Mapping Binaural Responses in the Brain: an fMRI Study

Artificial Intelligence Master thesis

July 2022

Student: E.A. Vast

First supervisor: Dr. B.M. Harvey

Second assessor: Dr. M.J. Mulder

Abstract

The neural mechanisms involved in determining the location of sound stimuli are a controversial subject with various theories offering different explanations. One theory speculates that binaural field maps may be responsible for processing the interaural level differences (ILD) as part of these mechanisms. In this study functional magnetic resonance imaging (fMRI) data on five subjects were analysed to identify potential binaural field maps, i.e., areas of the cortical surface tuned to ILD. Eight regions of interest (ROI) were identified in and outside the auditory cortex. These eight ROIs are the auditory cortex (AC), angular gyrus (AG), supramarginal gyrus (SMG), postcentral gyrus (PCG), posterior and anterior regions of Broca's area (BP and BA respectively), middle frontal gyrus (MFG), and anterior cingulate cortex (ACC). Three models using either ILD or (a combination of) contralateral or ipsilateral loudness level to predict neuronal responses were fit to the gathered data and compared. The model using ILD achieved significantly higher goodness of fit inside the defined regions, indicating that these regions are themselves tuned to ILD. A significant correlation between the preferred ILD of voxels inside the ROIs and their tuning widths was found, with tuning widths increasing as the preferred ILD moves further away from 0. This matches the observation that humans are more sensitive to changes in sound stimuli coming from directly in front of them. The neuronal responses inside the ROIs show repeatable spatial variations in ILD preferences, although the structure of these variations was not identified. These results provide evidence for the existence of binaural field maps and serve as the basis for further research into the existence of population receptive field maps for interaural time difference or the head-related transfer function.

Contents

1 Introduction

Determining the location of sound is something most of us do very frequently and without much thought. This skill is essential for the survival of most animals but how exactly the

brain processes auditory space remains an unresolved question. Answering this question requires an understanding of where auditory stimuli are processed and which areas of the brain perform what function in this processing. The auditory cortex (AC) is always recognized as an important area, but whether areas outside the AC play a part is controversial (King, Teki, & Willmore, 2018). One popular theory states that auditory space is encoded in the brain as a distributed representation in the two ACs (Ortiz-Rios et al., 2017). According to this theory, there is no direct representation of auditory space in our brain. Instead, both the left and right AC contain information that can be combined by other areas with higher level brain functions to determine the exact location of sounds. This theory was tested in research in which awake and anesthetized monkeys were exposed to auditory stimuli created specifically for each monkey using personalized, in-ear, binaural recordings of sounds played from 12 different directions. Neural responses of these monkeys were recorded in an fMRI scanner and analysed for patterns of activation. In line with their theory, the researchers found several areas in the AC that were broadly tuned to sounds from contralateral sectors and did not identify areas outside the AC that were tuned to such sound stimuli. However, other research did identify activation outside the AC and other theories have been put forward about how auditory space is processed.

In contrast to Ortiz-Rios et al., Rauschecker and Tian theorized that the auditory system is divided into two streams of processing, each consisting of several areas in the brain, with one being responsible for determining the location of sounds and the other for analysing the sound's content (Rauschecker & Tian, 2000). They concluded that these streams exist primarily within the AC, in line with Ortiz-Rios et al., but that some processing also occurs in areas outside the ACs. This theory was supported by research conduced partly in monkeys and partly in humans. In both groups, the researchers found that distinct areas were active when either exposed to directional audio or to non-directional audio like different important vocalizations for monkeys and speech-like stimuli for humans. Both this "two-stream" theory and the "distributed representation" theory provide a somewhat high-level view on processing sound location, but neither looks at specific features of sound and how the brain might process those to determine sound location.

Research from Panniello et al. provides a more low-level, detailed insight into how the brain might process sound location. They investigated the representation of interaural level difference (ILD) and sound frequency in the primary auditory cortex of mice using 2-photon calcium imaging (Panniello, King, Dahmen, & Walker, 2018). With this technique they were able to accurately map neural responses in small areas of the brain in living mice while exposing them to short sound stimuli that vary in either frequency or ILD. By looking at these responses it becomes possible to identify the neural structures that are responsible for processing certain features of sound. Panniello et al. found that approximately 65% of noise-responsive neurons respond more strongly to contralateral ILDs, i.e., sounds that appear to come from the opposite side. They also found neurons that prefer stimuli with a centered ILD ($\sim$10% of noise-responsive neurons) and neurons that prefer stimuli towards

the ipsilateral ear (∼24% of noise-responsive neurons). The authors did not find that these neurons had a spatial structure, instead reporting that they were distributed in a "largely unorganized manner". This research offers an in-depth view of how the primary auditory cortex might be tuned to ILD but does not rule out the existence of areas outside it, like those found by Rauschecker and Tian, nor does it give a definitive answer to how such areas might be structured in humans. It also provides a strong suggestion to look at the physical features of sound and how these are processed to further our understanding of this subject.

This idea of examining the auditory system by analysing specific features of sound and how the brain processes these has been around for over a decade. Based on research into the visual system and the way images are processed there, researchers have suggested that the auditory system may work in a similar fashion (King & Nelken, 2009). In the visual system, locations of objects are projected directly on to the retina. These projections are processed in the cortex through visual field maps; maps of the spatial properties of an image that are sensitive to certain aspects of it, such as a particular angle in the image (Wandell, Dumoulin, & Brewer, 2007). Through analysing functional magnetic resonance imaging (fMRI) data from human subjects many such visual field maps have been identified and shown to be located both within the visual cortex and outside it. The visual field maps have been shown to be transformed representations of the whole observed image as it appears on the retina, at least in the primary visual cortex (V1). After V1, the visual field maps have been shown to be sensitive to certain features or parts of the image shown. This approach to representing visual input raises the question whether the auditory system could not work in a similar fashion, representing features of audio input through binaural field maps which process different aspects of sound stimuli. If it does, this could explain the finding by Rauschecker & Tian of distinct areas responding to either sound content or sound location (Rauschecker & Tian, 2000). This theory also offers a more detailed version of the distributed representation suggested by Ortiz-Rios et al. (Ortiz-Rios et al., 2017).

A major contrast with the visual system that must be noted before continuing is that the sensory organ associated with hearing is not spatially structured, thus it may not be straightforward what features exactly could be represented in auditory field maps. On the retina, objects located higher in an image are directly shown as being higher simply by being picked up by lower parts of the retina (as visual input is inverted vertically after passing through the lens of our eyes). Our ears, on the other hand, can pick up a variety of auditory features, such as pitch or volume, but auditory space is not directly represented in the ear (Grothe, Pecka, & McAlpine, 2010). The receptor cells in the cochlea respond to different frequency bands, resulting in a tonotopic structure in both the ear and the cortex (King & Nelken, 2009). Instead of a direct observation of location the sound waves entering both ears must be compared, and the location of sounds determined using the differences found in that comparison. By looking at the features of sound stimuli coming into both ears, we can find three cues that are used to determine the location of sounds. The first is that sounds originating more from one side reach the ear on that side sooner,

the interaural time difference (ITD). Secondly, sounds coming from one side are louder in the ear on that side, the interaural level difference (ILD), also called the interaural intensity difference (IID). This is the feature that Panniello et al. investigated in mice (Panniello et al., 2018). Finally, auditory input sounds different between both ears due to the way it is transformed while moving through the skull and outer ears, this is called the head-related transfer function (HRTF) (Grothe et al., 2010). The HRTF results in the same stimulus sounding slightly different inside each ear as the path it took through the skull and outer ears will be different depending on which direction the sound originates from. Contrary to the ITD and ILD, which are binaural cues defined by comparing the sounds coming into the ears, the HRTF is a spectral cue that is different in both ears and not determined by comparing the stimulus between ears. These three features are not equally useful for each type of stimulus. For example, low-frequency sounds tend to travel around the head more easily, resulting in a reduced or nearly non-existent ILD, making the ILD more useful for higher frequency stimuli (Schnupp, Nelken, & King, 2011). ITDs, on the other hand, are mostly used for lower frequencies, approximately 800 Hz or lower. Finally, the HRTF is mostly used for determining the front-back location of stimuli, as the ILD and ITD do not change significantly based on whether the sound is coming from the front or back of the head.

The current research focused on finding areas tuned to ILD. We expect these binaural field maps to be similar to the various spatial field maps identified in the visual system. Some evidence for the existence of these auditory field maps has already been found in the master thesis this research is following up on, which was written by Fleur Bongaerts (Bongaerts, 2020). By exposing participants in an fMRI to a white noise stimulus with a changing ILD which gives the impression that it moves from one side to the other, and with a high or low pure tone presented intermittently throughout that participants had to respond to, she discovered several regions of interest (ROI) that seem to fit the expected behaviour of auditory field maps. Specifically, in her paper she identified up to nine separate ROIs in each hemisphere in which the neural populations appear to be tuned to ILD, primarily outside of the AC. Although the responses found within the AC were largely contralateral, each other region showed neural populations that responded to a larger and somewhat distinct range of ILD. In addition, the preferred ILD, i.e., the ILD the neural populations in that area most strongly respond to, appeared to be gradually change systematically across the cortical surface. These findings provide evidence for the existence of auditory field maps, at least for ILD. Aside from this master thesis, the tuning of areas in A1 found by Panniello et al. in mice also hints towards spatially-tuned responses, even though they could not discern a spatial structure in these areas (Panniello et al., 2018).

In the present study we continue where Bongaerts left off (Bongaerts, 2020). By slightly changing her experiment to have the pure tone (task target) change ILD along with the white noise we allowed the participant to attend to the ILD being presented with the goal of increasing the gain on neural responses to ILD. Furthermore, the preprocessing of fMRI data has been improved to reduce noise and clean up the data by adding a step that reduces

thermal noise. We also have been able to fit and validate a greater variety of models for the neural responses, improving certainty that an accurate parametrization of neural activity is found. Through this we aimed to find evidence for the existence of binaural field maps and the spatial structure of those field maps on the cortical surface.

2 Methods

For this research, fMRI data from five participants were used. These five participants consisted of subjects who took part in both the experiment conducted by Bongaerts (2020) and the current research, participants who only took part in the research by Bongaerts (2020), and subjects who only took part in the current experiment. To maintain anonimity, the subjects are referred to by their number as it was defined in the raw fMRI data, thus we have the following subjects:

- Subject 1: Both experiments
- Subject 13: Only Bongaerts' experiment
- Subject 14: Both experiments
- Subject 19: Only current experiment
- Subject 20: Only current experiment

Having this varied group of subjects allowed us to compare the results from both experiments and observe the effects of the changed experimental design which is described in the following sections.

Gathering and analysing the data to find evidence for the existence of auditory field maps tuned to ILD was done in three steps. First, participants were placed in an fMRI scanner where they were exposed to a white noise stimulus in which the ILD gradually changes from being louder in one ear to being louder in the other, giving the impression of the sound moving from one side to the other. To ensure participants paid attention to auditory input and reduce the risk of distractions, the room with the scanner was kept entirely dark and they were asked to respond to a pure tone that can be heard intermittently between the bursts of white noise. This pure tone was the target of the task, while the white noise was the auditory stimulus we were researching the neuronal response to. Secondly, after gathering the fMRI data, it needed to be cleaned up before analysis was possible. This preprocessing included coregistration of the functional data to a highly detailed anatomical scan of the brain, noise reduction by making use of phase data gathered while scanning, and preparing the data for averaging and modeling. Finally, parametric neural response models were created and compared to the gathered data to analyse how well they fit. These models always take as input the sound pressures in the two ears and output the predicted

activation of a neural population whose response followed a particular parametric function of these sound pressures. If a model explains the variance observed in an area better than competing response models, we can conclude that that area is tuned to the ILD. These three steps and the stimulus participants were exposed to are explained in greater detail in the subsections that follow.

2.1 Data Gathering

Before participants were placed in the scanner, the auditory stimulus had to be created. This process is described in the following section. After this, the design of the experiment is described in greater detail as is the scanning setup.

2.1.1 Auditory Stimulus

The auditory stimulus used in this experiment is mostly identical to the one used by Bongaerts in her experiment (Bongaerts, 2020), with the exception of the pure tone now changing ILD to match the ILD presentation. All stimuli used in this experiment were created using the Binaural Sound Creation Toolbox (Akeroyd, 2017) inside MatLab R2019b.

First, two white noise bursts were created, 50 ms and 200 ms long respectively. These two sounds were reused for all following steps so that the stimulus does not change during the experiment except for its ILD. 50 ms of silence was added to the end of both bursts, then a 1400 ms train of bursts was created by combining four of both sounds. These combinations were made in a controlled randomized order such that each train contains exactly four 'switches', meaning there are four occurrences of a short burst following a long one, or vice-versa, in a sequence. To find all patterns that meet this criterion, a list of all possible combinations of four short and four long bursts was generated. Then, all combinations that do not contain exactly four switches were discarded. For each TR (repetition time) in a scanning session, a pattern from this final list was selected. The selection was done in such a way that each pattern occurred approximately the same number of times and the patterns were distributed randomly throughout the scanning session. Finally, 100 ms silence was appended to each pattern so it fully filled the 1500 ms TR used in this experiment.

Having generated the patterns, the task for the participant needed to be added. For some of the sound trains the last 100ms of silence was replaced by a high (2000 Hz) or low (500 Hz) pure 'oddball' tone, followed by 50 ms of silence, this is the target for the participants to focus on. These oddballs were distributed across all runs and cycles such that each oddball (high and low) was presented exactly once per TR in the cycle and two oddballs were never presented directly after one another. The oddballs thus always had at least one TR without an oddball separating them. This step is taken to ensure the gathered data is balanced and the oddballs do not influence the final, average result that is

taken during the modeling stage. In the experiment by Bongaerts (Bongaerts, 2020), these oddballs were excluded from the ILD that got applied to the white noise sequence, but in this version the ILD also gets applied to the oddball. The steps up to here are shown in figure 1.

The final step was to apply ILD to the stimulus sequences. The target ILDs ranged from -22 dB to +22 dB, with a 0 ILD indicating an identical level in both ears and the midpoint of the ILD range. The ILD changes in log space with 16 steps between these extremes of the ILD range and +/-1 dB. An ILD value of +0.5 dB and -0.5 dB was introduced on both sides of the midpoint to smooth the transition from +/-1 dB to 0, combining to a total of 35 different possible ILD values. ILD was applied in such a way that the stimulus was always as loud as possible without going over the maximum loudness set by the participant. This means first subtracting half the maximum ILD (11 dB) from the maximum loudness set by the participant, then for the left ear half of the target ILD is subtracted from that value while for the right ear half of the target ILD is added to that value. For example, if the participant has set a maximum loudness of 110 dB and the target ILD is +6 dB, then for the left ear the volume will be 96dB ($110 - 11 - 3$). And the right ear volume will be 102dB ($110 - 11 + 3$). As such, the average sound pressure in the two ears will be constant (99 dB here) and independent of ILD: maximum loudness (110dB here) minus half the maximum ILD (11 dB).

The auditory stimuli were presented in cycles, where the ILD moved from -22 dB to +22 dB in logarithmic space and back again. The stimulus repeated four times at -22 dB and +22 dB in both directions. This approach resulted in a more dense sampling of ILDs in the center of auditory space (closer to 0 dB difference) compared to those further in the periphery (closer to -22 dB or +22 dB). A visualization of a single cycle can be seen in figure 2 A single scan run in this experiment consisted of four repeats of the full ILD cycle.

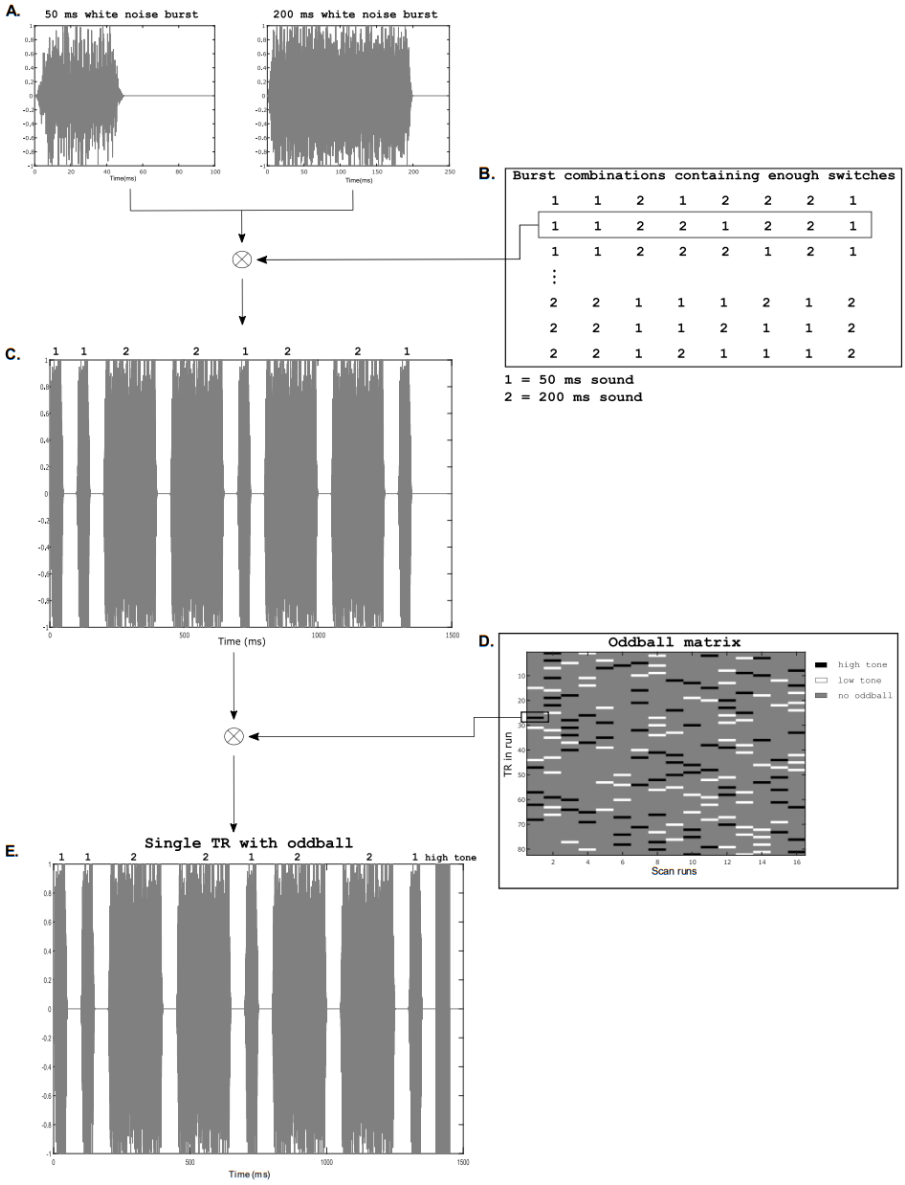

Figure 1: Creation of the white noise patterns and the oddballs in a single ILD presentation. (A) 50 ms and 200 ms white noise bursts. (B) The possible combinations of these bursts that include exactly four switches. (C) An example sequence of white noise bursts with 100ms silence at the end. (D) Different sequences of oddballs for different scan cycles. (E) An example sequence including a high pure tone at the end.
Figure taken from (Bongaerts, 2020).

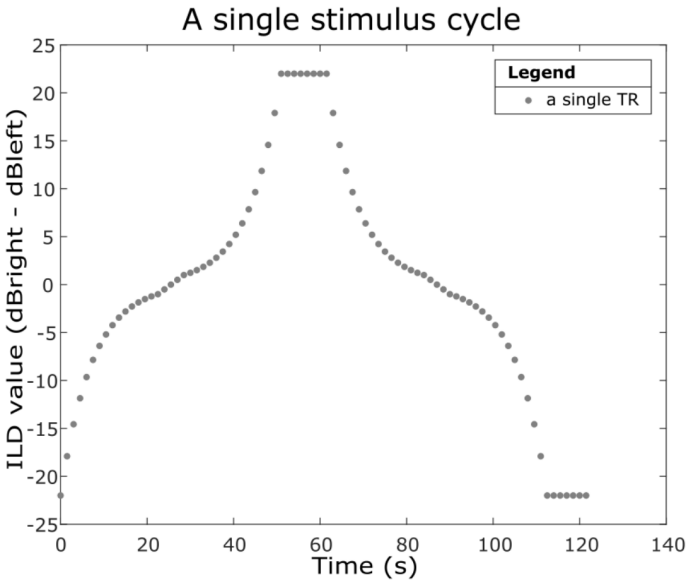

Figure 2: A single stimulus cycle with each dot representing a single TR of 1500 ms and the ILD value for both the white noise and any eventual oddball tones being shown on the left. The timing of oddball tones is not shown as that changes between cycles.
Figure taken from (Bongaerts, 2020).

2.1.2 Experimental Design

After the participant was placed in the scanner and they had set their preferred maximum level of loudness, the experiment could start. This part is again very similar to the design used by Bongaerts (Bongaerts, 2020), although for the new scans we were only interested in the binaural condition, meaning the stimulus was presented in both ears with the ILD applied as described above.

Data for participants was collected across multiple scanning sessions aiming for eight runs per session. The aim was to have at least two sessions in this new experiment both for new participants and those that already participated in the previous experiment by Bongaerts (Bongaerts, 2020).

2.1.3 Scanning Setup

As in the paper by Bongaerts, functional data were acquired at the Spinoza Centre for Neuroimaging in Amsterdam, Netherlands on a 7-Tesla (7T) Philips Achieva MRI scanner (Bongaerts, 2020). For each participant T1-weighted anatomical MRI data were acquired for use during the preprocessing. The blood oxygen level dependent (BOLD) responses during the experiment were acquired in the form of T2*-weighted functional images. For these images, a 32-channel head coil was used with a voxel size of 1.69x1.69x1.7mm and a

single TR was set to 1500 ms. 57 interleaved slices of 128x128 voxels each were captured, giving a field of view of 216x216x97 mm. Each functional run consists of 336 time frames, meaning it lasts 502 seconds. The first 12 seconds of frames of each scan were discarded to allow magnetization to reach a steady state with a person inside.

2.2 Preprocessing

Preprocessing refers to the process wherein we took the collected functional data and aligned it to the high-resolution anatomical scan for each participant, the coregistration of functional data. During this process, we also corrected for several sources of distortion in the functional data to minimize the influence of these sources during the analysis step. This process was broadly done in two steps. First, the individual steps required to coregister the functional data are calculated, each of which are explained in greater detail below. Next, the required transformations were combined into a single matrix which captures all these steps and then applied to the original functional data, giving us the final output for this process. This approach was taken as each spatial transformation and resampling of the data slightly blurs the data, reducing its effective spatial resolution. Combining the different transformations into a single step minimized this as the data are only resampled once from the original acquisition to the final anatomical space. After the individual sessions were preprocessed into this final anatomical space, the scans were imported into Vistasoft's mrVista framework (Wandell, 2014), which was used for analysis and model fitting. Preprocessing was done using AFNI (Cox, 1996), which is a package of computer programs created for analysis and visualization of three-dimensional human brain fMRI scans. The specific functions used for each step are given in the detailed descriptions below.

Aligning and denoising the functional data was done in four or five steps, depending on whether we collected phase data while scanning the participant. These steps are as follows:

- Thermal noise reduction (only if phase data was collected)
- Correcting for motion during scans
- Undoing distortions caused by the directionality of the magnetic field
- Aligning the functional data to the low-resolution anatomical scan
- Aligning the low-resolution to the high-resolution anatomical scan

2.2.1 Thermal Noise Reduction

The movement of molecules both inside the scanner and the scanned tissue introduces noise into the acquired data. If phase data has been collected during the scan, which was an option at the scanner we used, we were able to separate the noise from the functional

data. This noise then got subtracted from the functional data. As this is not a spatial transformation, this step was performed once and the denoised data was used for each following step. Compared to the previous research by Bongaerts (Bongaerts, 2020), this was an entirely new step that was not applied to any of the data collected during that project. The rest of the steps have remained identical.

2.2.2 Motion Correction

Next, we correct for participant motion during scans. As participants cannot lie perfectly still we expect slight movement to be visible within the data, between volumes of the scan run. Next to these minute movements, participants also shift over time in the scanner, causing the position of their head to change between runs collected in a single session. During this step we corrected for the first type of movement, which means that the goal for this step was that between volumes in a scan run the location of the head remains the same, but there could still be changes in this location between scan runs in the same session. The second type of movement is corrected for during the coregistration to high-resolution anatomical space. To achieve this, the function 3dvolreg was used to align the volumes in each scan run to the last volume of that scan run. This was also done to the top-up, which are quick anatomical scans collected using opposing magnetic gradients which will be used in the next step. Top-ups are very quick and thus motion of the participant was usually negligible, however the computational cost of 3dvolreg was also limited, so the top-ups were included here to minimize noise during the next step. The transformation for each volume calculated in this step was stored to be combined at the end into a single spatial transformation.

2.2.3 Distortion Reduction

Having corrected the scans for participant motion, the next step was to correct for the spatial warping caused by the direction of the magnetic field in the scanner. Before performing this step, the functional data looked as though it was stretched in one direction, which is a known phenomenon in MRI data (Doran, Charles-Edwards, Reinsberg, & Leach, 2005). To properly correct for this distortion, we made use of top-up scans. These are quick scans collected using an opposing gradient compared to the functional data, which means their shape was also distorted in the opposite direction. As motion correction has already happened here, we take the average of both the top-ups and functional scans over time for computational efficiency. Next 3dQWarp is applied to the scan average and top-up average. This function calculates the transformation required to move from the functional scan average to the top-up average and then returns the median between that and no transformation. This result then can be applied to give approximately an unwarped brain shape, as the midpoint between the top-up and the functional scan should give the unwarped spatial structure of both scan types. The resulting spatial transformation matrix

was stored and then applied to the motion corrected functional data using the function 3dNWarpApply, which returned the unwarped versions.

2.2.4 Coregistration

At this point the data was free of participant motion and spatial warping, so we could move on to coregistration of the functional data to the high-resolution anatomical scan. Calculating this transformation in a single step was prohibitively difficult as the functional data looks different from the anatomical scan, lies in an entirely different space, and was usually collected with different settings, thus resulting in a search space so large that there was no realistic solution to this problem. To resolve this, we first calculated the transformation required to align the functional data to the low-resolution anatomical scan (T1Matched), which was collected during the same session and using the same position, orientation, and resolution as the functional data. This means the transformation was much smaller and thus the problem was less complex. Solving this problem for each scan run individually would still be computationally demanding, so instead we first corrected for the motion between scan runs. This process was the same as described in the previous section on motion correction, except now the scan runs are compared to each other and aligned using 3dvolreg. Following this, the average of these scan runs is taken and the coregistration to T1Matched is calculated. Having found the required transformation for each scan run, these were stored, applied to each scan run and the final step could begin.

The last step was very similar to the previous one, only now the transformation we were interested in was going from the T1Matched to the high-resolution anatomical scan collected for each participant. These two scans are very similar in terms of their tissue contrasts and they are both anatomical scans undistorted by the magnetic gradient of the scanner. The transformation required thus was only a change of coordinates between nearly identical scans. After finding the transformation, the resulting spatial transformation matrix was again saved and could be applied to the output of the previous step to give the end result.

2.2.5 Finalizing Output

As you can tell, many transformations occur during this process, each of which slightly blurs the data. To minimize this blurring, a single large transformation matrix was created which combined all the steps described above except the thermal noise reduction. This final matrix was applied to either the original data, or the thermal noise reduced data, depending on which was available. This output was the functional data coregistered with the high-resolution, whole-brain anatomy, ready for analysis. The output of this process was visually checked to ensure no errors occurred and no low quality data was accidentally entered as input, for every subject and each scan run. If the end result was deemed to be properly aligned and free of unexpected distortions, it was then imported into the mrVista

framework (Wandell, 2014), where the modeling and analysis would take place.

A visual representation of the steps described above can be found in figure 3. The final step of combining the matrices and applying that has been excluded from that image for visual clarity.

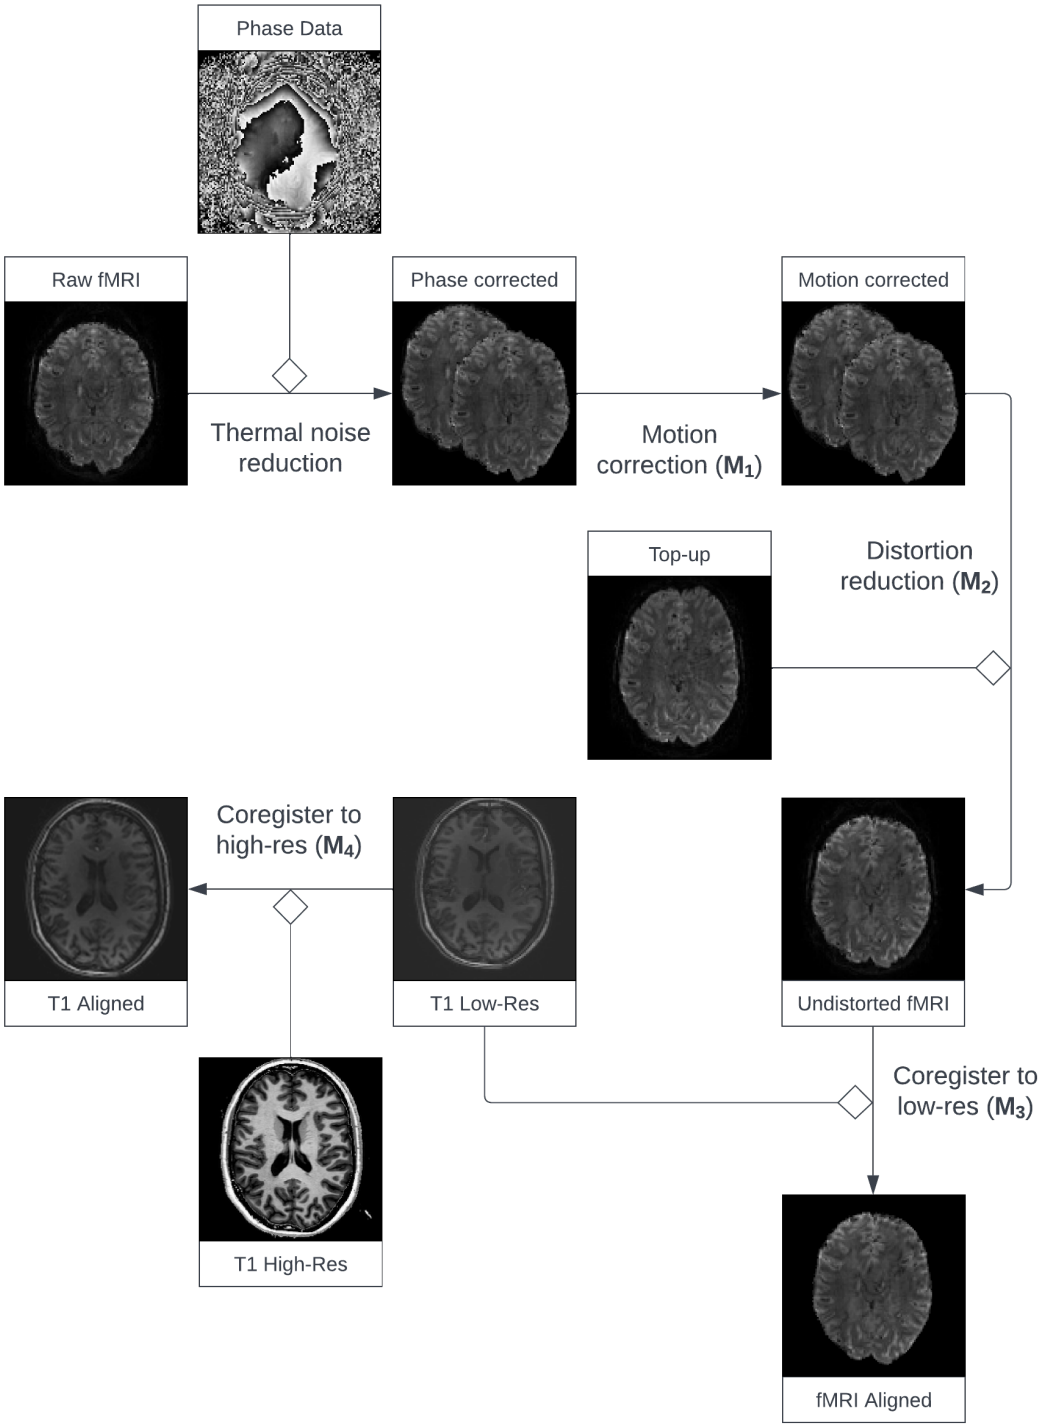

Figure 3: Steps taken while preprocessing fMRI data. Each arrow indicates a transformation of the source to the next step while a diamond head arrow indicates that the file is used without transforming in this step. M1 through M4 are the transformation matrices created in these steps which eventually get combined and applied to the raw fMRI data, this is not shown in this figure. The "Phase corrected" and "Motion corrected" images show two volumes from a single scan session to indicate that multiple volumes are used in these steps to correct for the motion of a participant during this step.

2.3 fMRI Analysis

The collected fMRI data was analysed by first modeling the ILD responses using a population receptive field (pRF) modeling approach as described in Harvey & Dumoulin (2017) and first used by Harvey et al. (2013), after which regions of interest (ROIs) were defined on the cortical surface and the existence of auditory field maps was confirmed using a similar approach as described in Harvey, Dumoulin, Fracasso, & Paul (2020).

2.3.1 Aggregating Data

Before modeling could begin, we first aggregated the various scan runs together. Three averages were created for both types of stimulus used, one with all the scan runs that match that stimulus type, one with the odd numbered scan runs matching that type, and one with the even scan runs matching that type. The stimulus type is defined by either having a static target for the participants or a target that changes ILD along with the white noise stimulus. An aggregate consists of the average response for each recording site over time in a single stimulus cycle given the scan runs that went into that aggregate, thus we averaged both over the stimulus cycles in each scan run and over the different scan runs. This meant that after creating the aggregates the subjects that took part in both Bongaerts' (2020) experiment and the current research had MovingBinaural and StaticBinaural aggregates, the participants that only took part in Bongaerts' (2020) experiment had StaticBinaural aggregates, and those that only took part in the current research had MovingBinaural aggregates. This process allowed us to crossvalidate the candidate models on unseen data as we fit models on each aggregate separately and the odd and even aggregates never overlap.

2.3.2 Defining Regions of Interest

One of the main goals of the current research was to find evidence for the existence of binaural field maps. To find this evidence, regions of interest (ROI) needed to be defined. These ROIs were determined using the linear Gaussian model, which is described in greater detail in section 2.4.3 (Fitting Models). We used this model as we had a strong hypothesis that this would best fit inside the binaural field maps we were looking for, as it is fit using the ILD level as input. These field maps would present as areas where the goodness of fit of the linear Gaussian model is noticeably higher than outside those areas. To identify the regions of interest (ROI), we first rendered the goodness of model fit for onto the cortical surface, this can be seen in figure 4 for a single subject. Based on this rendering, eight ROIs were defined that we wanted to investigate further. These 8 are the auditory cortex (AC), angular gyrus (AG), supramarginal gyrus (SMG), postcentral gyrus (PCG), posterior and anterior regions of Broca's area (BP and BA respectively), middle frontal gyrus (MFG), and anterior cingulate cortex (ACC). Note that these ROIs did not exactly overlap with the anatomical location they were named for, instead these names were given

for the proximity they held to the associated locations. BA and BP were named thusly as they were generally seen as a smaller area in front of and behind Broca's area.

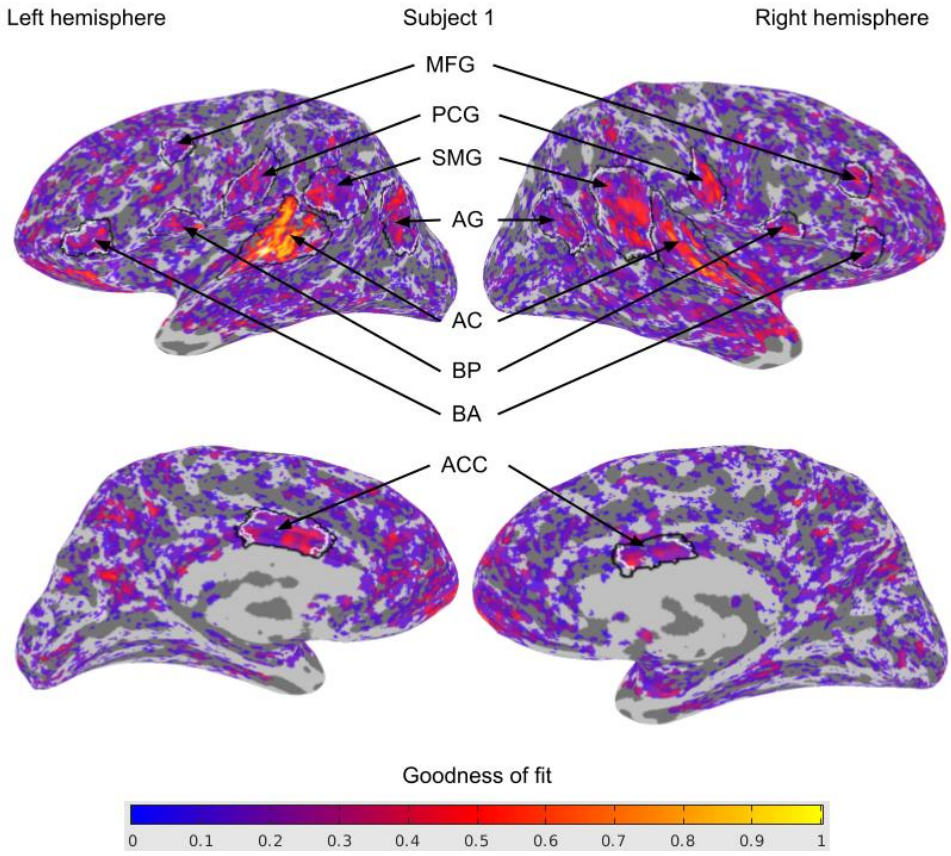

Figure 4: Goodness of fit for subject 1. The variance explained threshold for a voxel showing up is set to 0.1. The empty space on the medial side is where the corpus callosum would be when rendering both hemispheres simultaneously. The 8 ROIs can be seen here with their boundaries in black and white. These are the middle frontal gyrus (MFG), postcentral gyrus (PCG), supramarginal gyrus (SMG), angular gyrus (AG), auditory cortex (AC), Broca posterior (BP), Broca anterior (BA), and anterior cingulate cortex (ACC).

Aside from the location of the ROIs, we also investigated whether a structure existed in these areas, i.e., whether the preferred ILD changes from one side of the region to the other. This was done as Harvey & Dumoulin (2017) found that cortical maps often follow a linear progression of preferred values as you move from one edge to the other. To determine the direction of this structure, we projected the preferred ILD of the response model onto the cortical surface, this render for one subject can be seen in figure 5. For each of the ROIs identified above, except the AC, we then defined a high and low side of the boundary

of the area. This high and low side followed the values of ILD as defined in section 2.1.1 Auditory Stimulus. The AC was excluded from this step as it showed no structure in its response, instead the AC on both hemispheres displayed a strong contralateral response. The renders for goodness of fit and preferred ILD for the other participants can be found in Appendix A.

As we visually identified the regions where the variance explained was clearly higher than outside those regions, this process also provided evidence that the models fit better inside the ROIs compared to outside the ROIs.

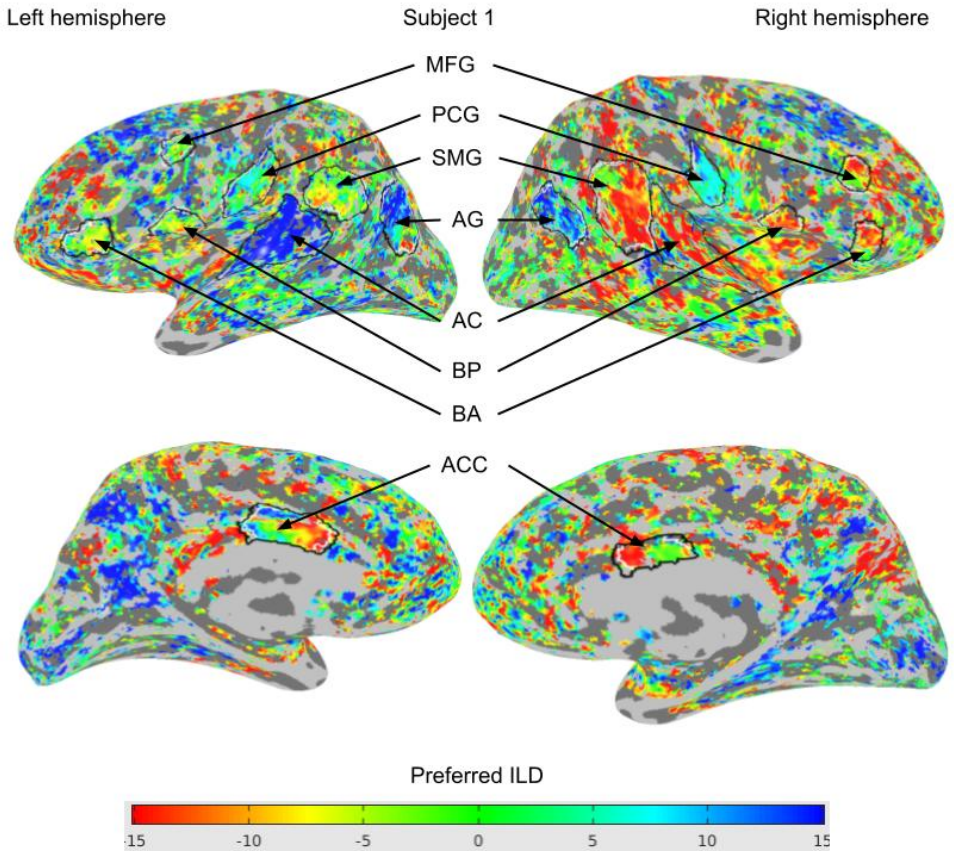

Figure 5: Preferred ILD for subject 1. The variance explained threshold for a voxel showing up is set to 0.1. The 8 ROIs can be seen here with their boundaries in black and white. The white boundaries mark their high (blue) or low (red) boundary which were later used to determine the structure of the ROI.

2.3.3 Fitting Models

As described in the paper by Harvey & Dumoulin (2017) we created pRF models, which describe the tuning of neural populations within each grey matter voxel of interest. Such models were created for each participant and aggregate of scan runs in those participants using forward modeling. The models predicted the neural response for the different neural populations at each time point in a stimulus cycle as it is shown in figure 2. Three candidate models were created for each aggregate which predict the BOLD responses in the neuronal population in each voxel using different relationships between neural response amplitude and the dB level in each ear. The linear Gaussian functions were characterized by: (1) a preferred ILD level, i.e., the Gaussian's mean, and (2) a tuning width, i.e., the Gaussian's standard deviation. Compressive monotonic functions used a variable exponent to predict neural response amplitudes from the contralateral dB level. Due to the nonlinear nature of these functions they had to be fit once on the dB level of the left ear for the right hemisphere and once on the dB level of the right ear for the left hemisphere. Finally, the compressive balance functions were based on the idea that in a single voxel, there might be neuronal populations that respond to contralateral dB level and populations that respond to ipsilateral dB level. To model this both the left and right ear loudness level were taken as input, then for the voxel a weighted sum of two compressive monotonic functions were taken. To fit these models, the exponent for both functions is assumed to be identical and then this exponent and the weight for the sum of the functions is altered to achieve the best fit.

Each of these functions represented a different hypothesis for how the observed BOLD responses came to pass, and whichever fit best across the different ROIs could be taken as the best explanation for the observed data. If the linear ILD Gaussian function best fit the data, this implied that the ILD level is indeed the input for the neuronal populations, thus providing evidence for integrated binaural responses. The compressive monotonic functions fitting best would imply that neural response amplitudes nonlinearly scaled with contralateral loudness level and not ILD. The compressive balance functions having the highest variance explained would then imply that the ROIs are made up of neuronal populations that respond to both the contralateral and ipsilateral dB level.

These models were fit by starting with a candidate model with a set of free parameters, such as the mean and standard deviation in the case of the Gaussian distributions. This candidate model was then given the ILD or dB level over a single stimulus cycle as an input. This step resulted in a predicted neural response time course for each combination of free parameters. This predicted neural response was convolved with a haemodynamic response function (HRF) to generate a candidate fMRI time course. The optimal scaling between this prediction and each voxel's response was then solved, allowing us to calculate the sum of squared errors between each scaled candidate fMRI time course and the aggregate measured neural response. These steps were repeated with a wide range of candidate parameters, after which the parameters with the lowest sum of squared errors were chosen

as the best fitting model for this neural population. A visualization of this process for the Gaussian models for a single voxel can be seen in figure 6. This fitting process was repeated for each participant and each aggregate.

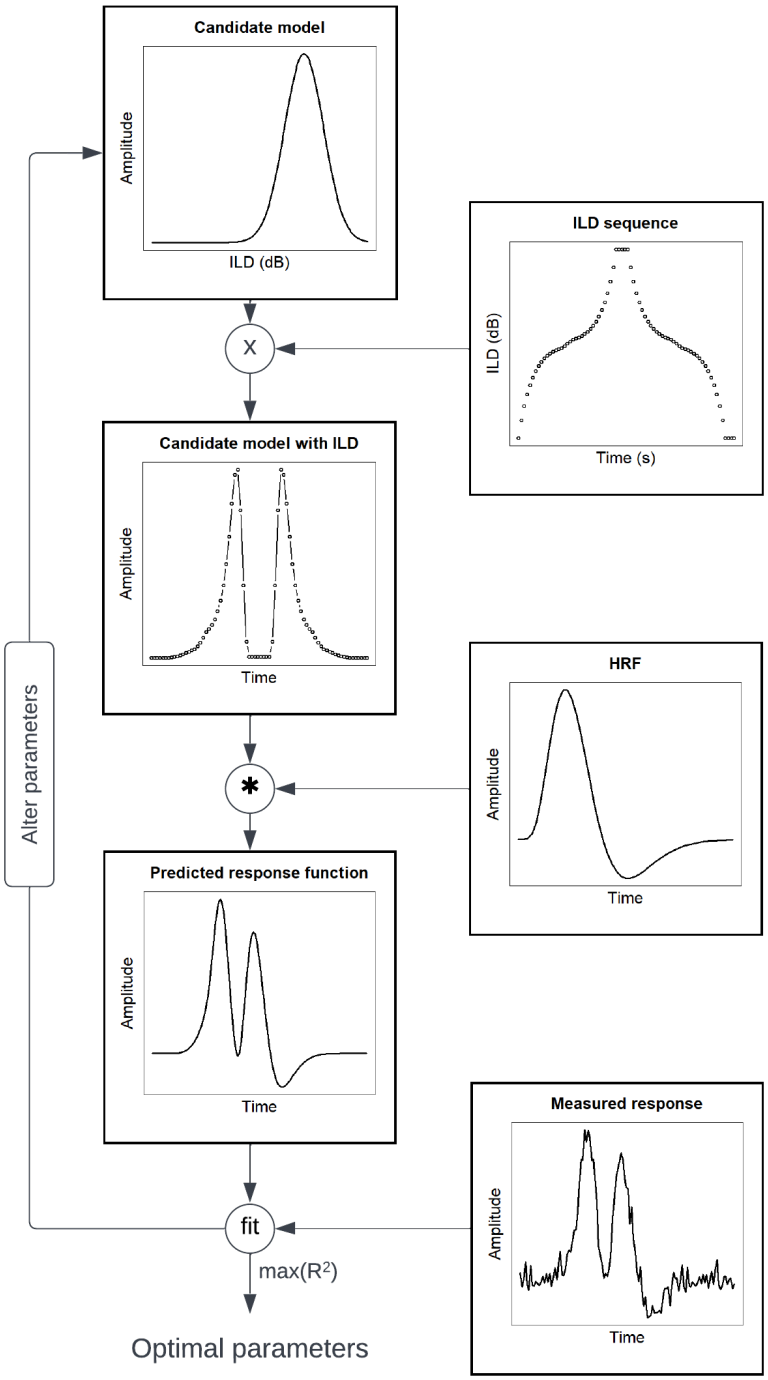

Figure 6: fMRI modeling process for linear Gaussian models. Starting with the candidate parametric neural response model at the top, the ILD sequence is used as input for this model. The resulting model is convolved with the haemodynamic response function (HRF), giving a predicted response function. This predicted response function gets compared to the measured response for the voxel in question. This entire process is repeated with a range of parameters (mean and standard deviation) and the model with the highest R^2 is chosen.

2.3.4 Comparing Models

To determine which of the candidate models best predicts the data we used crossvalidation. As described, we fit the various models on each participant and each aggregate in those participants, thus we had models fit to one half of the scans for a particulur stimulus type (e.g., even numbered scan runs) that could be evaluated on the complementary half of scans from the same stimulus (e.g., odd numbered scan runs). During this evaluation, the prediction had 0 degrees of freedom regardless of model complexity, thus allowing for direct comparison between the various models. For an unbiased choice of voxels to compare, we took all voxels inside the defined ROIs where any model explained more than 20% of the variance. We then calculated the mean variance explained for each model per ROI per odd and even split for each participant. Then we used a paired t-test to take all these values and see which models performed significantly better than others. This paired the fits for different models in the same data half, in the same ROI, and in the same participant to each other. Comparing all such pairs for each pair of candidate models gave a single t-value for this comparison.

As the candidate models were originally fit using identical HRF parameters and these parameters actually differ slightly between subjects (Harvey et al., 2013) we estimated subject-specific HRF parameters for the model that best explained the data and then re-fit the response models using these parameters. This was only done for the model with the highest goodness of fit as this is a computationally demanding and time consuming task.

Having found the best performing model and fitted the subject-specific HRF, all further analysis was done using that model.

2.3.5 Determining Preferred ILD and Tuning Width in ROIs

To investigate whether different ROIs were tuned to different ILD levels histograms were created of the preferred ILD of voxels inside each ROI. To ensure this analysis was limited to voxels with clear ILD tuning, a minimum variance explained threshold of 20% was used. Two sets of histograms were created, one for the StaticBinaural aggregate and one for the MovingBinaural aggregate. This allowed us to compare the preferred ILD and see what the effect of the altered experiment was on preferred ILD inside ROIs.

Aside from the preferred ILD, we also wished to investigate whether the tuning width was correlated to the preferred ILD. Finer tuning widths were expected with a close to central (0) ILD as that represents the area in front of us where we tend to be focused on. Smaller tuning widths imply an increased sensitivity of neural responses to small changes in ILD, so a more fine-grained processing of incoming stimuli. To investigate this, a linear function was fit to the absolute preferred ILD values and their associated tuning widths. For this function, the 95% confidence interval was calculated to show the spread of the data. A permutation test was performed to test whether the slopes found inside each ROI were significantly different from zero. For this test, a p-value smaller than 0.05 was

considered significant, after False Discovery Rate correction to account for the multiple tests conducted in different ROIs (Benjamini & Hochberg, 1995).

2.3.6 Determining Structure of ILD Preference

As Harvey & Dumoulin (2017) found that cortical maps often follow a linear progression from one end to the other, we also investigated whether this held true for the defined ROIs in our participants. The preferred ILD of ROIs and progression of preferred ILD in ROIs were analysed using a similar method to Harvey et al. (2020). The AC has been excluded from this analysis as no structure was visible in that ROI in any participants, so 71 ROIs were investigated for their structure.

First, we used the identified ROIs to determine how consistently a significant correlation could be found by comparing the preferred ILD values for the voxels inside each ROI for the odd and even aggregates. This allowed us to investigate for how many ROIs a repeatable spatial variation in ILD preferences exists, without identifying the spatial structure of that variation. For the participants with MovingBinaural aggregates, those were used, otherwise the StaticBinaural aggregates were used. For this analysis, the Pearson correlation coefficient (r) was used and a p-value smaller than 0.05 was used as the cut-off for whether the correlation was significant.

Next, we investigated the likelihood of finding a similar set of correlations if the data were randomly distributed. Randomly distributed data should show negative and positive correlations between odd and even aggregates equally frequently. Therefore we took all the correlation coefficients from the previous step and performed a paired t-test of their distribution against an r-value of 0. Here a single p-value was returned and if that value was smaller than 0.05 we could conclude that the data consistently showed a repeatable spatial variation.

Aside from figuring out whether a repeatable spatial variation of ILD preferences exists, we also wished to test the hypothesis that this preferred ILD changes linearly as you move from one boundary to the other. This was done by plotting the preferred ILD over the cortical distance from one boundary to the other. The resulting plots gave us an r-value and p-value for each ROI, where again it was determined to be a statistically significant correlation if the p-value was smaller than 0.05.

3 Results

3.1 Discovered Regions of Interest

Eight ROIs were identified per hemisphere, resulting in a total of 16 ROIs that were mostly consistently identified in each participant. These eight ROIs are the auditory cortex (AC), angular gyrus (AG), supramarginal gyrus (SMG), postcentral gyrus (PCG), posterior and anterior regions of Broca's area (BP and BA respectively), middle frontal gyrus (MFG), and

anterior cingulate cortex (ACC). In all five participants, the AC, AG, SMG, PCG, and MFG could consistently be identified in both hemispheres. There were four exceptions, which are the ACC, which was not identified in the right hemisphere of subject 13, BP, which was not identified in the left hemisphere of subject 19 or in either hemisphere of subject 20, and BA, which was not identified in the left hemisphere of subject 20. Therefore, we identified 76 out of 80 potential ROIs (5 participants x 16 ROIs) that could be used for further analysis. The resulting ROIs for subject 1 can be seen in figure 4 and figure 5, the other participants are included in Appendix A.

3.2 Modeling Comparisons

As described in the methods section, for each voxel in each aggregate three models were fit. These are the linear Gaussian model, the compressive monotonic model, and the compressive balance model. Note that the compressive monotonic model had two versions, one for the left hemisphere and one for the right.

The results of the paired t-test between the models over all participants can be seen in figure 7 for the StaticBinaural aggregate and in figure 8 for the MovingBinaural aggregate. In both situations the linear Gaussian model significantly outperformed the other models without fitting the HRF. Fitting the subject-specific HRF for the linear Gaussian significantly improved its performance compared to the version without the HRF fit, thus this model was used from here on out for further analysis.

The variance explained is consistently higher for the models with the MovingBinaural aggregates. This could be due to either the change in the stimulus compared to Bongaerts' (2020) experiment, as the stimulus used in the current research changed ILD along with the white noise, or due to the addition of the thermal noise reduction step in the preprocessing pipeline, or a combination of these two.

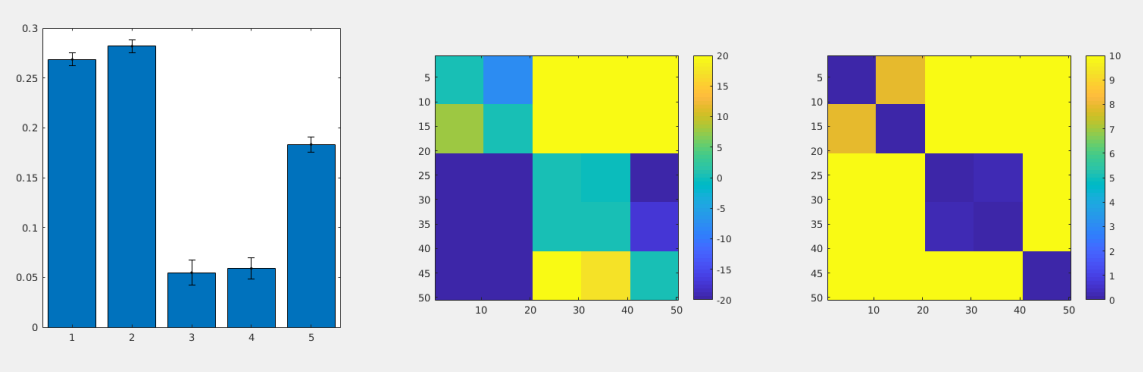

Figure 7: Performance metrics of the models for StaticBinaural aggregates. From left to right the graphs show the variance explained of each model with 95% confidence intervals, a matrix comparing each model against each other model with the t-value of the comparison, and a matrix comparing each model against each other model with the $-\log_{10}$ p-value of the comparison. The t-values were limited to [-20, 20] and any $-\log_{10}$ p-value greater than 10 was set to 10. The models shown are (1) the linear Gaussian model, (2) the linear Gaussian model with HRF fit, (3) compressive monotonic model for left hemisphere, (4) compressive monotonic model for right hemisphere, and (5) compressive balance model. All model comparisons reached significance at $p < 0.05$ except for the comparison between model 3 and 4, with a p-value of 0.67. This shows the linear Gaussian model outperforms all other base models for these aggregates and that it is significantly improved after fitting the subject-specific HRF.

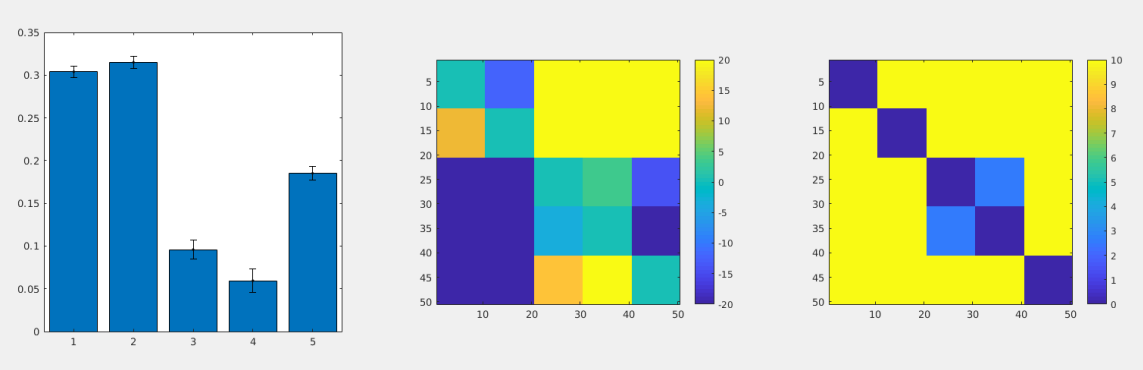

Figure 8: Performance metrics of the models for MovingBinaural aggregates. The images and model order are the same as in figure 7. All model comparisons reached significance at $p < 0.05$ including the comparison between model 3 and 4, with a p-value of 0.0022. This shows the linear Gaussian model outperforms all other base models for these aggregates and that it is significantly improved after fitting the subject-specific HRF.

3.3 Preferred ILD within ROIs

In figures 9 and 10 the histograms can be seen showing the number of voxels for each preferred ILD inside the different ROIs for the StaticBinaural and MovingBinaural aggregates respectively. Only voxels for which the model reached 20% variance-explained are included. In both cases, the preferred ILD of the AC is very strongly contralateral. For all other ROIs the responses are less clearly contralateral, with a larger fraction of non-contralateral or even ipsilateral preferred ILD values. Notably, the amount of voxels that are included for both the left and right ACC in the MovingBinaural aggregate. This seems to be the result of the changed stimulus compared to the experiment by Bongaerts (2020). Aside from this observation, it can clearly be seen that the various ROIs exhibit different preferred ILDs and that they are responsive to a wider range of ILD values, which is in line with the expected behaviour for binaural field maps.

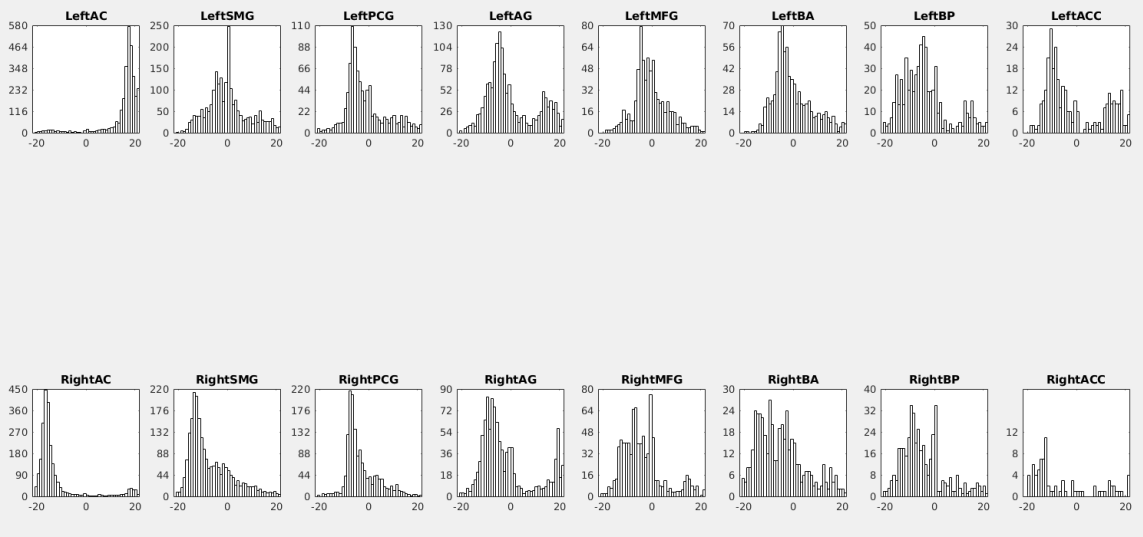

Figure 9: Preferred ILD for ROIs for the StaticBinaural aggregate. Only voxels for which the variance explained exceeded 20% were included in these charts. The responses in the AC are strongly contralateral, while the other ROIs display non-contralateral responses as well with some even showing a large fraction of ipsilateral responses such as the AG.

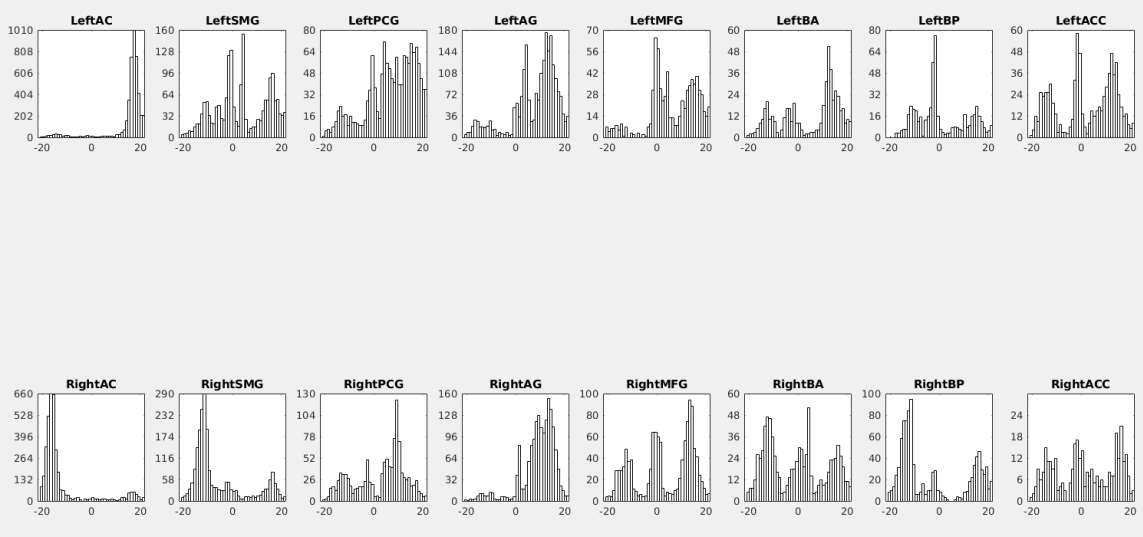

Figure 10: Preferred ILD for ROIs for the MovingBinaural aggregate. Only voxels for which the variance explained exceeded 20% were included in these charts. The responses in the AC are strongly contralateral, while the other ROIs display non-contralateral responses as well. Of note in this image is the increased number of voxels included in the ACC compared to the StaticBinaural aggregate and the balanced response visible for those ROIs. This seems to indicate that this region is more active in the current research due to the target changing ILD along with the white noise.

3.4 ROI Structure

Having determined the preferred ILD inside the ROIs, the next step was to analyse whether the preferred ILD inside each ROI is spatially structured. The first step in investigating the structure of the preferred ILD inside ROIs was to identify whether a significant correlation existed between the preferred ILD estimates from odd and even aggregates. The correlation coefficients found can be seen in table 1. The ROIs with a significant correlation between the odd and even aggregates have been coloured green. From this table we can see that a repeatable structure seems to exist in 22 out of 71 ROIs (28,9%). This appears to vary between ROIs: the right SMG and left PCG (near the primary auditory cortex) show a significant correlation relatively consistently, while the left MFG and left ACC (in the frontal lobe) only show significant correlation in at most one participant. The defined ROIs for subject 19 also show a significant correlation in 10 out of 14 cases. This result shows that there does seem to be some structure present in at least two ROIs, although a larger pool of participants would be needed to draw a stronger conclusion.

ROI	s01		s13		s14		s19		s20	
	r	p	r	p	r	p	r	p	r	p
Left SMG	0.180	0.131	0.386	0.055	-0.124	0.021	0.548	0.000	0.022	0.785
Right SMG	0.458	0.000	0.512	0.000	0.234	0.003	0.390	0.000	0.065	0.570
Left PCG	0.301	0.018	0.305	0.136	0.013	0.889	0.413	0.000	0.271	0.010
Right PCG	0.190	0.100	0.131	0.379	-0.094	0.346	0.325	0.002	0.331	0.001
Left AG	0.432	0.004	0.238	0.023	0.139	0.121	0.020	0.788	-0.069	0.243
Right AG	-0.406	0.046	-0.054	0.622	0.323	0.000	0.374	0.006	0.175	0.078
Left MFG	-0.229	0.668	-0.290	0.052	-0.121	0.254	-0.054	0.704	0.125	0.253
Right MFG	0.578	0.067	0.371	0.042	-0.158	0.076	0.226	0.002	0.124	0.198
Left BA	-0.028	0.922	-0.200	0.254	0.005	0.959	0.480	0.001	—	—
Right BA	0.367	0.104	-0.100	0.696	0.181	0.065	0.377	0.000	0.157	0.238
Left BP	0.730	0.000	0.484	0.242	-0.179	0.055	—	—	—	—
Right BP	0.254	0.190	0.448	0.117	-0.021	0.882	0.572	0.000	—	—
Left ACC	0.363	0.143	-0.1423	0.883	0.041	0.695	0.286	0.036	0.303	0.203
Right ACC	0.663	0.943	—	—	0.157	0.237	0.220	0.325	0.224	0.388

Table 1: Correlation coefficients and their associated probabilities for the comparison between odd and even aggregates inside each ROI. Blank ROIs were not identified in those participants. All cells where the correlation reached significance at $p < 0.05$ have been coloured green. 22 out of 71 ROIs showed a repeatable structure of preferred ILD.

Taking the r-values from the previous step, we performed a paired t-test against an r-value of 0. This returned a p-value of $2.1 * 10^{-5}$, therefore we conclude that the data consistently showed some repeatable spatial variation: positive correlations are more common than negative correlations. This conclusion does not imply a specific spatial structure to this variation.

Finally, we investigated whether the ROIs showed a linear progression in preferred ILD, as was found in Harvey & Dumoulin (2017) for their cortical maps of numerosity preferences. A significant correlation between preferred ILD and cortical distance was found in 13 out of 71 ROIs (18.3%). The results of this analysis for subject 1 can be seen in figure 11 and the results for the rest of the participants are included in Appendix B.

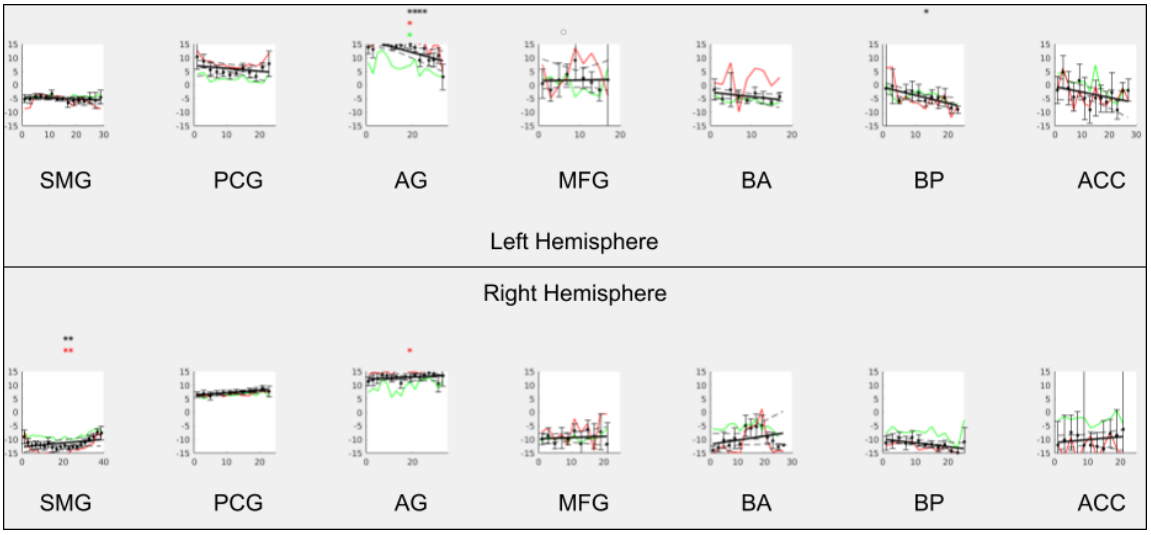

Figure 11: Preferred ILD plotted over cortical distance for subject 01. The stars over the plots show whether the correlation reached significance in the ROI in question. $* : p < 0.05$, $** p < 0.01$, $*** : p < 0.001$, $**** : p < 0.0001$.The red and green stars have the same meaning, but for the odd and even halves of that aggregate respectively.

3.5 Tuning Widths and Preferred ILD

Aside from the preferred ILD inside the ROIs, we also investigated the relationship between tuning width and preferred ILD. The results of this analysis can be seen in figures 12 and 13. The False Discovery Rate corected p-values for the correlations as found from a permutation test can be seen in table 2. The correlation between tuning width and preferred ILD reaches significance for each ROI in the MovingBinaural aggregate and only in 7 out of 16 ROIs for the StaticBinaural aggregate. This difference is likely due to the higher variance explained in the MovingBinaural aggregate, which was visible in figure 7 compared to figure 8. More voxels will have reached the 20% variance explained threshold in the MovingBinaural aggregate, leading to higher overall p-values. Looking at the plots, this correlation is also very clear for the MovingBinaural aggregate, with the expected V-shape appearing very clearly in all ROIs except the AC. This V-shape means that the tuning widths increase as the preferred ILD moves further away from the center (0 ILD). Comparing this to figure 12, the correlation plots for the StaticBinaural aggregate are much less clear, appearing more jumbled. This result shows that there is a correlation between tuning widths and preferred ILD. We cannot, however, conclude that the stronger correlation is due to the target in the experiment changing ILD along with the stimulus, or due to the addition of thermal noise reduction in the preprocessing pipeline.

	StaticBinaural	MovingBinaural
Left AC	0.030	0.0157
Right AC	0.024	0.001
Left SMG	0.000	0.000
Right SMG	0.000	0.000
Left PCG	0.509	0.000
Right PCG	0.709	0.004
Left AG	0.114	0.000
Right AG	0.854	0.000
Left MFG	0.061	0.000
Right MFG	0.000	0.000
Left BA	0.219	0.000
Right BA	0.141	0.000
Left BP	0.219	0.000
Right BP	0.000	0.000
Left ACC	0.867	0.000
Right ACC	0.024	0.000

Table 2: P-values for the correlation between tuning width and absolute ILD value. These values are the result of a permutation test intended to check whether slopes are significantly different from zero, and they are False Discovery Rate corrected. Changing the experiment to have the task target's ILD change along with the stimulus for the participants has made every correlation between preferred ILD and tuning width significant.

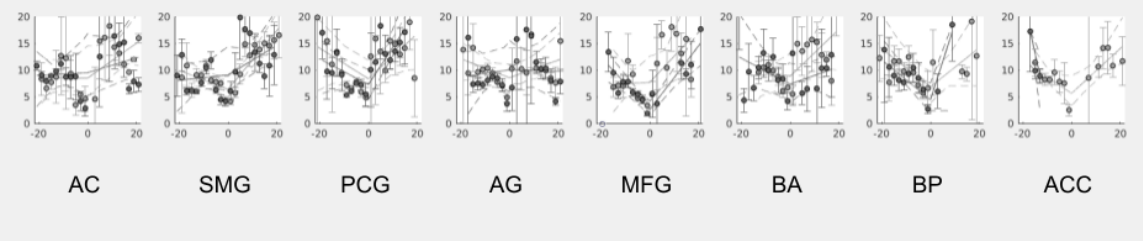

Figure 12: Relationship between tuning widths and preferred ILD for the StaticBinaural aggregate. The circles show the mean tuning widths for each bin in the left (light gray) and right (dark gray) hemispheres. Error bars are standard errors. A combination of linear functions of ILD and absolute ILD (solid line) was fitted to the tuning widths. The dashed lines show the upper and lower 95% confidence intervals to the fits. Only voxels for which the variance explained of the model reached 20% were included here. Note that these plots do not show the data used in the correlation itself, as that was calculated using unbinned data. The correlation between preferred ILD and tuning width is quite difficult to see in most ROIs, which is reflected in the table 2, as only half of the ROIs in the StaticBinaural aggregate reached significance.

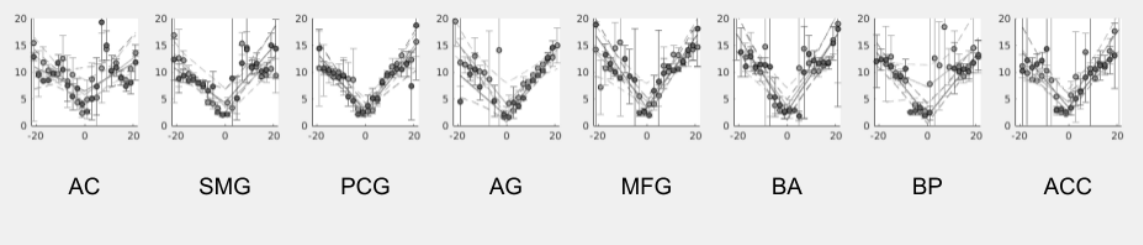

Figure 13: Relationship between tuning widths and preferred ILD for the MovingBinaural aggregate. The graphs were created in the same way as in figure 12. Only voxels for which the variance explained of the model reached 20% were included here. The correlation between preferred ILD and tuning width is very clear here, with the V-shape visible in each ROI. This indicates that the tuning width increases as preferred ILD moves further from 0. This correlation is more evident than in figure 12, which is in line with the p-values found in table 2 as all correlations in the MovingBinaural aggregate reached significance.

4 Discussion

This study aimed to find evidence for the existence of binaural field maps, i.e., areas of the cortex showing tuned responses to interaural level difference (ILD). As a secondary goal, this study aimed to investigate the spatial structure of these field maps (i.e., if a repeatable spatial variation could be identified) and then to see whether their preferred ILD changed as you move from one edge to the other, as was found by Harvey & Dumoulin (2017) for topographic numerosity maps. To achieve these goals fMRI data and population receptive field (pRF) modeling were used.

Using this approach, we provided evidence for tuned responses to ILD in several areas in the human cortex. 8 regions of interest (ROI) were identified per hemisphere that consistently showed a higher goodness of fit than surrounding areas when a model predicting neural responses from ILD was fit to the data. These 8 are the auditory cortex (AC), angular gyrus (AG), supramarginal gyrus (SMG), postcentral gyrus (PCG), posterior and anterior regions of Broca's area (BP and BA respectively), middle frontal gyrus (MFG), and anterior cingulate cortex (ACC). The ILD tuned linear Gaussian model significantly outperformed the two other models across the defined ROIs, showing that this model best reflected the measured data. Therefore, we concluded that these ROIs are indeed showing tuned responses to ILD values, and not to contralateral or ipsilateral sound levels or a weighted combination of those sound levels.

Starting out, we had no strong hypothesis for where the ROIs would be located, except that we expected the AC to show a strong contralateral response. Broadly speaking, the areas we identified are involved in either attention control or speech understanding and production. More specifically, the ACC and MFG are both associated with attention regulation (Posner, Rothbart, Sheese, & Tang, 2007; Japee, Holiday, Satyshur, Mukai,

& Ungerleider, 2015). Determining spatial location might be an important factor for this functionality, as it is a common situation where we might like to focus our attention only on sounds coming from a specific direction or guide visual attention to the sound's direction and make eye movements toward it. The ACC and MFG were also both more visibly active in the current study in which the attention of the participants is likely to shift along with the changing ILD to follow the task target (i.e., MovingBinaural), than in previous data where the target for the task remained at an ILD of 0 (i.e., StaticBinaural). This may support the idea that these areas are processing ILD to aid in attention regulation. The AG and SMG are both associated with multisensory integration of sound and vision (Oberman & Ramachandran, 2008; Yu, Li, & Sun, 2016), which is another task that benefits from determining the location sounds are coming from to match auditory and visual stimuli together. Broca's area is active while hearing speech or producing it (Flinker et al., 2015). Spatial location can play an important role in understanding speech in crowded situations, where speech may be coming from multiple directions and we need to separate out what we are trying to pay attention to. The PCG was a surprising find, as it is part of the somatosensory cortex and that tends to be activated only due to physical stimuli from our body. We, however, mostly observed activity in the lower region of the PCG, which is usually associated with sensory input from the tongue and mouth region (Picard & Olivier, 1983). As we make use of feedback from our body while speaking and Broca's area plays an important role in speech production, these two ROIs may be strongly linked, which could explain the activity we observed.

Evidence was also found for the preferred ILD of the ROIs not simply being contralateral. This behaviour is in line with what we expected from binaural field maps, with different regions being tuned to different ILDs. Akin to this, we found that tuning widths are correlated with preferred ILD, with smaller tuning widths being present in voxels with a preferred ILD closer to the center (i.e., 0). This correlation was more visible in the experiment with the target that changed ILD. This is again expected behaviour for binaural field maps, as it shows that, similar to the visual system being more fine-tuned for input coming from the center of our visual field, we have a more fine-grained processing for sounds coming from a central position in our binaural field.

Aside from showing the existence of tuned responses to ILD, this study also provided evidence for a repeatable spatial variation inside the identified ROIs, further showing that these may be the binaural field maps we set out to find. Evidence was found that some structure exists in 28.9% of the identified ROIs, with the right SMG and left PCG most consistently showing this structure. Ten out of 14 ROIs for subject 19 showed a significant correlation. This raises the question whether subject 19 is an outlier with a seemingly more consistently structured preferred ILD inside their ROIs or whether something is wrong with the defined ROIs or the data for our other participants. Taking the ILD preference correlations from repeated measures in all ROIs together, we found that the combined data were highly unlikely to come from a random distribution and instead are structured in some way. As for the question whether this structure is a linear progression of

ILD preferences across the cortical surface within the ROIs, as in the research by Harvey & Dumoulin (2017), this correlation between preferred ILD and cortical distance only reached significance in 18.3% of identified ROIs. Therefore, we conclude that a linear structure could not consistently be found, however some spatial structure does seem to exist for the identified ROIs. An interesting note here is that subject 13, which is the only participant for who the StaticBinaural aggregate was used, showed a significant correlation between preferred ILD and cortical distance in five out of 13 ROIs (38.5%), which is the most out all our participants. This could be the result of switching to a task target that changes ILD in the MovingBinaural aggregate for the other participants, which may have resulted in a weaker correlation between cortical distance and preferred ILD.

One drawback to the current study is the approach to identifying ROIs. The basis for the ROIs is the linear Gaussian model using ILD as input, which then turned out to be the model that best fit inside the ROIs. This resulted in a somewhat circular reasoning that could be avoided by drawing the ROIs once for each model. By identifying ROIs once for each model, one would be able to compare the model fit across all these ROIs and determine which model performs best across all situations. After determining this, the ROIs identified using the best fitting model could be used to conduct further research. However, most other models fit poorly and did not allow the identification of responsive regions of interest beyond the auditory cortex.

An argument could be made that the claim that the models fit better inside the ROIs than outside them needs to be statistically shown. To achieve this, one could fit the models to all white matter voxels, which are not responding to the stimulus presented. As these responses are essentially random, the variance explained for them will also vary. This distribution of variance explained reflects what we can expect when data drawn from a null distribution are used. For every value of variance explained in a responsive voxel one would then have a proportion of white matter voxels that have at least that same value. Thus, it is possible to determine the likelihood of finding this value by chance. This approach would provide statistical evidence for the claim that our models fit better inside the ROIs than expected by chance.

Aside from drawbacks specific to this study, some issues also exist with using an fMRI study for auditory research in a more general sense. Although precautions are taken to shield the participants from the background noise, an fMRI scanner simply produces a high level of noise. We repeated the stimulus multiple times and our participants wore headphones to limit the impact of this, but there is still the risk of the noise impacting the neuronal responses. Specifically, as the stimulus needs to be played at a high loudness to be heard over the background noise, the question remains what the neuronal responses would look like at lower volume levels with less background noise. This volume level is not what the human brain regularly deals with, thus it is possible that the responses would look significantly different at lower levels.

The current study offers some avenues for further research. Evidence was found for the existence of some spatial structure inside the identified ROIs, however this structure

did not appear to consistently be a linear correlation between preferred ILD and cortical distance. This raises the question what the shape of this structure could be and for which ROIs the structure consistently exists. To investigate this further, additional participants will be required to identify if a pattern exists in terms of which ROIs most consistently show the repeatable spatial variations. After identifying these, the shape of the structure could possibly be identified, starting with visual inspection of the ROIs and followed by statistical validation of the identified shape.

This research differed from the prior research by Bongaerts (2020) in the target presented to participants while the stimulus played, as the pure tone the participants have to respond to now changes ILD along with the white noise. This change was made to increase the gain on neural responses to ILD by allowing the participant to attend to the currently presented ILD. However it also seems to have improved model fit and increased the correlation between preferred ILD and tuning widths. To investigate whether these changes are due to the altered stimulus and not a consequence of the addition of thermal noise reduction in the preprocessing pipeline extra scan sessions could be performed in which phase data is collected, but the task does not change ILD. Alternatively, the preprocessing of the data gathered for the current study can be redone without the thermal noise reduction, after which the comparisons made in this study can be repeated to see if the differences remain. Investigating this change more thoroughly would help shed light on the impact of attention on these neural responses.

Having provided evidence for the existence of binaural field maps, the question remains how the other aspects of sound stimuli might be processed in the brain. The interaural time difference (ITD) and head-related transfer function (HRTF) are both used to determine the location of a sound stimulus (Grothe et al., 2010), but have been explicitly excluded in this research by only changing the ILD of the presented stimulus. It might be possible to identify similar field maps for these two features using the same type of experiment.

Taken together, our results seem to show the existence of binaural field maps outside the AC. These field maps appear to be tuned to ILD and exhibit a repeating spatial variation of some kind, providing support to the theory that they are somehow spatiotopically organized. They exist in the angular gyrus, supramarginal gyrus, postcentral gyrus, Broca's area, middle frontal gyrus, and anterior cingulate cortex. These areas are broadly associated with attention regulation or speech production and understanding, two activities that could benefit from sound localization. Although the AC shows primarily a strong contralateral response, the other maps vary in their preferred ILD. Inside the ROIs a correlation between preferred ILD and tuning width was identified, with the tuning widths increasing as the preferred ILD moves further away from zero. This shows that these regions are more sensitive to changes in ILD when these occur close to a central position. Further research is still needed to determine the shape of the discovered structure and corroborate the findings presented in this research.

References

Akeroyd, M. (2017). A binaural cross-correlogram toolbox for matlab. doi: http://doi.org/10.17639/nott.320

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the royal Statistical Society*, *57*, 289-300. doi: https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

Bongaerts, F. L. P. (2020). Neural tuning to and topographic representations of interaural level differences in the human brain. *Utrecht University*.

Cox, R. W. (1996). Afni: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*, 162-173. doi: https://doi.org/10.1006/cbmr.1996.0014

Doran, S. J., Charles-Edwards, L., Reinsberg, S. A., & Leach, M. O. (2005). A complete distortion correction for mr images: I. gradient warp correction. *Physics in Medicine and Biology*, *50*, 1343 - 1361. doi: http://doi.org/10.1088/0031-9155/50/7/001

Flinker, A., Korzeniewska, A., Shestyuk, A. Y., Franaszczuk, P. J., Dronkers, N. F., Knight, R. T., & Crone, N. E. (2015). Redefining the role of broca's area in speech. *Proceedings of the National Academy of Sciences of the United States of America*, *112*, 2871-2875. doi: https://doi.org/10.1073/pnas.1414491112

Grothe, B., Pecka, M., & McAlpine, D. (2010). Mechanisms of sound localization in mammals. *Physiological Reviews*, *90*, 983-1012. doi: https://doi.org/10.1152/physrev.00026.2009

Harvey, B. M., & Dumoulin, S. O. (2017). A network of topographic numerosity maps in human association cortex. *Nature Human Behaviour*, *1*. doi: https://doi.org/10.1038/s41562-016-0036

Harvey, B. M., Dumoulin, S. O., Fracasso, A., & Paul, J. M. (2020). A network of topographic maps in human association cortex hierarchically transforms visual timing-selective responses. *Current Biology*, *30*, 1424-1434. doi: https://doi.org/10.1016/j.cub.2020.01.090

Harvey, B. M., Klein, B. P., Petridou, N., & Dumoulin, S. O. (2013). Topographic representation of numerosity in the human parietal cortex. *Science*, *341*, 1123-1126. doi: https://doi.org/10.1126/science.1239052

Japee, S., Holiday, K., Satyshur, M. D., Mukai, I., & Ungerleider, L. G. (2015). A role of right middle frontal gyrus in reorienting of attention: a case study. *Frontiers in Systems Neuroscience*. doi: https://doi.org/10.3389/fnsys.2015.00023

King, A. J., & Nelken, I. (2009). Unraveling the principles of auditory cortical processing: can we learn from the visual system? *Nature Neuroscience*, *12*, 698-701. doi: https://doi.org/10.1038/nn.2308

King, A. J., Teki, S., & Willmore, B. D. B. (2018). Recent advances in understanding the auditory cortex. doi: https://doi.org/10.12688%2Ff1000research.15580.1

Oberman, L. M., & Ramachandran, V. S. (2008). Preliminary evidence for deficits in multisensory integration in autism spectrum disorders: The mirror neuron hypothesis. *Social Neuroscience*, *3*, 348-355. doi: https://doi.org/10.1080/17470910701563681

Ortiz-Rios, M., Azevedo, F. A. C., Kusmierek, P., Balla, D. Z., Munk, M. H., Keliris, G. A., ... Rauschecker, J. P. (2017). Widespread and opponent fmri signals represent sound location in macaque auditory cortex. *Neuron*, *93*, 971-983. doi: https://doi.org/10.1016/j.neuron.2017.01.013

Panniello, M., King, A. J., Dahmen, J. C., & Walker, K. M. M. (2018). Local and global spatial organization of interaural level difference and frequency preferences in auditory cortex. *Cerebral Cortex*, *28*, 350–369. doi: https://doi.org/10.1093/cercor/bhx295

Picard, C., & Olivier, A. (1983). Sensory cortical tongue representation in man. *Journal of Neurosurgery*, *59*, 781-789. doi: https://doi.org/10.3171/jns.1983.59.5.0781

Posner, M. I., Rothbart, M. K., Sheese, B. E., & Tang, Y. (2007). The anterior cingulate gyrus and the mechanism of self-regulation. *Cognitive, Affective, & Behavioral Neuroscience*, *7*, 391-395. doi: http://doi.org/10.3758/CABN.7.4.391

Rauschecker, J. P., & Tian, B. (2000). Mechanisms and streams for processing of "what" and "where" in auditory cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *97*, 11800-11806. doi: https://doi.org/10.1073/pnas.97.22.11800

Schnupp, J., Nelken, I., & King, A. J. (2011). *Auditory neuroscience: Making sense of sound*. The MIT Press.

Wandell, B. A. (2014). *Vision imaging science and technology lab.* `http://vistalab.stanford.edu/`. (Accessed: 2022-03-08)

Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, *56*, 366-383. doi: https://doi.org/10.1016/j.neuron.2007.10.012

Yu, H., Li, Q., & Sun, H. (2016). A task-irrelevant sound modulates the effects of simultaneous visual cue on visual discrimination: An fmri study. *2016 IEEE International Conference on Mechatronics and Automation*, 1965-1970. doi: https://doi.org/10.1109/ICMA.2016.7558867

A Regions of Interest per Participant

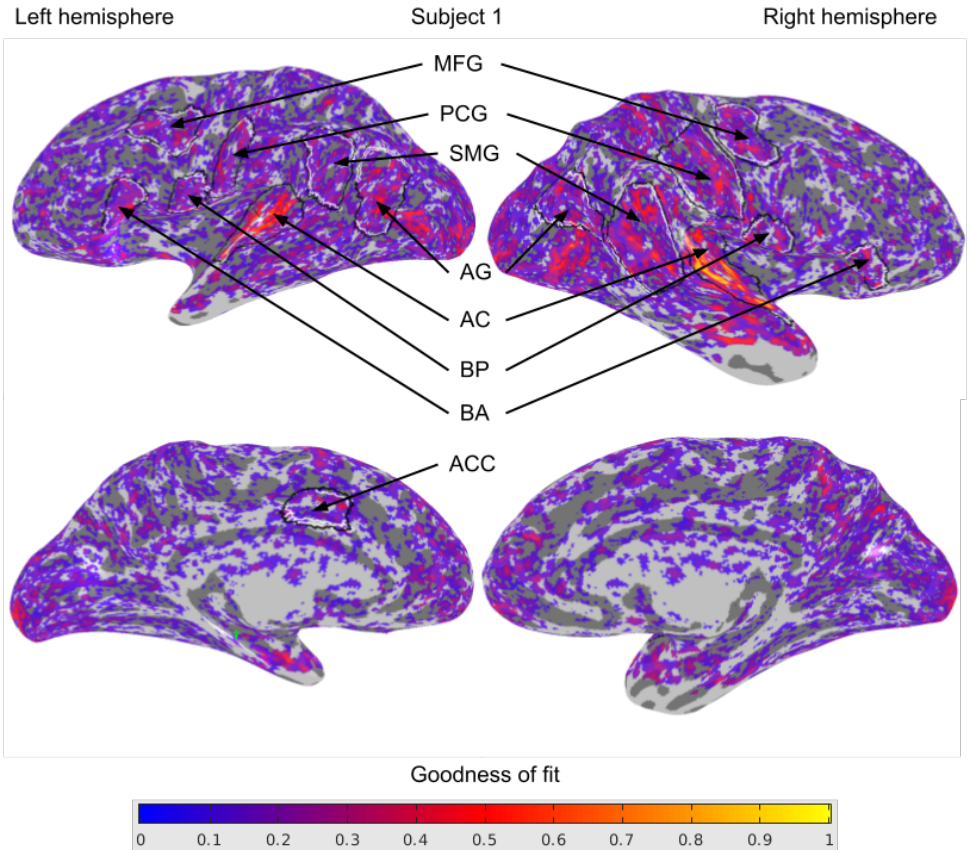

Figure 14: Goodness of fit for subject 13. The variance explained threshold for a voxel showing up is set to 0.1. The ROIs can be seen here with their boundaries in black and white. The right ACC was not identified in this participant.

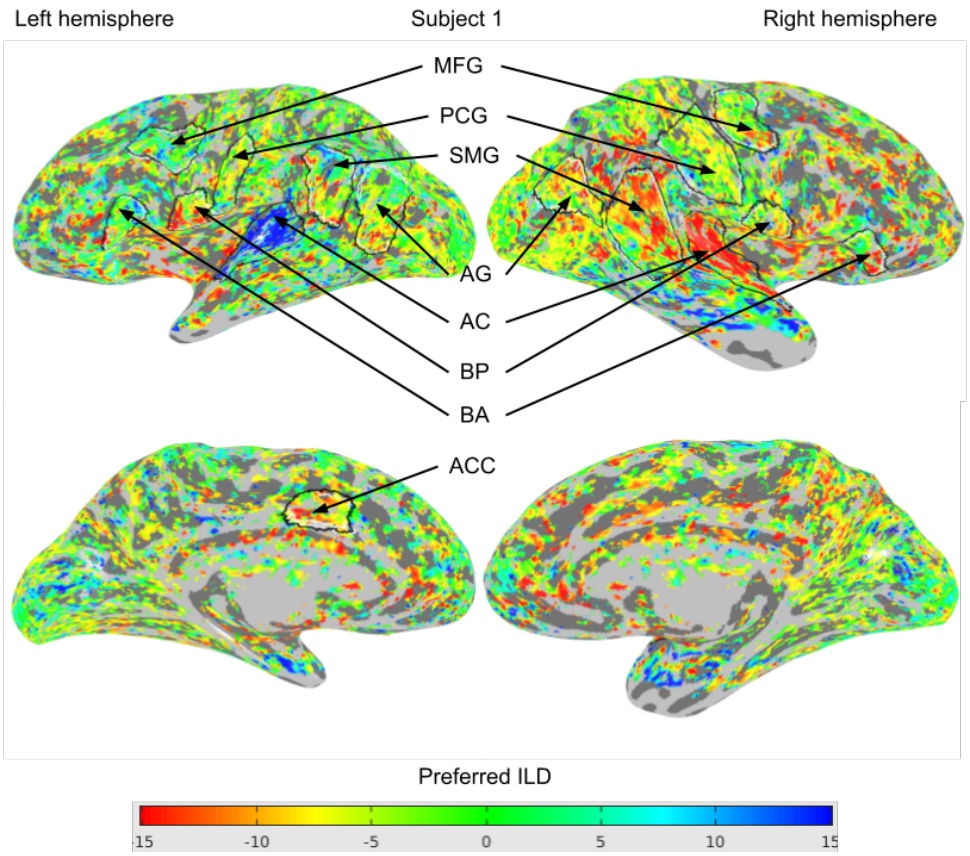

Figure 15: Preferred ILD for subject 13. The variance explained threshold for a voxel showing up is set to 0.1. The ROIs can be seen here with their boundaries in black and white. The white boundaries mark their high (blue) or low (red) boundary which were later used to determine the structure of the ROI. The right ACC was not identified in this participant.

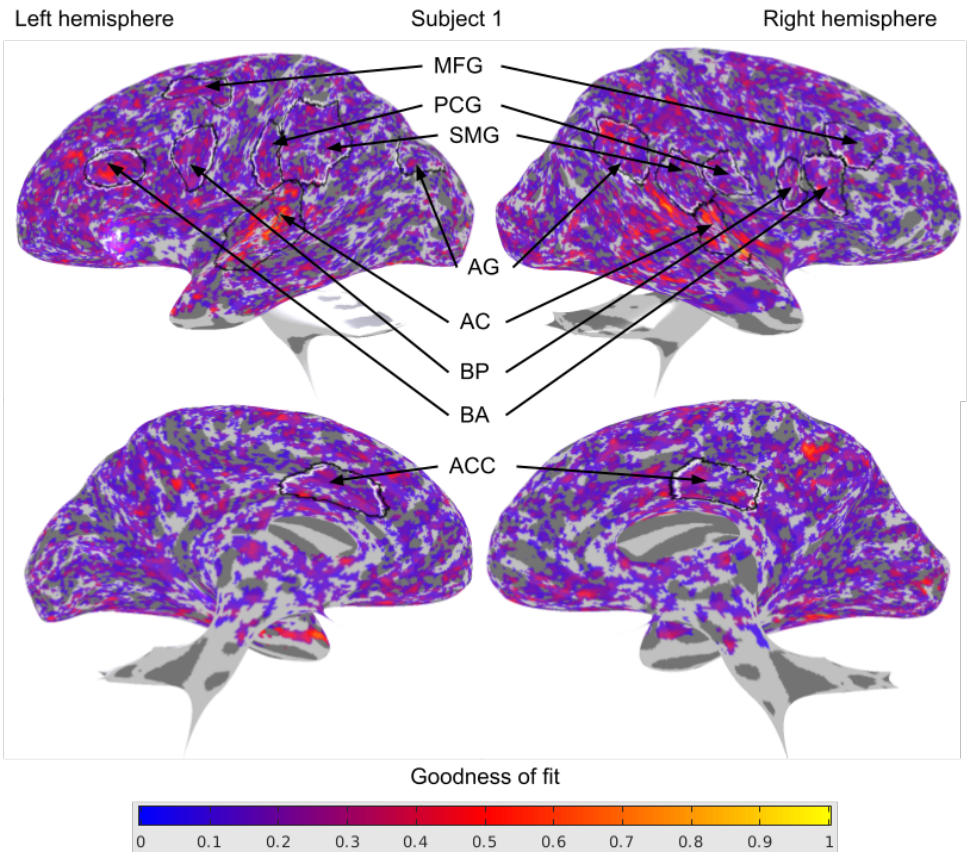

Figure 16: Goodness of fit for subject 14. The variance explained threshold for a voxel showing up is set to 0.1. The ROIs can be seen here with their boundaries in black and white.

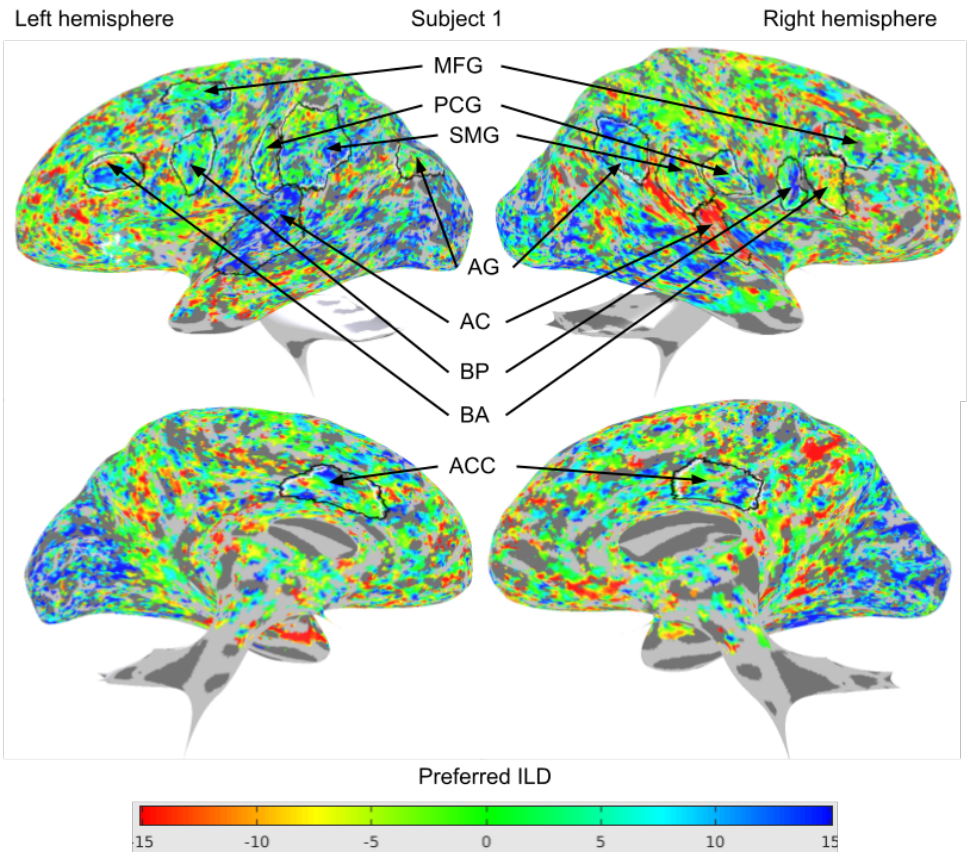

Figure 17: Preferred ILD for subject 14. The variance explained threshold for a voxel showing up is set to 0.1. The ROIs can be seen here with their boundaries in black and white. The white boundaries mark their high (blue) or low (red) boundary which were later used to determine the structure of the ROI.

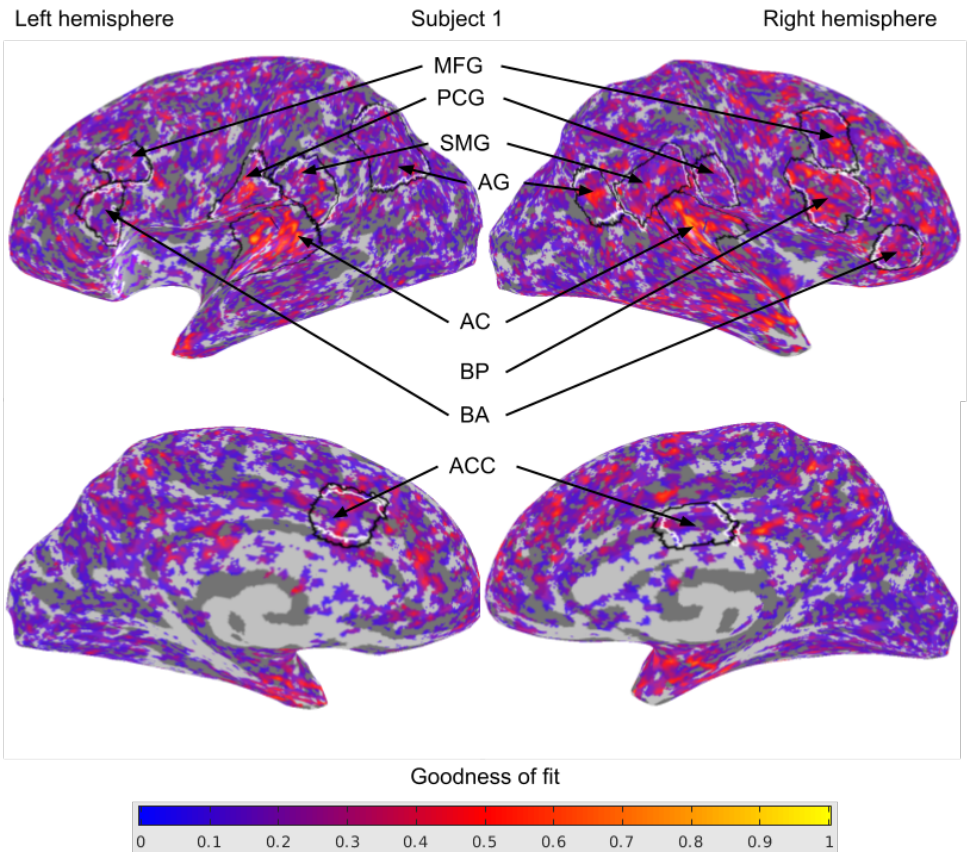

Figure 18: Goodness of fit for subject 19. The variance explained threshold for a voxel showing up is set to 0.1. The ROIs can be seen here with their boundaries in black and white. The left BP was not identified in this participant.

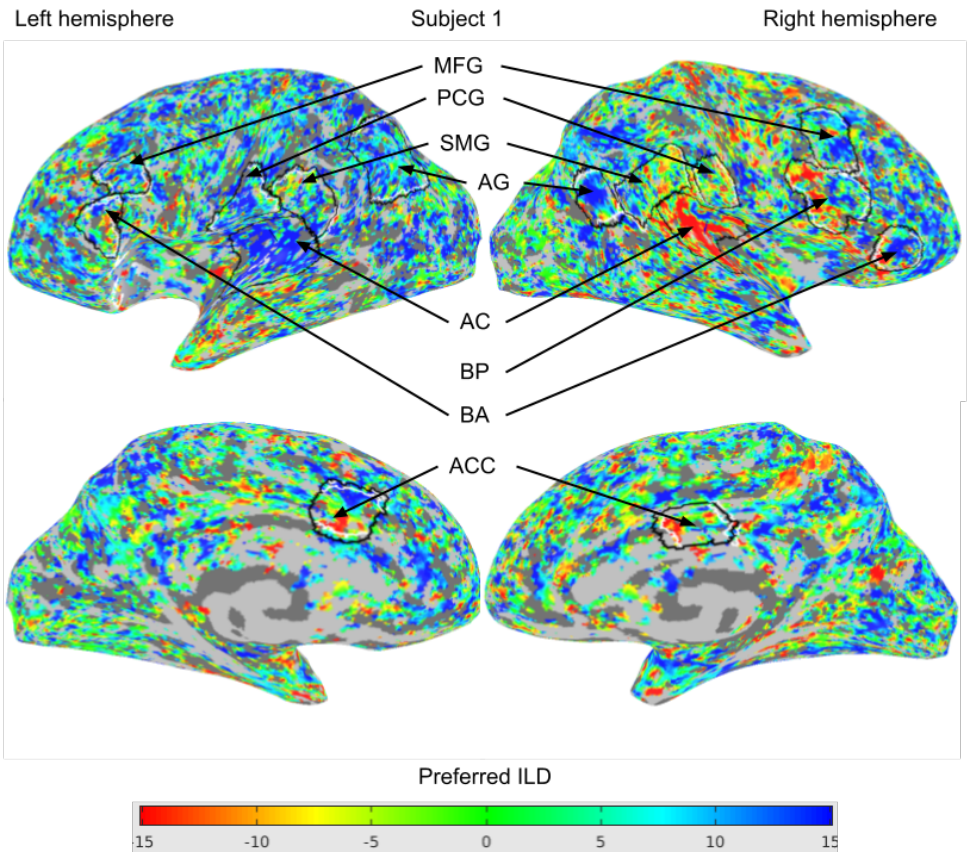

Figure 19: Preferred ILD for subject 19. The variance explained threshold for a voxel showing up is set to 0.1. The ROIs can be seen here with their boundaries in black and white. The white boundaries mark their high (blue) or low (red) boundary which were later used to determine the structure of the ROI. The left BP was not identified in this participant.

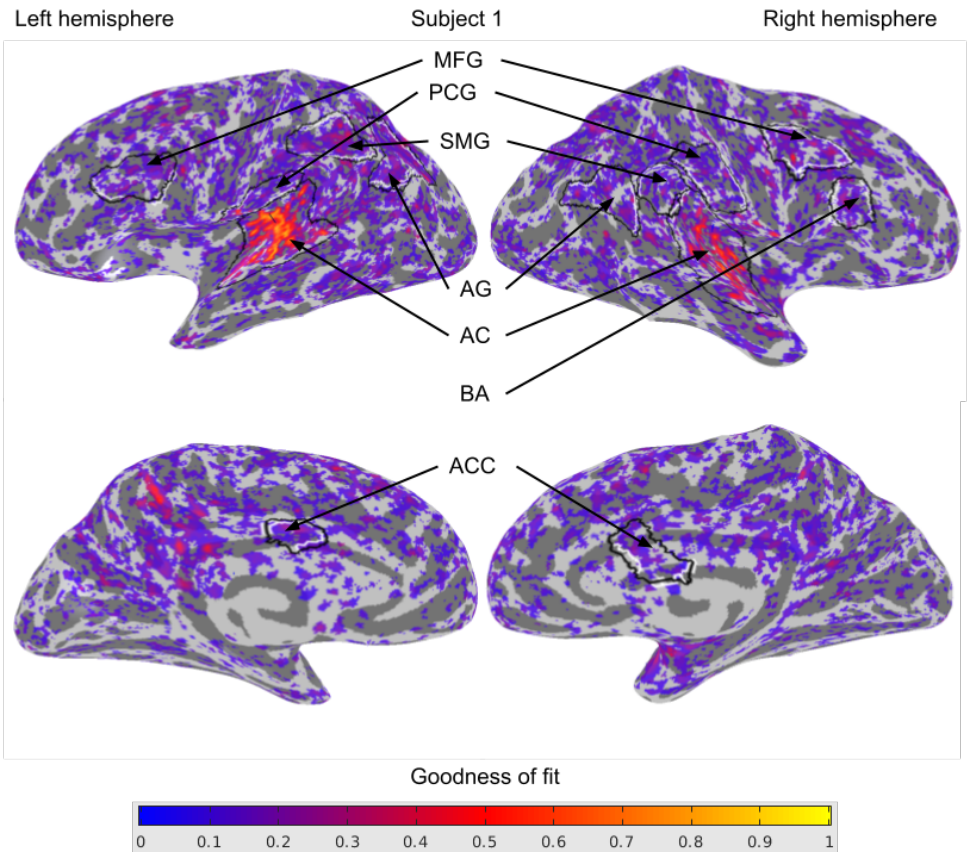

Figure 20: Goodness of fit for subject 20. The variance explained threshold for a voxel showing up is set to 0.1. The ROIs can be seen here with their boundaries in black and white. The left BA, left BP, and right BP were not identified in this participant.

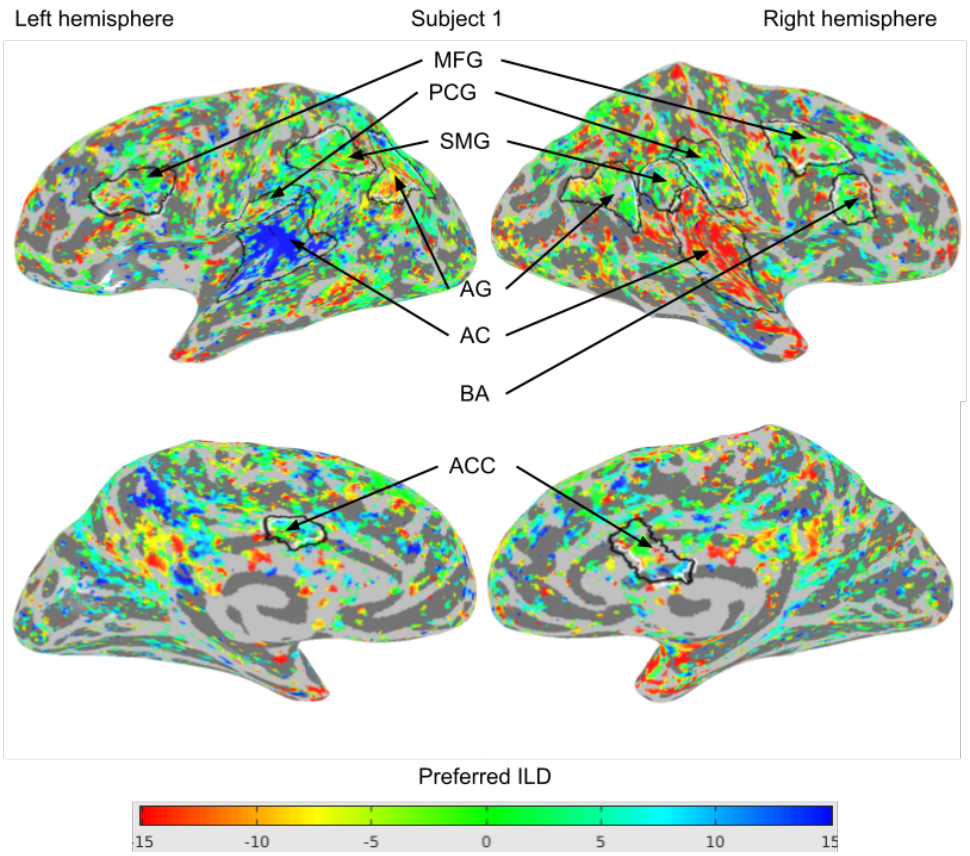

Figure 21: Preferred ILD for subject 20. The variance explained threshold for a voxel showing up is set to 0.1. The ROIs can be seen here with their boundaries in black and white. The white boundaries mark their high (blue) or low (red) boundary which were later used to determine the structure of the ROI. The left BA, left BP, and right BP were not identified in this participant.

B Preferred ILD Over Cortical Distance per Participant

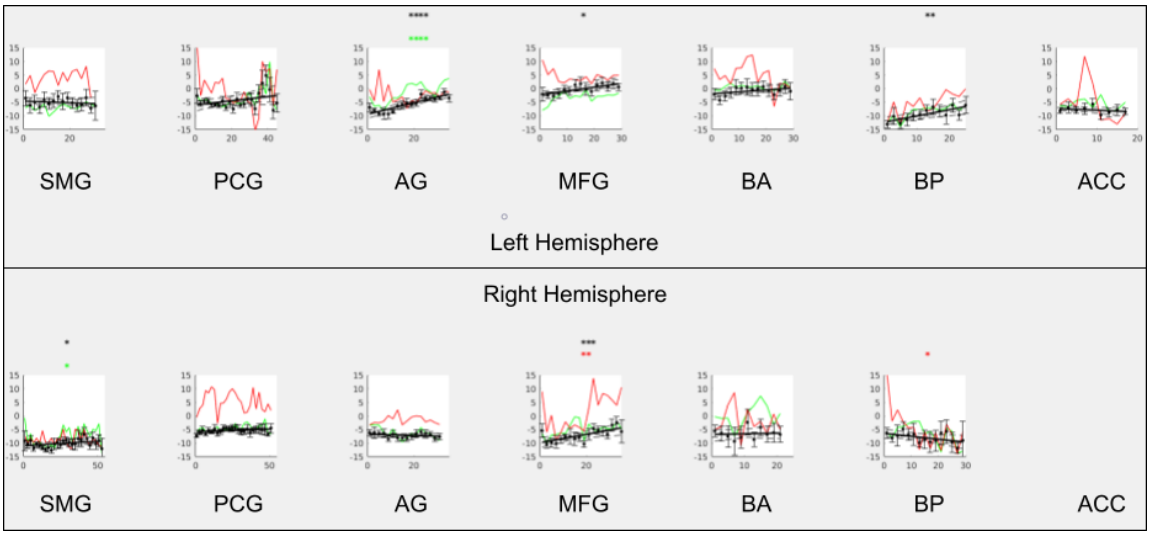

Figure 22: Preferred ILD plotted over cortical distance for subject 13. The stars over the plots show whether the correlation reached significance in the ROI in question. $* : p < 0.05$, $** p < 0.01$, $*** : p < 0.001$, $**** : p < 0.0001$.The red and green stars have the same meaning, but for the odd and even halves of that aggregate respectively.

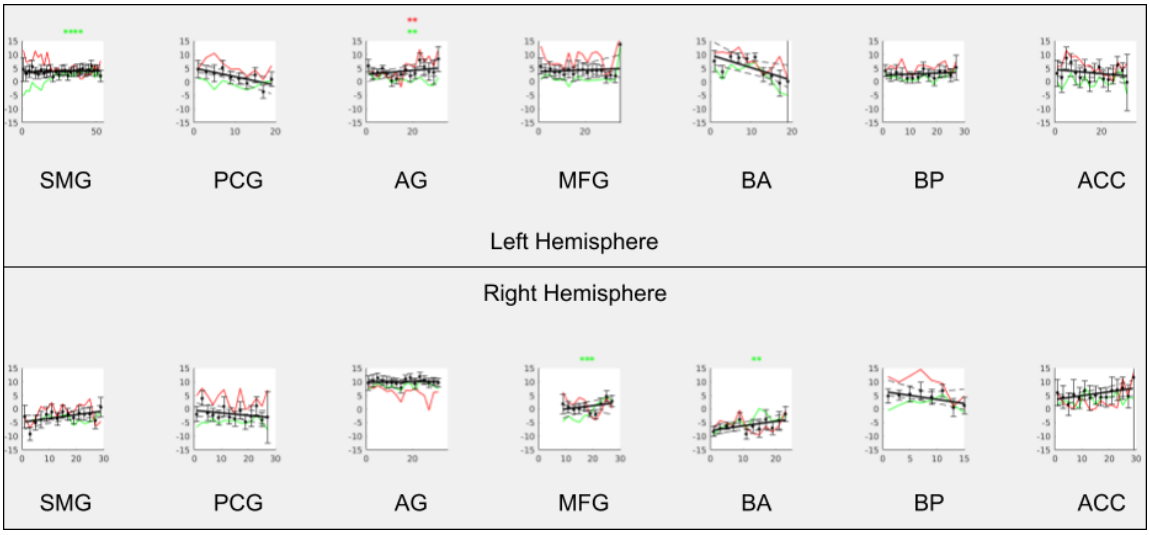

Figure 23: Preferred ILD plotted over cortical distance for subject 14. The stars over the plots show whether the correlation reached significance in the ROI in question. $* : p < 0.05$, $**p < 0.01$, $*** : p < 0.001$, $**** : p < 0.0001$.The red and green stars have the same meaning, but for the odd and even halves of that aggregate respectively.

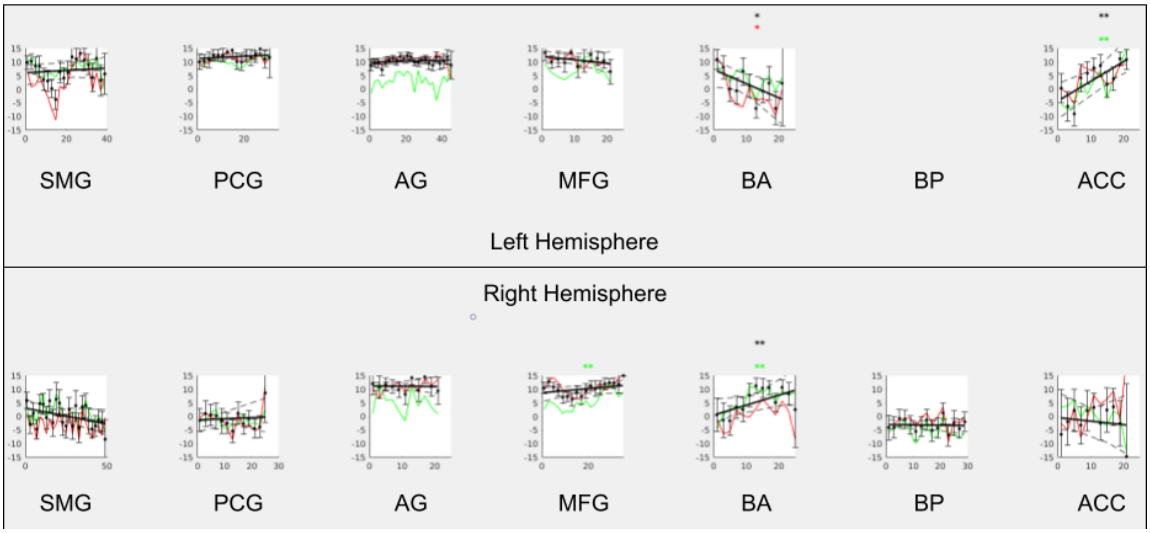

Figure 24: Preferred ILD plotted over cortical distance for subject 19. The stars over the plots show whether the correlation reached significance in the ROI in question. $* : p < 0.05$, $**p < 0.01$, $*** : p < 0.001$, $**** : p < 0.0001$.The red and green stars have the same meaning, but for the odd and even halves of that aggregate respectively.

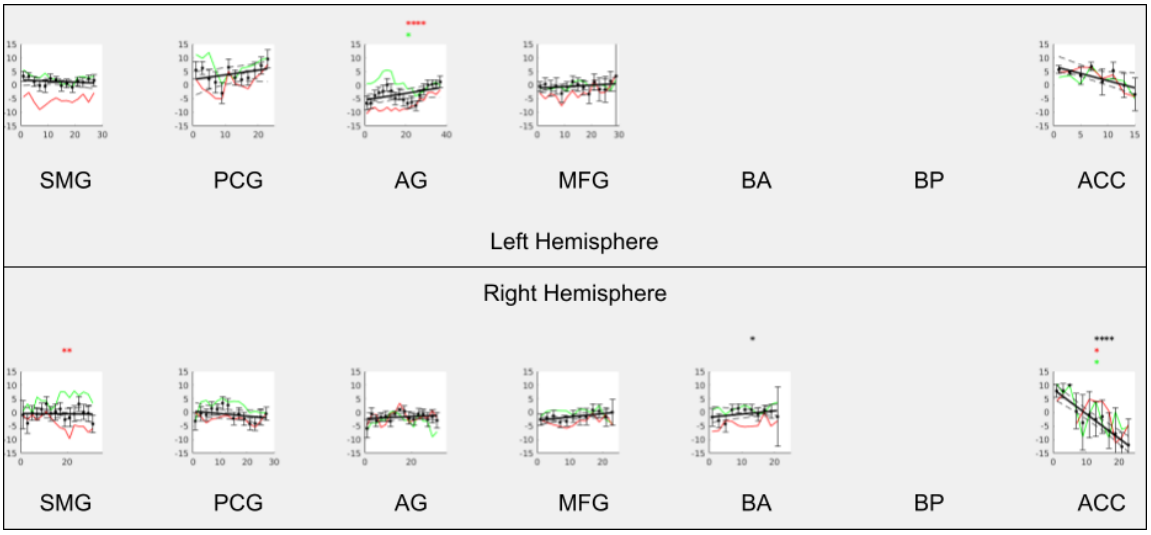

Figure 25: Preferred ILD plotted over cortical distance for subject 20. The stars over the plots show whether the correlation reached significance in the ROI in question. $* : p < 0.05$, $** p < 0.01$, $*** : p < 0.001$, $**** : p < 0.0001$.The red and green stars have the same meaning, but for the odd and even halves of that aggregate respectively.