Proof of concept simulations for validity solutions (D1.9 – SGA3)

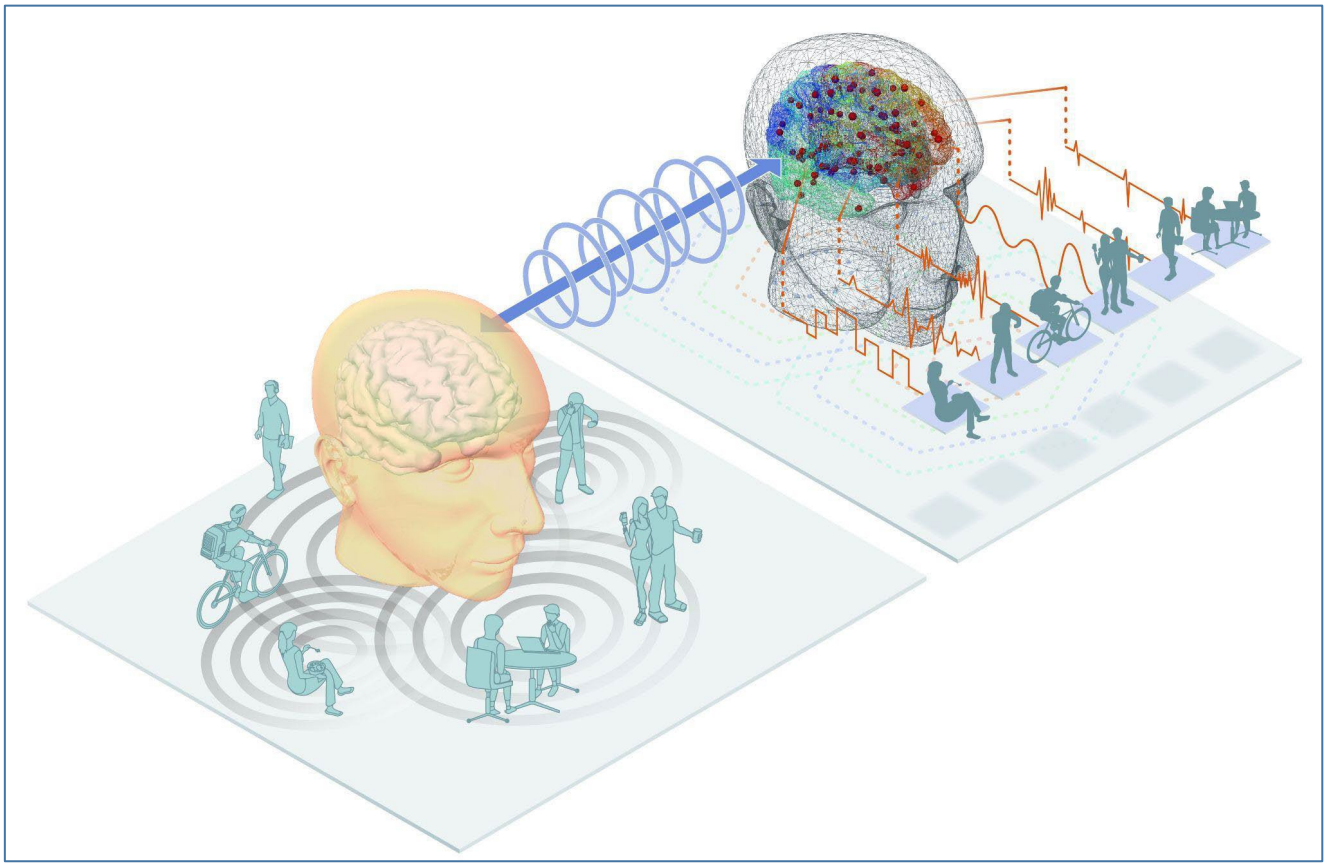

Figure 1: How to make the Virtual Brain think?

To develop steps towards functionally relevant information processing in the Virtual Brain, biologically realistic full brain network modelling is linked with machine learning approaches for pattern discrimination in the visual and auditory domain.

Project Number:	945539	Project Title:	HBP SGA3

Document Title:	Proof of concept simulations for validity solutions
Document Filename:	D1.9 (D118) SGA3 M42 SUBMITTED 230928.docx
Deliverable Number:	SGA3 D1.9 (D118)
Deliverable Type:	Demonstrator
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA3 M42 / 29 SEP 2023
Actual Delivery Date:	SGA3 M42 / 28 SEP 2023
Author(s):	Jan FOUSEK, Viktor JIRSA, AMU (P78)
Compiled by:	Jan FOUSEK, AMU (P78)
Contributor(s):	Jan FOUSEK, AMU (P78) contributed to all Sections Viktor JIRSA, AMU (P78) contributed to all Sections Alain DESTEXHE CNRS (P10) contributed to all Sections Mario SENDEN, UM (P117) contributed to all Sections Rainer GOEBEL, UM (P117) contributed to all Sections Maria V. SANCHEZ-VIVES, IDIBAPS (P93) contributed to all Sections Emre BASPINAR CNRS (P10) contributed to Section 3 Abdullah TURAN, CNRS (P10), contributed to Section 3 Salil BHAT, UM (P117) contributed to Section 2 Leonardo DALLA PORTA, IDIBAPS (P93) contributed to Section 2
WP QC Review:	Giovanna RAMOS QUEDA, AMU (P78), Pilar F. ROMERO, UPM (P68)
WP Leader / Deputy Leader Sign Off:	Viktor JIRSA, AMU (P78)
T7.4 QC Review:	N/A
Description in GA:	A software will be developed showing first traits of integrating task relevant brain function and task-related learning into activity-based models. The prototypes will be a first implementation of the concepts proposed in the blueprint (D3.19 at M40)

	and demonstrate the functionality of individual key elements relevant to model integration (such as plasticity via deep learning, Structured Flows on Manifolds (SFM) for Motion Description Language, or auditory discrimination). Tasks contributing to this Deliverable: T1.5, T1.11, T2.1, T2.4, T2.9 and T3.1
Abstract:	The Virtual Brain models successfully predict evolution of brain activity within the anatomical constraints and can be personalized in order to describe changes of subject's brain activity across brain states and a range of diseases. However, these models currently lack the capacity to represent functionally meaningful processes such as cognition. This report presents the outcome of WP1-WP2-WP3 interactions, culminating in prototypes of virtual brain models capable of executing cognitive tasks. These prototypes draw inspiration from visual and auditory discrimination tasks, and despite their early stages they offer the potential for experimental validation and deeper insights into underlying brain mechanisms.
Keywords:	Virtual brain, deep networks, auditory discrimination, visual discrimination
Target Users/Readers:	computational neuroscience community, cognitive neuroscientists, computer scientists, general public, HPC community, neuroimaging community, neuroinformaticians, researchers, students

Table of Contents

Table of Figures

1. Introduction

Today the virtual brain does not think. What this phrase is meant to communicate is the fact that no functionally meaningful processes such as cognition, memory, or pattern discrimination are represented in the brain activation patterns when simulated by the The Virtual Brain (TVB). What the latter does successfully is to predict the evolution of brain activity under the anatomical constraints imposed upon the brain model. Virtual brains can be personalised through the use of individual brain imaging data and describe a subject's brain activity changes under influence of variation of diseases such as epilepsy, stroke, neurodegeneration. This makes the virtual brain useful in clinical applications (WP1, WP2). It also is deeply rooted in physiology and explains the generation of well known brain states, such as details of the up and down states for neural populations and characteristics of ageing, sleep, or anaesthesia (WP2). But virtual brains do not address brain function in service of cognitive task performance. WP3 builds such brain network models that successfully simulate a range of brain functions with the goal to address the challenge of performing ecologically valid tasks in complex environments. This includes the generation of saccadic eye movements, in hand object manipulation and spatial navigation. Their models, while biologically inspired, are not realistic or extended enough to explain multiscale activity patterns and full brain dynamics. Some of the modelling activity in WP2 and WP3 takes an intermediate position in the range of biological realism, developing largely micro and mesoscale data-driven models addressing different brain states (physiological and pathological) and cognitive tasks such as object recognition.

This deliverable presents the outcome of WP1-WP2-WP3 interactions with the objective to tackle the issues that arise when integrating task relevant brain function and task-related learning into activity-based models (see Figure 1). We did not expect to provide a fully functional virtual brain model (this activity started only 9 months ago), but we aimed to identify the barriers and questions, when attempting this task. Taking a constructive approach, we worked together to develop a first blueprint or prototype of a virtual brain model capable of some form of brain function (D3.19 at M40), relying on the emergence of low-dimensional functionally meaningful brain activity dynamics (Structured Flows on Manifolds). In this endeavour, for concrete implementations, we chose visual and auditory stimulus discrimination tasks, for which experimental structural and functional brain imaging data are available.

For the visual task, the dataset consists of whole-brain, high-resolution fMRI measurements of healthy adult subjects. While viewing thousands of colour natural scenes, subjects were engaged in a continuous recognition task in which they reported whether they had seen each given image at any point in the experiment. These data constitute a massive benchmark dataset for computational models of visual representation and cognition and can support a wide range of scientific inquiry. When translating this data set into virtual brain modelling, several challenges arise including how to represent the visual stimulus and the corresponding processes leading to activation patterns in V1 in the virtual brain. A spatial dimension of the image in pixel space needs to be maintained, which requires a high-resolution virtual brain as developed in SGA3. Other pre-processing steps like early visual processing, such as cortical magnification, need to be addressed and there is a range of options that can be considered (integration in the cortical node, phenomenological ad-hoc pre-processing, or introduction of novel pathways). The arriving stimulus in V1 then activates through the connectome other brain areas and network propagation occurs, depending on the granularity of the encoding of the stimulus in V1. Such will result in different activation patterns, which are expressed in the corresponding fMRI signals for both, in silico and empirical data sets.

For the auditory task, we implemented a simple paradigm to simulate auditory discrimination using the virtual brain. We exploit a decision-making mechanism that was previously implemented using the AdEx mean-fields (Baspinar et al., 2023), and which forms the basis of the discrimination performed by the prefrontal cortex in this model (Turan et al., 2023). This very simple simulation shows that the activity evoked by different stimuli can be routed to the prefrontal cortex decision mechanism, which can trigger (or not) a motor action. Further work will be needed to develop this preliminary model, but one can already see a great advantage of using the TVB-AdEx: the model can predict the activity of excitatory and inhibitory neurons underlying this simple cognitive task. Therefore, it opens the route towards experimental testing of these predictions, where experiments

and theoretical models can cooperate to yield a deeper understanding of the underlying brain mechanisms.

We have made certain implementation choices, which are described in the subsequent sections, leading to a first prototype capable of entertaining a critical discussion of how to link connectome-based and artificial convolution-based brain networks.

2. Technical specification for the visual task

The current prototype consists of three integrated components: high-resolution brain network model, mean-field models capturing distinct brain states (awake, sleep), and a deep neural network processing of visual input. The virtual brain modelling workflow (Schirner et al 2022) comprises a high-resolution Brain Network Model (BNM) implemented in The Virtual Brain simulator, receiving inputs to the V1 mapped from the output of a CNN processing naturalistic scene images. Figure 2 details this workflow in a diagram and comprises the two streams of information processing using the virtual brain (lower branch) and artificial neural network (middle stream). The prototype is available here: https://wiki.ebrains.eu/bin/view/Collabs/sga3-d1-9-poc-simulations-for-validity

We make use of the Natural Scenes Dataset (NSD; http://naturalscenesdataset.org), which serves as an input to both streams. This dataset features fMRI data from 8 individuals, each having observed between 9,000 and 10,000 unique coloured natural scenes, resulting in 22,000–30,000 individual trials spanning 30–40 scanning sessions. The scanning utilised a 7T environment with full-brain gradient-echo EPI at a 1.8-mm precision and a 1.6-second interval between repetitions. The natural scenes were sourced from the MS COCO database (Lin 2014). A subset of 73,000 images were used which were then modified to be square-shaped and displayed at dimensions of 8.4° x 8.4°. A specific of 1,000 were viewed by all participants, while the others were exclusive to each. These scenes were showcased for a 3-second period, with 1-second pauses. Participants maintained central focus and engaged in an extended continuous image recognition activity. To process the fMRI outputs, there was a single temporal adjustment to manage slice time disparities and a spatial adjustment to manage head movement. Subsequently, a general linear approach was deployed to discern single-trial beta values. Using FreeSurfer, cortical surface renditions were formulated, leading to both volume and surface interpretations of the beta values. For our demonstration, we used the data (resting state and functional) from subject01.

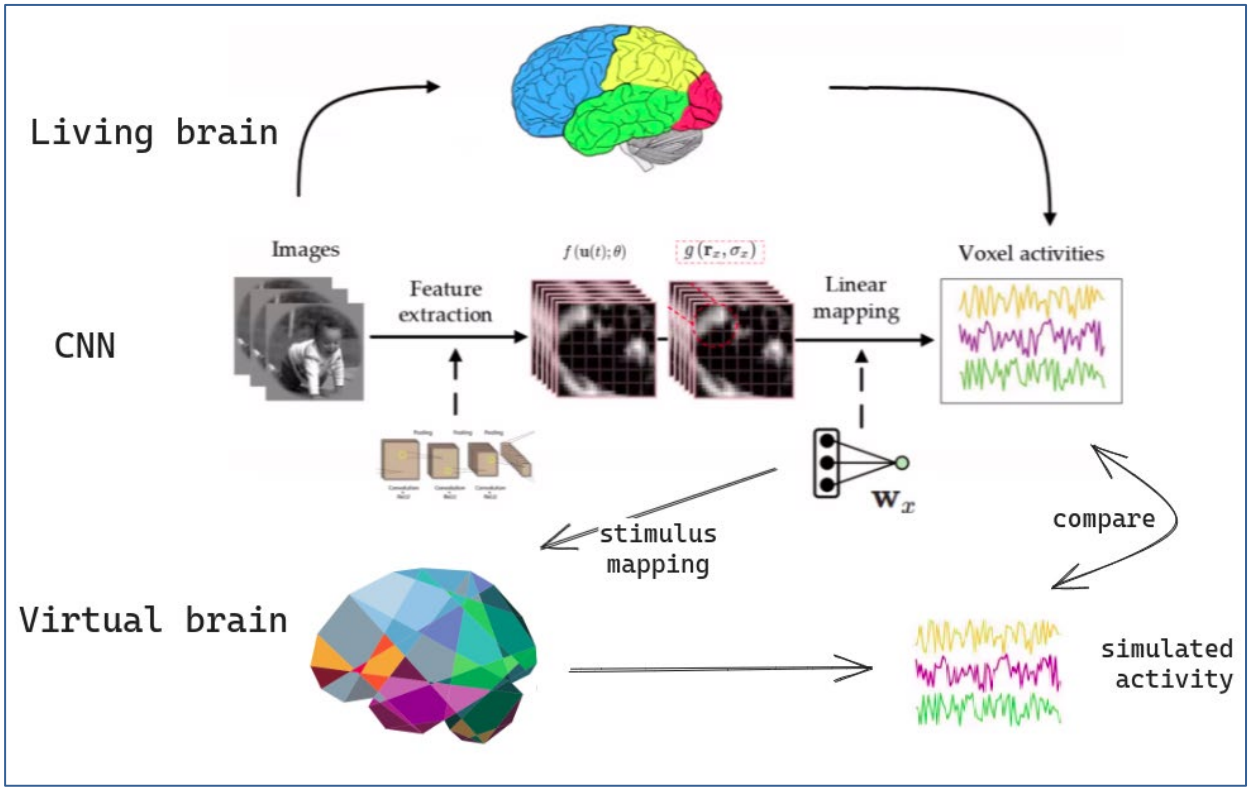

Figure 2: Workflow diagram of the demonstrator

The images from the Natural Scenes Dataset are passed through the convolutional neural network performing feature extraction, and the resulting activations are mapped against empirical fMRI recordings of the V1 area of a subject perceiving the same images. The resulting activation maps are used as a stimulus mask for the corresponding vertices of the high-resolution personalised virtual brain of the subject, generating a simulated evoked response.

2.1 Brain Network Model

The back-bone of whole brain network models is the personalised structural connectivity where the individual edges represent physical connections between neuronal populations represented as nodes (Ghosh et al 2008). The nodes are then in turn equipped with a mathematical model governing the dynamical evolution of the local state variables in time, given the network input, external stimulus and noise.

$$\dot{\Psi}(x,t) = \Lambda(\Psi(x,t)) + I(x,t) + \xi(x,t) + \int_\Gamma g_{local}(x,x\prime)S(\Psi(x\prime,t))dx\prime + \int_\Gamma g_{global}S(\Psi(x\prime,t-\frac{|x-x\prime|}{\nu}))dx\prime$$

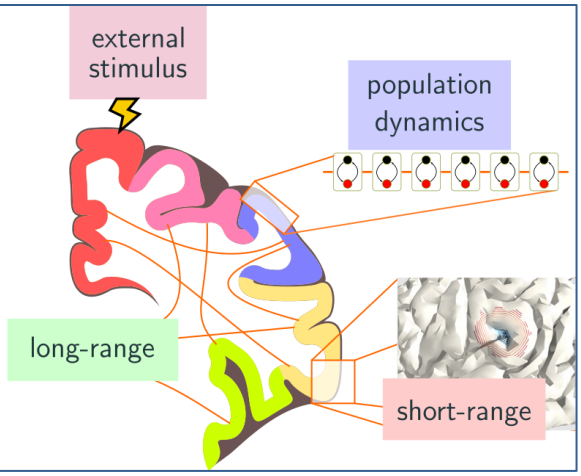

Figure 3: High-resolution brain network model

The system is defined by a network composed of local (short-range, cortico-cortical) and global (long-range, white matter) connectivity, connecting the nodes of the cortical mesh and subcortical areas. The node-level dynamics is governed by an appropriate neural mass model.

To accommodate the spatially detailed profile of the visual stimulus to the V1 area, this demonstrator employs a high-resolution BNM (Figure 3). In this case, the cortical surface is considered as a continuous domain, discretized into a triangular mesh. The mesh is a result of a pipeline adapted from Showcase 2, where the T1-weighted MRI is processed with Freesurfer package resulting in meshes with 27,099 vertices. Two connectivity matrices are then defined: long range connectome and the intra-cortical connectivity. The high-resolution connectome, g_{global}, is derived from the DWI by reconstructing the white-matter fibres and computing the intersection of the fibres with the cortical mesh, resulting in a vertex-to-vertex connectivity. The intra-cortical connectivity g_{local} is defined by a translationally invariant gaussian kernel on the geodesic distances on the cortical mesh. Both connectivity matrices are sparse yielding 2% non-zero elements for the intra-cortical connectivity, and 0.01 % for the long-range connectivity.

The nodes in the brain network mesh represent populations of neurons, whose activity is summarised by few state variables (e.g., mean firing rate, mean membrane potential...), and its evolution in time is prescribed by a system of differential equations. The demonstrator implements two mean field models to demonstrate the flexibility of the implementation and to integrate with efforts across work packages. Both mean field models show qualitatively the same behaviour (technically speaking, they have the same flow topology in state space). First, the Montbrio-Pazo-Roxin (Montbrió et al 2015) was chosen for the compactness of the model and the dynamical properties of bistability reflecting the up- and down-state and damped oscillatory dynamics in the up-state. Second model is the mean field AdEx (di Volo et al., 2019), which represents two excitatory and inhibitory sub-populations with adaptation. AdEx models are also chosen in the auditory discrimination task of this demonstrator and have been extensively employed in WP2 to simulate among others response to external stimulus and spontaneous dynamics in wake- and sleep-like states (Showcase 3, WP2).

As a baseline, both models were tuned to exhibit the spontaneous activity in the respective brain state. For the MPR model, this was achieved by setting the model parameters to a bistable regime and then scaling the coupling parameters of g_{global} and g_{local} such that fluid dynamics (defined by time-variable functional connectivity) are observed. In the case of the AdEx model, models for two brain states were prepared differing by the amount of level of adaptation of the excitatory population: low adaptation for wake-like, higher adaptation for sleep-like state. Again, scaling of coupling parameters for the two connectivity types was performed to achieve asynchronous dynamics in the wake-like state and synchronised slow-waves in the sleep-like state.

The external stimulus in the high-resolution BNM is applied with a given spatial mask defining the affected vertices, evoking local activity which then propagates through the network and interacts with the ongoing spontaneous dynamics (Spiegler 2020). In this demonstrator, we have applied a single pulse stimulus to selected vertices of V1. The weights for a particular image were determined as a distance-weighted sum of the activation values of the output layer of the CNN, which was spatially mapped to V1 mesh.

2.2 CNN scene Processing

To model the external stimulus to the V1, we process the input image (scene) through a shallow Convolutional Neural Network (CNN). While the layers of a typical CNN display characteristics of visual areas, they do not fully capture the biological properties of early visual processing, such as cortical magnification. In order to improve the biological plausibility of the CNN, the input image is preprocessed by distortion (Figure 4A) according to the ganglion cell distribution (da Costa et al., 2023). The radius of distortion is kept the same as the field of view used in NSD (8.4°). The images in MS COCO dataset do not have a single label per image. Instead, an image is characterised by the presence of multiple objects. Hence, the last layer of our CNN has 80 nodes corresponding to 80 object categories and each label is represented as a n-hot encoded binary vector. This CNN is trained on 73,000 images that were used in NSD.

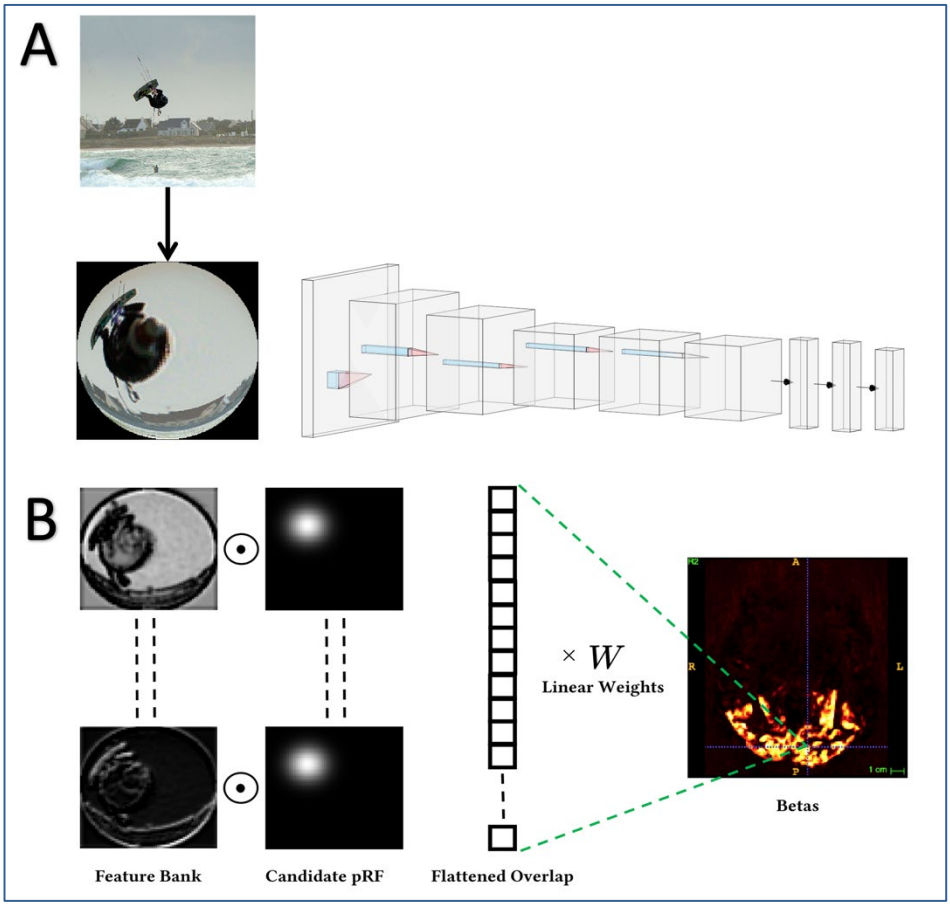

Figure 4: Processing of visual input using CNN and retinotopic mapping

(A) Each image is first distorted using a biologically plausible cortical magnification (da Costa et al, 2023) and then used to train a shallow CNN. (B) For every image, a feature bank is extracted from the CNN. For every feature, an overlap (element-wise product) is computed between that feature and an appropriate pRF candidate. These overlaps, across all the features, are then flattened out as a single vector which is used to predict the Betas of V1 voxels.

The purpose of training such a CNN is to extract meaningful features from a (distorted) image. A bank of features is obtained from the last convolutional layer. In order to perform retinotopic mapping, the overlap between the features and a population receptive field (pRF) is computed. The pRF is modelled as a 2D isotropic Gaussian and the overlap is computed as an element-wise product between a feature and the corresponding pRF. We select an appropriate pRF for each image by searching over a grid of pRF parameters. Note that the same pRF is used for all the features (per image).

The overlaps between the features and the pRFs are flattened out as a single vector. This single vector is then used to predict an activation (scalar value) in V1. We posit a straightforward linear relationship between the predictor (flattened vector) and its corresponding activation. The weights of this linear model are calculated using ridge regression. The activations are single-trial Betas, which are computed per image using GLMsingle (Prince et al., 2022). This process is described in Figure 4B.

In summary, for every image that was presented, we extract a feature bank using a CNN. Then, for every V1 voxel, we estimate an appropriate pRF and a set of linear weights that would estimate a Beta value for that image. This pipeline is implemented in order to follow a retina-LGN-V1 like processing of the visual input.

2.3 Brain network response to the visual stimulus

In the final step of the demonstrator the V1 activation maps defined by the single-trial Beta values were mapped from the MRI voxel space to the corresponding vertices of the high-resolution brain

network model, and used as a spatial mask to define the stimulus *I(x,t)*. The temporal profile of the stimulus was a single short pulse defined as a Gaussian kernel of width of 6 ms. This stimulus was applied to the excitatory firing rate of the AdEx mean-field tuned to the wake-like asynchronous regime. In response, the evoked activity propagated through the network carried both by the local connection within the cortical mesh, and the long-range connectivity. The propagation pattern is shown in Figure 5. Next step is the systematic parameter exploration of the weights for the local connectivity, global connectivity and stimulus to identify configurations where the model propagates the perturbation from the V1 without being locked in the hyper excited high-firing rate regime (a property of the AdEx mean field). This work is currently being performed using the EBRAINS HPC infrastructure.

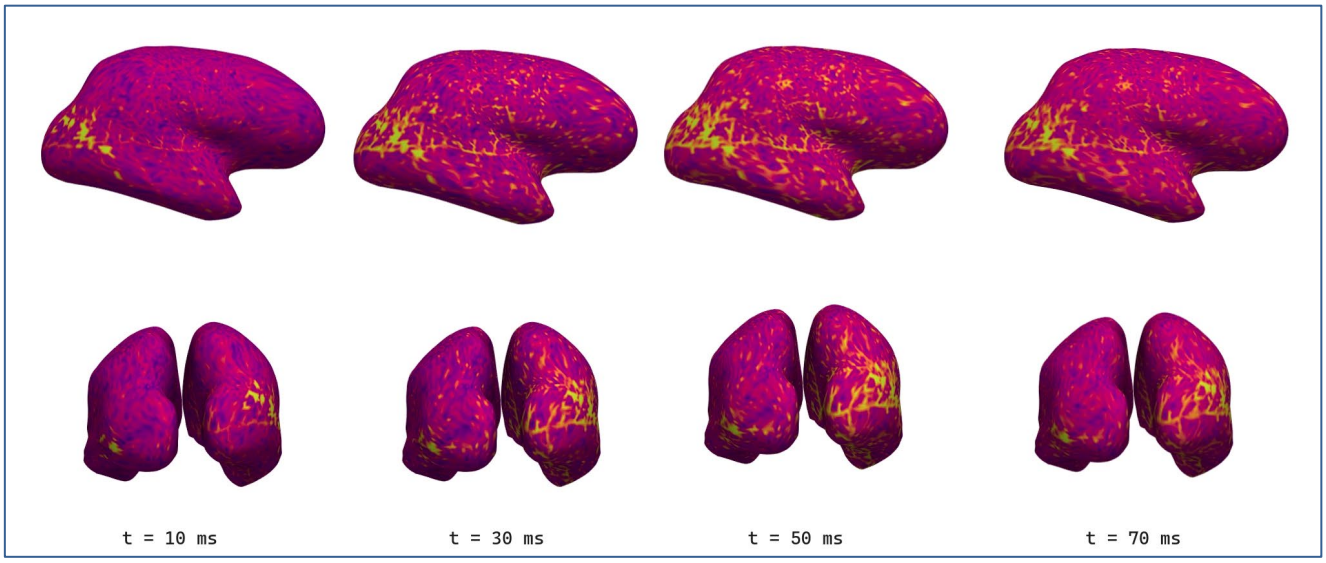

Figure 5: BNM activation map evoked by the stimulus in V1

Snapshots of the simulated activation map of the BNM evoked by the stimulus that is derived from the spatial profile of the Beta value distribution in V1 and plotted over time.

3. Technical specification for the auditory discrimination task

Our aim is to build a whole-brain model to simulate a paradigm involving multiple brain areas, for instance auditory discrimination, using a paradigm designed for the macaque cortex. We are interested in the neuronal dynamics at the whole-brain scale and perform the simulations by employing The Virtual Brain (TVB) simulation environment. Adaptive Exponential (AdEx) neuronal population models (di Volo et al., 2009) are used to describe each node's dynamics. For the connectivity, we use open-access CoCoMac connectivity dataset (Bakker et al., 2012), which is a matrix containing the connection weights between the nodes. We focus on a cognitive task that mainly involves the prefrontal cortex (PFC). In the auditory discrimination task, our pipeline starts from the primary auditory cortex stimulated by the auditory signals, it is then modulated in the PFC so that the stimulus discrimination occurs. Finally, it ends in the primary motor cortex which outputs the neuronal activity determining the motor action. The aim of this study is to demonstrate the use of whole-brain models to investigate simple cognitive paradigms, such as auditory discrimination in the macaque brain.

It is important to note that a detailed report is available (Turan et al., 2023), where all information such as the model structure, are given in detail. This detailed report describes two different tasks of auditory discrimination, of which we describe here only the second task.

This auditory discrimination task is described in Figure 6. The participant is provided simultaneously with two different auditory stimuli of different pitch and is instructed to push the button to indicate that it perceived the stimuli as different. Otherwise, the participant does not push the button,

indicating that it perceived the stimuli as the same (Figure 6A). The corresponding routing of information in the TVB model is shown in Figure 6B. The two stimuli (S1, S2) activate different areas of primary auditory cortex (A1, Ra and Rb), which then provide input to prefrontal cortex (PFC), where a decision is made, which will activate (or not) the motor cortex and thus the motor action (pushing the button). The TVB simulation does not include the subcortical sensory input, nor the motor command.

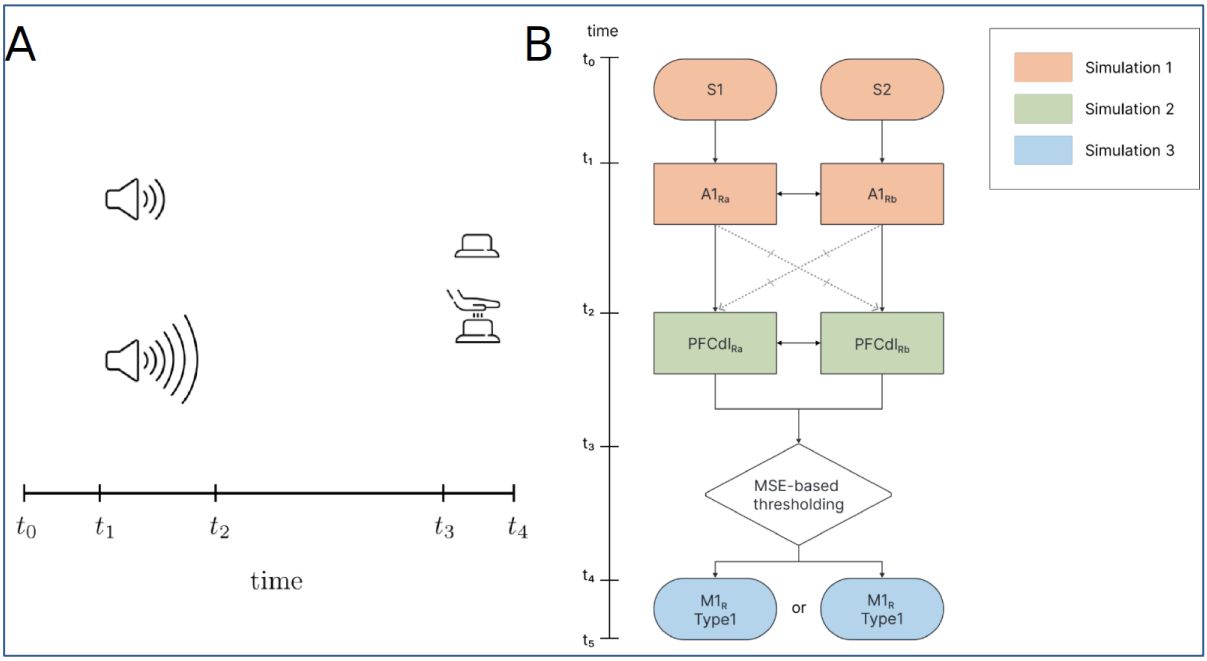

Figure 6: Scheme of the auditory discrimination task

(A) Temporal organisation of the task. (B) Information routing from auditory cortex to motor cortex, with a decision-making process occurring in the prefrontal cortex.

Importantly, the key of the simulation is the decision mechanism implemented in the prefrontal cortex. To do this, we implemented the decision-making model conceived in WP2, which was implemented in the AdEx mean-field model, with two competing AdEx cortical columns (Baspinar et al, 2023). This model was implemented in the TVB-AdEx framework (Goldman et al, 2020), which was conceived in the framework of Showcase 3 and with the help of WP1. Since both the TVB-AdEx and the decision-making model use the AdEx mean-fields, they are compatible and this compatibility was essential to the present integration.

The behaviour of the model is depicted in Figure 7. Two stimuli of different pitch were activating A1 (Figure 7A). The stimulus values were 0.001 and 0.02 here, and the output firing rates of A1 regions are indicated. Depending on the output of each A1 region, a strong or weak stimulus is routed to the corresponding PFC region (PFC firing rates are shown in Figure 7B). Here, the two competing PFC columns, PFCdl-Ra and PFCdl-Rb, receiving input from the two stimuli, maintain a different firing, showing that the model perceives the stimuli as different. Those outputs are compared to each other via the MSE thresholding. If MSE value is high, it triggers a strong increase of activity in M1, meaning that the decision is to push the button, otherwise a weak stimulus is provided to M1. The output of M1 is shown in Figure 7C.

In the case the two stimuli are the same, or very close, there is little difference between the different brain regions, and the decision-making mechanism classifies the stimuli as “not different”. In this case, no motor output is generated (not shown; see Turan et al., 2023).

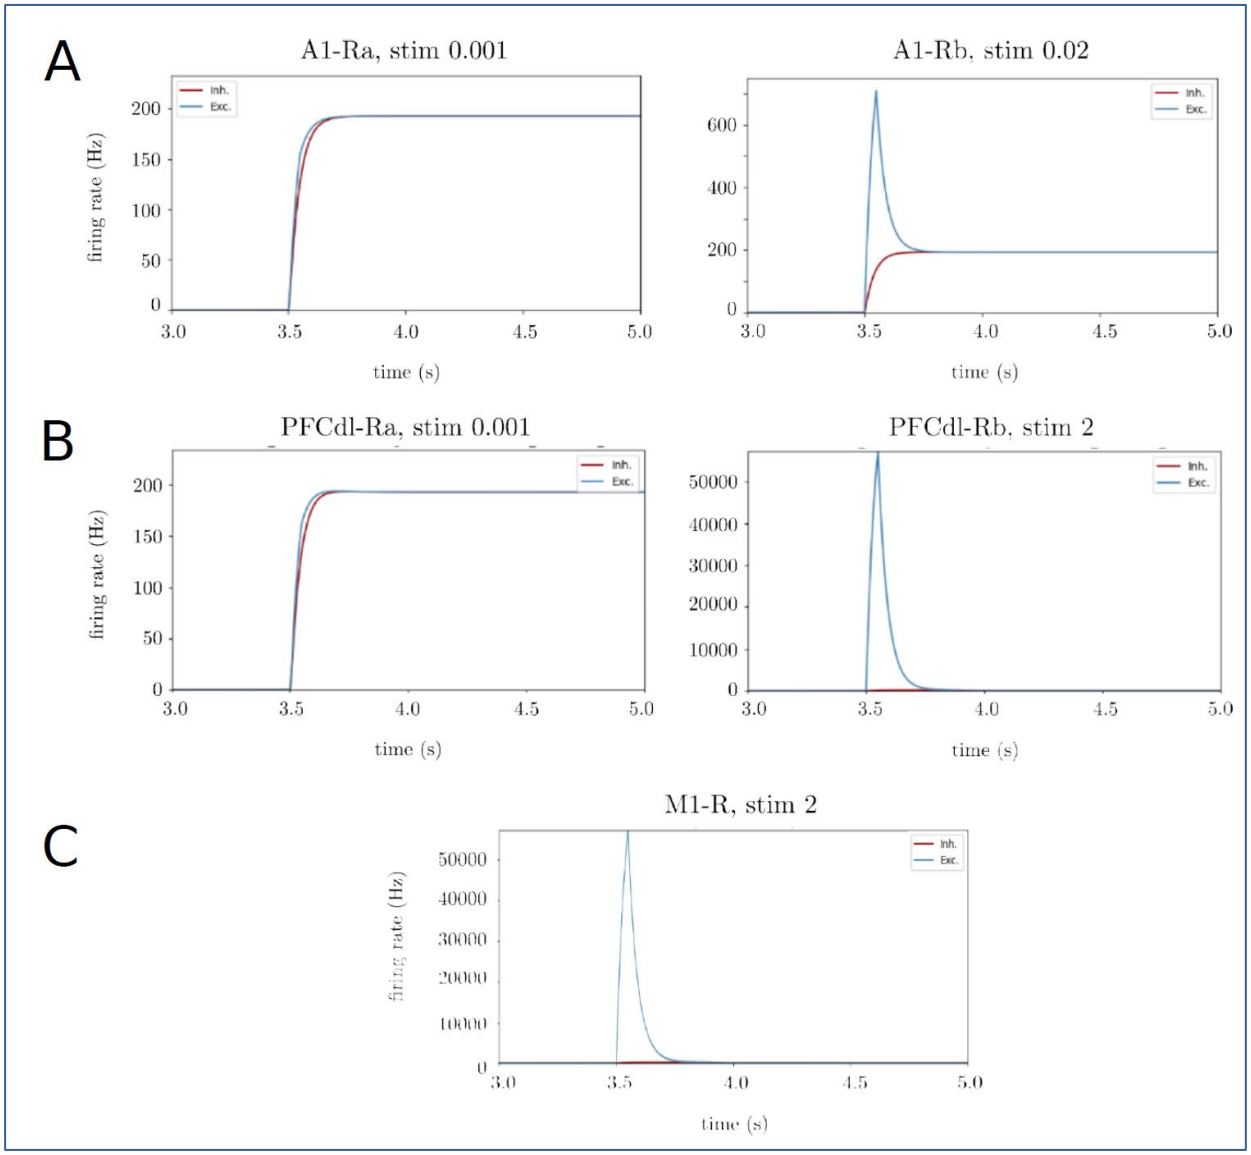

Figure 7: Auditory discrimination task simulated using TVB

(A) Activity in the auditory cortex. (B) Activity in prefrontal cortex. (C) Motor output.

In conclusion, we have shown here the TVB implementation of a simple paradigm to simulate auditory discrimination. We exploited a decision-making mechanism that was previously implemented using the AdEx mean-fields (Baspinar et al., 2023), and which form the basis of the discrimination performed by the prefrontal cortex in this model (Turan et al., 2023). This very simple simulation shows that the activity evoked by different stimuli can be routed to the prefrontal cortex decision mechanism, which can trigger (or not) a motor action. Further work will be needed to develop this preliminary model, but one can already see a great advantage of using the TVB-AdEx: the model can predict the activity of excitatory and inhibitory neurons underlying this simple cognitive task. Therefore, it opens the route towards experimental testing of these predictions, where experiments and theoretical models can cooperate to yield a deeper understanding of the underlying brain mechanisms.

4. Looking forward

Our prototypes described in this deliverable set the ground for discussion and developments. Our objective was to propose a first set of integrated virtual brain models that can be evaluated in cognitive tasks against empirical data. Although the current implementations clearly have many shortcomings, they nevertheless provide functional prototypes with all relevant ingredients including

external naturalistic stimulation, task context, neural representation, connectome based brain activity, physiological brain states, and realistic source-to-sensor mapping. As such, they represent first stepping stones towards a cognitive virtual brain and make progress in advancing our understanding of the connection between brain states and cognition - two routes that often advance in parallel and that we have connected for the first time. The use of the AdEx mean field model offers a systematic avenue for investigating the brain states' dynamics (di Volo et al., 2019). By integrating this model with connectome-based data, we can obtain the emergence of intricate spatiotemporal patterns. Notably, this model has demonstrated promising outcomes when simulating two distinct brain states: wakefulness and NREM (Non-Rapid Eye Movement) sleep (Goldman et al., 2022). During wakefulness, spontaneous brain dynamics exhibit irregular temporal patterns. Upon external stimulation, in a specific cortical region, complex spatiotemporal patterns emerge. In this state, information is processed across various spatial scales. For example, during a visual task, following the activation of early visual areas, activity propagates and amplifies in a top-down manner, engaging different cortical regions and facilitating cognitive processing. Conversely, during NREM sleep (also known as slow wave sleep) or deep anaesthesia, the brain's dynamics self-organise into slow oscillations (Sanchez-Vives et al., 2017). This state entails alternating periods of neuronal activity known as up-states and periods of relative silence referred to as down-states. These slow oscillations manifest as highly synchronised events, propagating as travelling waves predominantly in an anteroposterior direction (Massimini et al., 2004). In this state, information processing is disrupted. Perturbations in the system lead to stereotypical activation of cortical areas followed by periods of silence, disrupting propagation and resulting in a less complex spatiotemporal regime, where cognition is impaired (Casali et al., 2013). It's important to note that wakefulness and NREM sleep represent two extremes within a high-dimensional space (Harris & Thiele, 2011). In other words, various other brain states can manifest, contingent upon behavioural and experimental conditions. The significance of these states extends beyond unravelling the mechanisms of the healthy brain; it also pertains to understanding the pathological brain, particularly in cases of consciousness disorders where the brain may engage in slow oscillations (Rosanova et al., 2018). The AdEx mean field model integrated into a connectome-based dataset opens the possibility to address these questions. All its properties were maintained in the two discrimination tasks described here, and in fact exploited for the emergence of the network's spatiotemporal dynamics. The model not only enables us to replicate sleep and awake dynamics, but it also opens the possibility to link brain dynamics to intrinsic neuronal properties, reproducing the regional variability observed in the human brain (Hawrylycz et al., 2012). For instance, the neurotransmitter acetylcholine is known to be involved in the awake-sleep transitions, as well cognitive processes (Hasselmo, 2006). With the recent studies of receptors map (Hawrylycz et al., 2012; Hansen et al., 2022), it becomes feasible to bind the model with receptors densities observed in the human brain. This not only makes the model more biologically plausible, but it also improves the similarity between model and data, as shown in (Deco et al. 2021).

Considering these insights, a brain model that integrates cognitive tasks with spatiotemporal patterns opens opportunities to explore the relationship between cognition and brain states. This approach allows us to investigate under what conditions cognitive processes are impaired and their corresponding brain states correlates.

5. References

P4209: A.Turan, E. Baspinar and A. Destexhe. A whole-brain model of auditory discrimination. Biorxiv (2023) https://doi.org/10.1101/2023.09.23.559095

P3995: da Costa, D., Kornemann, L., Goebel, R., & Senden, M. (2023). Unlocking the Secrets of the Primate Visual Cortex: A CNN-Based Approach Traces the Origins of Major Organizational Principles to Retinal Sampling. bioRxiv, 2023-04.

P3132: Deco, G., Kringelbach, M. L., Arnatkeviciute, A., Oldham, S., Sabaroedin, K., Rogasch, N. C., ... & Fornito, A. (2021). Dynamical consequences of regional heterogeneity in the brain's transcriptional landscape. Science Advances, 7(29).

P3903: E. Baspinar, G. Cecchini, M. DePass, M. Andujar, P. Pani, S. Ferraina, R. Moreno-Bote, I. Cos, and A. Destexhe, A biologically plausible decision-making model based on interacting cortical columns, bioRxiv, (2023).

P3023: Goldman, J. S., Kusch, L., Aquilue, D., Yalçınkaya, B. H., Depannemaecker, D., Ancourt, K., Nghiem, T.-A. E., Jirsa, V., & Destexhe, A. (2022). A comprehensive neural simulation of slow-wave sleep and highly responsive wakefulness dynamics. Frontiers in Computational Neuroscience, 16, 1058957.

P2681: J. S. Goldman, L. Kusch, B. H. Yalcinkaya, D. Depannemaecker, T.-A. E. Nghiem, V. Jirsa, and A. Destexhe, Brain-scale emergence of slow-wave synchrony and highly responsive asynchronous states based on biologically realistic population models simulated in the virtual brain, bioRxiv, (2020).

P3023: J. S. Goldman, L. Kusch, D. Aquilue, B. H. Yalc ̧ınkaya, D. Depannemaecker, K. Ancourt, T.-A. E. Nghiem, V. Jirsa, and A. Destexhe, A comprehensive neural simulation of slow-wave sleep and highly responsive wakefulness dynamics, Frontiers in Computational Neuroscience, 16 (2023).

P1864: M. Di Volo, A. Romagnoni, C. Capone, and A. Destexhe, Biologically realistic mean-field models of conductance-based networks of spiking neurons with adaptation, Neural Computation, 31 (2019), pp. 653–680.

P1516: Rosanova, M., Fecchio, M., Casarotto, S., Sarasso, S., Casali, A. G., Pigorini, A., ... & Massimini, M. (2018). Sleep-like cortical OFF-periods disrupt causality and complexity in the brain of unresponsive wakefulness syndrome patients. Nature communications, 9(1), 4427.

P1186: Sanchez-Vives, M. V., Massimini, M., & Mattia, M. (2017). Shaping the default activity pattern of the cortical network. Neuron, 94(5), 993-1001.

P2973: Schirner, M., Domide, L., Perdikis, D., Triebkorn, P., Stefanovski, L., Pai, R., Prodan, P., Valean, B., Palmer, J., Langford, C., Blickensdörfer, A., van der Vlag, M., Diaz-Pier, S., Peyser, A., Klijn, W., Pleiter, D., Nahm, A., Schmid, O., Woodman, M., … Ritter, P. (2022). Brain simulation as a cloud service: The Virtual Brain on EBRAINS. NeuroImage, 251, 118973.

P2593: Spiegler, A., Abadchi, J. K., Mohajerani, M., & Jirsa, V. K. (2020). In silico exploration of mouse brain dynamics by focal stimulation reflects the organization of functional networks and sensory processing. Network Neuroscience (Cambridge, Mass.), 4(3), 807–851.

P1864: Volo, M. di, Romagnoni, A., Capone, C., & Destexhe, A. (2019). Biologically Realistic Mean-Field Models of Conductance-Based Networks of Spiking Neurons with Adaptation. Neural Computation, 31(4), 653–680.

Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K. R., ... & Massimini, M. (2013). A theoretically based index of consciousness independent of sensory processing and behavior. Science translational medicine. DOI: 10.1126/scitranslmed.3006294

Ghosh, A., Rho, Y., McIntosh, A. R., Kötter, R., & Jirsa, V. K. (2008). Noise during rest enables the exploration of the brain's dynamic repertoire. PLoS Computational Biology, 4(10), e1000196. https://doi.org/10.1371/journal.pcbi.1000196

Hansen, J.Y., Shafiei, G., Markello, R.D. et al. Mapping neurotransmitter systems to the structural and functional organization of the human neocortex. Nat Neurosci 25, 1569–1581 (2022). https://doi.org/10.1038/s41593-022-01186-3

Harris KD, Thiele A. Cortical state and attention. Nat Rev Neurosci. 2011 Aug 10;12(9):509-23. doi: 10.1038/nrn3084. PMC3324821

Hasselmo ME. The role of acetylcholine in learning and memory. Curr Opin Neurobiol. 2006;16(6):710-715. doi:10.1016/j.conb.2006.09.002 PMC2659740

Hawrylycz MJ, Lein ES, Guillozet-Bongaarts AL, et al. An anatomically comprehensive atlas of the adult human brain transcriptome. Nature. 2012;489(7416):391-399. doi:10.1038/nature11405 PMC4243026

Massimini, M., Huber, R., Ferrarelli, F., Hill, S., & Tononi, G. (2004). The sleep slow oscillation as a traveling wave. Journal of Neuroscience, 24(31), 6862-6870. https://doi.org/10.1523/JNEUROSCI.1318-04.2004

Montbrió, E., Pazó, D., & Roxin, A. (2015). Macroscopic Description for Networks of Spiking Neurons. Physical Review X, 5(2), 021028. https://doi.org/10.1103/PhysRevX.5.021028

Prince, J.S., Charest, I., Kurzawski, J.W., Pyles, J.A., Tarr, M., Kay, K.N. Improving the accuracy of single-trial fMRI response estimates using GLMsingle. eLife (2022). https://doi.org/10.7554/eLife.77599