Optimising a computational model of human auditory cortex with an evolutionary algorithm

Ewelina Tomana[a], Nina Härtwich[b], Adam Rozmarynowski[a], Reinhard König[b], Patrick J. C. May[c], Cezary Sielużycki[a]

[a] *Department of Biomedical Engineering, Wrocław University of Science and Technology, Wyb. Wyspiańskiego 27, 50-370, Wrocław, Poland*
[b] *Research Group Comparative Neuroscience, Leibniz Institute for Neurobiology, Brenneckestr. 6, 39118, Magdeburg, Germany*
[c] *Department of Psychology, Lancaster University, LA1 4YR, Lancaster, United Kingdom*

Abstract

We demonstrate how the structure of auditory cortex can be investigated by combining computational modelling with advanced optimisation methods. We optimise a well-established auditory cortex model by means of an evolutionary algorithm. The model describes auditory cortex in terms of multiple core, belt, and parabelt fields. The optimisation process finds the optimum connections between individual fields of auditory cortex so that the model is able to reproduce experimental magnetoencephalographic (MEG) data. In the current study, this data comprised the auditory event-related fields (ERFs) recorded from a human subject in an MEG experiment where the stimulus-onset interval between consecutive tones was varied. The quality of the match between synthesised and experimental waveforms was 98%. The results suggest that neural activity caused by feedback connections plays a particularly important role in shaping ERF morphology. Further, ERFs reflect activity of the entire auditory cortex, and response adaptation due to stimulus repetition emerges from a complete reorganisation of AC dynamics rather than a reduction of activity in discrete sources. Our findings constitute the first stage in establishing a new non-invasive method for uncovering the organisation of the human auditory cortex.

Keywords: auditory cortex, computational modelling, event-related field, evolutionary algorithms, MEG, optimisation

1. Introduction

Magnetoencephalographic (MEG) signals are generated mainly by primary currents running in the apical dendrites of a large number of synchronously active pyramidal cells in the cortex (Hämäläinen et al., 1993). The cortical response to a single stimulus is very weak and, thus, it is buried in brain activity unrelated to stimulus processing and in the noise of the MEG measurement. Hence, the standard approach to obtain a detectable event-related field (ERF) is to use a large number of stimulus repetitions (typically > 100) and average the artefact-corrected trials of the MEG signals. Apart from this technical constraint, the geometric properties of the network of the active pyramidal cells themselves play a decisive role in the generation and measurability of ERFs: The apical dendrites of the pyramidal cells are locally aligned with each other, and they are oriented perpendicularly to the cortical surface. The total current running in these dendrites reflects the size of the population generating the ERF and therefore the size of the active cortical area.

MEG is particularly sensitive to signals originating from sensory cortical areas such as the visual or auditory cortex. In this work, we focus on auditory cortex (AC). The anatomy of AC is characterised by a hierarchical core-belt-parabelt structure, which is similar across mammals, with cross-species differences appearing as variations in the number of cortical fields and the pattern that reflects the connectivity between them. In general, the core fields connect to each other and to the surrounding belt fields, which in turn are connected to nearest-neighbour belt fields and to parabelt fields (Kaas and Hackett, 2000b; Hackett et al., 2014). Information processing in the AC is mainly determined by two streams of signal flow. Feedforward connections drive the signal in the core→belt→parabelt direction and contribute to an ERF deflection with a polarity such as that shown by the P1m response. Feedback connections represent signal flow in the opposite direction, from parabelt to belt to core (Hackett et al., 2014). This results in ERF deflections with opposite polarity as exhibited, for example, by the N1m response (Ahlfors et al., 2015).

The auditory ERF is characterised by a series of waves. These are generally categorised with respect to their peak latency into early- (< 8 ms), mid- (8–40 ms) and late-latency (> 40 ms) responses, with the latter being the most prominent. The general structure of late auditory ERFs to almost any transient auditory stimulus comprises the sequence P1m, N1m, and P2m. These waves have alternating polarities (indicated by P for positive and N for negative), and their approximate peak latencies in relation to stimulus onset are 50 ms, 100 ms, and 150–180 ms, respectively. The peak amplitudes and latencies of these responses are sensitive to the intensity, duration, sound source location, and frequency composition of the stim-

ulus as well as to the stimulus-onset interval (SOI; for a review, see May and Tiitinen, 2010).

The neural origin of the late-latency ERFs with regard to the connectivity of fields within AC is not fully understood. Fundamental information about the connectivity of fields must usually be obtained from invasive recordings in animals, predominantly in non-human primates and Mongolian gerbils, since, for ethical reasons, comparable invasive measurements in humans are only permitted in some clinical cases. Hence, the transfer of knowledge about animal AC to human AC is a true challenge, as this requires the linking of different types of recordings. On the one hand, there are the electrophysiological measurements on the single cell- and circuit level in animals. On the other, there are the MEG measurements taken from humans, where a large population of neurons is required to obtain detectable ERFs.

Computational modelling has the potential to merge the knowledge gained at the different levels of observation in humans and animals. May and Tiitinen (2010, 2013), May et al. (2015), and Westö et al. (2016) developed such a computational model of signal processing in AC. The model takes into account the idiosyncratic architecture of AC by mimicking the connectivity pattern between AC fields. On the microscopic scale, the model yields firing rates for excitatory and inhibitory neuron populations, which are organised into cortical columns, which are the basic units of the model. On the macroscopic scale, it yields ERFs as the spatial sum of the synaptic inputs to the excitatory populations.

This AC model has dynamics evolving on two time scales. On the faster time scale, mean population firing rates are translated into synaptic input currents, which in turn feed into determining the firing rate. This kind of dynamics can be approximated as a set of normal modes that spatially overlap with each other and extend over the entire auditory cortex (Hajizadeh et al., 2019, 2021, 2022). On the slower time scale, the excitatory connections are modulated by short-term synaptic depression (STSD) which operates as a function of the pre-synaptic firing rate. STSD results in adaptation, the reduction of the response to a repeated stimulus, and it also underlies the ability of the model to mimic the selective responsiveness of AC neurons to complex stimulation with a temporal and spectral structure. Further, STSD is the main mechanism which allows the model to reproduce a wide variety of mismatch negativity responses of the auditory ERF (May, 2021). Notably, since the anatomy of the AC among mammals has a comparable basic structure, it is relatively simple to modulate the structure of the model to mirror that of AC of different species, thus allowing for cross-species comparisons of function.

Our modelling approach differs from those of other groups by including, as key element, the detailed core-belt-parabelt structure of the AC of mammals. Other models investigate auditory processing at different levels of observation, such as a single AC field (see, for example, Loebel et al., 2007; Noto et al., 2016; Wang and Knösche, 2013), or in terms of large-scale dynamic networks of distributed brain areas, where cortical areas are treated as nodes (Dynamic Causal Modelling, David et al., 2006) or the entire AC is considered a single node (Virtual Brain project, Ritter et al., 2013). These approaches, however, are not intended to provide mechanistic explanations of how AC processes and represents sound.

The purpose of computational modelling is to unravel the complex mechanisms underlying experimental observations. Such modelling is often realised by reducing the complexity of the modelled system by projecting hidden information onto a well-defined parameter space. These model parameters can then be interpreted in terms of quantitative information about the biological system under investigation. Based on our previous modelling work (Hajizadeh et al., 2019, 2021, 2022; Turczak, 2022), we hypothesise that the still elusive anatomical structure of the human AC is reflected in the morphology of ERFs recorded by MEG. Here, we use the original nonlinear version of the AC model developed by May and Tiitinen (2013) and May et al. (2015), where the model parameters were chosen on a trial-and-error basis. This was largely because the knowledge base does not yet exist, especially in the case of the human brain, that would allow one to look up parameter values based on biological data. In the present work, we demonstrate that parameter values can be derived by employing the advanced optimisation method of the evolutionary algorithm (EA, Michalewicz, 1996). Specifically, consecutive iterations of the EA lead to the ERF generated by the model to match that measured in an MEG experiment on a human participant. As a result, we arrive at estimates of relative connection strengths between distinct AC fields.

In Sect. 2, we briefly introduce the AC model. Sect. 3 provides a description of the MEG data used for the optimisation of the model. Sect. 4 presents the building blocks of the EA used in this work. In Sect. 5, we show results of the EA optimisation of the AC model. Finally, Sect. 6 provides some discussion and sketches perspectives for future research.

2. Computational modelling of auditory cortex

2.1. Architecture of the model

We are working with the AC model developed by Patrick May and colleagues; see, for example, May and Tiitinen (2010, 2013) and May et al. (2015). Key features of the model are the specific anatomical structure of the auditory cortex of mammals and the inclusion of STSD. In general terms, the AC of mammals is divided into three hierarchically organised areas, the core, belt, and parabelt. Each of these main areas is further subdivided into fields, the number of which strongly differs across species. Since the exact number of fields in the human AC is unknown, we have turned to the AC of primates (in particular of

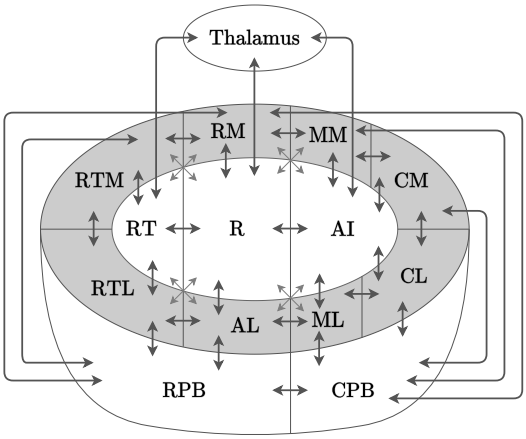

Figure 1: Schematic representation of the core-belt-parabelt structure of the AC of the macaque monkey (May et al., 2015), along with the subcortical thalamus which provides the afferent input to AC. Dark and light arrows indicate high and low density of connections between two fields, respectively.

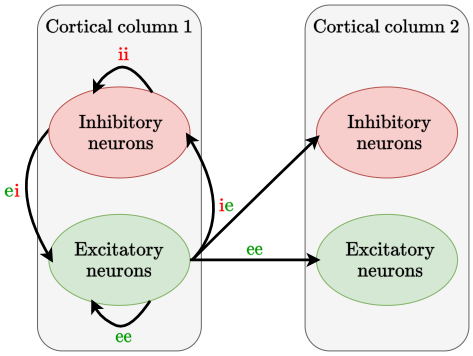

Figure 2: Possible connections between populations of excitatory (label e) and inhibitory (label i) neurons of two cortical columns of the model.

macaque monkeys) as an animal model of the human AC. Based on cytoarchitectonic, anatomical, and physiological differences, one can identify thirteen distinct fields in the macaque AC, which are interconnected according to a specific pattern (Romanski and Averbeck, 2009; Kaas and Hackett, 2000a). Fig. 1 summarises the experimental results regarding the field parcellation and connectivity pattern of the macaque AC. The core area consists of three fields, the primary auditory field (AI), the rostral field (R), and the rostral temporal field (RT), which receive parallel input from subcortical regions. The core is surrounded by eight belt fields, the anterolateral (AL), the middle lateral (ML), the caudolateral (CL), the caudomedial (CM), the mediomedial (MM), the rostromedial (RM), the medial rostrotemporal (RTM), and the lateral rostrotemporal (RTL) field. Laterally, the belt is bounded by the parabelt, which, in the macaque monkey, is divided into two fields: the rostral parabelt (RPB) and the caudal parabelt (CBP) field.

In the computational model, the interactions between fields are described using a 14×14 matrix. This includes the thirteen AC fields as well as a subcortical field representing thalamus, specifically the ventral division of the medial geniculate nucleus (MGN). The MGN is used as an input field which relays activation to the cortex. This means that there are a total of 196 possible connections between fields, including intra-field connections. Each AC field is further subdivided into 16 cortical columns, each composed of an excitatory and an inhibitory neuron population. The different connection types between excitatory and inhibitory populations are shown in Fig. 2. Thus, per connection type, there are $(14 \cdot 16)^2$ potential network connections between columns. The number of actual connections present in the model network is bounded by the experimentally determined connectivity pattern of the macaque AC.

In the model, the interactions between neuronal populations are represented by the elements of weight matrices denoted by W, reflecting the strength of a given type of connection. The W matrices are separated into four types, according to the direction of interactions occurring within and between the populations of excitatory and inhibitory neurons. They are denoted as follows:

- W_{ee} for excitatory populations targeting excitatory populations,
- W_{ei} for inhibitory populations targeting excitatory populations,
- W_{ie} for excitatory populations targeting inhibitory populations,
- W_{ii} for inhibitory populations targeting inhibitory populations.

2.2. Model dynamics

The dynamics of the model are based on the works of Wilson and Cowan (1972) and Hopfield and Tank (1986). They are described by two coupled nonlinear differential equations comprising the time-dependent state variable $\boldsymbol{u} = [u_1(t), \ldots, u_N(t)]^{\mathrm{T}}$ for the excitatory and $\boldsymbol{v} = [v_1(t), \ldots, v_N(t)]^{\mathrm{T}}$ for the inhibitory population, with N being the total number of columns in the network and $^{\mathrm{T}}$ denoting transpose:

$$\tau_{\mathrm{m}} \frac{\mathrm{d}\boldsymbol{u}}{\mathrm{d}t} = W_{\mathrm{ee}}\, Q(t)\, g(\boldsymbol{u}(t)) - W_{\mathrm{ei}}\, g(\boldsymbol{v}(t)) - \boldsymbol{u}(t) + \boldsymbol{i}_{\mathrm{aff}}(t), \tag{1a}$$

$$\tau_{\mathrm{m}} \frac{\mathrm{d}\boldsymbol{v}}{\mathrm{d}t} = W_{\mathrm{ie}}\, Q(t)\, g(\boldsymbol{u}(t)) - W_{\mathrm{ii}}\, g(\boldsymbol{v}(t)) - \boldsymbol{v}(t). \tag{1b}$$

Here, τ_{m} denotes the membrane time constant, and $\boldsymbol{i}_{\mathrm{aff}}(t)$ is the afferent current reflecting the input due to auditory stimulation. The weight matrices W_{ee}, W_{ei}, W_{ie}, and W_{ii} have been introduced in Sect. 2.1. The output of the excitatory and inhibitory populations are represented

by the mean firing rates $g(\boldsymbol{u})$ and $g(\boldsymbol{v})$, respectively; see Eq. (2) for $g(\boldsymbol{u})$:

$$g(u_n) = \begin{cases} \tanh\left(\frac{2}{3}\,(u_n - \theta)\right) & \text{for} \quad u_n \geq \theta, \\ 0 & \text{for} \quad u_n < \theta, \end{cases} \tag{2}$$

with $n \in \{1, \ldots, N\}$. The firing rates are nonlinear, monotonically increasing functions of the state variables $u_n(t)$ and $v_n(t)$ that take values in the $[0, 1]$ interval. Initially, as $u_n(t)$ and $v_n(t)$ increase, the values of $g(u_n)$ and $g(v_n)$ remain zero, until the state variables exceed a certain threshold θ, at which point the populations of neurons begin to fire. In our simulations, we used $\theta = 0.05$.

To reflect STSD, the weights in W_{ee} and W_{ie} are scaled via the time-dependent diagonal matrix $Q(t)$,

$$Q(t) = \text{diag}(\boldsymbol{q}(t)), \tag{3a}$$

$$\frac{\mathrm{d}\boldsymbol{q}}{\mathrm{d}t} = -\frac{\boldsymbol{q}(t) \circ g(\boldsymbol{u}(t))}{\tau_{\text{o}}} + \frac{\mathbf{1} - \boldsymbol{q}(t)}{\tau_{\text{rec}}}, \tag{3b}$$

where, for each time point t, diag in (3a) transforms the synaptic efficacy $\boldsymbol{q}$ into Q with elements of $\boldsymbol{q}$ on the leading diagonal, and $\circ$ in (3b) denotes element-wise multiplication (Hadamard product). Eq. (3b) is a simplified, continuous-time description of STSD (c.f. Tsodyks and Markram, 1997; Loebel et al., 2007). In this view, the depression of synaptic activity between pre- and post-synaptic excitatory cell populations depends only on the firing rate of the pre-synaptic excitatory cell population, $g(u_n(t))$, and STSD is characterised by two stages: a fast drop-off in synaptic efficacy determined by the time constant τ_{o} occurring directly after stimulus onset when the firing rate $g(u_n(t))$ is non-zero, and a slow recovery determined by the time constant τ_{rec}, when the firing rate $g(u_n(t))$ returns to zero after stimulus offset. The values $q_n(t)$ range from 1 (no STSD) to 0 (connection fully deactivated due to STSD). In our simulations, τ_{o} and τ_{rec} were set to 0.100 s and 1.600 s, respectively, for the AC, whereas, for the thalamus, they were set to 0.020 s and 0.100 s. The order-of-magnitude difference of the time constants of STSD in the different parts of the model reflects the different time scales of the dynamics observed experimentally along the lemniscal part of the auditory pathway (Ulanovsky et al., 2004; Pérez-González and Malmierca, 2014).

The computation of a synthetic MEG signal based on activity in the model network relies on the summation of the excitatory input currents to the excitatory neuron populations, i.e., the first term on the right-hand side of Eq. (1a). This reflects the fact that, in-vivo, the MEG signal is, to a large extent, generated by dendritic currents flowing in the apical dendrites of synchronously activated pyramidal neurons in the cortex (Hämäläinen et al. (1993); see Hajizadeh et al. (2019) for a detailed discussion of the computation of the MEG response in the AC model). In the traditional view, inhibitory synapses compared to excitatory ones are thought to contribute only little to dendritic current and therefore to the generation of the MEG (Bartos et al., 2007; Mitzdorf, 1985). This is because the reversal potentials of inhibitory synapses are relatively close to the resting membrane potential. However, the situation changes when the membrane potential is elevated, and this is indeed supported by experimental results (e.g., Buzsáki et al., 2012). Nevertheless, the actual magnitude of these effects in MEG generation remains to be explored, and is likely to be small when the pyramidal neuron is close to its resting state. For this reason, we opted for the parsimonious approach of ignoring the contribution of inhibitory synapses in order to minimise the number of model parameters. We note that this approximation will be better when the pyramidal population is responding in an adapted way, with short SOIs, when the deviations from the resting state are smaller.

The MEG signal $R(t)$ of the model is computed as:

$$R(t) = \sum_{i,j} k_{i,j}\, w_{i,j}\, q_j(t)\, g(u_j(t)), \tag{4}$$

where the indices i and j range from 1 to N (the number of cortical columns in the network) and identify post- and presynaptic columns respectively, $w_{i,j}$ represents elements of the weight matrix W_{ee}, and $k_{i,j}$ is a scaling factor. For connections located in thalamus, $k_{i,j} = 0$, because thalamic activation does not contribute to the auditory ERF. For intra-field (i.e., for $i = j$) as well as for feedback connections, $k_{i,j} = 1$, and for feedforward connections, $k_{i,j} = -1$ (May et al., 2015). The opposite polarity of the scaling factor k for feedforward and feedback connections reflects the fact that the different connection types target the apical dendrites of pyramidal neurons at different cortical depths. Feedforward connections tend to form in the middle cortical layers, predominantly in layer IV, targeting proximal locations of the pyramidal dendrites, and thus result in a dendritic current directed towards the cortical surface. In contrast, active feedback connections result in a dendritic current directed away from the cortical surface. This is because they target the upper layers I and II, and thus the distal locations on the apical dendrite (see, for example, Ahlfors et al., 2015; Douglas and Martin, 2004; Kohl et al., 2022; Schroeder and Foxe, 2002). Fig. 3 maps out the non-zero k values for the different connection types. Connections beyond the macaque AC connectivity pattern have zero weights and the corresponding scaling factors $k_{i,j}$ were set to zero too.

The MEG signal depends both on the strength and orientation of the primary current, and therefore the curvature of the cortical surface due to cortical folding will affect the MEG (Hämäläinen et al., 1993). We decided to ignore this curvature effect for two reasons. First, the effect of cortical folding is mathematically simple enough to include as a location-specific scaling factor, as was done in our previous work (Hajizadeh et al., 2021). However, its inclusion is not a practical concern in this stage of our current work. It would require a priori knowledge — not

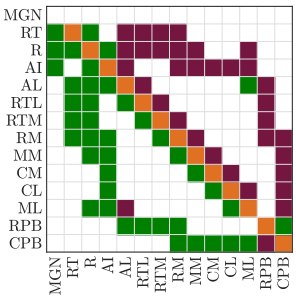

Figure 3: Distribution of the scaling factor $k_{i,j}$. Feedback connections with $k_{i,j} = 1$ are shown in burgundy, feedforward connections with $k_{i,j} = -1$ in green, and intra-field connections with $k_{i,j} = 1$ in orange. The scaling factor for the remaining elements shown in white is zero.

currently available — of the subject-specific cortical topography onto which each field of the core, belt, and parabelt are projected. Second, the simplification can be partly justified by considering the macroanatomy of the human AC, which is located in the Sylvian fissure, on the supratemporal plane. The AC is characterised by two flat areas, the planum temporale and the planum polare, surrounding Heschl's gyrus (HG). Core areas are located along HG, and, in the left hemisphere, the parabelt may be located on the external part of PT, with the belt areas presumably situated in between (Moerel et al., 2014). Thus, it is likely that the majority of cortical columns in AC are oriented in a direction roughly perpendicular to the supratemporal plane — although we note that this is unlikely to be true for all belt fields.

In this work, we simplified the architecture of the AC model by reducing the parameter space and represented each field by a single column to focus on connectivity patterns across fields. The main aim of this work was to optimise the values of two real-valued 14×14 matrices of connection weights shown in Fig. 4. All computations were done in MATLAB (MathWorks, USA).

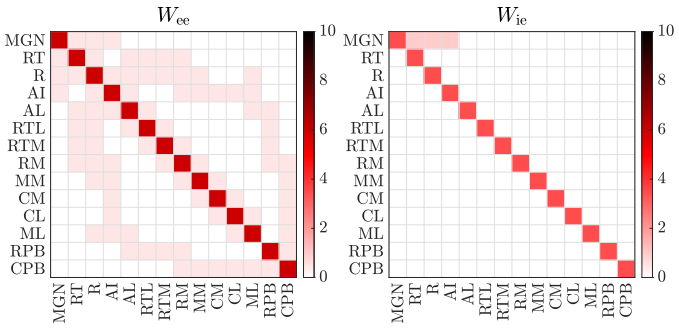

Figure 4: Weight values in the W_{ee} and W_{ie} matrices that were used in the AC model (May et al., 2015; May, 2021) before we started with our EA optimisation work. For W_{ee}, non-zero (non-white) weights are equal to 6 on the leading diagonal and 0.5 elsewhere. For W_{ie}, leading-diagonal weights are equal to 3.5, and the connections between the three core areas (RT, R, AI) and MGN were set to 1. These values are used to create the initial population for our EA implementation; see Sect. 4.2.

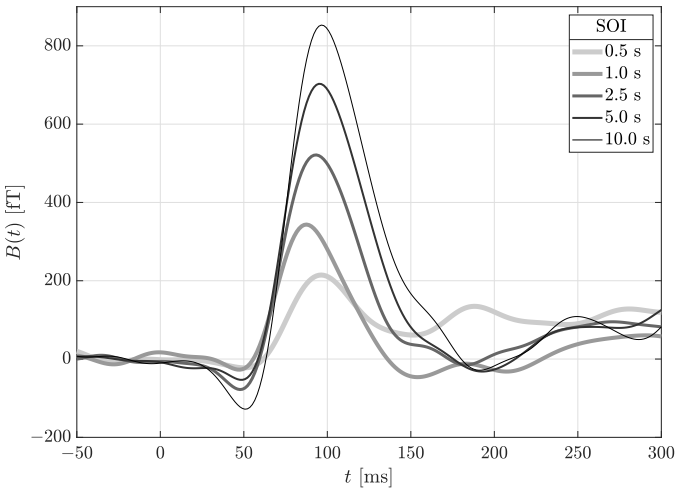

Figure 5: Trial-averaged time courses from an MEG measurement with a single subject. Tones were presented in five discrete blocks which differed in SOIs ranging from 0.5 s to 10 s. With increasing SOI, a gradual increase of the N1m-peak amplitude and of the corresponding N1m-peak latency is observed. $t = 0$ ms reflects stimulus onset.

3. MEG data

MEG recordings were acquired from a subject who gave written informed consent to participate in the MEG study, which was approved by the Ethics Committee of the Otto von Guericke University in Magdeburg, Germany. The subject was passively listening to sequences of pure tones (audio frequency: 1.5 kHz, duration: 100 ms, sound-pressure level: 80 dB), which were presented in five separate blocks (111 stimuli per block) of a regular-SOI paradigm. Each block was characterised by the constant SOI between two consecutive stimuli. The SOI intervals were 0.5 s, 1.0 s, 2.5 s, 5.0 s, and 10.0 s, which translates to stimulation rates of 2.0 Hz, 1.0 Hz, 0.4 Hz, 0.2 Hz, and 0.1 Hz, respectively (Zacharias et al., 2012). The presentation of the blocks was randomised.

Fig. 5 shows the SOI-dependence of the artefact-corrected trial-averaged magnetic fields $B(t)$ for the five SOIs. Recordings taken from the MEG channel with the largest N1m signal above the left hemisphere are shown. The important features of the waveforms are the increase of the N1m peak amplitude with increasing SOI, the common rising slope of all five waveforms towards the N1m peak, and the different falling slopes after the peak.

4. Evolutionary algorithm

Evolutionary algorithms (EAs) are a group of algorithms that search the space of alternative solutions to find the optimal one. These algorithms are called evolutionary because they are based on mechanisms known from genetics and evolutionary theory. The operation of an EA starts with the initiation of an initial population, which consists

of a given number of specimens. Each such specimen represents a specific solution to a given problem. An EA operates in a space called the environment, which is created by the problem to be solved. The adaptation of specimens to the environment is measured through a fitness function, which assigns to each specimen a numerical value indicating the quality of the solution it offers. As in biology, each specimen is assigned a genotype consisting of (the so-called) chromosomes and these are composed of elementary units called genes (Arabas, 2001; Michalewicz, 1996).

4.1. Assumptions and representation of chromosomes

So far, the values of the weight-matrix elements of our AC model have been chosen by trial-and-error; they have not been optimised using advanced optimisation methods. Hence, we decided to investigate how such an optimisation would improve the MEG response generated by the model in comparison to the corresponding ERF recorded in an MEG experiment with a human subject. In the current approach, we have limited ourselves to the optimisation of the matrices W_{ee} and W_{ie}. The weights in the other two matrices W_{ei} and W_{ii} remained fixed at the values reported in previous publications (May et al. (2015); May (2021); ei-connections on the leading diagonal with a weight of 3.5 and equal to zero otherwise; no non-zero ii-connections (omitted for simplicity)).

If, in the context of our EA optimisation, a vectorized connection matrix were to be presented as a chromosome, the computational and memory complexity of the EA algorithm would become unnecessarily large, because only a fraction of elements is assumed to be non-zero. These elements are defined by the pattern of field connections known from the macaque monkey (see elements in different shades of red in Fig. 4). Specifically, no additional connections beyond the known pattern were allowed to be generated. Thus, in our implementation, the chromosome that is the vector representing the W_{ee} matrix has a size of 1×92 and the chromosome representing the W_{ie} matrix has a size of 1×17. We restricted the weight values in the chromosomes to the range $[0, 10]$. Since one of the defining characteristics of the cortical column is the coupling of the neurons it comprises (for a review, see, for example, Gerstner et al., 2014), we allowed one exception: the range for the leading-diagonal elements of W_{ee} (representing weights intra-column connections) had a lower bound of 10^{-3} instead of zero.

4.2. Initial population

The operation of the EA begins with the generation of the initial population. We constrained the solution space such that the initial specimens were created using the weight values shown in Fig. 4, i.e. in accordance with May and Tiitinen (2013); May et al. (2015). However, in order to increase the diversity of the initial population (which is a desired feature in an EA design), we decided to add a random value from the uniform distribution $\mathcal{U}(-0.5, 0.5)$ to each gene in the chromosome of each specimen of the initial population. This approach resulted in good convergence (see the fitness function introduced in Sect. 4.3 and Fig. 9) at an acceptable time cost (see Sect. 6).

4.3. Fitness function

Once created, the initial population is evaluated by examining how well a specimen x has adapted to the environment. This is done for each specimen of the population. Since the time-independent AC structure reflected by W_{ee} and W_{ie} is assumed not to depend on temporal properties of the stimulus, we ran the EA optimisation for all SOIs simultaneously, as the same optimised W_{ee} and W_{ie} apply for all SOIs. For this purpose, the fitness function $\Phi(x)$ was used, which we defined as the dot product of the concatenated form of the true MEG signal (see Sect. 3), $\mathrm{MEG}^{\mathrm{true}}_{\mathrm{conc}} = \left[\mathrm{MEG}^{\mathrm{true}}_{1}, \ldots, \mathrm{MEG}^{\mathrm{true}}_{S}\right]$ with $S = 5$ denoting the number of considered SOIs, and its ℓ^2 normalised counterpart generated by the model using the optimised weight values of the $W^{\mathrm{opti}}_{\mathrm{ee}}$ and $W^{\mathrm{opti}}_{\mathrm{ie}}$ matrices offered by specimen x:

$$\Phi(x) = \left\langle \mathrm{MEG}^{\mathrm{true}}_{\mathrm{conc}}, \frac{\mathrm{MEG}^{\mathrm{opti}}_{\mathrm{conc}}(x)}{\left\|\mathrm{MEG}^{\mathrm{opti}}_{\mathrm{conc}}(x)\right\|_2} \right\rangle. \tag{5}$$

Normalisation of the MEG signal generated by the optimised model provides all specimens with equal chances in the competition for the best fit in terms of signal morphology. The value of $\Phi(x)$ is maximised during subsequent iterations of EA. This means that a specimen with a greater value of $\Phi(x)$ is better than those with smaller values. The value of $\Phi(x)$ constitutes the scaling factor for the normalised synthetic signal to obtain optimal fits with respect to the reference signal. In order not to complicate the notation by introducing a new name or label, $\mathrm{MEG}^{\mathrm{opti}}$ will denote this scaled synthetic signal in later sections.

To constrain the fitness function $\Phi(x)$ given by Eq. (5) to $[0, 1]$ for visualisation purposes, we normalised it by its upper bound, that is, by

$$\left\langle \mathrm{MEG}^{\mathrm{true}}_{\mathrm{conc}}, \frac{\mathrm{MEG}^{\mathrm{true}}_{\mathrm{conc}}}{\left\|\mathrm{MEG}^{\mathrm{true}}_{\mathrm{conc}}\right\|_2} \right\rangle = \left\|\mathrm{MEG}^{\mathrm{true}}_{\mathrm{conc}}\right\|_2. \tag{6}$$

Hence, the normalised fitness function $\Phi_{\mathrm{n}}(x)$ — see Fig. 9 — is given by

$$\Phi_{\mathrm{n}}(x) = \frac{\Phi(x)}{\left\|\mathrm{MEG}^{\mathrm{true}}_{\mathrm{conc}}\right\|_2}. \tag{7}$$

In order to reduce computational burden, we restricted the MEG time signal to the $[-50, 200]$-ms time window, which suffices for a good optimisation with respect to the N1m waveform (see Fig. 5).

4.4. Selection

The purpose of selecting the best parents for reproduction is to obtain best possible offspring. There are many methods to model the selection process. Preliminarily, in the process of designing our algorithm, the roulette method was used for selection, but it did not prove successful due to too little selection pressure, i.e. poor differentiation of specimens with a similar value of the fitness function. Therefore, the ranking method was eventually used (Rutkowska et al., 1997). This method consists in arranging specimens in a temporary population according to the value of the fitness function and assigning each specimen to the rank R. If we define the size of the temporary population as N, then the worst specimen in this population is assigned rank 1, and the best one rank N. To increase selection pressure in the ranking method, the reproduction probability for a specimen to become a parent is calculated according to the formula:

$$P_{\mathrm{r}}(x) = \frac{R(x)}{\sum_{z \in p_{\mathrm{t}}} R(z)}, \tag{8}$$

where $R(x) \in \{1, \ldots, N\}$ is the rank value of a given specimen, and $\sum_{z \in p_{\mathrm{t}}} R(z)$ is the sum of the rank values of all specimens belonging to the current temporary population p_{t}. The operation of the ranking method after defining the probability according to the above formula is still similar to the roulette wheel method, and therefore to the game of roulette, yet, it offers an increased selection pressure. If we imagine that each specimen is assigned an area on the roulette wheel proportional to the calculated probability of reproduction $P_{\mathrm{r}}(x)$, then the choice of the specimen depends on what area the ball will land on. Better specimens will have a larger $P_{\mathrm{r}}(x)$, so they will occupy a larger area on the roulette wheel, which means they will have a better chance of becoming a parent (Arabas, 2001; Michalewicz, 1996; Rutkowska et al., 1997).

4.5. Reproduction

After choosing a pair of parents from the temporary population, reproduction takes place, i.e. the process of creating offspring. In our implementation, four types of arithmetic crossover were used: one-point, two-point, single-point, and whole, which are presented in Fig. 6. Each of these four operators has an equal chance of being used, but only one of them works for a single pair of parents. The task of crossovers in EAs is to explore new solution spaces, because as a result of these operations, offspring differ from their parents. Via an arithmetic crossover of two parental chromosomes, x_1 and x_2, two offspring chromosomes, y_1 and y_2, are created, which are linear combinations of those from the parents according to the formula:

$$y_1 = \alpha x_1 + (1 - \alpha) x_2, \tag{9a}$$

$$y_2 = (1 - \alpha) x_1 + \alpha x_2. \tag{9b}$$

Depending on the problem to be solved, different approaches to the value of the α parameter are used. For example, α may have a value that does not change over the course of the EA, or it changes with each iteration or crossover. Our implementation uses a technique in which the value of α is randomly selected from the range $(0, 1)$ in each crossover. Values 0 and 1 are excluded, since they could result in a replication of a specimen (Michalewicz, 1996).

4.6. Mutation

After the creation of offspring, they are subject to random mutations, which do not cause changes as large as for arithmetic-crossover operations, because the mutated specimens resemble their previous version. In our EA, the possibility of the lack of mutation in a specimen was allowed, which resulted in an additional parameter that can be adjusted, that is, the mutation probability P_{m}. The mutation was simulated by three operators: insertion, deletion, and inversion, whose operation is illustrated in Fig. 7. Only one of these three operators acts on a specimen randomly selected for mutation, but each of the three has equal chance of occurrence. Initially, we used uniform mutation, but since it resulted in too large changes in the chromosomes, we eventually used non-uniform mutation, which was designed to counteract the excessive randomness present in the uniform mutation. The principle of the non-uniform mutation at a randomly selected position i in the chromosome is to randomly draw $\gamma \in \{0, 1\}$, which translates to deriving a new value g_{new} of a gene at position i, taking into account the previous value g_{old} of that gene, according to the formula:

$$g_{\mathrm{new}} = \begin{cases} g_{\mathrm{old}} + \Delta(t, UB - g_{\mathrm{old}}) & \text{if } \gamma = 0, \\ g_{\mathrm{old}} - \Delta(t, g_{\mathrm{old}} - LB) & \text{if } \gamma = 1. \end{cases} \tag{10}$$

Here, LB is the lower bound for g_{old}, equal to 10^{-3} for genes located on the leading diagonal of W_{ee} and 0 for the other genes, UB is the upper bound for g_{old}, equal to 10 for all genes, and t is the EA iteration number. Note that $\Delta(t, \rho)$ attains values in $[0, \rho]$, approaching zero with increasing EA iterations, when its changes become smaller. At these last iterations, the EA explores the solution space more locally than in the initial phase of operation when it explores the space in a more "courageous" fashion. The function $\Delta(t, \rho)$ is defined as:

$$\Delta(t, \rho) = \rho \left(1 - r^{(1 - t/T)b}\right), \tag{11}$$

where r is a random number from $[0, 1]$, t is the current EA iteration, T is the maximum number of iterations, and b is a parameter that determines the degree of dependency on the iteration number (Michalewicz, 1996).

4.7. Succession and stopping criterion

After producing a certain number of offspring, the value of the fitness function is calculated for each of them. Then the set of parent specimens and the set of offspring specimens are merged, and this new set is sorted according

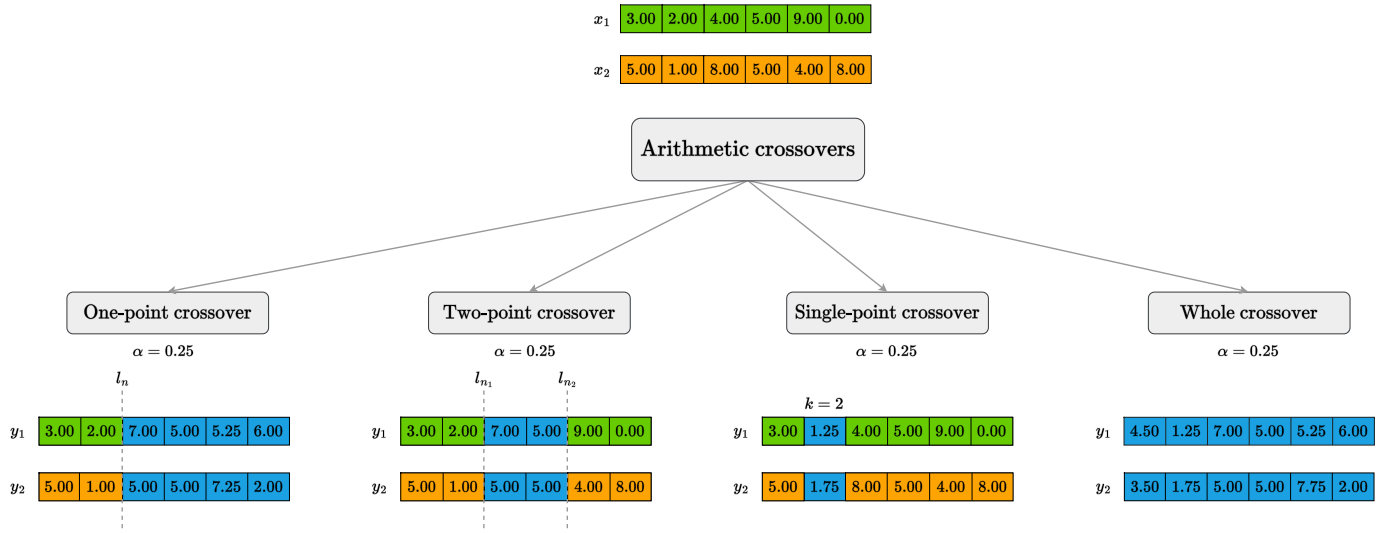

Figure 6: The four types of arithmetic crossover implemented in the EA. In all cases, two offspring y_1 and y_2 are produced from two parents x_1 and x_2. **One-point crossover** involves drawing a cut locus l_n denoting the cut between genes n and $n+1$, where $n \in \{1, \ldots, N-1\}$ and N is the length of the chromosome. The first offspring y_1 is then formed by transcribing the values of the genes from the chromosome of the first parent x_1 located at positions 1 to n, while operations specified by Eq. (9a) are performed at positions $n+1$ to N. Similarly, the second offspring y_2 is formed by transcribing the values of genes with indices 1 to n from the chromosome of parent x_2, while Eq. (9b) operates at positions $n+1$ to N. In a **two-point crossover** two loci, l_{n_1} and l_{n_2}, are drawn, where $n_1 < n_2$. Offspring y_1 is formed by rewriting parental genes at positions 1 to n_1 and n_2+1 to N from parent x_1, whereas (9a) operates at positions n_1+1 to n_2. Similarly, offspring y_2 inherits unchanged genes from parent x_2 at positions 1 to n_1 and n_2+1 to N, whereas (9b) operates at positions from n_1+1 to n_2. **Single-point crossover** involves drawing a position $k \in \{1, \ldots, N\}$. Offspring y_1 is then formed by inheriting all but one genes from parent x_1, that is, all except the gene at position k, where Eq. (9a) holds. Similarly, offspring y_2 is formed from parent x_2, with the crossover given by Eq. (9b) operating at the kth gene. **Whole crossover** produces two offspring, y_1 and y_2, according to Eqs (9a) and (9b) at all gene locations.

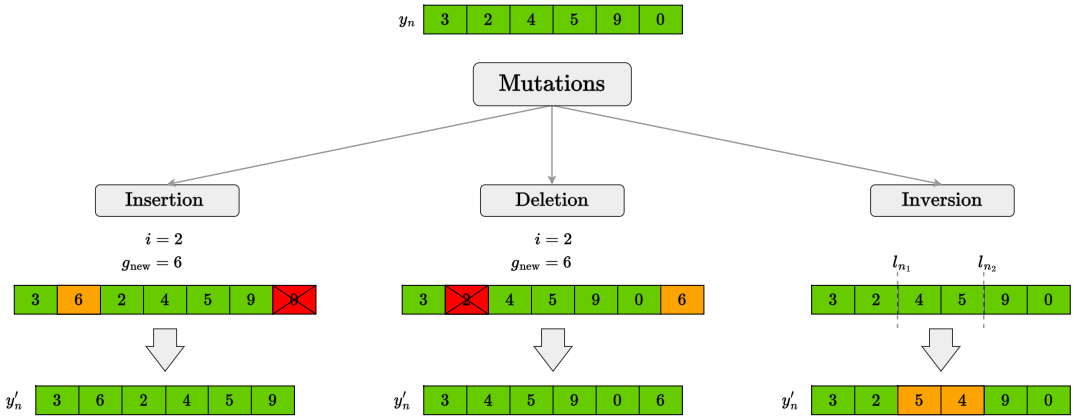

Figure 7: The three types of mutation implemented in the EA. **Insertion** and **Deletion** operate at a randomly selected location $i \in \{1, \ldots, N\}$ in the y_n chromosome, where g_{new} is derived according to Eq. (10). In **Insertion**, a new gene is inserted at the position preceding i, this location becoming the new position i, and this is followed by deletion at the former position N (now $N+1$) in order to preserve the original length of the chromosome equal to N. In **Deletion**, the randomly drawn index i indicates the position of the gene to be deleted from the y_n chromosome, resulting in a chromosome of length $N-1$. In order to preserve the length of the chromosome, g_{new} derived according to Eq. (10) is then appended at the end of the chromosome. In **Inversion**, two loci, l_{n_1} and l_{n_2}, are randomly drawn, with $l_{n_1} < l_{n_2}$, and then the order of genes in the chromosome fragment between these two loci is reversed.

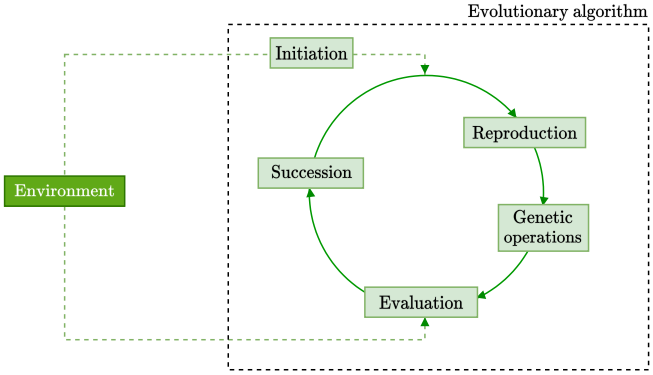

Figure 8: The principle of operation of an evolutionary algorithm.

to the value of the fitness function of each of its members. Subsequently, succession takes place, i.e. a certain number of the best specimens pass to the next generation. Later, reproduction, mutation, and succession are performed again, in a loop, as depicted in Fig. 8. There are two possible criteria for stopping this loop: (1) the number of iterations predetermined by the programmer or (2) the fitness function reaching a certain value. Our implementation uses the first criterion.

4.8. In search of optimal EA parameters

Once the process of designing and implementing the EA has been completed, another task must be solved, which is tuning the parameters of the algorithm, such as the number of specimens or probability of mutation. These parameters were eventually assigned values listed in Tab. 1. The first parameter in the table determines the number of generations after which the EA will stop. Then there is *Number of specimens* which denotes how many specimens were created in the initial population. It was assumed that this number would be kept constant in consecutive generations. Next comes *Selection method*, specifying the strategy chosen for selecting specimens to become parents. Then there is also a parameter called *Mutation*, which indicates the probability that the given offspring will undergo one of the three implemented mutations (see Sect. 4.6). Due to the fact that the values on the leading diagonal of the W_{ee} matrix cannot be equal to zero, another parameter was created and called *Minimum diagonal value of* W_{ee}. The last parameter in the table, *Initial population*, describes the method of creating specimens that will be included in the initial population. The setting shown in the table indicates that all specimens carried a similar pattern of connection strengths across fields, which differed slightly due to the addition of random numbers from a uniform distribution $\mathcal{U}(-0.5, 0.5)$.

Table 1: Summary of the parameters of the evolutionary algorithm that were used to obtain the optimisation results.

Parameter	Setting
Number of iterations	10^3
Number of specimens	10^2
Selection method	ranking method
Mutation	non-uniform
Probability of mutation	0.9
Minimum diagonal value of W_{ee}	10^{-3}
Initial population	identical $+\,\mathcal{U}(-0.5, 0.5)$

5. Results from optimising the AC model

5.1. Connection matrices and waveforms

Due to the fact that EAs belong to the group of heuristic algorithms, we decided to run the EA implementation ten times with the parameter settings listed in Tab. 1. Fig. 9 shows, as an example, results of the EA run with the largest value of the normalised fitness function $\Phi_{\mathrm{n}}(x)$ — see Eq. (7) — which happened to be the first of the ten runs; see Tab. 2. The results for the nine remaining runs can be found in *Supplementary material*. The panels in the top row of Fig. 9 show, for each SOI, three different waveforms. The black trace represents the experimental $\mathrm{MEG}_s^{\mathrm{true}}$ signal with $s \in \{1, \ldots, 5\}$ indexing across the five SOIs, the blue trace represents the scaled (see Sect. 4.3) $\mathrm{MEG}_s^{\mathrm{init}}$ signal generated by the AC model for the initial $W_{\mathrm{ee}}^{\mathrm{init}}$ and $W_{\mathrm{ie}}^{\mathrm{init}}$ matrices, and the red time course represents the model-generated $\mathrm{MEG}_s^{\mathrm{opti}}$ signal for the optimised $W_{\mathrm{ee}}^{\mathrm{opti}}$ and $W_{\mathrm{ie}}^{\mathrm{opti}}$ matrices. For each SOI, we make two distinct observations: 1) $\mathrm{MEG}_s^{\mathrm{opti}}$ matches the time course of $\mathrm{MEG}_s^{\mathrm{true}}$ much better than $\mathrm{MEG}_s^{\mathrm{init}}$ does. 2) There is a good agreement between $\mathrm{MEG}_s^{\mathrm{opti}}$ and $\mathrm{MEG}_s^{\mathrm{true}}$. This holds even for the case of the 0.5-s SOI, where the match is likewise acceptable, given the small magnitude of the $\mathrm{MEG}_1^{\mathrm{true}}$ signal and the overall complex morphology of its entire time course between 0 ms and 200 ms.

The bottom row of Fig. 9 depicts the progression of the normalised fitness function $\Phi_{\mathrm{n}}(x)$ with growing number of iterations as well as the W_{ee} and W_{ie} matrices in their initial and EA-optimised forms. As a result of the optimisation process, the MEG response generated by the model significantly improves from $\mathrm{MEG}_s^{\mathrm{init}}$ to $\mathrm{MEG}_s^{\mathrm{opti}}$. This improvement manifests itself in the monotonic increase of $\Phi_{\mathrm{n}}(x)$. Further, in the course of EA optimisation, the distance between the best and the worst specimen understood as the difference between the corresponding $\Phi_{\mathrm{n}}(x)$ values declined steadily. This behaviour is characteristic of EA algorithms, and it is observed in each of the ten runs. Tab. 2 lists the respective $\Phi_{\mathrm{n}}(x)$ values obtained from the final of the 10^3 iterations for all ten EA runs.

The excellent agreement between true and optimised waveforms is attributed to the transformation of the connections matrices W_{ee} and W_{ie} during the optimisation process, as indicated in the panels of the bottom row of Fig. 9. The homogeneity of the initial matrix $W_{\mathrm{ee}}^{\mathrm{init}}$ is characterised by two different main values, one on the leading diagonal and the other one above and below it. These only vary slightly due to the addition of random numbers from a uniform distribution $\mathcal{U}(-0.5, 0.5)$ to generate a genetically diversified initial population (see Sect. 4.2). Over the

Table 2: Values of $\Phi_{\mathrm{n}}(x)$ — see Eq. (7) — for the ten EA runs together with arithmetic-mean (AM) and standard-deviation (SD) values across the runs.

EA run	1	2	3	4	5	6	7	8	9	10	AM	SD
$\Phi_{\mathrm{n}}(x)$	0.983	0.978	0.981	0.981	0.981	0.977	0.981	0.981	0.981	0.981	0.981	0.002

course of the optimisation, the homogeneity of weight values across the leading diagonal as well as of weight values above and below it was transformed to a pattern of weights in $W_{\mathrm{ee}}^{\mathrm{opti}}$, in which the difference of the weight values between intra-field and feedforward or feedback connections was modified compared to $W_{\mathrm{ee}}^{\mathrm{init}}$. Further implications of these findings are discussed in Sect. 5.2. The changes from $W_{\mathrm{ie}}^{\mathrm{init}}$ to $W_{\mathrm{ie}}^{\mathrm{opti}}$ were more subtle and did not entail a similarly large change in the intra-field weights as obtained for W_{ee}.

Fig. 10 displays the MEG time courses from all ten EA runs. The true MEG signals ($\mathrm{MEG}_s^{\mathrm{true}}$) are plotted in black, whereas the vivid-blue trace displays the arithmetic mean of the ten initial MEG signals ($\mathrm{MEG}_s^{\mathrm{init}}$, shown in pale blue). The pale-red traces ($\mathrm{MEG}_s^{\mathrm{opti}}$) are the ten MEG signals computed with the ten EA-optimised connection matrices, and their arithmetic mean is shown in vivid red. Note that for each run, the EA does not return identical results due to EA's stochastic properties. Nevertheless, Figs 9 and 10 show that the MEG signals produced by the AC model after EA optimisation of the two connection matrices W_{ee} and W_{ie} are very similar to the true MEG, especially in the vicinity of the N1m waveform.

Fig. 11 shows the arithmetic-mean and standard-deviation representations of the $W_{\mathrm{ee}}^{\mathrm{opti}}$ and $W_{\mathrm{ie}}^{\mathrm{opti}}$ matrices computed across all ten EA runs. Notably, standard deviations are relatively small compared to the means, which speaks for a decent reproducibility of the $W_{\mathrm{ee}}^{\mathrm{opti}}$ and $W_{\mathrm{ie}}^{\mathrm{opti}}$ estimates across EA runs.

5.2. Asymmetry in feedforward and feedback connections

As shown in the bottom row of Fig. 9, the optimisation approach resulted in an asymmetry of weight values between feedforward and feedback connections in the EA-optimised W_{ee} matrices. Mean feedback connection weights from parabelt to belt, $\overline{w}_{\mathrm{b\leftarrow p}}^{\mathrm{opti}} = 1.307$, between belt fields, $\overline{w}_{\mathrm{b\leftarrow b}}^{\mathrm{opti}} = 1.592$, and from belt to core, $\overline{w}_{\mathrm{c\leftarrow b}}^{\mathrm{opti}} = 0.800$, are stronger than the respective feedforward connection weights of $\overline{w}_{\mathrm{b\rightarrow p}}^{\mathrm{opti}} = 1.247$, $\overline{w}_{\mathrm{b\rightarrow b}}^{\mathrm{opti}} = 1.394$, and $\overline{w}_{\mathrm{c\rightarrow b}}^{\mathrm{opti}} = 0.729$. This asymmetry is also reflected in the corresponding waveforms shown in the upper row of Fig. 12. Here, for each SOI of the run whose waveforms and matrices are presented in Fig. 9, the $\mathrm{MEG}_s^{\mathrm{opti}}$ waveforms (red curves) are broken down into feedforward (green curves), feedback (burgundy curves) and intra-field contributions (orange curves). Both feedback and intra-field connections provide positive contributions to the overall MEG response, with the feedback contributions being roughly 1.5 times as large as the intra-field contributions. In general, the dominance of feedback over intra-field contributions holds true for all SOIs and across all ten EA runs, with feedback contributions being between 1.5 times and twice as large as intra-field contributions. The feedforward contributions are of opposite polarity, and, thus, reduce the sum of the large positive feedback and intra-field contributions to the $\mathrm{MEG}_s^{\mathrm{opti}}$ response. The difference between the absolute values of the feedback and feedforward waveforms is indicated by the grey curves. The initial dominance of feedforward contributions around the P1m can be explained by the delayed onset of feedback activity relative to feedforward activity. Subsequently, the feedback contribution becomes dominant during the N1m. Again, this holds true for all SOIs and all ten EA runs (see bottom row of Fig. 12). This confirms that the stronger feedback connections in the time-independent $W_{\mathrm{ee}}^{\mathrm{opti}}$ matrix, described earlier, also results in a difference at the level of the dynamic response of the model.

Tab. 3 summarises means across feedforward, feedback, and intra-field connections calculated using the classification depicted in Fig. 3. Note that the ratio of the mean feedback weight to the mean feedforward weight in $W_{\mathrm{ee}}^{\mathrm{opti}}$ — that is, $\overline{w}_{\mathrm{FB}}^{\mathrm{opti}}/\overline{w}_{\mathrm{FF}}^{\mathrm{opti}}$ — ranges from 1.073 to 1.338, compared to the respective range of $\overline{w}_{\mathrm{FB}}^{\mathrm{init}}/\overline{w}_{\mathrm{FF}}^{\mathrm{init}}$ spanning from 0.625 to 1.162 for the initial conditions reflected by $W_{\mathrm{ee}}^{\mathrm{init}}$. Some small asymmetries in the weights of the initial W_{ee} matrices are caused by the aforementioned diversification of the initial population by adding random numbers from the uniform distribution $\mathcal{U}(-0.5, 0.5)$ to the fixed values inherited from May et al. (2015); see Sect. 4.2 and Fig. 4. Also, the feedback connections appear to be stronger than previously assumed by May et al. (2015) and May (2021); see Fig. 4. Moreover, the earlier assumption of identical connection weights across the entire leading diagonal of both W_{ee} and W_{ie} turns out not to hold, as revealed by the EA-optimised versions of these matrices. For clarity, data for the ten EA runs from Tab. 3 are visualised in Fig. 13.

5.3. Spatial distribution of optimised waveforms

In the traditional view, ERFs are considered to arise from spatially discrete, local (point-like) sources. The inclusion of the anatomy of the auditory cortex as an essential part of our AC model provides an alternative perspective on the generation of ERFs. In this view, there is only one underlying process distributed throughout the entire network: ERFs are generated by the network properties of the whole AC (Hajizadeh et al., 2019, 2022). Specifically, all AC fields contribute to the ERF, which, as a corollary, can be broken down into individual contributions from core, belt, and parabelt. An example of this decomposition of the optimised ERF is shown in Fig. 14.

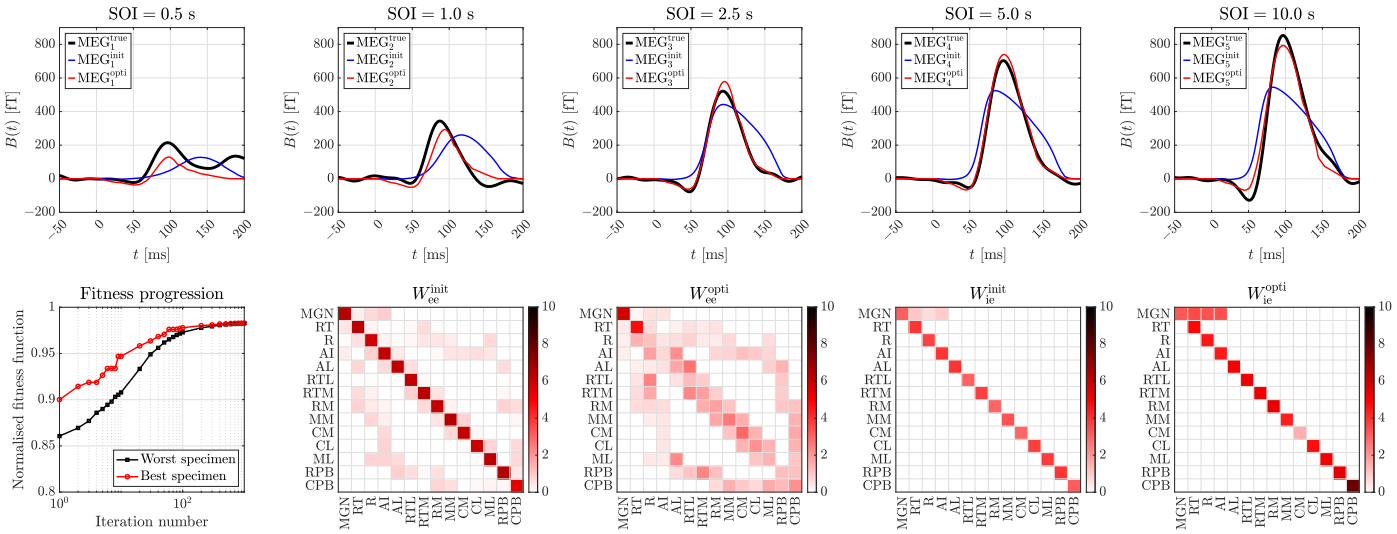

Figure 9: Results for the one EA run (out of ten) that resulted in the largest value of the fitness function in the last iteration. The top row shows MEG waveforms, with black traces representing the true MEG, blue traces the MEG output of the AC-model for the best initial specimen, and red traces the MEG output of the AC-model for EA optimised W_{ee} and W_{ie}. The left panel in the bottom row shows the progression of the normalised fitness function (see Eq. (7)), while the panels to the right depict the initial and optimised W_{ee} as well as the initial and optimised W_{ie} matrices.

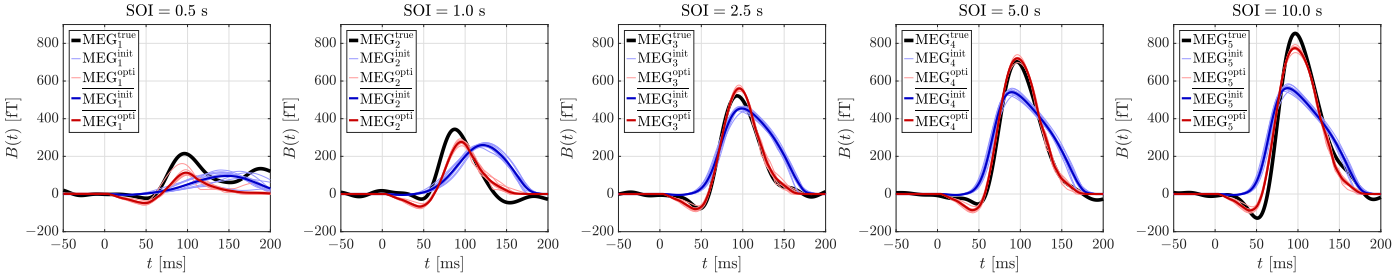

Figure 10: EA optimisation results for the five SOIs. In each panel, the black curve represents the true MEG signal, the vivid-blue curve the MEG signal created by averaging the ten MEG signals (shown in pale blue) simulated with the AC model using the initial W_{ee} and W_{ie} matrices, and the vivid-red curve the MEG signal computed by averaging the MEG signals (in pale red) simulated with the AC model using the EA-optimised W_{ee} and W_{ie}.

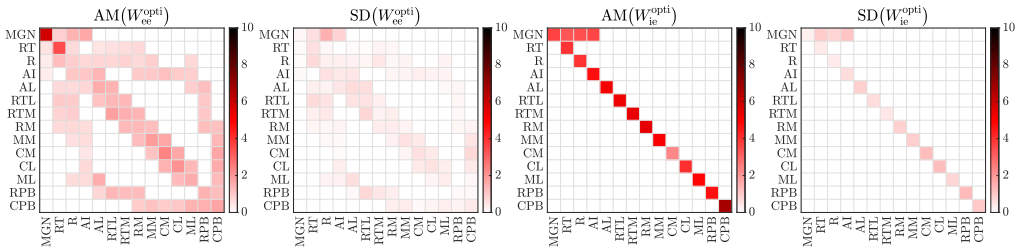

Figure 11: Arithmetic-mean (AM) representations of the optimised weight matrices W_{ee} (first panel) and W_{ie} (third panel) computed across the ten EA runs along with the corresponding standard deviations (SD, second and fourth panel, respectively). The standard deviations are small relative to the respective mean weight values.

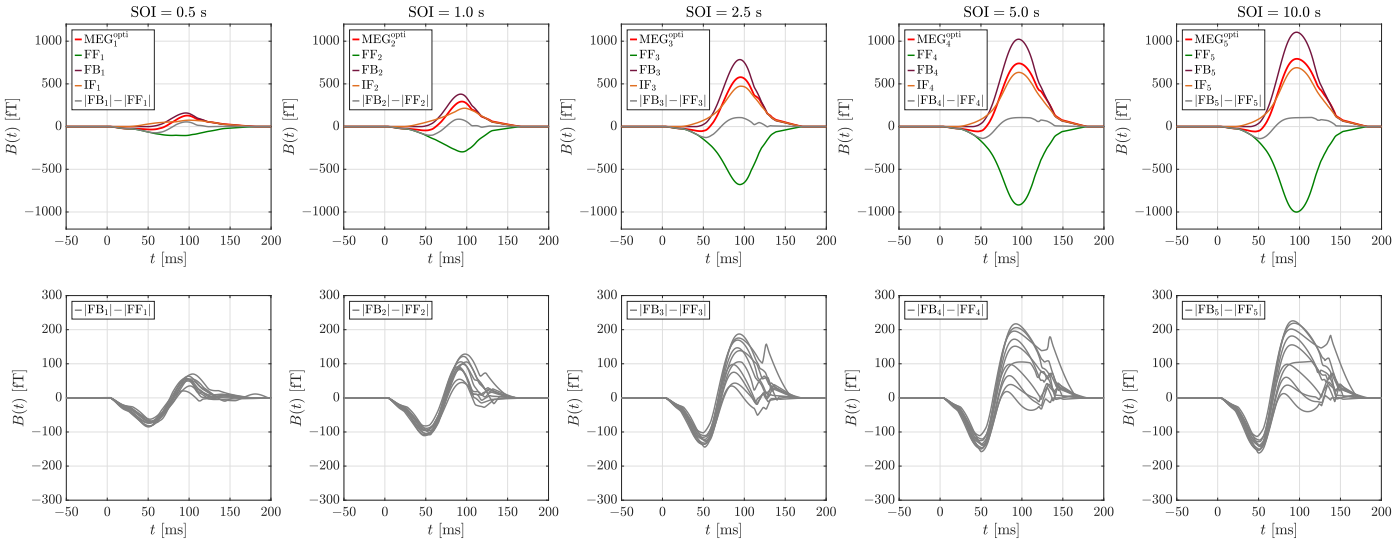

Figure 12: Decomposition of $\mathrm{MEG}_s^{\mathrm{opti}}$, that is, the net simulated ERF into contributions from the different connection types. Akin to Fig. 10, the top row shows results for the one EA run (out of ten) that resulted in the largest value of the fitness function in the last iteration. Here, in each of the five panels the net ERF (red curve) is decomposed into contributions from feedforward connections (green curve), feedback connections (burgundy curve) and intra-field connections (orange curve) for a given SOI. Additionally, the difference between the absolute values of the time courses for feedback and feedforward contributions is plotted in grey. In the bottom row, this time course difference is plotted across all ten EA runs of the EA and for each SOI respectively.

Table 3: Arithmetic means of $w_{i,j}$ values (see Eq. (4)) of the initial and EA-optimised W_{ee} matrices across feedforward (FF), feedback (FB), and intra-field (IF) connections, respectively (see Fig. 3). The last two rows show the arithmetic mean (AM) and standard deviation (SD) of these mean values across the ten EA runs.

EA run	$\overline{w}_{\mathrm{FF}}^{\mathrm{init}}$	$\overline{w}_{\mathrm{FB}}^{\mathrm{init}}$	$\overline{w}_{\mathrm{IF}}^{\mathrm{init}}$	$\overline{w}_{\mathrm{FF}}^{\mathrm{opti}}$	$\overline{w}_{\mathrm{FB}}^{\mathrm{opti}}$	$\overline{w}_{\mathrm{IF}}^{\mathrm{opti}}$
1	0.479	0.430	6.178	1.001	1.137	2.027
2	0.433	0.470	5.981	0.982	1.245	1.600
3	0.477	0.428	5.905	0.945	1.095	1.835
4	0.426	0.495	5.823	0.985	1.318	1.581
5	0.456	0.402	5.956	0.953	1.230	1.868
6	0.491	0.454	5.928	0.935	1.191	1.540
7	0.428	0.459	5.921	1.053	1.172	1.685
8	0.467	0.339	6.118	1.023	1.327	1.606
9	0.536	0.335	6.033	0.928	0.996	2.018
10	0.454	0.430	5.832	0.986	1.080	1.961
AM	0.465	0.424	5.968	0.979	1.179	1.772
SD	0.034	0.053	0.115	0.040	0.106	0.191

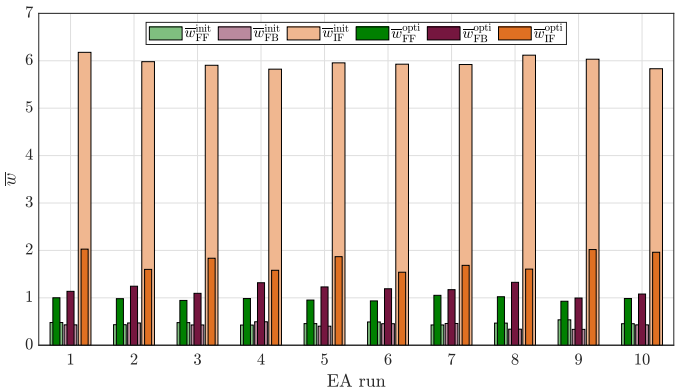

Figure 13: Visual representation of the data from Tab. 3 for the ten EA runs.

Here, the dominant contribution to the overall ERF signal (red curves) stems from the summation across the eight belt fields (cyan curves), which clearly exceeds the summed contribution from the three core fields (lime curves). This observation applies for all SOIs and also holds true for a total of six out of the ten EA runs. In the remaining four runs, contributions from core and belt were roughly equal across all SOIs. The contributions from the two parabelt fields (purple curves) are the weakest across all SOIs. This was the case across all EA ten runs. We also noticed a clear shift of the parabelt waveform to higher N1m-peak latency values compared with the core and belt waveforms.

Adaptation, the reduction of a neuronal response due to stimulus repetition, has a lifetime which is traditionally estimated through using the regular-SOI paradigm. In this approach, the amplitude of the response increases as a function of SOI, and this dependence is described by fitting the peak amplitudes with an exponentially saturating function (Lü et al., 1992; Lu et al., 1992). The steepness of the fitted function provides the adaptation time constant τ_{SOI}. For the $\mathrm{MEG}^{\mathrm{opti}}$ signal, the mean adaptation time constant across the ten EA runs was $\overline{\tau}_{\mathrm{SOI}} = 2.029$ s (SD = 0.054 s). The decomposition of the waveforms into core, belt, and parabelt contributions enabled the computation of area-specific τ_{SOI} values: The shortest $\overline{\tau}_{\mathrm{SOI}} = 1.841$ s (SD = 0.128 s) was found in the belt area, whereas the means for core and parabelt barely differed from each other, with $\overline{\tau}_{\mathrm{SOI}} = 2.459$ s (SD = 0.151 s) for the core and $\overline{\tau}_{\mathrm{SOI}} = 2.431$ s (SD = 0.242 s) for the parabelt.

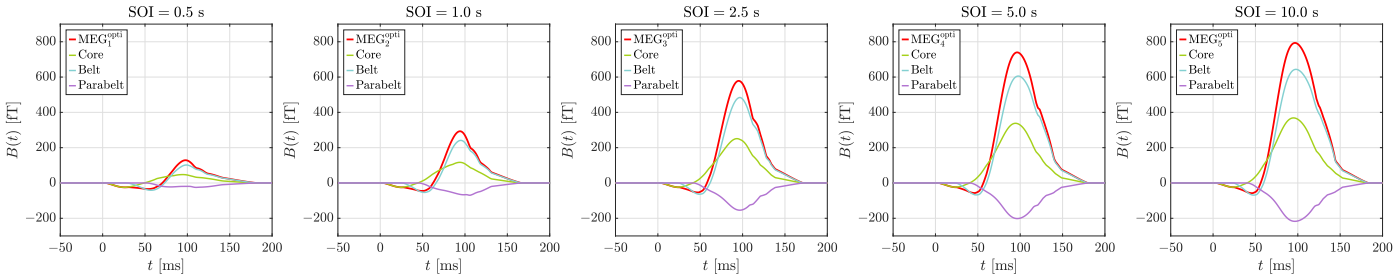

Figure 14: Decomposition of $\text{MEG}_s^{\text{opti}}$, that is, the net simulated ERF into contributions from the three different AC areas for the one EA run (out of ten) that resulted in the largest value of the fitness function in the last iteration. In each panel, the net ERF (red curve) is decomposed into contributions from core (lime curve), belt (cyan curve), and parabelt (purple curve) for a given SOI.

6. Discussion

The human auditory cortex is poorly understood in terms of its anatomical structure, that is, in terms of its parcellation into fields and the connectivity between fields. It is also poorly understood how the interaction between these fields shapes the auditory ERF. Here, we demonstrated that computational modelling combined with advanced optimisation methods could be useful for tackling these questions. We propose an EA optimisation of the AC model by May et al. (2015) and May (2021) so that the model is able to synthesise auditory ERFs measured with a variety of SOIs. In this approach, the model parameters to be optimised are the connection strengths between the various fields of the model. The optimisation resulted in a strong improvement of the synthesised MEG outputs of the model in terms of their fits to true MEG signals. In particular, the N1m waveforms elicited by stimuli presented at various SOIs could be synthesised with a very good fit to the experimental data. In earlier work of May et al. (2015) and May (2021), feedforward connection weights were assumed to be symmetrical with feedback weights, and intra-field connection weights expressed on the leading diagonals of the weight matrices were assumed to be identical across the fields (Fig. 4). These assumptions did not survive EA optimisation (see Fig. 9): In the optimised model, feedback connections were stronger than feedforward ones, and the intra-field connections varied from field to field.

The function and significance of cortical feedback connections are still unclear, with the literature containing various views. Crick and Koch (1998) and Sherman and Guillery (2011) reported that feedback connections are weak and modulatory. Bastos et al. (2012) assumed that they are functionally inhibitory, supporting suppressive top-down prediction signals in the predictive coding framework. Our results indicate otherwise: feedback connections are neither weak nor functionally inhibitory. This result is consistent with that of Hajizadeh et al. (2019) who investigated the differential contributions to the synthesised MEG signal made by active feedforward and feedback connections. The weights of these contributions are expressed in the k terms in Eq. (4). By systematically varying the weights of the feedforward contributions to the MEG signal while keeping the feedback contributions fixed, and vice versa, they found that the P1m reflects primarily feedforward activation, whereas feedback activation predominantly drives the N1m and P2m. In the current optimised model, therefore, the prominence of the N1m response of the auditory ERF is linked to the prominence of feedback connections over feedforward ones.

We inspected ERF generation and adaptation in the three areas of AC separately. The core, the belt, and the parabelt all contributed to ERF generation across the whole time course of the response. The ERF, therefore, does not arise out of spatially discrete, localised sources but, rather, reflects the activity of the entire AC. The contribution to the ERF from each area adapts, and thus adaptation emerges from a complete reorganisation of AC dynamics rather than a reduction of activity in discrete sources (see also Hajizadeh et al., 2019, 2021, 2022). Further, the lifetime of adaptation is area-specific, with the belt recovering more quickly than the core and the parabelt. This offers a prediction that might be testable through non-invasive measurements.

An inspection of the results across the ten EA runs reveals a diversity both in the optimised forms of W_{ee} and W_{ie} (see *Supplementary material*) as well as in the synthetic MEG waveforms, due to the stochastic nature of EAs. The diversity in the connection matrices might potentially indicate the risk that the mapping $M : \{W_{\text{ee}}, W_{\text{ie}}\} \to \text{MEG}$ is non-injective surjective.[1] However, since we have never obtained two identical MEG time courses from any pair of independent EA runs optimising the AC model, we identify no formal objections to recapitulate the hypothesis stated in Sect. 1 that the anatomical structure of the human AC is reflected in the signatures of auditory evoked responses recorded by MEG (Hajizadeh et al., 2019, 2021, 2022). Further, noting the diversity of the synthetic MEG waveforms, we derived their respective arithmetically averaged versions, as depicted in Fig. 10.

Calculations using EAs can be demanding. In order to decrease time cost, we parallelised the pipeline wherever it was possible so that efficient performance could be

[1] https://www.wolframalpha.com/examples/mathematics/mathematical-functions/injectivity-and-surjectivity

achieved using multi-thread processing. Using a computer with two 24-core 2.9-GHz Intel Xeon Platinum 8268 CPUs, a single EA iteration takes only about 10 s. Further reduction of the total time can be achieved by distributing the computations across multiple nodes of a computational cluster or by using GPU architectures.

In future work, it would be worth exploring other optimisation methods such as gradient-based algorithms. These tend to approach the optimal solution faster than evolutionary algorithms when the objective function is differentiable and has an unambiguous extremum. Nevertheless, evolutionary algorithms, due to their stochastic nature and ability to explore the solution space, can avoid getting stuck in local minima more easily and thus search for a global solution more effectively, although this does require more time (Yu and Gen, 2010). An effective method that could also be used in advanced optimisation of parameters of computational models is the Sequential Neural Posterior Estimation (SNPE), which provides both optimised parameters and their distributions (Gonçalves et al., 2020). However, Gonçalves et al. (2020) state that scaling SNPE to optimise more than 30 parameters is challenging, as estimating full posterior scales exponentially with the dimensionality determined by the number of parameters. Note that in our work, we are optimising 109 weight values (92 constituting genes in the chromosome reflecting the non-zero entries in W_{ee} and 17 constituting genes in the chromosome reflecting the non-zero entries in W_{ie}). Regardless, the SNPE method could be used in our problem to adjust other parameters than weight values. In the future, we intend to explore alternative optimisation methods to determine if the optimisation results are reproducible with acceptable variance and free from bias.

The current work represents a promising first step in developing a new tool for exploring the connection patterns of human auditory cortex. This relies on MEG measurements being used as the basis for optimising the structure of a computational model of AC. The main limitation of the current approach is the a priori parcellation of the model into fields. This field structure is poorly understood in humans, and we therefore used the organisation of the monkey auditory cortex (Kaas and Hackett, 2000b; Hackett et al., 2014) as the template. In future work, we look forward to extending our approach so that the field structure itself can be optimised.

7. Acknowledgements

This research was supported by the Deutsche Forschungsgemeinschaft through project KO1713/12-1. Calculations have been carried out at Wrocław Centre for Networking and Supercomputing (wcss.pl), grant No. 596.

References

Ahlfors, S.P., Jones, S.R., Ahveninen, J., Hämäläinen, M.S., Belliveau, J.W., Bar, M., 2015. Direction of magnetoencephalography sources associated with feedback and feedforward contributions in a visual object recognition task. Neuroscience Letters 585, 149–154.

Arabas, J., 2001. Wykłady z algorytmów ewolucyjnych. Wydawnictwo Naukowo-Techniczne. In Polish.

Bartos, M., Vida, I., Jonas, P., 2007. Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. Nature Reviews Neuroscience 8, 45–56.

Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J., 2012. Canonical microcircuits for predictive coding. Neuron 76, 695–711.

Buzsáki, G., Anastassiou, C.A., Koch, C., 2012. The origin of extracellular fields and currents: EEG, ECoG, LFP and spikes. Nature Reviews Neuroscience 13, 407–420.

Crick, F., Koch, C., 1998. Constraints on cortical and thalamic projections: the no-strong-loops hypothesis. Nature 391, 245–250.

David, O., Kiebel, S.J., Harrison, L.M., Mattout, J., Kilner, J.M., Friston, K.J., 2006. Dynamic causal modeling of evoked responses in EEG and MEG. NeuroImage 30, 1255–1272.

Douglas, R.J., Martin, K.A.C., 2004. Neuronal circuits of the neocortex. Annual Review of Neuroscience 27, 419–451.

Gerstner, W., Kistler, W.M., Naud, R., Paninski, L., 2014. Neuronal dynamics: From single neurons to networks and models of cognition. Cambridge University Press.

Gonçalves, P.J., Lueckmann, J.M., Deistler, M., Nonnenmacher, M., Öcal, K., Bassetto, G., Chintaluri, C., Podlaski, W.F., Haddad, S.A., Vogels, T.P., Greenberg, D.S., Macke, J.H., 2020. Training deep neural density estimators to identify mechanistic models of neural dynamics. eLife 9, e56261.

Hackett, T.A., de la Mothe, L.A., Camalier, C.R., Falchier, A., Lakatos, P., Kajikawa, Y., Schroeder, C.E., 2014. Feedforward and feedback projections of caudal belt and parabelt areas of auditory cortex: refining the hierarchical model. Frontiers in Neuroscience 8, 1–21.

Hajizadeh, A., Matysiak, A., Brechmann, A., König, R., May, P.J.C., 2021. Why do humans have unique auditory event-related fields? Evidence from computational modeling and MEG experiments. Psychophysiology 58, e13769.

Hajizadeh, A., Matysiak, A., May, P.J.C., König, R., 2019. Explaining event-related fields by a mechanistic model encapsulating the anatomical structure of auditory cortex. Biological Cybernetics 113, 321–345.

Hajizadeh, A., Matysiak, A., Wolfrum, M., May, P.J.C., König, R., 2022. Auditory cortex modelled as a dynamical network of oscillators: Understanding event-related fields and their adaptation. Biological Cybernetics 116, 475–499.

Hämäläinen, M., Hari, R., Ilmoniemi, R.J., Knuutila, J., Lounasmaa, O.V., 1993. Magnetoencephalography—theory, instrumentation, and applications to non-invasive studies of the working human brain. Reviews of Modern Physics 65, 413–497.

Hopfield, J.J., Tank, D.W., 1986. Computing with neural circuits: A model. Science 233, 625–633.

Kaas, J., Hackett, T., 2000a. Subdivisions of auditory cortex and processing streams in primates. Proceedings of the National Academy of Sciences of the United States of America 97, 11793–11799.

Kaas, J.H., Hackett, T.A., 2000b. Subdivisions of auditory cortex and processing streams in primates. Proceedings of the National Academy of Sciences of the United States of America 97, 11793–11799.

Kohl, C., Parviainen, T., Jones, S.R., 2022. Neural mechanisms underlying human auditory evoked responses revealed by Human Neocortical Neurosolver. Brain Topography 35, 19–35.

Loebel, A., Nelken, I., Tsodyks, M., 2007. Processing of sound by population spikes in a model of primary auditory cortex. Frontiers in Neuroscience 1, 197–209.

Lü, Z.L., Williamson, J., Kaufman, L., 1992. Human auditory primary and association cortex have differing lifetimes for activation traces. Brain Research 572, 236–241.

Lu, Z.L., Williamson, S.J., Kaufman, L., 1992. Behavioral lifetime of human auditory sensory memory predicted by physiological measures. Science 258, 1668–1670.

May, P.J.C., 2021. The adaptation model offers a challenge for the predictive coding account of mismatch negativity. Frontiers in Human Neuroscience 15, 1–10.

May, P.J.C., Tiitinen, H., 2010. Mismatch negativity (MMN), the deviance-elicited auditory deflection, explained. Psychophysiology 47, 66–122.

May, P.J.C., Tiitinen, H., 2013. Temporal binding of sound emerges out of anatomical structure and synaptic dynamics of auditory cortex. Frontiers in Computational Neuroscience 7, 1–15.

May, P.J.C., Westö, J., Tiitinen, H., 2015. Computational modelling suggests that temporal integration results from synaptic adaptation in auditory cortex. European Journal of Neuroscience 41, 615–630.

Michalewicz, Z., 1996. Genetic Algorithms + Data Structures = Evolution Programmes. Springer, Berlin, Heidelberg.

Mitzdorf, U., 1985. Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. Physiological Review 65, 37–100.

Moerel, M., Martino, F.D., Formisano, E., 2014. An anatomical and functional topography of human auditory cortical areas. Frontiers in Neuroscience 8, article 225.

Noto, M., Nishikawa, J., Tateno, T., 2016. An analysis of nonlinear dynamics underlying neural activity related to auditory induction in the rat auditory cortex. Neuroscience 318, 58–83.

Pérez-González, D., Malmierca, M.S., 2014. Adaptation in the auditory system: an overview. Frontiers in integrative neuroscience 8, 19.

Ritter, P., Schirner, M., Mclntosh, A.R., Jirsa, V.K., 2013. The virtual brain integrates computational modeling and multimodal neuroimaging. Brain Connectivity 3, 121–145.

Romanski, L., Averbeck, B., 2009. The primate cortical auditory system and neural representation of conspecific vocalizations. Annual Review of Neuroscience 32, 315–346.

Rutkowska, D., Piliński, M., Rutkowski, L., 1997. Sieci neuronowe, algorytmy genetyczne i systemy rozmyte. Wydawnictwo Naukowe PWN. In Polish.

Schroeder, C.E., Foxe, J.J., 2002. The timing and laminar profile of converging inputs to multisensory areas of the macaque neocortex. Cognitive Brain Research 14, 187–198.

Sherman, S.M., Guillery, R.W., 2011. Distinct functions for direct and transthalamic corticocortical connections. Journal of Neurophysiology 106, 1068–1077.

Tsodyks, M., Markram, H., 1997. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. Proceedings of the National Academy of Sciences of the United States of America 94, 719–723.

Turczak, E., 2022. Investigating signal processing in the human auditory cortex. Master's thesis. Wrocław University of Science and Technology.

Ulanovsky, N., Las, L., Farkas, D., Nelken, I., 2004. Multiple time scales of adaptation in auditory cortex neurons. Journal of Neuroscience 24, 10440–10453.

Wang, P., Knösche, T.R., 2013. A realistic neural mass model of the cortex with laminar-specific connections and synaptic plasticity-evaluation with auditory habituation. PLoS One 8.

Westö, J., May, P.J.C., Tiitinen, H., 2016. Memory stacking in hierarchical networks. Neural Computation 28, 327–253.

Wilson, H.R., Cowan, J.D., 1972. Excitatory and inhibitory interactions in localized populations of model neurons. Biophysical Journal 12, 1–24.

Yu, X., Gen, M., 2010. Introduction to Evolutionary Algorithms. Springer, London.

Zacharias, N., König, R., Heil, P., 2012. Stimulation-history effects on the M100 revealed by its differential dependence on the stimulus onset interval. Psychophysiology 49, 909–919.