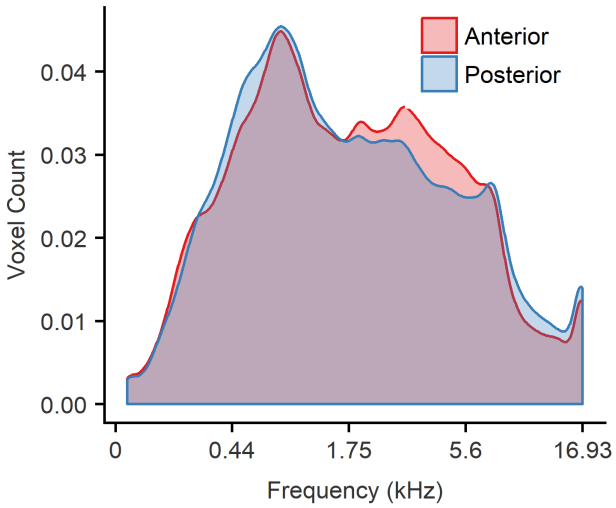


Figure 4.6: Distribution of preferred frequency in anterior and posterior AC gradients

Distribution of preferred frequency on a nDLF scale from anterior and posterior AC gradients from all subjects hemispheres. No significant differences between anterior and posterior gradients can be seen.

point image from estimates of the cortical magnification factor and tuning width. As expected, the linear transformation was a much poorer fit than any of the other functions. All other 3 functions fitted reasonably well, meaning that the frequency range represented by a given cortical area increases more or less linearly along the tonotopic gradient, see figure 4.3 panel A. This is also expected given what is know about the general distribution of preferred frequencies along the cochlea (Glasberg and Moore 1990). The cortical magnification function predicted from the behavioural frequency discrimination threshold function is the best-fitting model for the observed cortical magnification function. This suggests that the cortical magnification function is related to frequency discrimination. The amount of cortex dedicated to different frequencies, as predicted by the discrimination function under the assumption that behavioural performance is directly proportional to cortical surface, seems

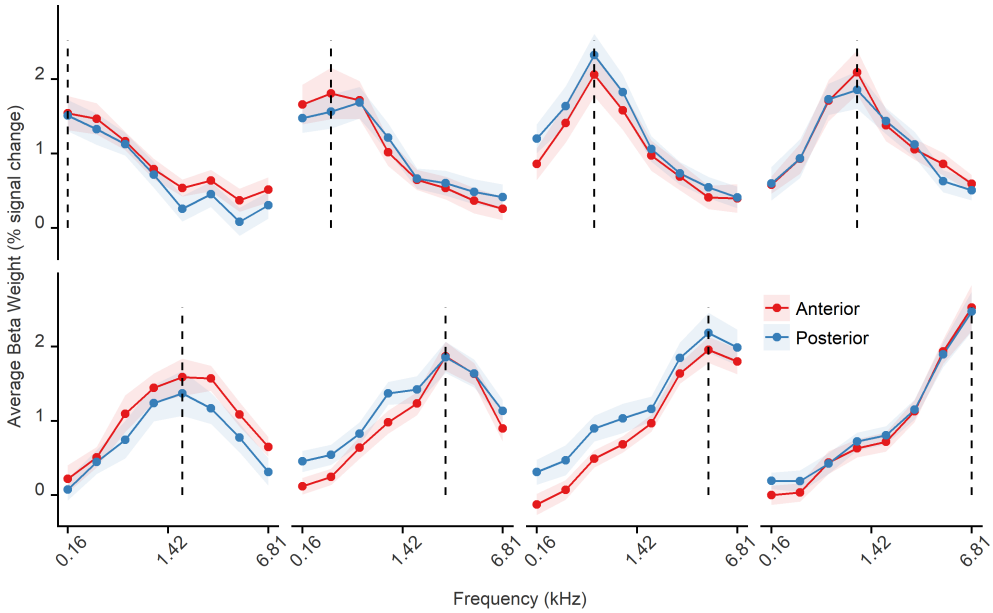


Figure 4.7: Average voxel frequency response profiles of the anterior and posterior AC gradients

Average frequency response profiles were derived by averaging voxelwise response profiles after grouping AC gradient ROI voxels in 8 preferred frequency bins. Voxelwise preferred frequencies and resulting average frequency response profiles were derived using a split-half method to avoid artificially narrow frequency response profiles (for details, see General Methods). The frequency profiles estimated from continuous data are less consistent in shape that the sparse data. Points connected by solid lines represent the average beta weight associated with each stimulus frequency bin across all hemispheres from all subjects (n =16) and the ribbon represent the corresponding standard error. No significant differences between anterior and posterior gradients can be seen.

to correspond to the amount of cortex that is actually observed to respond to these frequencies, see figure 4.4.

These results suggest that frequency discrimination may be determined by cortical representation of frequency in the primary auditory cortex. This agrees with the findings from both visual and somatosensory modalities, where the cortical magnification of primary sensory areas has been found to correlate with within-subject measures of Vernier acuity and tactile hyperacuity respectively (Duncan, Boynton, and Jolla 2003; Duncan and Boynton 2007). These studies compared the cortical magnification function to the behavioural functions of individual subjects, rather than the generic behavioural functions used in the presented study. Looking at within-subject correlations between auditory cortical (tonotopic) magnification and frequency discrimination thresholds would provide a stronger experimental test of this relationship.

The tonotopic cortical magnification factor, the cortical surface distance between two points representing a constant change in frequency, decreases as a function of preferred frequency, see figure 4.4 panel C. The linearity of the relationship between preferred frequency and cortical distance breaks down below roughly 0.4 kHz, which makes the cortical magnification at low frequencies is difficult to interpret. Cortical frequency tuning widths are narrowest at the centre region of the human audible range (0.4-1 kHz) see figure 4.4 panel B, suggesting that frequency selectivity is greatest for frequencies within the range of human speech (Everest and Pohlmann 2009d). Point image is the product of cortical magnification and tuning width. The near inverse relationship between these, see figure 4.4 panel D, leads to point image being roughly constant across frequency above preferred frequencies of 0.4 kHz, akin to the constant point image found in the primary visual cortex (Harvey and Dumoulin 2011). Point

image is greater for low frequencies it isn't clear is this is a reflection of cortical properties or inadequacy in the estimation of low prefered frequencies.

No meaningful differences were ascertained between anterior and posterior tonotopic gradients in terms of preferred frequency distribution (figure 4.6), frequency response profiles (figure 4.7, cortical distance as a function of preferred frequency (figure 4.5), and, therefore, cortical magnification and point image. It may be that further parcellation of the auditory cortex is required. In the non-human primates, tonotopic gradient reversals do not deliminate all boundaries between auditory cortical areas and further functional and structures are needed to fully parcellate the auditory core (Kaas and Hackett 2000). It is possible that more than one distinct cortical area is present in a single tonotopic gradient and therefore using only gradient reversals to delineate cortical areas will result in a number of distinct cortical regions being selected together. This is unavoidable until methods are established to further delineate the auditory cortex. It is reasonable to assume that cortical areas specialising in different tasks would have different cortical magnification factors, for example, the frequency range important for effective intelligibility of human speech, are different to those important for optimum auditory spatial localisation (Everest and Pohlmann 2009d; Moore 2012).

Chapter 5

How can we reliably detect hearing loss induced cortical tonotopic reorganisation in humans?

5.1 Introduction

Peripheral hearing loss occurs following damage or degradation of hair cells in the inner ear, which convert the mechanical movement of the basilar membrane into neuronal activity. Loss or damage of hair cells reduces the sensitivity of the ear, requiring sounds to be louder before perception by the listener. Peripheral hearing loss not only results in reduced sensitivity to quiet sounds, but also leads to significant problems in supra-threshold hearing, such as difficulty with understanding speech in noise (Dubno, Dirks, and Morgan 1984), and is often associated with tinnitus (Norena et al. 2002), the phantom perception of sound in the ab-

sence of an external sound source (Baguley 2002). Animal research suggests that such impairments could be a consequence of central processes in response to reduced peripheral output (Eggermont and Roberts 2004). In animals, peripheral hearing loss has been shown to cause tonotopic reorganisation in auditory cortex, where neurons that would normally respond to frequencies within the hearing loss, usually higher frequencies, shift their tuning towards lower, and therefore less affected, frequencies (Noreña and Farley 2013). It has been suggested that this cortical tonotopic reorganisation and over-representation of frequencies around the hearing loss edge, leads to tinnitus (see Eggermont and Roberts (2014) for a review). Determining if similar tonotopic reorganisation occurs in humans with hearing impairment would be of great clinical benefit, possibly opening avenues to treatment.

Sereda, Hall, et al. (2011) found that the pitch of tinnitus percept is generally within the area of hearing loss, supporting the reorganisation model of tinnitus. In some patients with steeply sloping high-frequency hearing loss, where the transition between normal hearing and hearing loss is abrupt, the tinnitus pitch is located near the edge of the hearing loss region and has a narrow bandwidth, suggesting that the properties and shape of the hearing loss profile are important (Sereda, Hall, et al. 2011; Sereda, Edmondson-Jones, and Hall 2015). The tonotopic reorganisation model of tinnitus suggests that the tinnitus pitch arises at the edge of the hearing loss region as a result of over-representation of neurons tuned to the frequency of transition between normal hearing and hearing loss (Eggermont and Roberts 2004).

There is indirect evidence of possible tonotopic reorganisation from human behavioural measures of frequency discrimination. Enhanced frequency discrimination has been found near the edge of hearing loss for

individuals with steeply sloping hearing loss (Moore and Vinay 2009; Thai-Van et al. 2003), in contrast to the degraded frequency selectivity often associated with sensorineural hearing loss (Glasberg and Moore 1986). These changes fit with the reorganisation model of chronic tinnitus, as an increased number of neurons responding to hearing loss edge frequencies would be expected to result in an increased ability to discriminate these frequencies.

Measuring tonotopic maps is sensitive to changes in sensation level, which is a problem when comparing tonotopic organization between hearing-impaired and normal-hearing subjects. The sensation level of auditory stimuli presented to a subject with hearing loss is decreased at the hearing loss frequencies and this will decrease the amplitude of the cortical response at these frequencies compared to unaffected frequencies. This, in turn, will modify the apparent tonotopic organization, even in the absence of true tonotopic reorganisation. If not accounted for it will be impossible to determine if differences detected in tonotopic maps from normal and hearing loss subjects are due to changes in sensation level or genuine cortical tonotopic reorganisation.

Figure 5.1 depicts how high-frequency hearing loss may bias voxel estimates of preferred frequency, the frequency that elicits the greatest response, and tuning width, the range of frequencies that elicit a response above baseline brain activity. In a normal hearing condition, the frequency that elicits the largest response will be attributed to the preferred frequency of the voxel. Hearing loss will reduce the energy presented to the central auditory system from stimuli within the hearing loss frequency range, but not outside of it. If this is not accounted for in the analysis, the predicted preferred frequency and tuning width will be biased. For example, in high-frequency hearing loss, responses to frequen-

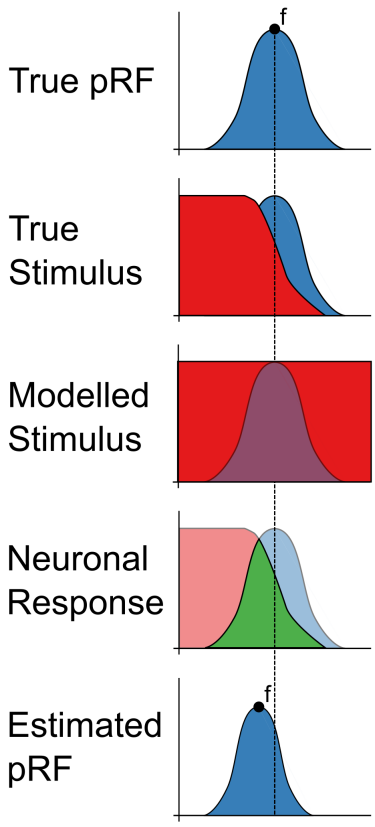


Figure 5.1: Impact of high-frequency hearing loss on estimates of pRF properties

From the top to bottom panel: *True pRF* - The first panel depicts the true population receptive field as a blue Gaussian function. *True stimulus* - The true stimulus sensation level for a subject with hearing loss is depicted in the second panel in red. *Modelled Stimulus* - The profile of modelled stimulus by a tradition analysis if the hearing loss isn't accounted for. *Neuronal Response* - The true stimulus will elicit an activity profile, represented by green, that is the product of the true pRF and the true stimulus (the first two panels). *Estimated pRF* - Fitting a model which doesn't take into account the true stimulus will assign a pRF that is lower in frequency and narrower in tuning width, rather than the true pRF.

cies within the hearing loss region will be reduced and estimated tonotopic maps will be biased toward lower frequency. This systematic bias may erroneously be attributed to changes in cortical properties and cortical tonotopic reorganisation, rather than a measurement bias arising from variation in stimulus sensation level due to peripheral hearing loss.

Studies have attempted to mitigate this bias in the study of tonotopic reorganisation by recruiting patients with tinnitus but clinically normal hearing, in either both (Langers, Kleine, and Dijk 2012) or one ear (Ghazaleh et al. 2017). Despite using high-resolution fMRI, Langers, Kleine, and Dijk (2012) found no evidence of over-representation of frequencies matching subjects' tinnitus pitches, in comparison to normal-hearing controls without tinnitus. However, the lack of evidence for tonotopic reorganisation in the tinnitus subjects recruited for this study isn't conclusive evidence that tonotopic reorganisation isn't a possible mechanism in tinnitus generation or a possible side effect of hearing loss. Tinnitus is a heterogeneous auditory deficit (Baguley 2002) and there may be several mechanisms that lead to its generation. The patients recruited in the study by Langers, Kleine, and Dijk (2012) may have had tinnitus due to another tinnitus generation mechanism or the reorganisation may have been too subtle to detect using the experimental procedure.

Another high-resolution fMRI study, by Ghazaleh et al. (2017), recruited patients with chronic tinnitus and unilateral hearing loss. Auditory stimuli presented to the normal hearing and matched in perceived loudness within-subjects (80 phons), but not across. This allowed them to present auditory stimuli and image the auditory cortex of both hemispheres without the confounding effects of hearing loss but at different sound pressure level between subjects. This study found a greater number of voxels responsive to low frequencies in the patients, suggesting potential

reorganisation. However, this study employed a continuous scanning protocol. Patients would have heard less of the scanner noise, due to hearing loss, than the normal hearing controls, and therefore the sensation level of stimuli would be different. As a consequence of this, cortical responses would have been enhanced in the scanner noise frequency range, which matches well with the frequency range where Ghazaleh et al. (2017) found the enhancement and brings the findings into question.

Wolak et al. (2017) recruited patients with symmetrical steeply sloping high-frequency hearing loss but did not report the prevalence of tinnitus within the group to assess tonotopic reorganisation. Steeply sloping hearing loss has been suggested to lead to tonotopic reorganisation in animal studies, making the patients ideal candidates (Irvine, Rajan, and Mcdermott 2000; Sereda, Hall, et al. 2011). The patients were found to have reduced cortical responses to high-frequency stimuli and increased low-frequency cortical representation compared to normal hearing controls. Wolak et al. (2017) took this to suggest evidence of tonotopic reorganisation but as discussed previously, the confounding effects of peripheral hearing loss would account for these results without cortical reorganisation.

To study this "ideal" patient population, we propose to simulate hearing loss in normal-hearing controls to match the hearing loss profiles of the patients. Peripheral deficit simulation has been successfully used in visual fMRI to study patient populations (Baseler et al. 2011), akin to hearing loss simulation in the auditory domain (Desloge et al. 2012). Furthermore, we suggest employing a method to minimise biases in the estimation of tonotopic properties caused by variation in sensation level. Binda et al. (2013) demonstrated that modelling a visual peripheral deficit simulation in the definition of the population receptive field model stimulus

description could elevate the bias resulting from the deficit (Binda et al. 2013).

This chapter presents the results of mapping the tonotopic organisation of auditory cortex in normal-hearing subjects, in the presence and absence of simulated steeply sloping high-frequency hearing loss. Patients with steeply sloping high-frequency hearing loss are likely candidates for the tonotopic organisation model of tinnitus (Sereda, Hall, et al. 2011; Sereda, Edmondson-Jones, and Hall 2015), making them ideal subjects for an fMRI tonotopic mapping study of tonotopic organisation. Tonotopic organisation was estimated from BOLD-fMRI responses at 3T to narrow-band noise (NBN) stimuli at 32 different centre frequencies and hearing loss was simulated using threshold elevating noise, a method which Desloge et al. (2012) demonstrated to be an effective simulation of hearing loss by the assessment of psychoacoustic measurements. Scanning was performed at 3T to determine if changes in tonotopic properties could be observed at a field strength readily available in clinic environments.

We predict that high-frequency hearing loss simulation will bias voxel-wise preferred frequency estimates towards lower frequencies, causing an underrepresentation of high frequency responding voxels and an overrepresentation of low to mid frequencies in the resulting tonotopic maps. The level of cortical activity in response to NBN stimuli is dependent on the stimulus sensation level, therefore hearing loss simulation will reduce cortical activity to stimuli with reduced sensation level, and in turn, this will bias estimates of preferred frequency and selectivity.

This study develops a better control for future studies of tonotopic reorganisation in subjects with hearing loss, and its potential as a mechanism

for supra-threshold hearing deficits and chronic tinnitus, which is not addressed directly in the current literature. First, it will be determined if hearing loss simulation biases tonotopic maps in the manner expected due to variation in sensation level with frequency by comparing the tonotopic estimates between conditions. Second, a method to correct for this bias will be evaluated against the true organisation. To account for the variation in sensation level when estimating cortical frequency tuning properties, a modified pRF fitting procedure with the additional step of modelling the peripheral deficit - the simulated hearing loss, has been fitted to the data collected in the hearing loss simulation condition (Binda et al. 2013).

5.2 Methods

Seven normal-hearing participants (5 male; Mean age = 28.71, STE = 7.95) voluntarily took part in the present study. Normal-hearing was defined using clinical criteria (hearing thresholds $\leq$ 20 dB HL, threshold asymmetry $\leq$ 10 dB HL between ears) for frequencies between 250 Hz to 8 kHz, as assessed using pure tone audiometry.

5.2.1 Data collection

As in the previous two studies, BOLD responses to amplitude-modulated NBN stimuli centred at thirty-two different frequencies between 0.1 and 8 kHz and presented in random order at 75 dB SPL were measured using a sparse acquisition protocol (See Chapter 2 - General Methods for further details). Scanning was performed on a Philips Ingenia 3T MRI system using a volume transmit coil and a 32-channel receiver coil. For each subject, all data were collected within a single session.

Cortical responses to the different NBNs were measured in two conditions; a normal-hearing condition, where stimuli are presented at equal sensation levels across frequencies, and a simulated hearing loss condition, in which the sensation level varied as a function of frequency. Sensation level was controlled using masking noise, the level of which varied as a function of frequency in the simulated hearing loss condition. The assumption is made that data collected in the normal hearing condition will allow for the true representation to be derived, or as close as we are able to determine, of the tonotopic properties of the auditory cortex. Statistical analysis of tonotopic properties was restricted to voxels within the AC gradients ROI (see General Methods for details).

5.2.2 Hearing Loss Simulation

In the normal-hearing control condition, stimuli at all frequencies were presented at a constant sensation level (SL). In the simulated hearing loss condition, the sensation level of different stimuli varied as a function of the stimulus centre frequency to mimic high-frequency hearing loss (high-frequency stimuli were presented at lower SL).

To achieve either a constant SL at all frequencies or a decrease in SL at high frequencies, stimuli were presented in the presence of a continuous masking (or threshold-elevating) noise that differed in its frequency spectrum. In the normal-hearing condition, the masking noise was the same as used in the previous two studies and was presented at 25 dB SPL/ERB, resulting in a constant sensation level of 50 dB for all 32 stimuli.

In the simulated hearing loss condition, the threshold elevating noise was modified to increase sensation level at high frequencies. Controlling sen-

sation level using masking level allowed for all stimuli to be presented at a constant decibel sound pressure level (dB SPL), of 75 dB SPL. The threshold-elevating noise in the simulated hearing loss condition was identical to that of the normal-hearing condition up to 3 kHz and then increased to mimic a high-frequency hearing loss with a slope of 36.5 dB HL/oct between 3 and 8 kHz. The masking noise level in the simulated-hearing condition was limited to 75 dB SPL at any frequency to protect the subjects' hearing. The desired hearing loss simulation was first defined on a dB Hearing Level (dB HL) scale. Hearing loss dB HL values were then converted to dB SPL to allow the creation of the threshold elevating noise to simulate hearing loss. 0 dB HL is the average threshold of hearing as a function of frequency and therefore the SPL level of 0 dB HL is frequency-dependent. The values of which have been determined behaviourally using pure tone audiometry at discrete standardised values and are documented in BS EN ISO 389-2:1997 and ISO 389-5:2006. A continuous function was interpolated from these data sets. The summation of the threshold of hearing and the threshold of hearing loss determines the spectral level of simulated hearing loss in dB SPL.

$$sHL(f) = HL(f) + HT(f) \tag{5.1}$$

Where sHL(f) is the level of the masking noise for the hearing loss simulation in dB SPL at frequency f, HL(f) is the hearing loss in dB HL as defined above (0 dB HL until 3 kHz and increasing by 36.5 dB HL/oct between 3 and 8 kHz) and HT(f) is the normal threshold of hearing in dB SPL.

5.2.3 Data Analysis

Pre-processing of functional and structural data was performed as described in Chapter 2 - General Methods. The preferred frequency and the tuning of frequency selectivity of voxels were estimated using both the population receptive field modelling (pRF) described previously (for details see Chapter 2 - General Methods) and a modified pRF procedure detailed below. Subsequent data analysis (to produce across-subject statistics) was restricted to AC gradients ROIs. Voxel-wise GLMs were fitted to estimate the BOLD response to the presented auditory stimulus, the parameters of which are detailed in Chapter 2 - General Method. GLM and pRF models were fitted independently to the normal hearing and simulated hearing loss conditions. The impact of simulated hearing loss was assessed by comparing estimates between conditions and restricted to voxels within AC gradients ROIs.

5.2.4 Modified population receptive field analysis

The modified pRF analysis took into account the hearing loss simulation in an attempt to estimate the true tonotopic properties of voxels, devoid of the bias introduced by decreased sensation levels at certain frequencies. To achieve this, the model's event matrix, each row of which is an indicator function of when a stimulus was presented, was weighted by a hearing threshold function, such that frequencies with decreased sensation level contributed relatively less to the modelled response. This function was created by subtracting the hearing loss simulation level from the stimulus presentation level to get the sensation level and then normalised by dividing by the maximum stimulus sensation level. The method presented by Binda et al. (2013) has been developed here by applying the

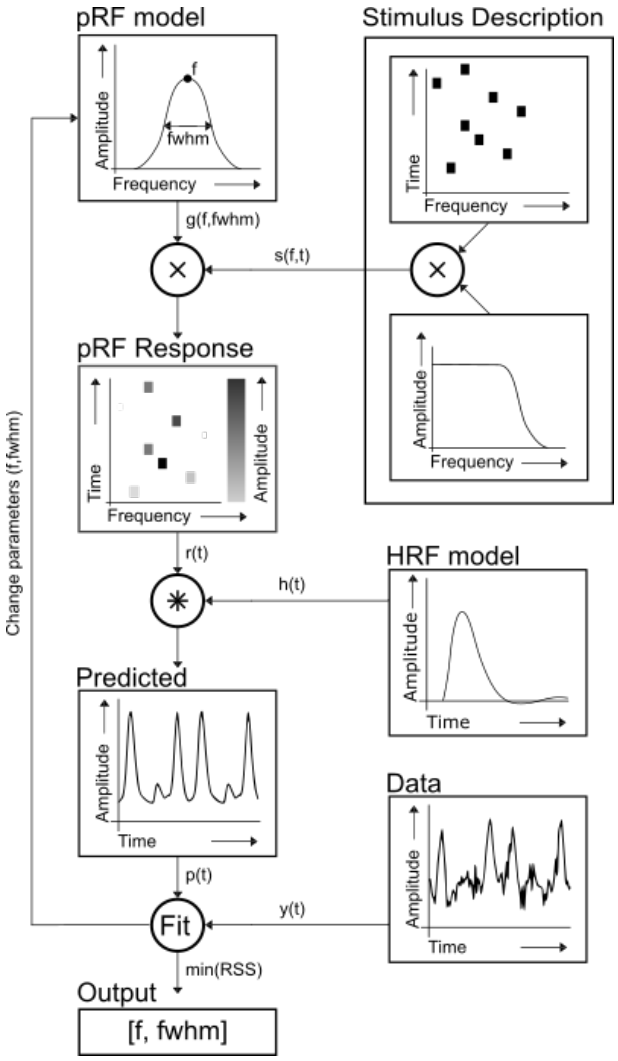


Figure 5.2: A flow chart describing the modified population receptive field (pRF) modelling procedure

The modified population receptive field (pRF) modelling followed the same fitting procedure as a traditional pRF analysis (see General Methods - 2.3.2) but the model's event matrix, each row of which is an indicator function of when a stimulus was presented, was weighted by a hearing threshold function, such that frequencies with decreased sensation level contributed relatively less to the modelled response.

technique to a deficit that varies the sensation level on a continuous scale from the baseline level to the maximum attenuation rather than a binary model of the deficit. The rest of the pRF analysis was implemented as described in Chapter 2 - General methods. See figure 5.2 for a schematic of the analysis.

5.2.5 Data Visualisation

Data visualisation was employed using the methods described in Chapter 2 - General methods. The creation of average tuning curves was modified

as follows. Average tuning curves for simulated hearing loss condition were calculated by using the voxels binning determined by the normal hearing conditions frequency response profiles.

5.3 Results

5.3.1 Impact of hearing loss simulation on preferred frequency estimates

Preferred frequency estimates were derived from data acquired using stimuli presented in the presence and absence of simulated hearing loss. In addition, a modified pRF analysis was applied to data collected in the presence of hearing loss simulation. Mirror tonotopic gradients, with a low-frequency response focus located on the central-lateral part of HG and two mirror-reversed tonotopic gradients, in a V like shape, connecting two elongated high-frequency foci, one anterior-medial and the other posterior to HG, were reliably found on all hemispheres in both conditions and analyses (see figure 5.3 for a selection of subjects).

From tonotopic maps (see figure 5.3), it can be seen that the preferred frequencies of voxels located in the high-frequency loci in the normal hearing conditions are consistently shifted towards lower preferred frequencies. The size of the high-frequency foci decreases, suggesting that regions of the cortex representing high frequencies were shifted to low frequencies. Tonotopic maps estimated from the hearing loss simulation data using the modified pRF analysis appear to reflect the normal hearing condition more than when estimated using a standard pRF analysis. The modified pRF analysis leads to an increase in voxels with high-frequency preference in cortical areas corresponding to the high-frequency

foci in the normal hearing condition.

Hearing loss simulation reduces the number of voxel with a high-frequency preference, therefore increasing the number of lower-frequency preferring voxels (see figure 5.4 and figure 5.3). Accounting for sensation level in the modified pRF analysis procedure partially alleviate this shift in preferred frequency distribution. This correction is not perfect, however, as there remains an over-representation of low frequencies ($>$ 0.44 kHz), mid-high frequencies (1.75-5.5 kHz) become under-represented compared to both normal hearing and hearing loss simulation conditions and an exaggeration of high-frequencies at the limit of the frequency range.

5.3.2 Impact of hearing loss simulation on average voxel frequency response profiles

Average voxel frequency response profiles were derived using the beta weights of the GLM analysis and calculated independently for each condition from all voxels located within the AC gradient ROIs of all subjects. Simulated hearing loss attenuates average percentage signal change in response to frequencies with reduced sensation level due to hearing loss simulation, see figure 5.5. Interestingly, the suppression extends beyond the region of hearing loss simulation and the response level gradually returns to match the normal condition. The extension of suppression below the edge of the simulated hearing loss is also visible in figure 5.6 which depicts the difference between the average beta weight from each condition and the sensation level of stimuli in the hearing loss simulation, as a function of frequency for all subjects. The cortical responses from 0.5 kHz and above is attenuated whereas the reduction in sensation level only impacts frequencies above 3 kHz. The group average percentage sig-

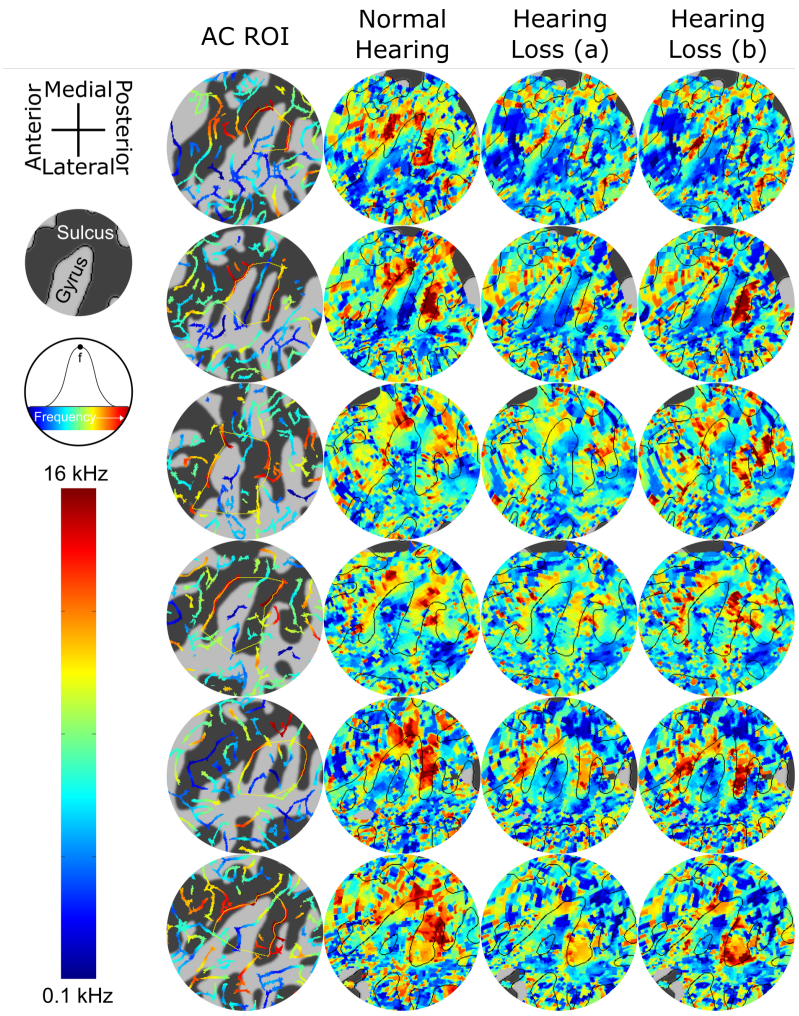


Figure 5.3: Spatial distribution of voxel-wise preferred frequency estimates from normal-hearing and hearing loss simulation conditions

From left to right: AC ROI; sulci/gyri boundaries and frequency gradient reversals, Normal hearing; maps of voxel-wise preferred frequencies estimated from the normal-hearing condition, Hearing loss (a); the hearing loss simulation condition, Hearing loss (b); the hearing loss simulation condition using the modified pRF analysis. Hearing loss simulation reduced the number of high-frequency responding voxels, which could be miss interpreted as tonotopic reorganisation. Implementing the modified pRF analysis increases the number of voxels with high-frequency preference but tonotopic organisation is not exactly the same as the normal hearing condition. Each row is an individual hemispheres, with ones from the right flipped horizontally to align the orientation of HG across all hemispheres. The low-frequency gradient reversal on and high-frequency surrounding Heschl'S Gyrus were used to delineate an auditory core ROI (depicted by yellow line in the first column).

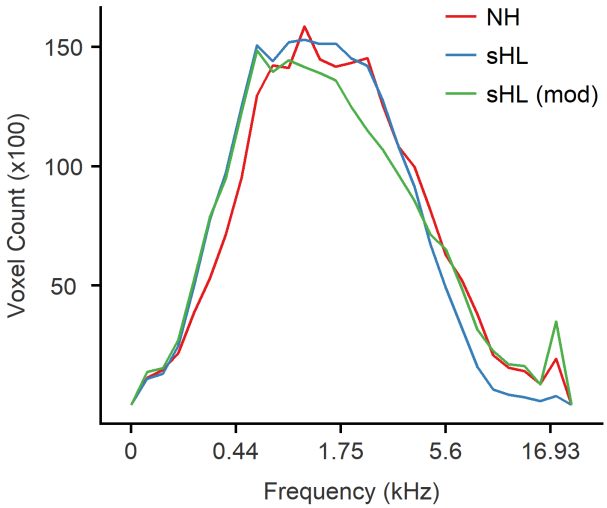


Figure 5.4: Distribution of preferred frequency estimates from normal-hearing and hearing loss simulation conditions

Hearing loss simulation reduces the number of voxel with a high-frequency preference, therefore increasing the number of lower-frequency preferring voxels. Accounting for sensation level in the modified pRF analysis procedure partially alleviate this shift in preferred frequency distribution but doesn't correct the over-representation of low frequencies ($>$ 0.44 kHz) or under-represented mid-high frequencies (1.75-5.5 kHz), and leads to an exaggeration of high-frequencies at the limit of the frequency range in comparison to the normal hearing condition. Preferred frequency is plotted in kHz on an nERB scale. The red line depicts estimates from the normal-hearing condition, blue and green lines are estimates from the hearing- loss simulation condition. Red and blue were derived using a standard pRF analysis, whereas the green line was derived using a modified pRF analysis that included the hearing loss simulation in the models' stimulus description.

nal change in response to all stimuli was significantly different (Welch Two Sample t-test p=9.458e-11) reduced by the hearing loss simulation, from 0.81 to 0.66 percent.

5.3.3 Impact of hearing loss simulation on average voxel frequency response profiles and tuning width

The average voxel frequency response profiles are Gaussian-like in shape on an ERB scale, for all frequency bins from data collected in both normal-hearing and simulated hearing loss conditions, see figure 5.7. The simulated hearing loss causes an overall attenuation of responses elicited from voxels with a preferred frequency within the frequency range of hearing loss simulation. Voxels with preferred frequencies beyond this area exhibited artificially narrow frequency selectivity due to the attenuation of activity in response to high-frequency stimuli within the hearing loss simulation frequency range.

Hearing loss simulation significantly narrows the average tuning width of voxels within the AC gradient ROI of all subjects, from 5.48 nERB in the normal-hearing condition to 4.39 nERB (Welch Two Sample t-test, p-value $< 2.2e-16$). Tuning widths estimates from the modified pRF fitting procedure remain significantly narrower than in the normal-hearing condition, at 4.76 nERB (Welch Two Sample t-test, p-value $< 2.2e-16$). Estimates of tuning width upon HG appear to be relativity consistent along the middle of HG and anteriorly, whereas posterior cortical HG areas have larger tuning widths. Extremely narrow tuning widths, at the lower limit of the fitting procedure parameters, are estimated for a majority voxels beyond HG and its adjacent anterior and posterior sulci.

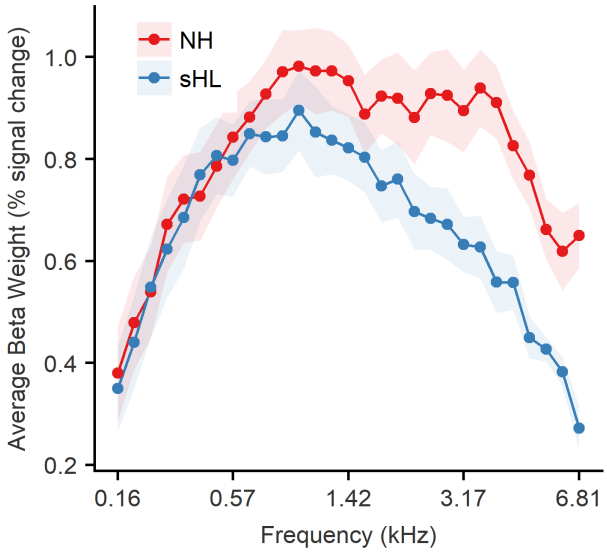


Figure 5.5: Average percentage signal change in response to stimuli at different centre frequencies in the presence and absence of hearing loss simulation

Impact of hearing loss simulation on average BOLD responses in primary auditory cortex as a function of stimulus frequency. Average BOLD response as a function of frequency was computed by averaging voxelwise GLM response profiles (beta weights) across voxels with the AC gradient ROI. Blue represents estimates from data acquired in the hearing loss simulation condition and red from data acquired in the normal-hearing condition. The solid line is the mean for all hemispheres (n =14) from all subjects AC gradients ROIs and the ribbons represent the standard error. The frequency specific attenuation of BOLD responses due to high-frequency can be seen by the difference (significant) between the frequency profiles.

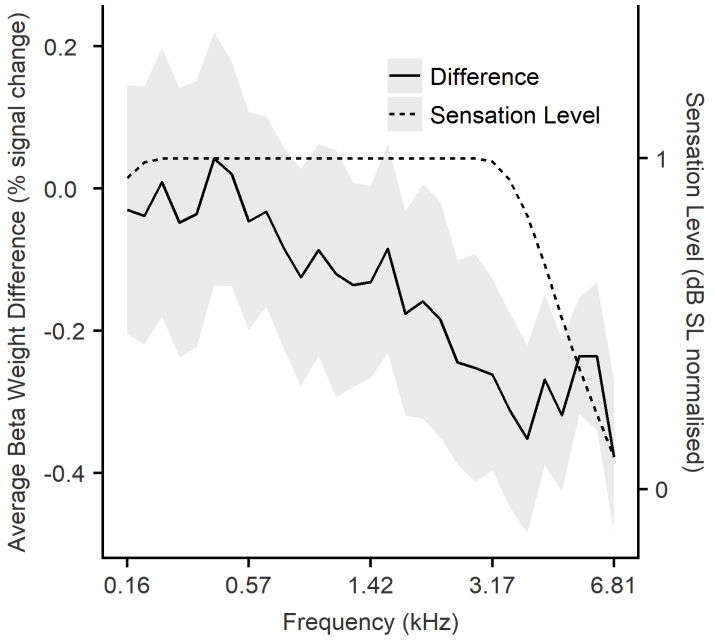


Figure 5.6: Difference between average beta weights collected in the presence and absence of hearing loss

The difference between the mean beta weights of each condition (solid black line) was calculated by subtracting the average betas weights of the normal hearing condition from the simulated hearing loss condition. The grey ribbon depicts the sum of the standard error of both conditions. The dashed black line is the normalised sensation level of stimulus frequencies in the hearing loss condition. The difference in average BOLD response between the two conditions demonstrates how the attenuation follows the hearing loss simulations frequency profile but extends beyond to frequencies with equal sensation levels in both conditions.

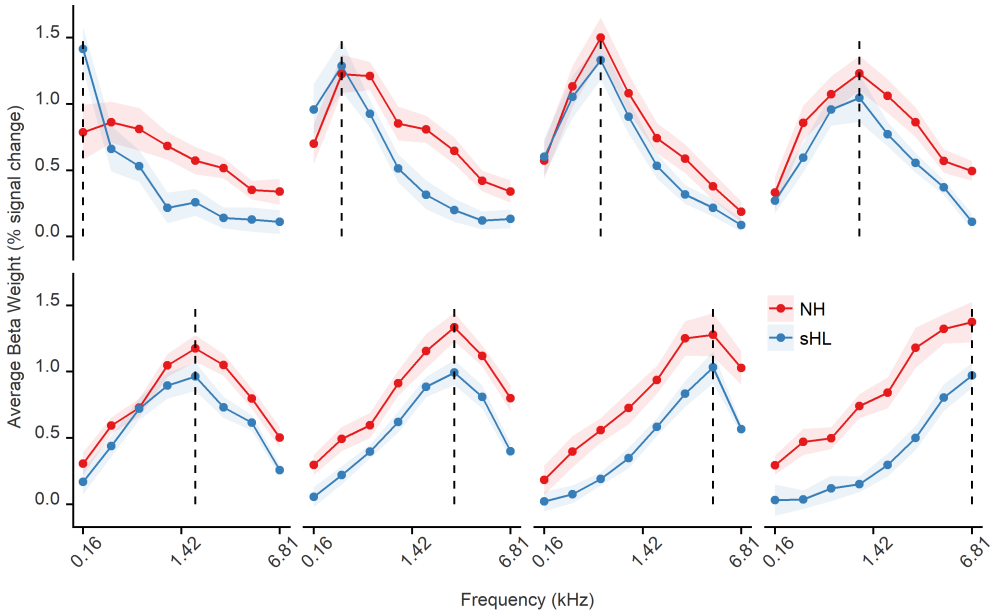


Figure 5.7: Impact of hearing loss simulation on average voxel frequency response profiles averaged

Frequency response profiles averaged according to voxel's preferred frequency for data collected in the presence and absence of hearing loss simulation. Average frequency response profiles were derived by averaging voxelwise response profiles after grouping AC gradient ROI voxels in 8 preferred frequency bins. Voxelwise preferred frequencies and resulting average frequency response profiles were derived using a split-half method to avoid artificially narrow frequency response profiles (for details, see Methods and Chapter 2 - General Methods). Points connected by solid lines represent the average beta weight associated with each stimulus frequency bin across all hemispheres from all subjects (n=14) and the ribbon represent the corresponding standard error. Hearing loss simulation results in attenuation of responses to all frequencies for voxels with a preference to frequencies within the hearing loss range (high-frequencies). In addition, the frequency profiles of voxels with preferred frequencies outside of this range are narrower than in the normal-hearing condition, with the impact being the greatest for frequencies the furthest from the hearing loss range (low-frequencies).

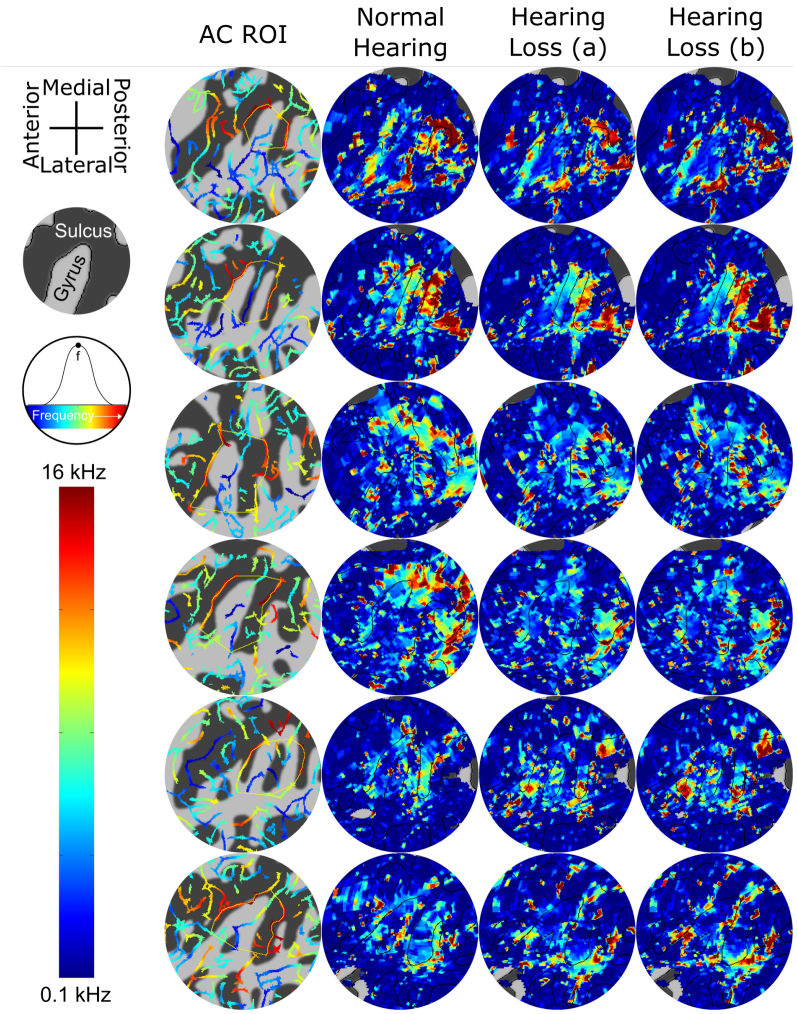


Figure 5.8: Spatial distribution of voxel-wise frequency selectivity estimates from normal-hearing and hearing loss simulation conditions

From left to right: AC ROI; sulci/gyri boundaries and frequency gradient reversals, Normal hearing; maps of frequency selectivity estimates from the normal-hearing condition, Hearing loss (a); the hearing loss simulation condition, Hearing loss (b); the hearing loss simulation condition using the modified pRF analysis. In the hearing loss condition there is a reduction in the number of broadly tuned voxels but the maps are hard to interpret due to poorly fixed voxels being assigned extremely narrow tuning widths. Each row is an individual hemispheres, with ones from the right flipped horizontally to align the orientation of HG across all hemispheres. The low-frequency gradient reversal on and high-frequency surrounding Heschl'S Gyrus were used to delineate an auditory core ROI (depicted by yellow line in the first column).

5.4 Discussion

The presented study measured the impact of simulated hearing loss on estimates of tonotopic properties from the human auditory cortex, using auditory fMRI. It was found that hearing loss simulation in normal-hearing subjects biases the distribution of preferred frequency estimates in auditory cortex to lower frequencies, compared to a normal hearing condition, despite the absence of tonotopic re-organization. This bias can be seen in individual tonotopic maps (see figure 5.3) as a reduction in the cortical area covered by high-frequency regions and the group average distribution of preferred frequency (see figure 5.4) as a decrease in the number of high frequency responding voxels. The distribution of preferred frequency also demonstrates an increase in voxels with preferred frequencies just below that of the hearing loss simulation. The bias in estimates can be explained by the impact of hearing loss simulation on the BOLD activity elicited by auditory stimuli. There is a significant reduction in BOLD percentage signal change to stimuli within the frequency range of the simulated hearing loss, seen in figure 5.5. Grouping voxels by preferred frequency shows that this attenuation impacts all voxels, irrespective of frequency preference. This bias in voxel profiles leads to under-estimating both preferred frequency and overestimating frequency selectivity (underestimating tuning width), see figure 5.8. Desloge et al. (2012) demonstrated that behavioural measures of frequency selectivity from hearing loss subjects could be reproduced in normal hearing subjects using threshold elevating noise hearing loss simulation. The systematic biases in estimates observed in the present study appear to simulate what would be expected from hearing loss subjects due to reduced sensitivity to sound. The findings presented here suggest that threshold elevating noise can also be implemented as an effective method of sim-

ulating hearing loss in auditory fMRI experiments. The use of hearing loss simulation allows the creation of normal-hearing controls for patients with hearing profiles expected to, based on animal models, lead to tonotopic reorganisation without the confound of different stimulus sensation levels between groups.

Binda et al. (2013) demonstrated that modelling a peripheral deficit in the analysis of visual fMRI data can minimise its impact on estimates of functional cortical properties. The presented study has developed this method by applying the technique to a deficit that varies in sensation level on a continuous scale, requiring greater consideration than the binary deficit modelled in the original study (Binda et al. 2013). The sensation level was accounted for in the description of auditory stimulus in the pRF analysis. Bias in preferred frequency was partially alleviated, with estimates from the hearing loss simulation using the modified analysis being in greater agreement with estimates from normal hearing data - in terms of the distribution estimates, see figure 5.4 and the correlation been datasets. However, the correction for the bias is only partial. The distribution of preferred frequency estimates is still biased towards lower frequencies compared to estimates from normal hearing datasets and average tuning widths remain narrower. This may be due to inadequately modelling the impact of threshold elevating noise on cortical responses to auditory stimuli measured using BOLD fMRI. The model could have failed to account for a number of unknown factors.

The presented methods offer a valuable contribution to future auditory fMRI studies investigating tonotopic reorganisation in clinical populations. The findings will allow for the creation of matched hearing loss controls and the opportunity to minimise the effect of hearing loss on the estimates of voxel functional properties. This will allow subject recruit-

ment to be focused on the potential for reorganisation, rather than to mitigate hearing loss in stimulus presentation.

Chapter 6

General Discussion

6.1 Accurate measurement of tonotopic properties

The first study presented in this thesis, Chapter 3, aimed to determine how to best measure the tonotopic properties of the human auditory cortex using fMRI. To understand the functional properties of the auditory cortex, it is vital that the limitations of the measures employed are understood, and its impact is characterised and well defined. Investigating the human auditory cortex using BOLD fMRI poses unique challenges compared to the other sensory cortices. The acoustic noise generated by the scanner when collecting data is sufficiently loud to mask acoustic stimuli, despite the use of hearing protection. The use of sparse acquisition protocols, where stimuli are presented in the silence between data acquisition, may avoid this issue but at the expense of sampling resolution and number of observations possible in a given time.

The first part of this study assessed the impact of acquisition protocol on

estimates of tonotopic properties. Acoustic scanner noise was found to bias estimates of preferred frequency and selectively, observed in the distribution of preferred frequencies and the average tuning curves of voxels from the auditory cortex. This bias appears to be due to the scanner noise masking the acoustic stimuli, and as a result, attenuating cortical responses to them. The resulting attenuation of cortical response was frequency-dependent in correspondence with the frequency spectrum of the scanner noise, suggesting that cortical activity is dependent on stimulus sensation level. This result is corroborated by the results from the last study presented in chapter 6 demonstrated, where stimulus sensation level was controlled, the sensation level of a stimulus directly impacts the cortical activity in response to it. A direct relationship between sensation level and cortical activation has also been observed in a study by Langers, Van Dijk, et al. (2006), which varied the sensation level by altering the sound pressure level of stimuli. Their results and the results presented in this thesis suggest the strength of cortical activity in response to auditory stimuli reflects the sensation level of the stimulus (Langers, Van Dijk, et al. 2006). To present stimuli in silence using a sparse acquisition protocol requires a longer time between data acquisition and therefore a reduction in the number of data points acquired. In the presented study, this resulted in nearly one-fifth of the data compared to the continuous dataset. However, this didn't result in a significant reduction in the reliability of estimates from the data. The results of this study strongly suggest that sparse acquisition protocols should be used for the measurement of cortical frequency tuning properties.

The first study also assessed statistical analysis methods to infer cortical frequency properties from BOLD fMRI time-course data. The results of this study suggest using multiple analysis methods to overcome short-

falls of any one technique. Population receptive field modelling appears to yield estimates of preferred frequency with less bias due to the stimulus frequency range than estimates based on the centroids of beta weights from a GLM analysis. However, GLM beta weights allow for voxel frequency response profiles to be investigated with minimal assumptions made about the shape of population receptive fields of the auditory cortex.

The aim of the first study is to determine how to best measure cortical frequency tuning properties to create methods sufficiently sensitive to measure tonotopic reorganisation. The results from the final study suggest that the methods presented are sufficiently sensitive to detect changes due to hearing loss, with the impact of hearing loss simulation seen in the distribution of preferred frequency and frequency response profiles. The results of the first study demonstrate that the methodology employed has the potential to bias results. Inference of the true cortical properties from indirect measures of populations of neurons, such as BOLD fMRI, must be done in the context of experimental and analytical methodologies employed and the potential impact that each method may have on estimates.

6.2 Cortical tonotopic magnfiction

The study presented in chapter 4 details the measurement of the cortical magnification of the auditory cortex. The results of this study show that the average cortical magnification function is consistent with the frequency discrimination function across frequency. This suggests that frequency discrimination is dependent on the tonotopic magnification of the human auditory cortex, consistent with findings in the visual and so-

matosensory modalities where the cortical magnification of primary sensory areas has been found to correlate with within-subject measures of Vernier acuity and tactile hyperacuity respectively (Duncan, Boynton, and Jolla 2003; Duncan and Boynton 2007). In the visual and tactile domain, it has been shown that inter-individual difference in the absolute value of the cortical magnification correlates with the inter-individual differences in discrimination ability (Duncan, Boynton, and Jolla 2003; Duncan and Boynton 2007). To determine if the same is true in the auditory domain would require behavioural measurement of frequency discrimination from subjects to be compared to measures of their cortical magnification, which would be a very interesting topic for a future study.

6.3 Measuring hearing loss induced tonotopic reorganisation

Animal research has suggested that peripheral hearing loss may lead to changes in central auditory processing arising as a consequence of a reduction in peripheral auditory output. In animals, steeply sloping hearing loss has been shown to cause large-scale tonotopic reorganisation in auditory cortex (Irvine, Rajan, and Mcdermott 2000; Sereda, Hall, et al. 2011), whereby neurons responsive to frequencies affected by the hearing loss (typically higher frequencies), shift their tuning towards less affected (lower) frequencies (Noreña and Farley 2013). To detect changes in tonotopic reorganisation requires the recruitment of subjects with steeply sloping hearing loss, the characteristics suggested to lead to reorganisation. To determine if reorganisation occurs, the tonotopic maps of these subjects need to be compared to normal-hearing subjects measured using equivalent stimuli. The results of the study presented in chapter 5

suggest that hearing loss simulation through threshold evolving noise achieves this. Implementing threshold elevating noise to simulate hearing loss allows for stimuli to be presented at the same sensation level and sound pressure level to hearing loss subjects and normal hearing controls. The biases observed in tonotopic estimates from data collected normal hearing subjects in the presence of threshold elevating noise are consistent with what would be expected with peripheral hearing loss and it has been shown to effectively stimulate behavioural changes observed in hearing loss subject (Desloge et al. 2012).

6.4 Future Directions

The next step will be to apply the methods developed in this thesis to the study of patients with hearing loss characteristics consistent with the expectations to lead to tonotopic reorganisation based on animal models, steeply sloping hearing loss and tinnitus (Irvine, Rajan, and Mcdermott 2000; Sereda, Hall, et al. 2011; Noreña and Farley 2013). The results of the studies presented in this thesis suggest the following. Collecting BOLD fMRI responses to auditory stimuli using a sparse imaging paradigm to mitigate the effects of acoustic scanner noise and using threshold elevating noise to simulate hearing loss in normal-hearing controls. Analysing the data using a number of analysis methods to allow for a holistic interpretation of the data. A pRF analysis will produce the least biased (of methods assessed in this thesis) estimates of preferred frequency. From these estimates, cortical magnification and point image can be calculated and compared between groups and related to the characteristics of hearing loss, i.e. the hearing loss profile or the properties of a tinnitus percept. Comparing measured cortical magnification functions and frequency discrimination thresholds in patients to normal hearing

controls may help detect tonotopic reorganisation. Cortical responses to auditory stimuli estimated using a GLM analysis would allow for the assessment of cortical frequency response profiles and cortical tuning properties, which help with the interpretation preferred frequency and selectivity estimates. Hopefully, the work presented in this thesis will provide methodologies that improve the study of hearing-loss-induced tonotopic reorganisation in human auditory cortex.

Bibliography

Ahveninen, Jyrki et al. (2016). "Intracortical depth analyses of frequency-sensitive regions of human auditory cortex using 7T fMRI HHS Public Access". In: *Neuroimage* 143, pp. 116–127. DOI: 10.1016/j.neuroimage.2016.09.010. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5124525/pdf/nihms817209.pdf.

Alho, Kimmo et al. (2014). "Stimulus-dependent activations and attention-related modulations in the auditory cortex: A meta-analysis of fMRI studies". In: *Hear. Res.* 307, pp. 29–41. ISSN: 03785955. DOI: 10.1016/j.heares.2013.08.001. URL: https://linkinghub.elsevier.com/retrieve/pii/S0378595513001913.

Baguley, David M. (2002). "Mechanisms of tinnitus". In: *Br. Med. Bull.* 63, pp. 195–212. ISSN: 00071420. DOI: 10.1093/bmb/63.1.195.

Baseler, Heidi A et al. (2011). "Large-scale remapping of visual cortex is absent in adult humans with macular degeneration". In: *Nat. Publ. Gr.* 14.5, pp. 649–655. ISSN: 1097-6256. DOI: 10.1038/nn.2793. URL: http://dx.doi.org/10.1038/nn.2793.

Baumann, Simon et al. (2013). "A unified framework for the organization of the primate auditory cortex". In: *Front. Syst. Neurosci.* DOI: 10.3389/fnsys.2013.00011. URL: www.frontiersin.org.

Békésy, Georg v. (1949). "The Vibration of the Cochlear Partition in Anatomical Preparations and in Models of the Inner Ear". In: *J. Acoust. Soc. Am.* 21.3, pp. 233–245. ISSN: 0001-4966. DOI: 10.1121/1.1906502. URL: http://asa.scitation.org/doi/10.1121/1.1906502.

Besle, Julien et al. (2019). "Is Human Auditory Cortex Organization Compatible With the Monkey Model? Contrary Evidence From Ultra-High-Field Functional and Structural MRI". In: *Cereb. Cortex* 29, pp. 410–428. DOI: 10.1093/cercor/bhy267. URL: https://academic.oup.com/cercor/article-abstract/29/1/410/5144239.

Binda, Paola et al. (2013). "Minimizing biases in estimating the reorganization of human visual areas with BOLD retinotopic mapping." In: *J. Vis.* 13.7, p. 13. ISSN: 1534-7362. DOI: 10.1167/13.7.13.doi. URL: http://www.journalofvision.org/content/13/7/13.full.

Bloch, F (1946). *Nuclear Induction*. Tech. rep. URL: https://journals.aps.org/pr/pdf/10.1103/PhysRev.70.460.

Brodmann, Korbinian and Laurence J. Gary (1909). *Brodmann's localization in the cerebral cortex : the principles of comparative localisation in the cerebral cortex based on cytoarchitectonics*. Springer, p. 298. ISBN: 0387269193.

Buračas, Giedrius T. and Geoffrey M. Boynton (2002). "Efficient Design of Event-Related fMRI Experiments Using M-Sequences". In: *Neuroimage* 16.3, pp. 801–813. ISSN: 10538119. DOI: 10.1006/nimg.2002.1116. URL: http://linkinghub.elsevier.com/retrieve/pii/S105381190291116X.

Da Costa, Sandra et al. (2011). "Human Primary Auditory Cortex Follows the Shape of Heschl's Gyrus". In: DOI: 10.1523/JNEUROSCI.2000-11.2011. URL: www.psychtoolbox.org.

Dale, Anders (1999). "Optimal Experimental Design for Event-Related fMRI". In: *Hum. Brain Mapp.* 8, pp. 109–114.

Daniel, P. M. and D. Whitteridge (1961). "The representation of the visual field on the cerebral cortex in monkeys". In: *J. Physiol.* 159.2, pp. 203–221. ISSN: 14697793. DOI: 10.1113/jphysiol.1961.sp006803.

De Martino, Federico et al. (2015). "Frequency preference and attention effects across cortical depths in the human primary auditory cortex". In: *Proc. Natl. Acad. Sci.* 112.52. DOI: 10.1073/pnas.1507552112. URL: https://www.pnas.org/content/pnas/112/52/16036.full.pdf.

Desloge, Joseph G et al. (2012). "Auditory-Filter Characteristics for Listeners With Real and Simulated Hearing Impairment". In: *Trends Amplif.* 16.1. DOI: 10.1177/1084713812445510.

Dick, Frederic et al. (2012). "In Vivo Functional and Myeloarchitectonic Mapping of Human Primary Auditory Areas". In: 32.46, pp. 16095–16105. DOI: 10.1523/JNEUROSCI.1712-12.2012.

Dubno, Judy R., Donald D. Dirks, and Donald E. Morgan (1984). "Effects of age and mild hearing loss on speech recognition in noise". In: *J. Acoust. Soc. Am.* 76.1, pp. 87–96. ISSN: 0001-4966. DOI: 10.1121/1.391011. URL: http://asa.scitation.org/doi/10.1121/1.391011.

Dumoulin, Serge O. and Brian A. Wandell (2008). "Population receptive field estimates in human visual cortex". In: *Neuroimage* 39.2, pp. 647–660. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2007.09.034. URL: http://www.ncbi.nlm.nih.gov/

pubmed/17977024%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3073038%20http://linkinghub.elsevier.com/retrieve/pii/S1053811907008269.

Duncan, R. O. and G. M. Boynton (2007). "Tactile Hyperacuity Thresholds Correlate with Finger Maps in Primary Somatosensory Cortex (S1)". In: *Cereb. Cortex* 17.12, pp. 2878–2891. ISSN: 1047-3211. DOI: 10.1093/cercor/bhm015. URL: https://academic.oup.com/cercor/article-lookup/doi/10.1093/cercor/bhm015.

Duncan, Robert O, Geoffrey M Boynton, and La Jolla (2003). "Cortical Magnification within Human Primary Visual Cortex Correlates with Acuity Thresholds". In: *Neuron* 38, pp. 659–671.

Eggermont, Jos J. and Larry E. Roberts (2004). "The neuroscience of tinnitus". In: *Trends Neurosci.* 27.11, pp. 676–682. ISSN: 01662236. DOI: 10.1016/j.tins.2004.08.010.

— (2014). "Tinnitus: animal models and findings in humans". In: *Cell Tissue Res.* ISSN: 0302-766X. DOI: 10.1007/s00441-014-1992-8. URL: http://link.springer.com/10.1007/s00441-014-1992-8.

Engelien, A (2002). "Physiological Mapping of Human Auditory Cortices with a Silent Event-Related fMRI Technique". In: *Neuroimage* 16.4, pp. 944–953. ISSN: 10538119. DOI: 10.1006/nimg.2002.1149. URL: https://linkinghub.elsevier.com/retrieve/pii/S1053811902911493.

Everest, F. Alton and Ken C. Pohlmann (2009a). "Fundamentals of Sound". In: *Master Handb. Acoust.* 5th. McGraw-Hill, pp. 1–17.

— (2009b). *Master handbook of acoustics*. McGraw-Hill. ISBN: 0071603336. URL: https://books.google.co.uk/books?id=6tiJ1cwnwxoC%7B%5C&%7Dprintsec=frontcover%7B%5C&%7Ddq=master+

handbook+of+acoustics+5th%7B%5C&%7Dhl=en%7B%5C&%7Dsa=X%7B%5C&%7Dved=0ahUKEwjRqvqU9Z3kAhW9VBUIHQl9BP0Q6AEIKjAA%7B%5C#%7Dv=onepage%7B%5C&%7Dq=master%20handbook%20of%20acoustics%205th%7B%5C&%7Df=false.

— (2009c). "Sound in the Free Field". In: *Master Handb. Acoust.* Pp. 33–38.

— (2009d). "The Perception of Sound". In: *Master Handb. Acoust.* Pp. 39–66.

Formisano, Elia et al. (2003). *Mirror-Symmetric Tonotopic Maps in Human Primary Auditory Cortex*. Tech. rep., pp. 859–869. URL: https://ac.els-cdn.com/S089662730300669X/1-s2.0-S089662730300669X-main.pdf?%7B%5C_%7Dtid=b07c5d2d-bf5c-46ed-9f60-ada9e68cb155%7B%5C&%7Dacdnat=1539181704%7B%5C_%7D.

Gander, Phillip E. et al. (2019). "Direct electrophysiological mapping of human pitch-related processing in auditory cortex". In: *Neuroimage* 202, p. 116076. ISSN: 10959572. DOI: 10.1016/j.neuroimage.2019.116076.

Gardner, Justin L. et al. (2018a). "mrTools: Analysis and visualization package for functional magnetic resonance imaging data". In: DOI: 10.5281/ZENODO.1299483. URL: https://zenodo.org/record/1299483%7B%5C#%7D.XWLWv-hKiHs.

— (2018b). *mrTools: Analysis and visualization package for functional magnetic resonance imaging data*. Version 4.7. mrTools provides a set of Matlab tools to analyze functional magnetic resonance imaging data. It can do basic analyses like correlation analyses used in retinotopy experiments, event-related, population RF and GLM analyses. It can display the results of analyses on inplane anatomies, flat maps and surfaces. It is designed

to make it easy to write your own script and programs in Matlab to analyze your data. DOI: 10.5281/zenodo.1299483. URL: https://doi.org/10.5281/zenodo.1299483.

Geisler, Wilson S (1984). *Physical limits of acuity and hyperacuity.* Tech. rep. 7.

Ghazaleh, Naghmeh et al. (2017). "High-Resolution fMRI of Auditory Cortical Map Changes in Unilateral Hearing Loss and Tinnitus". In: *Brain Topogr.* 30.5, pp. 685–697. ISSN: 0896-0267. DOI: 10.1007/s10548-017-0547-1. URL: http://link.springer.com/10.1007/s10548-017-0547-1.

Glasberg, B R and B C Moore (1986). "Auditory filter shapes in subjects with unilateral and bilateral cochlear impairments." In: *J. Acoust. Soc. Am.* 79.4, pp. 1020–1033. ISSN: 00014966. DOI: 10.1121/1.393374.

— (1990). "Derivation of auditory filter shapes from notched-noise data". In: *Hear. Res.* 47.1-2, pp. 103–138. ISSN: 03785955. DOI: 10.1016/0378-5955(90)90170-T.

Gowland, P A and M 0 Leach (1991). *A Simple Method for the Restoration of Signal Polarity in Multi-image Inversion Recovery Sequences for Measuring T1.* Tech. rep., pp. 224–247. URL: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910180123.

Hall, Deborah A et al. (1999). "" Sparse " Temporal Sampling in Auditory fMRI". In: *Hum. Brain Mapp.* 223, pp. 213–223.

Harvey, Ben M and Serge O Dumoulin (2011). "The Relationship between Cortical Magnification Factor and Population Receptive Field Size in Human Visual Cortex : Constancies in Cortical Architecture". In: *J. Neurosci.* 31.38, pp. 13604–13612. DOI: 10.1523/JNEUROSCI.2572-11.2011.

Hawkins, J. E. and S. S. Stevens (1950). "The Masking of Pure Tones and of Speech by White Noise". In: *J. Acoust. Soc. Am.* 22.1, pp. 6–13. ISSN: 0001-4966. DOI: 10.1121/1.1906581. URL: http://asa.scitation.org/doi/10.1121/1.1906581.

Humphries, Colin, Einat Liebenthal, and Jeffrey R Binder (2010). "Tonotopic organization of human auditory cortex". In: *Neuroimage* 50, pp. 1202–1211. DOI: 10.1016/j.neuroimage.2010.01.046. URL: http://surfer.nmr.mgh.harvard..

Irvine, Dexter R F, Ramesh Rajan, and Hugh J Mcdermott (2000). "Injury-induced reorganization in adult auditory cortex and its perceptual consequences". In: *Hear. Res. 147* 147.1-2, pp. 188–199. URL: http://ac.els-cdn.com/S0378595500001313/1-s2.0-S0378595500001313-main.pdf?%7B%5C_%7Dtid=56a36822-7c2f-11e7-aa90-00000aacb362%7B%5C&%7Dacdnat=1502193022%7B%5C_%7D.

Kaas, Jon H and Troy A Hackett (2000). "Subdivisions of auditory cortex and processing streams in primates". In: *Proc. Natl. Acad. Sci.* 97.22, pp. 11793–11799. DOI: 10.1073/pnas.97.22.11793. URL: www.pnas.org.

Langers, Dave R M (2014). "Assessment of tonotopically organised subdivisions in human auditory cortex using volumetric and surface-based cortical alignments". In: *Hum. Brain Mapp.* 35.4, pp. 1544–1561. ISSN: 10970193. DOI: 10.1002/hbm.22272.

Langers, Dave R M and Pim Van Dijk (2012). "Mapping the Tonotopic Organization in Human Auditory Cortex with Minimally Salient Acoustic Stimulation". In: September, pp. 2024–2038. DOI: 10.1093/cercor/bhr282.

Langers, Dave R M, Emile de Kleine, and Pim van Dijk (2012). "Tinnitus does not require macroscopic tonotopic map reorganiza-

tion." In: *Front. Syst. Neurosci.* 6, p. 2. ISSN: 1662-5137. DOI: 10.3389/fnsys.2012.00002. URL: http://www.ncbi.nlm.nih.gov/pubmed/22347171%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3269775.

Langers, Dave R M, Katrin Krumbholz, et al. (2014). "Neuroimaging paradigms for tonotopic mapping (I): The in fl uence of sound stimulus type". In: *Neuroimage* 100, pp. 650–662. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2014.07.044. URL: http://dx.doi.org/10.1016/j.neuroimage.2014.07.044.

Langers, Dave R M, Rosa M Sanchez-panchuelo, et al. (2014). "Neuroimaging paradigms for tonotopic mapping (II): The in fl uence of acquisition protocol". In: *Neuroimage* 100, pp. 663–675. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2014.07.042. URL: http://dx.doi.org/10.1016/j.neuroimage.2014.07.042.

Langers, Dave R M, Pim Van Dijk, et al. (2006). "fMRI activation in relation to sound intensity and loudness". In: DOI: 10.1016/j.neuroimage.2006.12.013. URL: www.elsevier.com/locate/ynimg.

Levitt, H (1971). "Transformed Up-Down Methods in Psychoacoustics". In: *J. Acoust. Soc. Am.* 49, pp. 467–477.

Moerel, M, Federico De Martino, and Elia Formisano (2012). "Processing of Natural Sounds in Human Auditory Cortex : Tonotopy , Spectral Tuning , and Relation to Voice Sensitivity". In: 32.41, pp. 14205–14216. DOI: 10.1523/JNEUROSCI.1388-12.2012.

Moore, Brian C. J. (2012). *An introduction to the psychology of hearing.* Emerald. ISBN: 1780520387. URL: https://books.google.co.uk/books?hl=en%7B%5C&%7Dlr=%7B%5C&%7Did=LM9U8e28pLMC%7B%5C&%7Doi=fnd%7B%5C&%7Dpg=PP1%7B%5C&%7Ddq=an+

introduction+to+the+psychology+of+hearing+reference%7B%5C&%7Dots=L2WnlXUIDf%7B%5C&%7Dsig=6QQ%7B%5C_%7DcAyBb5vc8-9tp4QTpMxPwfY%7B%5C#%7Dv=onepage%7B%5C&%7Dq=an%20introduction%20to%20the%20psychology%20of%20hearing%20reference%7B%5C&%7Df=false.

Moore, Brian C J and S. N. Vinay (2009). "Enhanced discrimination of low-frequency sounds for subjects with high-frequency dead regions". In: *Brain* 132.2, pp. 524–536. ISSN: 00068950. DOI: 10.1093/brain/awn308.

Moore, Brian C.J. and Brian R. Glasberg (1987). "Formulae describing frequency selectivity as a function of frequency and level, and their use in calculating excitation patterns". In: *Hear. Res.* 28.2-3, pp. 209–225. ISSN: 03785955. DOI: 10.1016/0378-5955(87)90050-5. URL: https://linkinghub.elsevier.com/retrieve/pii/0378595587900505.

Morel, A., P. E. Garraghty, and J. H. Kaas (1993). "Tonotopic organization, architectonic fields, and connections of auditory cortex in macaque monkeys". In: *J. Comp. Neurol.* 335.3, pp. 437–459. ISSN: 0021-9967. DOI: 10.1002/cne.903350312. URL: http://doi.wiley.com/10.1002/cne.903350312.

Nelson, D A, M E Stanton, and R L Freyman (1983). "A general equation describing frequency discrimination as a function of frequency and sensation level". In: *J. Acoust. Soc. Am.* 73.6. ISSN: 00014966. DOI: 10.1121/1.381278.

Nestares, Oscar and David J Heeger (2000). *Robust Multiresolution Alignment of MRI Brain Volumes*. Tech. rep. URL: https://www.cns.nyu.edu/heegerlab/content/software/imgregsoft/Nestares-MRM2000.pdf.

Noreña, Arnaud J. and Brandon J. Farley (2013). "Tinnitus-related neural activity: Theories of generation, propagation, and centralization". In: *Hear. Res.* 295, pp. 161–171. ISSN: 03785955. DOI: 10.1016/j.heares.2012.09.010. URL: http://dx.doi.org/10.1016/j.heares.2012.09.010.

Norena, Arnaud et al. (2002). "Psychoacoustic characterization of the tinnitus spectrum: Implications for the underlying mechanisms of tinnitus". In: *Audiol. Neuro-Otology* 7, pp. 358–369. ISSN: 14203030. DOI: 10.1159/000066156.

Nourski, Kirill V et al. (2012). "Spectral Organization of the Human Lateral Superior Temporal Gyrus Revealed by Intracranial Recordings". In: DOI: 10.1093/cercor/bhs314. URL: https://academic.oup.com/cercor/article-abstract/24/2/340/347678.

Overath, Tobias et al. (2012). "Sensitivity to temporal modulation rate and spectral bandwidth in the human auditory system: fMRI evidence." In: *J. Neurophysiol.* 107.8, pp. 2042–56. ISSN: 1522-1598. DOI: 10.1152/jn.00308.2011. URL: http://www.ncbi.nlm.nih.gov/pubmed/22298830%7B%5C%%7D5Cnhttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3331610.

Oxenham, Andrew J (2013). "Mechanisms and mechanics of auditory masking". In: *J Physiol* 591, p. 2375. DOI: 10.1113/jphysiol.2013.254490. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3678028/pdf/tjp0591-2375.pdf.

Papanikolaou, Amalia et al. (2015). "Nonlinear population receptive field changes in human area V5/MT+ of healthy subjects with simulated visual field scotomas". In: DOI: 10.1016/j.neuroimage.2015.06.085. URL: http://ac.els-cdn.com/S105381191500590X/

1-s2.0-S105381191500590X-main.pdf?%7B%5C_%7Dtid=8391ef90-7386-11e7-ba46-00000aacb35f%7B%5C&%7Dacdnat=1501240904%7B%5C_%7D.

Patterson, R D (1976). "Auditory filter shapes derived with noise stimuli." In: *J. Acoust. Soc. Am.* 59.3, pp. 640–54. ISSN: 0001-4966. DOI: 10.1121/1.380914. URL: http://www.ncbi.nlm.nih.gov/pubmed/1254791.

Pickles, James O. (2012). *An introduction to the physiology of hearing*. Emerald Group Publishing. ISBN: 1780521669. URL: https://books.google.co.uk/books?id=mMTFWiLLw-cC%7B%5C&%7Ddq=an+introduction+to+the+physiology+of+hearing+reference%7B%5C&%7Dhl=en%7B%5C&%7Dsa=X%7B%5C&%7Dved=0ahUKEwjcx5Wn9Z3kAhWsSRUIHRzRBZsQ6AEIMDAB.

Pienkowski, Martin and Jos J. Eggermont (2011). "Cortical tonotopic map plasticity and behavior". In: *Neurosci. Biobehav. Rev.* 35.10, pp. 2117–2128. ISSN: 01497634. DOI: 10.1016/j.neubiorev.2011.02.002. URL: http://dx.doi.org/10.1016/j.neubiorev.2011.02.002.

Pugh, Kenneth R. et al. (1996). "Auditory Selective Attention: An fMRI Investigation". In: *Neuroimage* 4.3, pp. 159–173. ISSN: 10538119. DOI: 10.1006/nimg.1996.0067. URL: https://linkinghub.elsevier.com/retrieve/pii/S1053811996900671.

R Core Team (2013). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. URL: http://www.R-project.org/.

Scarff, Carrie J. et al. (2004). "The effect of MR scanner noise on auditory cortex activity using fMRI". In: *Hum. Brain Mapp.* 22.4, pp. 341–349. ISSN: 1065-9471. DOI: 10.1002/hbm.20043. URL: http://doi.wiley.com/10.1002/hbm.20043.

Schira, Mark M., Christopher W. Tyler, et al. (2009). "The Foveal Confluence in Human Visual Cortex". In: *J. Neurosci.* 29.28, pp. 9050–9058. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.1760-09.2009. URL: http://www.ncbi.nlm.nih.gov/pubmed/19605642%20http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC6665445.

Schira, Mark M, Alex R Wade, and Christopher W Tyler (2007). "Two-dimensional mapping of the central and parafoveal visual field to human visual cortex." In: *J. Neurophysiol.* 97.6, pp. 4284–95. ISSN: 0022-3077. DOI: 10.1152/jn.00972.2006. URL: http://www.ncbi.nlm.nih.gov/pubmed/17360817.

Schönwiesner, Marc, Peter Dechent, et al. (2015). "Parcellation of Human and Monkey Core Auditory Cortex with fMRI Pattern Classification and Objective Detection of Tonotopic Gradient Reversals". In: *Cereb. cortex* 25.10, pp. 3278–3289. ISSN: 14602199. DOI: 10.1093/cercor/bhu124.

Schönwiesner, Marc, D Yves Von Cramon, and Rudolf Rü (2002). "Is It Tonotopy after All?" In: DOI: 10.1006/nimg.2002.1250. URL: https://ac.els-cdn.com/S1053811902912504/1-s2.0-S1053811902912504-main.pdf?%7B%5C_%7Dtid=d6463c74-e392-4234-ab31-9a6aecbcc1df%7B%5C&%7Dacdnat=1537778771%7B%5C_%7D.

Sereda, Magdalena, Mark Edmondson-Jones, and Deborah A Hall (2015). "Relationship between tinnitus pitch and edge of hearing loss in individuals with a narrow tinnitus bandwidth". In: *Int. J. Audiol.* 54.4, pp. 249–256. ISSN: 1708-8186. DOI: 10.3109/14992027.2014.979373. URL: http://www.tandfonline.com/action/journalInformation?journalCode=iija20.

Sereda, Magdalena, Deborah a Hall, et al. (2011). "Re-examining the relationship between audiometric profile and tinnitus pitch." In: *Int. J. Audiol.* 50.5, pp. 303–312. ISSN: 1499-2027. DOI: 10.3109/14992027.2010.551221.

Talavage, Thomas M. et al. (2000). "Frequency-dependent responses exhibited by multiple regions in human auditory cortex". In: *Hear. Res.* 150.1-2, pp. 225–244. ISSN: 03785955. DOI: 10.1016/S0378-5955(00)00203-3.

Tamè, Luigi et al. (2019). "Reconstructing neural representations of tactile space". In: *bioRxiv*, p. 679241. DOI: 10.1101/679241. URL: https://www.biorxiv.org/content/10.1101/679241v1.abstract.

Thai-Van, Hung et al. (2003). "Enhanced frequency discrimination near the hearing loss cut-off: A consequence of central auditory plasticity induced by cochlear damage?" In: *Brain* 126.10, pp. 2235–2245. ISSN: 00068950. DOI: 10.1093/brain/awg228.

Thomas, Jessica M. et al. (2015). "Population receptive field estimates of human auditory cortex". In: *Neuroimage* 105, pp. 428–439. ISSN: 10959572. DOI: 10.1016/j.neuroimage.2014.10.060. URL: http://dx.doi.org/10.1016/j.neuroimage.2014.10.060.

Tootell, R. B H et al. (1988). "Functional anatomy of macaque striate cortex. II Retinotopic organization". In: *J. Neurosci.* 8.5, pp. 1531–1568. ISSN: 02706474. DOI: 10.1523/jneurosci.08-05-01531.1988.

Turner, R. et al. (1993). "Functional Mapping of the Human Visual Cortex at 4 and 1.5 Tesla Using Deoxygenation Contrast EPI". In: *Magn. Reson. Med.* 29.2, pp. 277–279. URL: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910290221.

Tzourio-Mazoyer, Nathalie and B. Mazoyer (2017). "Variations of planum temporale asymmetries with Heschl's Gyri duplications and association with cognitive abilities: MRI investigation of 428 healthy volunteers". In: *Brain Struct. Funct.* 222.6, pp. 2711–2726. ISSN: 1863-2653. DOI: 10.1007/s00429-017-1367-5. URL: http://link.springer.com/10.1007/s00429-017-1367-5.

Wallace, Mark N., Peter W. Johnston, and Alan R. Palmer (2002). "Histochemical identification of cortical areas in the auditory region of the human brain". In: *Exp. Brain Res.* 143.4, pp. 499–508. ISSN: 0014-4819. DOI: 10.1007/s00221-002-1014-z. URL: http://link.springer.com/10.1007/s00221-002-1014-z.

Wandell, Brian A, Serge O Dumoulin, and Alyssa A Brewer (2007). "Review Visual Field Maps in Human Cortex". In: 1893, pp. 366–383. DOI: 10.1016/j.neuron.2007.10.012.

Wandell, Brian A and Stelios M Smirnakis (2009). "Plasticity and stability of visual field maps in adult primary visual cortex". In: *Nat. Rev. Neurosci.* 10. DOI: 10.1038/nrn2741. URL: http://dx.doi.org/10.1038/nrn2741.

Wasserthal, Christian et al. (2014). "NeuroImage Localizing the human primary auditory cortex in vivo using structural MRI". In: *Neuroimage* 93, pp. 237–251. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2013.07.046. URL: http://dx.doi.org/10.1016/j.neuroimage.2013.07.046.

Westheimer, G (1975). "Editorial: Visual acuity and hyperacuity." In: *Invest. Ophthalmol.* 14.8, pp. 570–2. ISSN: 0020-9988. URL: http://www.ncbi.nlm.nih.gov/pubmed/1150397.

Wickham, Hadley (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN: 978-3-319-24277-4. URL: https://ggplot2.tidyverse.org.

Wier, C C, W Jesteadt, and D M Green (1977). "Intensity discrimination as a function of frequency and sensation level." In: *J. Acoust. Soc. Am.* 61.1, pp. 169–177. ISSN: 00014966. DOI: 10.1121/1.381278.

Wolak, Tomasz et al. (2017). "Tonotopic organisation of the auditory cortex in sloping sensorineural hearing loss". In: *Hear. Res.* ISSN: 18785891. DOI: 10.1016/j.heares.2017.09.012.