Modelling Spatiotemporal Variability of Brain Responses in Functional Ultrasound

Erol, A.

DOI
10.4233/uuid:75f86071-0133-4491-9033-397e71c29e8b

Publication date
2024

Document Version
Final published version

Citation (APA)
Erol, A. (2024). *Modelling Spatiotemporal Variability of Brain Responses in Functional Ultrasound*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:75f86071-0133-4491-9033-397e71c29e8b

Modelling Spatiotemporal Variability of Brain Responses in Functional Ultrasound

MODELLING SPATIOTEMPORAL VARIABILITY OF BRAIN RESPONSES IN FUNCTIONAL ULTRASOUND

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op
dinsdag 3 december 2024 om 10:00 uur

door

Aybüke EROL

Master of Science in Electrical and Electronics Engineering
geboren te Ankara, Turkije.

Dit proefschrift is goedgekeurd door de promotoren.

Samenstelling promotiecommissie bestaat uit:

Rector Magnificus,	voorzitter
Prof. dr. ir. A. J. van der Veen	Technische Universiteit Delft, promotor
Dr. B. Hunyadi	Technische Universiteit Delft, copromotor
Dr. ir. P. Kruizinga	Erasmus Medisch Centrum, copromotor
Onafhankelijke leden:	
Prof. dr. ir. G.J.T. Leus,	Technische Universiteit Delft
Prof. dr. M. Carandini,	University College London, UK
Prof. dr. T. Adali,	University of Maryland, Baltimore County, USA
Dr. D. Maresca,	Technische Universiteit Delft
Prof. dr. ir. W.A. Serdijn,	Technische Universiteit Delft, reservelid

Keywords: Functional Ultrasound, Hemodynamic Response Modelling, Spatiotemporal Variability

Front & Back: Cover art capturing the unintended parts of the work.

ISBN 000-00-0000-000-0

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

CONTENTS

1

INTRODUCTION

"The biology of mind bridges the sciences - concerned with the natural world - and the humanities - concerned with the meaning of human experience."

— Eric Kandel

Part of this chapter is published as: A. Erol and B. Hunyadi (2022). Tensors for neuroimaging: A review on applications of tensors to unravel the mysteries of the brain. Tensors for Data Processing, 427-482.

1

UNderstanding how the brain works has long posed a challenge for humankind. Neurosurgery is believed to be the oldest medical specialty, with its first signs dating back to ancient times [1]. Our fascination with the brain stems from both clinical and cognitive interests, where the former focuses on developing diagnostic and treatment tools for brain disorders and the latter is concerned with the neuronal basis of mental processes such as perception, action, language, attention, and memory. As our quest to unravel the brain's mysteries continues, we have reached many breakthroughs in the diagnosis and treatment of neurological or psychosomatic diseases. These breakthroughs include early detection of Alzheimer's disease even before clinical symptoms appear [2], deep stimulation of brain areas involved in motor control to improve the mobility of individuals with Parkinson's disease [3], and brain-computer interfaces to facilitate the communication of paralysed patients with the outside world [4].

The ongoing advances in neuroscience are achieved thanks to the joint progresses in neuroimaging and signal processing techniques. While neuroimaging enables the visualization of brain anatomy and real-time monitoring of brain activity, signal processing delineates the behavior of brain signals, such as how they vary while performing certain tasks or under pathological conditions. On the other hand, learning more about brain's anatomy and functioning has also inspired numerous engineering innovations. For instance, deep neural networks mimic the human nervous system for artificial intelligence tasks, and are used in various fields including medicine, nanotechnology, telecommunications, autonomous vehicles, art and finance [5].

To improve our understanding of the brain, we are in a continuous search for better neuroimaging modalities and finer models for processing the acquired data. Each neuroimaging technique comes with its own set of limitations: while some offer high temporal but low spatial resolution, others provide the opposite. Additional challenges include invasiveness, lack of portability and high costs, all of which make it challenging to image the brain with the accuracy and convenience that we desire. There are still many open questions regarding the optimal acquisition of neuroimaging data and development of models and algorithms that can reliably interpret and analyze this data. In this thesis, we focus on a relatively recent yet promising neuroimaging technique called functional ultrasound, and propose new ways to process the data. We start with a brief overview of existing neuroimaging techniques.

1.1. BRIEF HISTORY OF NEUROIMAGING

Neuroimaging techniques are employed to image the structure and function of the nervous system. Structural imaging deals with the analysis of anatomical properties of the brain and is useful for diagnosing intracranial lesions such as tumors. The journey towards obtaining a complete image of the brain began with X-rays in the early 20th century. For many years, computed tomography (CT) was the primary method for structural imaging of the brain, which utilizes X-rays taken from multiple angles to create detailed cross-sectional images [6]. However, since its introduction in 1978, magnetic resonance imaging (MRI) has largely replaced CT. MRI uses powerful magnets instead of ionizing radiation and offers greater contrast between normal and abnormal brain tissue [7].

On the other hand, functional imaging is used to identify brain areas and processes

that are associated with performing a particular cognitive or behavioral task. Information flow in the brain while processing a task is controlled by the firing of neurons via both electrical and chemical signals [8]. This electrical activity can be detected directly from isolated neurons in a localized manner, or imaged from larger neuronal populations in the form of aggregated post-synaptic potentials via electroencephalography (EEG) and magnetoencephalography (MEG). EEG is the oldest functional brain imaging technique, with the first reported human EEG recording taking place in 1929. EEG detects the electrical activity of neurons via electrodes placed along or below (intracranial EEG or electrocorticography, ECoG) the scalp [9], while MEG records the magnetic field produced by this electrical activity using magnetometers [10]. Neuronal activity can also be monitored via optical imaging using voltage-sensitive dyes or fluorescent calcium indicator dyes for detecting the calcium influx induced by neuronal firing [11].

Fluctuations in neuronal activity result in associated changes in blood dynamics. Activation of a brain region leads to heightened consumption of glucose and oxygen in that area. These metabolic demands prompt an augmented blood flow to the region. Neurovascular coupling (NVC) describes this interaction between changes in local neuronal activity and cerebral blood flow (CBF) and volume (CBV) [12]. NVC forms the basis of many functional neuroimaging techniques including positron emission tomography (PET), functional near-infrared spectroscopy (fNIRS), functional magnetic resonance imaging (fMRI) and functional ultrasound (fUS) (Fig. 1.1).

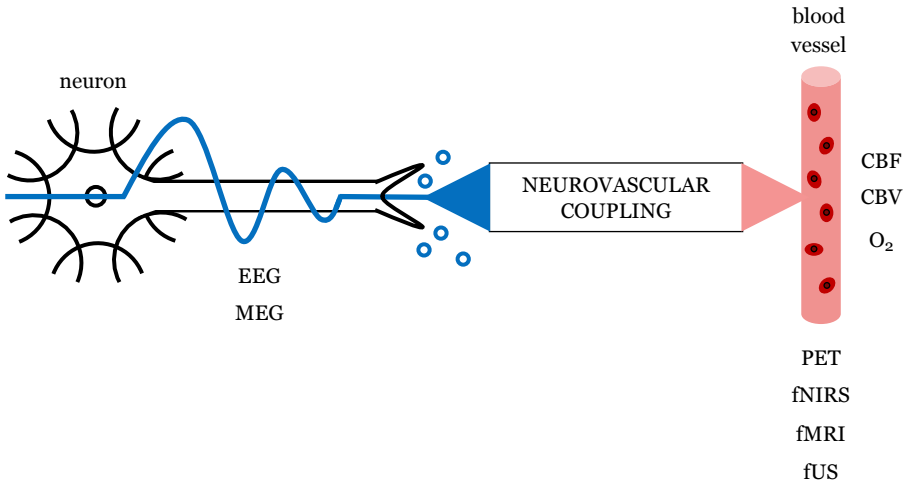

Figure 1.1: Functional neuroimaging modalities, in relation to neuronal and hemodynamic activity. Neurovascular coupling describes the relationship between neuronal activity and resulting changes in the blood flow. EEG and MEG directly measure neuronal activity whereas PET, fNIRS, fMRI and fUS provide an indirect measure through neurovascular coupling.

PET measures the alterations in glucose level in response to metabolic activity by injection of radioactive tracers to the brain which are attached to glucose and absorbed by the bloodstream [13]. Meanwhile, the changes in oxygenation of hemoglobin in

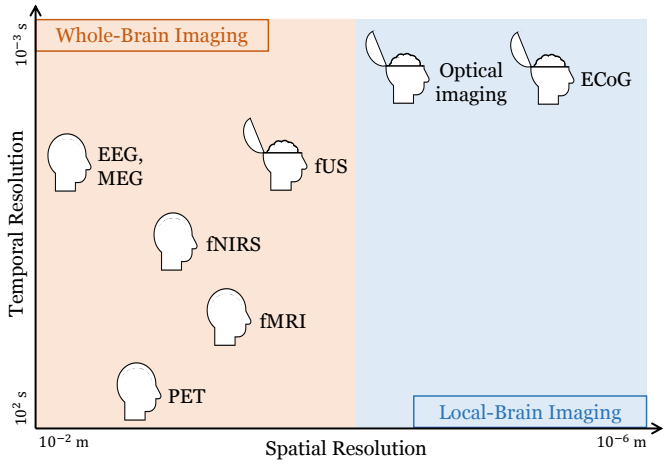

Figure 1.2: Functional neuroimaging modalities with their temporal and spatial resolutions, where an opened skull refers to intracranial imaging. Image adapted from [16].

red blood cells can be detected by fNIRS and fMRI. In fNIRS, near-infrared light is used to track hemodynamic changes based on the differences in optical properties of hemoglobin states [14]. The magnetic properties of hemoglobin are affected as well by the amount of oxygen that the cells carry, resulting in the blood-oxygen-level-dependent (BOLD) signal detected by fMRI via electromagnets. Since the early 1990s, fMRI has come to dominate brain mapping research due to its non-invasive nature (requiring no injections or surgery) and high spatial resolution [15].

1.2. FUNCTIONAL ULTRASOUND

Ultrasound imaging is one of the most used imaging techniques in medicine due to its low-cost and portability, as well as the harmless nature of ultrasonic waves. Although medical ultrasound imaging traces its roots back to over half a century ago, it was almost a decade ago that it entered the field of neuroimaging, thanks to the advances in ultrafast ultrasound imaging. These advances significantly enhanced the technique's sensitivity to subtle variations in blood volume, ultimately establishing ultrasound as a new modality for functional neuroimaging [17]. Offering an unparalleled spatiotemporal resolution for whole-brain imaging (Fig. 1.2), fUS has been used in both cognitive and clinical studies since its development. Although at present fUS imaging is commonly performed in an intracranial setting to avoid attenuation of ultrasonic waves by the skull, ongoing improvements in its technology steer toward transcranial applications [18].

1.2.1. PRINCIPLES OF FUNCTIONAL ULTRASOUND

For fUS imaging, ultrasonic plane waves are transmitted to the brain through a cranial window, which are then backscattered by moving red blood cells (RBCs) or tissue. In

order to enhance the image quality and resolution, plane waves are steered at different angles and the reconstructed echo images are coherently summed, resulting in one compound image. The compound images obtained over time encompass information from two components: blood flow and tissue motion [19]. For separating the blood signal of interest from the slower-moving tissue component, initially high-pass temporal filtering was employed. However, it is not always straightforward to pinpoint the exact frequency that divides the two subspaces, since there can be overlaps in their spectra, as shown in Fig. 1.3. Recently, high-pass filtering has been replaced by singular value decomposition (SVD), which defines the subspaces based on their spatiotemporal coherence and yields better separation than high-pass filtering [20].

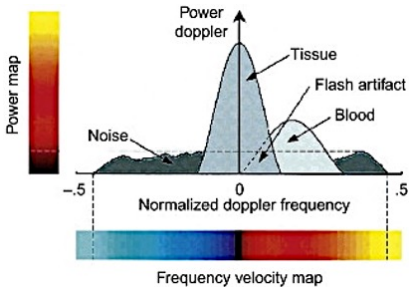

Figure 1.3: Doppler spectrum of blood flow and tissue signals. Flash artifact refers to the incorrect mapping of moving blood onto tissue regions. Image taken from [21].

The amplitude of each pixel (x, y) of the cleaned echo images, $p(x, y, t)$, varies under the Doppler effect due to the movement of RBCs. The spectrum of $p(x, y, t)$ reveals a received central frequency f_R, given by $f_R = f_T + f_D$, where f_T is the transmitted frequency. The Doppler shift f_D is given by $\frac{-2v}{c} f_T \cos\theta$, where v is the velocity of the red-blood scatterer, c is the speed of sound and θ is the angle between the ultrasound beam and the scatterer [22].

While RBC velocity and direction can be inferred from the spectrum of $p(x, y, t)$ using the Doppler shift, the integral of this spectrum (i.e., power-Doppler) has been observed to be a more suitable metric for imaging microvascular dynamics [19]. Indeed, power-Doppler is directly proportional to the number of moving RBCs, hence to the changes in CBV [17]. The last step of fUS imaging is the computation of power-Doppler images (PDIs) as the mean voxel intensities along ensembles of compound images. In mathematical terms, the power-Doppler intensity at (x, y) over an ensemble defined by the window $t_i, i = 1, \ldots, N$ is given by [23]:

$$I(x, y) = \frac{1}{N} \sum_{i=1}^{N} |p(x, y, t_i)|^2. \tag{1.1}$$

Choosing a larger ensemble increases the sensitivity of the estimation of $I(x, y)$, hence the ensemble size sets a trade-off between the signal-to-noise ratio (SNR) and temporal resolution of the PDIs. Compound images can be acquired every millisecond, whereas PDIs are commonly calculated from hundreds of compound images. PDI acquisition steps are illustrated in Fig. 1.4.

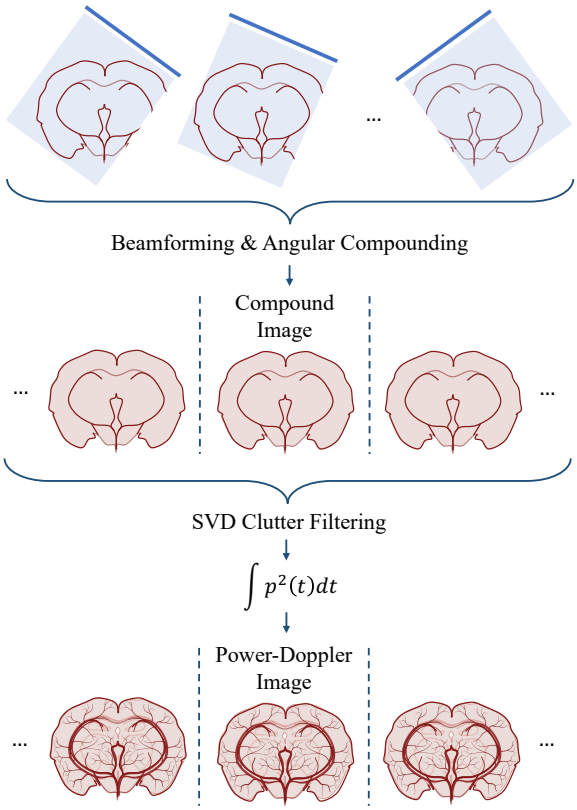

Figure 1.4: PDI acquisition pipeline. First, the brain is illuminated with planar ultrasonic waves at multiple angles. The echo images of different angles are beamformed and coherently compounded into one compound image. Compound images are subjected to SVD clutter rejection in ensembles. Then, a PDI is obtained from each ensemble by computing the per-voxel power.

Although initially developed for acquisition of a single brain slice, various 3D-volumetric extensions of fUS imaging have been employed since then. One such extension is concatenation of 2D fUS images of different brain slices obtained by repeating the same experiment at each slice [24]. Alternatively, a 2D array transducer can be used for ultrasound transmission [25], which requires expensive hardware and acquired volumes suffer from lower sensitivity. Recently, another solution known as swept-3D fUS has been proposed, which we use in this thesis. In swept-3D fUS, the ultrasound probe is moved continuously back-and-forth over the craniotomy during the experiment. It is worth emphasizing that due to the movement of the probe, swept-3D fUS acquires slices at different time instants.

1.3. SCOPE AND CONTEXT OF THIS THESIS

Neuroscience research has revealed that the brain operates through intricate networks of interconnected regions that exhibit synchronized activity during specific tasks or resting state. Understanding such processes of the brain unravels the mysteries of cognition and behavior, and sheds light on discrepancies seen in neurological disorders. Offering an impressive spatiotemporal resolution, fUS has the potential to help us uncover new insights regarding functional brain regions and their temporal dynamics.

In this thesis, we focus on stimulus-evoked fUS response of the mouse brain. While delving deeper into the spatial extent and temporal characteristics of the fUS response, we propose novel signal processing approaches that can as well be adapted to other hemodynamics-based neuroimaging modalities.

Functional ultrasound provides an indirect measure of neural activity through NVC by detecting fluctuations in local blood dynamics. The most common approach for addressing NVC in the literature is the linear time-invariant (LTI) model, expressing the voxel time-series as the convolution between neuronal activity and the impulse response of the underlying neurovascular system, also known as the hemodynamic response function (HRF) (Fig. 1.5), which varies across brain regions and subjects. When neuronal activity is unknown, the stimulus time course is used as a proxy for the input signal.

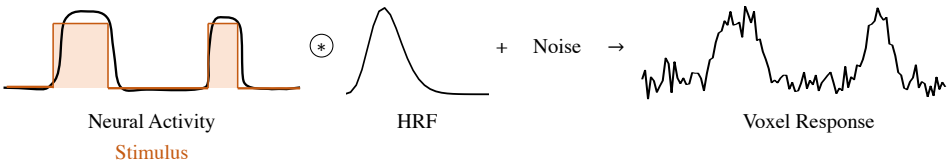

Figure 1.5: Diagram of the LTI model. When neural activity is unavailable, it is replaced with the stimulus time course.

By convolving a fixed HRF with the binarized stimulus signal, we can construct a modeled time course that stands for the expected response. The modeled time course is later used to calculate the voxel-wise Pearson correlation coefficients, or in a linear regression framework to estimate spatial activation maps of the brain, highlighting the brain areas that respond to the stimulus. While such fixed characterizations remain popular due to their good performance and simplicity, experimental evidence suggests that the brain exhibits non-LTI characteristics. For instance, brain signals may violate the superposition principle under certain conditions [26], or may exhibit varying response magnitudes and delays to the same stimulus upon repeated presentations; while the amount of this variability can also differ across brain regions [27]. Given the dynamic nature of the measured responses, we need more complex models than pre-determined input signals (stimuli) and impulse responses (HRFs).

RESEARCH OBJECTIVES

Our main objective in this thesis is to achieve a flexible characterization of the fUS signal while accounting for the spatiotemporal variability of the brain. We break down this

objective into three parts, and start by examining the non-LTI characteristics of brain response as measured by fUS:

Q1. *How well does the LTI model fit to the measured fUS response? Are there certain conditions (i.e., stimulus or brain region-dependent) under which the LTI model falls short in explaining the fUS response?*

In order to answer Q1, we can use the conventional approach of time-locking the fUS response to the stimulus onsets. First, we estimate HRFs using both linear and nonlinear convolution under different stimulus durations to determine if or when linear convolution remains plausible for characterizing the fUS signal. Subsequently, for each repetition of a stimulus of fixed duration, we investigate the changes in the activation strength of responsive regions to identify any potential dynamic traits in their responses. Together, these two investigations navigate us in identifying the weak aspects of existing techniques, designing our experiments, and proposing new methods for a more accurate depiction of brain activity.

For our next research question, we aim at expanding the definition of the input signal beyond binarized stimuli in order to capture the nuanced patterns of brain activity, even for task-based experiments. On the other hand, due to the variability of HRF across events, brain regions and subjects, there is no universally applicable HRF shape either [28]. Therefore, we need methods that can estimate both the HRFs and the input "source" signals. We use the term source signal to describe the actual (unknown) inputs of the hemodynamic system, which can absorb the dynamic nature of brain responses to stimuli. Our second research question is as follows:

Q2. *Can we jointly estimate source signals and HRFs?*

Although a joint estimation of the HRFs and sources would indeed provide great flexibility for revealing the region-dependent variations of the HRF and time-varying aspects of brain activity respectively, certain assumptions still need to be made to obtain physiologically meaningful results. Commonly used assumptions include parametrization of the HRF shapes (to resemble forms that were consistently observed in prior studies) and sharing of sources across the brain, in accordance with the functionally connected organization of the brain. Even so, both for handling the large number of unknowns and maintaining a reasonable subset of responses for which commonality of sources can be assumed, prior selection of regions of interest is essential.

Finally, we seek for a semi-flexible solution, where prior information of the stimulus can be included in the problem to solely *facilitate* the analysis, leading us to our last research question:

Q3. *Can we use the known stimulus time course to identify the evoked activity within the whole brain without absolutely relying on it as the input of the LTI system generating the fUS response?*

The answer to Q3 suggests finding a balance between the two approaches mentioned above, such that we neither directly equate the input signal of the LTI model to the stimulus time course, nor completely disregard the stimulus information and estimate ev-

erything from scratch. This way, we can construct a framework for an informed decomposition of the evoked activity in the whole brain while incorporating spatiotemporal variability.

1.4. OUTLINE

This thesis as organized as follows. To start with, we provide an exploratory analysis of fUS data and highlight the drawbacks of the common LTI model with a fixed input, i.e. the stimulus time course. Next, we propose two approaches to handle the spatiotemporal variability of brain responses. In the first approach, we perform a semi-blind deconvolution of selected regions of interest by only assuming a parametric model for the HRFs and no available prior information regarding the input sources, except that they are uncorrelated. For the second approach, we propose a framework for whole-brain decomposition by using the stimulus information as a guiding factor. A brief summary for each chapter is provided below, while a flowchart of the outline is given in Fig. 1.6.

Chapter 2: This chapter provides an overview of prevalent analysis techniques for functional neuroimaging data. These approaches can be broadly categorized into hypothesis-driven methods, where the HRF is modelled and stimuli serve as inputs to the hemodynamic system, and data-driven methods, which operate without relying on prior information. Finally, we refer to the existing fUS analysis methods and corresponding findings from the literature.

Chapter 3: We start our fUS data analysis by investigating Q1, i.e., the conditions for which the fUS signal can be well-approximated under the assumption of an LTI system. To that end, we train first-order and second-order HRF kernels to predict output responses and define a plausible range of stimulus durations for which the standard first-order kernels produce a satisfactory prediction. Subsequently, we demonstrate the variability of the fUS signal in spatial and temporal domains using linear regression of activation coefficients, estimated as the degree of fit to a voxel-specific HRF-based model.

Chapter 4: To incorporate the spatiotemporal variability of the brain in our model, we express the fUS signal as the output of a multiple-input multiple-output system, that is driven by unknown input sources which are shared across the selected regions of interest (ROIs). This enables us to estimate both the shared sources and region-specific HRFs, as described in Q2, by applying block-term decomposition on the tensor formed with lagged output correlations, assuming uncorrelated sources.

Chapter 5: This chapter addresses Q3 as we propose a new approach for decomposing the evoked activity of the brain by exploiting the stimulus information as a guiding factor in a regularized optimization problem. By adjusting the regularization strength, we control the influence of prior information in the estimated time courses and corresponding spatial activation maps of stimulus-induced brain response. Furthermore, we apply the proposed method on a new 3D fUS imaging technique known

as swept-3D fUS to uncover the evoked activity at the full resolution, referring to the different slice acquisition time instants during the probe's motion.

Chapter 6: In this chapter, we discuss the findings of previous chapters in the light of how multimodal fusion of fUS signals and neuronal processes can enhance our understanding. Accordingly, we state possible research directions while outlining a number of open questions regarding the Power-Doppler data acquisition pipeline.

Chapter 7: Lastly, we provide conclusions driven from the findings and scientific contributions presented in this dissertation.

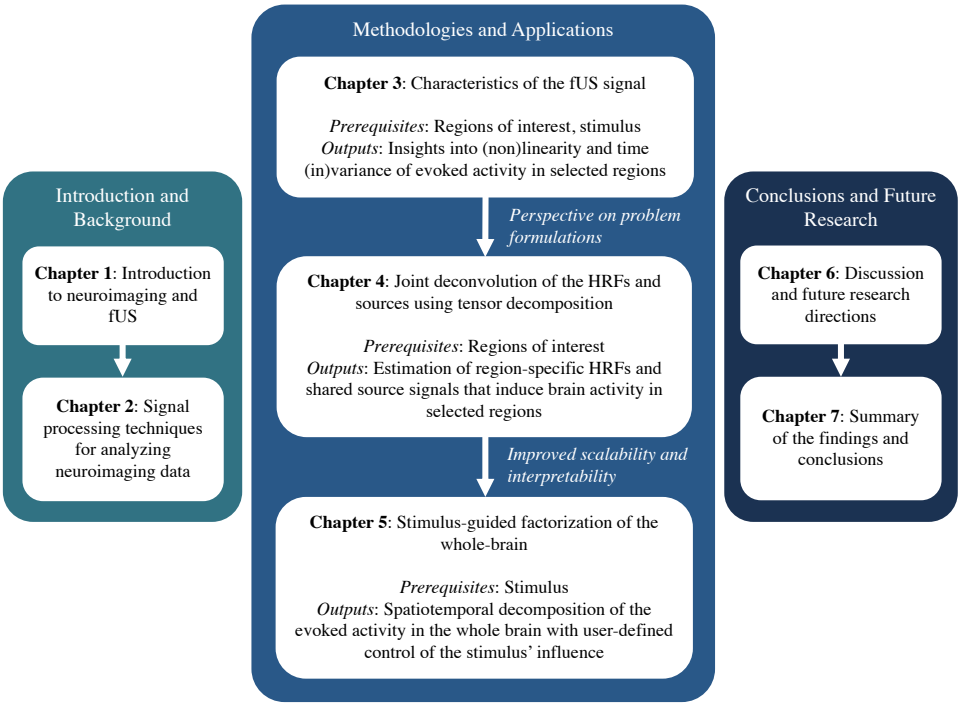

Figure 1.6: Thesis outline. After the introductory chapters, we investigate the characteristics of the fUS signal. Based on these insights, we propose a convolutive mixture model for the multivariate fUS response while assuming unknown sources and HRFs. This approach requires a prior selection of regions of interest that can be assumed to be driven by same sources. We then propose a technique that exploits the known stimulus time course up to a user-defined extent, facilitating the interpretability and scalability of evoked activity analysis while maintaining spatiotemporal flexibility. At last, we summarize our findings and suggest several future research directions.

1.5. LIST OF CONTRIBUTIONS

Finally, we provide an overview of the contributions made to the scientific literature during the work on this Ph.D. thesis.

Journal Papers

J1 **A. Erol**, C. Soloukey, B. Generowicz, N. Van Dorp, S. Koekkoek, P. Kruizinga, and B. Hunyadi (2023). Deconvolution of the Functional Ultrasound Response in the Mouse Visual Pathway Using Block-Term Decomposition. Neuroinformatics, 21(2), 247-265.

J2 **A. Erol**, B. Generowicz, P. Kruizinga and B. Hunyadi (2024). Evoked Component Analysis (ECA): Decomposing the Functional Ultrasound Signal with GLM-Regularization. IEEE Transactions on Biomedical Engineering. doi: 10.1109/TBME.2024.3395154.

Book Chapters

B1 **A. Erol** and B. Hunyadi (2022). Tensors for neuroimaging: A review on applications of tensors to unravel the mysteries of the brain. Tensors for Data Processing, 427-482.

Conference Papers

C1 **A. Erol**, S. Van Eyndhoven, S. Koekkoek, P. Kruizinga, and B. Hunyadi (2020, November). Joint estimation of hemodynamic response and stimulus function in functional ultrasound using convolutive mixtures. In 2020 54th Asilomar Conference on Signals, Systems, and Computers (pp. 246-250). IEEE.

C2 **A. Erol**, B. Generowicz, P. Kruizinga and B. Hunyadi (2023, June). GLM-Regularized Low-Rank Factorization For Extracting Functional Response From Swept-3D Functional Ultrasound. In 2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW) (pp. 1-5). IEEE.

C3 S. E. Kotti, **A. Erol** and B. Hunyadi (2023, June). Modeling Nonlinear Evoked Hemodynamic Responses in Functional Ultrasound. In 2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW) (pp. 1-5). IEEE.

C4 **A. Erol**, P. Kruizinga and B. Hunyadi (2024). Analyzing Trial-to-Trial Variability in the Mouse Visual Pathway using Functional Ultrasound. In 2024 IEEE 21st International Symposium on Biomedical Imaging (ISBI). IEEE.

2

PRINCIPLES OF FUNCTIONAL NEUROIMAGING DATA ANALYSIS

"No research without action, no action without research."

— Kurt Lewin

THE brain is the most intricately structured organ within the body. To gain a deeper understanding of how the brain operates, there is an ever-increasing demand for neuroimaging techniques with higher sensitivity (detecting subtle changes in blood flow or electrical activity, etc.), better contrast (such as differentiating between different types of tissue or activity based on their physical or biochemical properties) and larger coverage of the brain. Following the advancements in imaging technologies, the scale and complexity of neuroimaging data continue to expand. In order to sift through this vast amount of data and extract informative subsets containing spatial and temporal patterns of interest, various approaches have been explored in the literature. In this thesis, we focus on those that are most commonly applied on hemodynamics-based neuroimaging techniques, including functional ultrasound (fUS).

2.1. PHYSIOLOGY OF THE BRAIN

The brain is part of the central nervous system, orchestrating countless physiological and cognitive functions. The average adult human brain comprises approximately 86 billion neurons which communicate through networks of synaptic connections. This communication is enabled by electrical impulses known as action potentials, initiating the release of neurotransmitters upon reaching the synaptic terminals. This release activates ion channels on the postsynaptic neuron, facilitating the transmission of signals across the synapse [29].

Neurons depend on oxygen and glucose for proper functioning. As such, when a brain region becomes active, it requires more oxygen and nutrients. The hemodynamic response refers to the physiological process in the brain that controls the rapid delivery of blood to the active area, ensuring that the active neurons receive the resources needed to sustain their activity. Neurovascular coupling (NVC) defines this relationship between neuronal activity and the associated changes in blood flow. NVC plays a crucial role in regulating the hemodynamic responses for maintaining the supply-demand balance within the brain, and forms the basis for many neuroimaging techniques including fUS, which maps brain activity by detecting changes in cerebral blood volume.

The brain's structure is divided into several regions, each specializing in different functions. For instance, the cerebrum, the largest part, is responsible for cognitive skills and sensory interpretation, while the cerebellum focuses on coordination. On the other hand, the cortex is involved in higher processes such as memory, reasoning, problem-solving, emotions, and consciousness (Fig. 2.1(a)). In this thesis, we particularly focus on the visual processing pathway of the mouse brain, which consists of several key regions analogous to those found in the human brain: the Superior Colliculus (SC), Lateral Geniculate Nucleus (LGN), and Visual Cortex (VIS).

The vast majority of information about visual stimuli are conveyed via the retinal ganglion cells (RGCs) to the downstream, subcortical (i.e., beneath the cortex) targets LGN and SC, before being relayed to VIS. LGN and SC receive both similar and distinct visual stimulus input information from RGCs. SC is a laminated midbrain structure, meaning it consists of layers with distinct functions. It serves as a vital hub for integrating sensory information from various modalities, including vision. Receiving projections from 85-90% of RGCs [30], SC neurons function as "feature detectors", meaning

that subsets of SC neurons will respond differently depending on the stimulus type that was presented within its receptive field. SC is therefore thought to be useful for the swift detection of visual features that indicate potential threats (such as flashing, moving or looming spots [31, 32]).

LGN is a laminated thalamic centre that acts as a relay station for visual information traveling from the retina to the visual cortex. In both humans and mice, LGN is responsible for segregating and relaying visual signals to different layers of the visual cortex, allowing for the processing of visual stimuli with distinct spatial and temporal characteristics. The Visual Cortex (VIS), located in the occipital lobe, is where the majority of visual processing takes place. For both species, VIS is responsible for encoding visual features such as color, motion, and form, and integrating this information to construct a coherent visual representation of the environment. The mouse visual cortex has been reported to include up to 16 interconnected, retinotopically organized areas [33]. Different areas of the visual cortex are preferentially activated by different stimulus features [34, 35].

Overall, SC, LGN and VIS carry pivotal functions in the mouse brain's visual processing system, which parallel their counterparts in the human brain as well. A summary of the connections between these regions is depicted in Fig. 2.1(b).

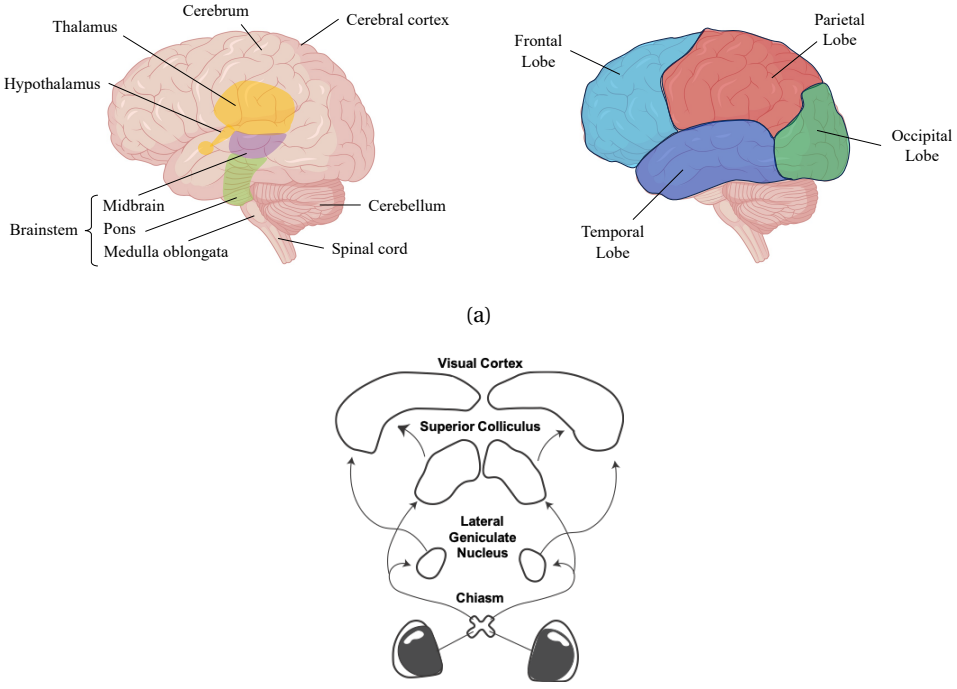

Figure 2.1: Schematic diagram of the anatomy of the human brain (a) and information flow between the eye and the brain, highlighting SC, LGN and VIS (b).

2.2. FUNCTIONAL NEUROIMAGING DATA ANALYSIS

This section investigates functional brain analysis in terms of its objectives and commonly used techniques, divided into two parts as data-driven and hypothesis-driven methods.

2.2.1. OBJECTIVES

Understanding the complex dynamics of brain function is a central goal in neuroscience research. In this section, we delve into fundamental aspects of functional brain analysis: functional connectivity, resting-state activity and task-evoked activity.

Functional Connectivity

Functional connectivity (FC) explores the statistical interdependence between the activity time courses of different anatomical regions. These interdependencies may give insight into the intrinsic organization of the brain activity and how spatially remote brain regions co-operate. FC may be studied in resting state or during task execution, by discovering networks that are commonly represented with nodes corresponding to specific brain regions and edges that describe their pairwise associations. Recently, there has been growing interest in dynamic networks, spurred by evidence of changes in statistical interdependence across brain areas over time [36]. FC patterns exhibit notable differences between patients with neurological disorders and healthy controls, such as reduced global synchrony in schizophrenic patients [37].

Resting-State Activity

Resting-state activity refers to the spontaneous neural activity that occurs in the absence of explicit tasks or stimuli. Studying resting-state activity allows researchers to explore the brain's intrinsic functional organization, as well as define a baseline and control state, which is inherently different from task-induced states. Indeed, studies have shown that a specific set of brain regions are more active during rest or mind-wandering than during task-oriented activities, known as the default mode network [38]. It is possible to examine the resting brain from the perspective of functional connectivity, or by detecting the intensity of low-frequency oscillations, which are shown to arise from spontaneous neural activity in the brain areas within the default mode network [39]. Resting-state networks are as well affected by neurological diseases. For instance, [40] shows that the connectivity within the default mode network is significantly weaker for stroke patients compared to healthy controls.

Task-Evoked Activity

Task-evoked activity corresponds to the brain's response to specific stimuli or tasks. Through various experimental paradigms, researchers can investigate how different brain regions are involved in a variety of neural mechanisms including sensory processing, motor control, cognitive functions and emotional responses. Similar to resting-state networks, task-evoked activity differs in patients with neurological disorders compared to healthy controls, often manifesting as reduced responses in specific brain areas [41].

With prior knowledge of activity-inducing stimuli, it is possible to formulate

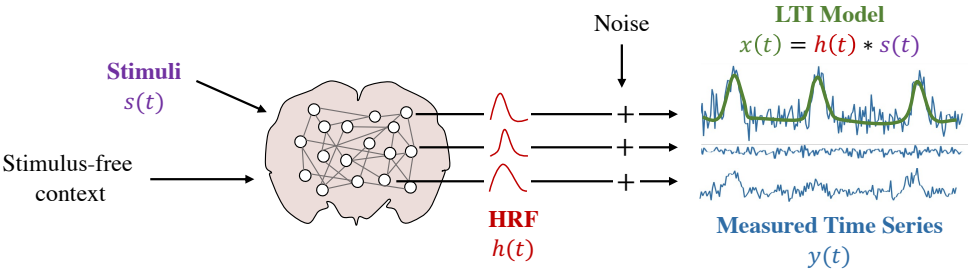

Figure 2.2: Schematic of the underlying signals comprising the fUS data. Brain activity can be induced by external stimuli, yet, it can also be driven by the brain's own intrinsic processes — referred as stimulus-free context in the figure. The resulting neuronal activity gets convolved with an HRF, which can vary from one region to another. Last but not least, the measured time series will be affected by noise, including both physiological and instrumentation noise. The part that is of interest in this thesis is the task-relevant content of the fUS signal, which is commonly modelled as the convolution of stimuli and HRF. Notice that while some regions may not be responsive to the stimuli, even the ones that respond might differ in the way they do, exhibiting different degrees of variability to repeated stimuli.

hypotheses regarding expected responses. In what follows, we will explain hypothesis-driven and data-driven methods for analyzing task-evoked activity of the brain.

2.2.2. HYPOTHESIS-DRIVEN METHODS

In hypothesis-driven methods, a model is specified based on the expected hemodynamic responses to experimental conditions. As mentioned in the previous chapter (Fig. 1.5), the modelled time course, representing the expected response of a voxel i, is typically expressed as:

$$x_i(t) = h_i(t) * s(t), \tag{2.1}$$

where $x_i(t)$ is the modelled time course, $h_i(t)$ is the voxel's hemodynamic response function (HRF), $*$ denotes convolution and $s(t)$ is the stimulus time course, given by a boxcar function indicating when the stimulus is on or off. The estimation of HRF is a topic of interest on its own, and existing models for the HRF will be discussed thoroughly in Section 2.2.2. Note that the voxel index i is dropped in the subsequent equations for notational simplicity. A schematic diagram of this model is provided in Fig. 2.2. In this diagram, brain activity that is not dependent on external stimuli is also depicted (i.e., stimulus-free context), although the conventional LTI model only takes interest in the stimulus-evoked content.

Subsequently, the temporal synchrony between the modelled time course $x(t)$ and measured voxel time series $y(t)$ is examined for each voxel. This way, an activation coefficient is obtained per voxel, with a higher value suggesting a stronger link between the voxel's response and the stimulus. In fUS, the voxel time series corresponds to its power variation across the Power-Doppler Image (PDI) stream.

A commonly used measure to quantify the match between $x(t)$ and $y(t)$ is by computing the Pearson Correlation Coefficient (PCC) r_{xy} between them:

$$r_{xy} = \frac{\sum_{t=1}^{N}[x(t)-\bar{x}][y(t)-\bar{y}]}{\sqrt{\sum_{t=1}^{N}[x(t)-\bar{x}]^2}\sqrt{\sum_{t=1}^{N}[y(t)-\bar{y}]^2}}, \tag{2.2}$$

where N is the total number of time samples during the experiment and $\bar{(\cdot)}$ denotes the sample mean of a time series [42]. In this equation, the expected value was replaced by the sample mean under the assumption of ergodicity; meaning that time series data sampled over a sufficiently long period can represent the true statistical properties of the underlying signal [43]. In the literature, it is also common practice to replace $x(t)$ with a lagged version of the stimulus such that $x(t) = s(t-\tau)$, in which case the value of the PCC defined in Eq. 2.2 will be a function $r_{sy}(\tau)$ of the time lag, τ. Due to the fact that brain regions might respond to the stimulus at different delays, the introduced lag is then swept in a range of plausible values to find the optimal value of PCC [44]. For the HRF-based definition of the modelled time course, such effects and more can be integrated via the HRF [45] by adjusting its parameters such as its peak latency or width. Both the HRF or the optimal value of the time lag incorporated directly on the stimulus can vary depending on the voxel.

The most widespread hypothesis-driven analysis tool used in hemodynamics-based neuroimaging modalities is the general linear model (GLM) [46, 47]. In GLM, modelled time courses are used as regressors (known as design variables) on the measured time series for a linear prediction of activation coefficients. The activation coefficient estimated for each voxel and regressor specifies the contribution of the regressor on the time series response of that particular voxel. Mathematically, GLM approximates a voxel time series $\mathbf{y} \in \mathbb{R}^{N\times 1}$ as:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \tag{2.3}$$

where $\mathbf{X} \in \mathbb{R}^{N\times R}$ is the design matrix consisting of R regressors, $\boldsymbol{\beta} \in \mathbb{R}^{Rx1}$ is the vector of activation coefficients and $\boldsymbol{\epsilon} \in \mathbb{R}^{N\times 1}$ represents noise. There can be multiple regressors of interest, calculated via Eq. 2.1, referring to stimulus onsets of different conditions $s(t)$ or different HRFs $h(t)$. While the former lets activations be estimated per experimental condition, the latter allows for a more precise characterization of hemodynamic latency throughout the brain [48]. Nuisance factors can as well be added as regressors to the design matrix, such as motion artifacts or low-frequency drifts. A schematic of GLM is provided in Fig. 2.3. When solved for all voxels, GLM provides spatial activation maps for each regressor.

Besides the linear approximations, there has also been an interest towards characterizing the nonlinear aspects of brain signals. Notably, the Balloon model has gained significant attention in functional magnetic resonance imaging (fMRI) research. The Balloon model delineates each physiological process as a state variable, dynamically influenced by others, ultimately contributing to the blood-oxygen-level-dependent (BOLD) signals detected in fMRI. This nonlinear relationship between the stimulus signal and output response is shown to be well captured by Volterra series. Volterra series have proven useful particularly when modelling the impact of prior input on closely-spaced

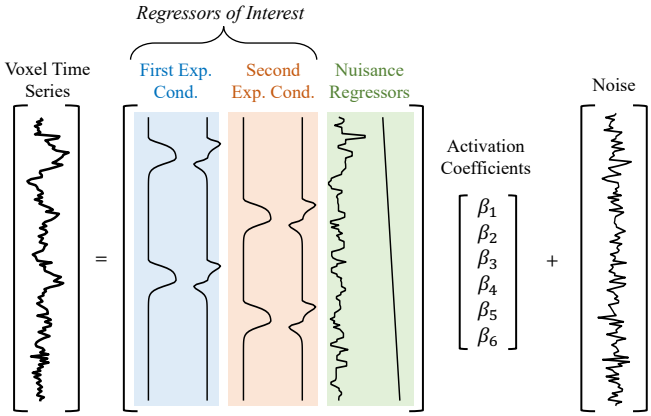

Figure 2.3: Example schematic of GLM, where the design matrix contains 4 regressors of interest (2 experimental conditions, each convolved with 2 HRFs), one regressor showing motion and one regressor for the linear drift. The HRFs used in this example are picked as the canonical HRF and its first-order derivative. The activation coefficients determine the weighting factor of each regressor on the voxel time series.

stimuli [49]. The success of Volterra series in such cases arises from its ability to precisely capture memory effects beyond what linear convolution can achieve, as it allows for modelling of higher-order interactions and nonlinear dependencies among different combinations of past and present input samples.

Notice that all of the aforementioned methods presume that the input signal is known and equal to the stimulus, leaving only one unknown in the convolution: the HRF [50]. However, considering that the connection between stimuli and brain response may feature nonlinear and time-varying traits, the reliability of the binarized stimulus representation comes into question. Moreover, not all stimuli that elicit a brain response can be pre-determined. For example, mental imagery can induce responses that are as intense as those triggered by real perception [51]. As such, there have also been attempts to predict the sources by assuming a known and fixed HRF [52] [53]. These methods fall short of depicting the HRF variability.

Lastly, there are several methods that follow the LTI model, but assume unknown input(s) and HRF(s). For example, [54], applies homomorphic filtering based on the assumption that the HRF is a predominantly low frequency signal while the unknown input (neuronal response) is a comparatively higher frequency signal. On the other hand, [55] first estimates a spike-like source signal by thresholding the fMRI data and selecting the time points where the activation begins, and subsequently fits a GLM using the estimated source signal to find the HRF. Both of these methods are univariate, whereas [56] proposes a multivariate approach to accomplish deconvolution of the whole-brain. This approach expands the LTI signal model to represent neuronal activation as a low-rank matrix - constructed by a certain (low) number of temporal activation patterns and corresponding spatial maps encoding functional networks - linked with the observed fMRI signals via region-specific HRFs. These regions characterized by unique HRFs are deter-

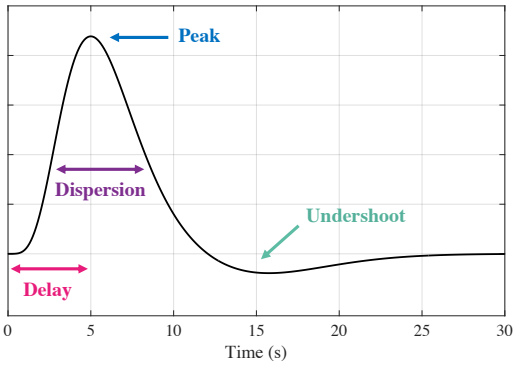

Figure 2.4: Canonical HRF.

mined *a priori* by anatomical parcellation. However, accurately parcellating the brain poses its own challenges, given the concordance problem associated with brain atlases and inter-subject variations [57].

HRF MODELS

To formulate the hypothesis for the aforementioned methods, HRF should be specified. In the literature, canonical HRF stands out as the prevailing mathematical representation to relate hemodynamic response to neural events or stimulus. Following many empirical observations, the canonical model approximates the HRF as the difference of two gamma distributions [58, 59]:

$$h(t;\boldsymbol{\theta}) = \theta_1 \frac{t^{\theta_2-1}\theta_3^{\theta_2}\mathrm{e}^{-\theta_3 t}}{\Gamma(\theta_2)} - \theta_4 \frac{t^{\theta_5-1}\theta_6^{\theta_5}\mathrm{e}^{-\theta_6 t}}{\Gamma(\theta_5)}. \tag{2.4}$$

In this model, θ_1, θ_2 and θ_3 control the amplitude, delay and dispersion of the HRF's peak response, whereas θ_4, θ_5 and θ_6 control those in the same respective order for the undershoot response. $\Gamma(\cdot)$ denotes the gamma function and it serves as a normalization parameter. While the canonical HRF, provided in Fig. 2.4, uses a fixed set of parameters $\boldsymbol{\theta} = [\theta_1,\theta_2,\theta_3,\theta_4,\theta_5,\theta_6]^{\mathrm{T}}$ in Eq. 2.4, HRF is known to vary across brain region and subjects. In fact, HRF has shown potential as a biomarker for healthy aging [60] or pathological brain functioning; examples of which include obsessive-compulsive disorder [61], mild traumatic brain injury [62], Alzheimer's disease [63], epilepsy [64] and severe psychosocial stress [65]. As such, HRF estimation has become one of the most predominant research topics amongst imaging modalities that depend on neurovascular coupling.

One approach for HRF estimation is to use the model in Eq. 2.4 (or its modified versions) and estimate the optimal set of parameters $\boldsymbol{\theta}$ that minimizes the difference between voxel time series and the HRF-modelled time course [66, 67].

Another option is to define a basis set of HRFs:

$$h(t) = \sum_{i=1}^{B} \beta_i g_i(t), \tag{2.5}$$

where $g_i(t)$ is the ith basis function. Notice that the β_i weights here are the same as the activation coefficients of the corresponding task-relevant design variables in GLM. For simplicity, let's focus on the GLM formulation for one experimental condition:

$$y(t) = h(t) * s(t) + \epsilon(t) = \sum_{i=1}^{B} \beta_i g_i(t) * s(t) + \epsilon(t) = \sum_{i=1}^{B} \beta_i x_i(t) + \epsilon(t), \tag{2.6}$$

where $x_i(t) = g_i(t) * s(t)$ is the design variable constructed with the ith HRF basis $g_i(t)$. Canonical HRF and its first and second-order derivatives have been a popular choice for the basis set in the HRF literature. Nevertheless, it is possible to generate different sets with more basis functions using parametric models, such as using Eq. 2.4 or half-cosines [68].

2.2.3. DATA-DRIVEN METHODS

Using a data-driven approach is unavoidable when no prior knowledge is available on the expected activation time courses, such as in case of resting-state experiments [69]. However, even if prior information is available to form a hypothesis, one might choose not to do so in order to flexibly capture (i.e., without forcing any modelling constraints) the nuanced patterns within the brain's evoked response. Data-driven methods utilized in functional neuroimaging data analysis range from matrix and tensor factorizations to clustering methods, graph-theoretical approaches and deep learning techniques [70].

The goal of matrix decompositions is to express the input data matrix $\mathbf{Y} \in \mathbb{R}^{N \times M}$, where N is the number of time samples and M is the number of voxels, as the product of factor matrices $\mathbf{U}$ and $\mathbf{V}$:

$$\mathbf{Y} \approx \mathbf{U}\mathbf{V}^{\mathrm{T}}. \tag{2.7}$$

Columns of $\mathbf{U} \in \mathbb{R}^{N \times K}$ and $\mathbf{V} \in \mathbb{R}^{M \times K}$ (called signatures) store the estimated time courses and spatial activation maps, respectively. The column size of $\mathbf{U}$ and $\mathbf{V}$, denoted by K, corresponds to the rank of the factor matrices, i.e. number of components used to approximate $\mathbf{Y}$. Note that Eq. 2.7 is equivalent to the following:

$$\mathbf{Y} \approx \sum_{k=1}^{K} \mathbf{u}_k \circ \mathbf{v}_k, \tag{2.8}$$

where $\mathbf{u}_k$ and $\mathbf{v}_k$ represent the temporal and spatial signatures placed in the kth column of $\mathbf{U}$ and $\mathbf{V}$ respectively, and $\circ$ denotes the outer product.

Singular value decomposition (SVD) achieves an exact factorization of $\mathbf{Y}$ such that $\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathrm{T}}$ by extracting orthogonal factor matrices ($\mathbf{U}$ and $\mathbf{V}$) and a diagonal matrix $\mathbf{\Sigma}$ of singular values, where each singular value determines the significance of the corresponding component over the data [71]. SVD is closely related to the concept of Principal Component Analysis (PCA). PCA aims to linearly transform the input data onto a new orthogonal coordinate system that captures the largest variances within the data. Usually, only the first few components where most of the variation occurs are considered to be of interest. As such, PCA is used for dimensionality reduction, feature extraction, denoising, or as an exploratory tool for revealing underlying spatial and temporal patterns within the data [72]. PCA is often computed by eigendecomposition of the data covariance matrix or SVD of the data matrix [73].

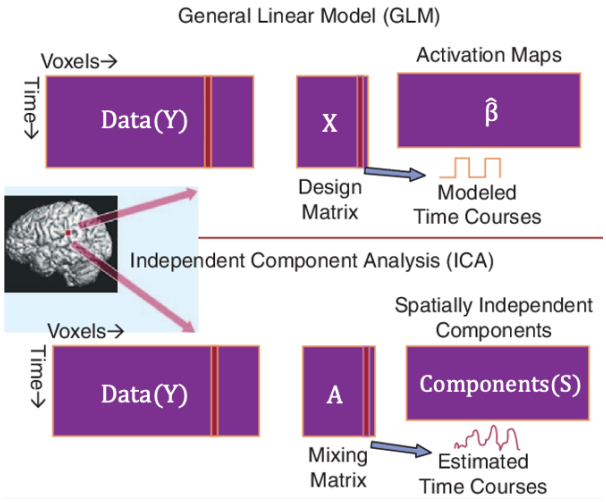

Figure 2.5: Comparison of GLM and ICA. In GLM, temporal signatures are modelled based on prior knowledge on experimental conditions, whereas in ICA they are estimated. Image adapted from [74].

On the other hand, ICA assumes that the observations in $\mathbf{Y}$ are linear mixtures of independent sources. In other words, while PCA seeks orthogonality in both of the estimated factor matrices, ICA assumes independence along a single mode, which is typically selected as space. ICA refers to $\mathbf{V}^T$ as the "source matrix" $\mathbf{S}$, which is optimized to contain statistically independent spatial maps in its rows. Each spatial map is associated to a time course contained in the respective column of $\mathbf{U}$, commonly called the "mixing matrix" $\mathbf{A}$, such that $\mathbf{Y} = \mathbf{AS}$. Owing to its intuitive nature and versatility in characterizing brain function, ICA has emerged as one of the most prevalent methods in neuroimaging literature [74]. Note that the formulation presented here corresponds to spatial ICA, which is the standard method of ICA for neuroimaging modalities that encompass a much larger spatial dimension than temporal dimension, such as fMRI and fUS. Temporal ICA is preferred, for example, in EEG studies, where $\mathbf{Y}$ gets transposed such that $\mathbf{S}$ contains independent time courses, whose associated spatial activation patterns are stored, unconstrained, in $\mathbf{A}$ [75]. In the rest of this thesis, ICA refers to spatial ICA, unless specified otherwise.

A comparison of ICA and GLM is provided in Fig. 2.5. In summary, while GLM specifies the temporal signatures of the extracted components (i.e., columns of the factor matrix in time), ICA estimates them based on the assumption that their spatial counterparts are independent, without requiring a priori modelling of neural processes. ICA has also proven effective in removal of artifacts [76], extracting functional connectivity networks [77], or for multimodal integration [78].

Due to the fact that functional brain data is inherently multidimensional (3D space and time), tensor factorizations have also been employed for a variety of neuroimaging applications [79]. Tensor factorizations are generalizations of matrix factorizations for higher-order data. For example, canonical polyadic decomposition (CPD) expresses an

input tensor $\mathcal{Y} \in \mathbb{R}^{I_1 \times I_2 \times \ldots \times I_D}$ as a sum of outer products (Fig. 2.6(a)):

$$\mathcal{Y} \approx \sum_{k=1}^{K} \mathbf{u}_k^{(1)} \circ \mathbf{u}_k^{(2)} \circ \ldots \circ \mathbf{u}_k^{(D)}, \tag{2.9}$$

where each term $u_k^{(d)}, d = 1, 2, \ldots, D$ represents the signature along mode d of the kth component. The factor matrices are given by $\mathbf{U}^{(d)} = [\mathbf{u}_1^{(d)} \ldots \mathbf{u}_K^{(d)}]$. Notice that Eq. 2.9 is a generalization of Eq. 2.8 to arrays of any order. CPD can also be viewed as an extension of SVD to higher orders, with the difference that factor matrices are not necessarily orthogonal for CPD. CPD, similar to aforementioned methods, can be used for source separation, localization, or artifact removal. An important advantage of CPD over many other factorizations, including ICA, is that CPD provides a unique decomposition of the data tensor under mild constraints [80].

CPD expresses the input tensor as a sum of K rank-1 terms, meaning that each component is defined via a single vector along each mode. Block-term decomposition (BTD), on the other hand, decomposes the input tensor in terms that are of low (≥ 1) multilinear rank (Fig. 2.6(b)):

$$\mathcal{Y} = \sum_{k=1}^{K} \mathcal{S}_k \times_1 \mathbf{U}_k^{(1)} \times_2 \mathbf{U}_k^{(2)} \times_3 \ldots \times_D \mathbf{U}_k^{(D)}, \tag{2.10}$$

where $\mathcal{S}_k$ is the core tensor of the kth component and $\times_d$ is the d-mode product. The d-mode product stands for the multiplication between a tensor by a matrix along mode d. For example, the d-mode (matrix) product of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \ldots \times I_D}$ and a matrix $\mathbf{U} \in \mathbb{R}^{J \times I_n}$ is given by a tensor $\mathcal{Z} = \mathcal{X} \times_d \mathbf{U}$, whose elements satisfy:

$$z_{i_1 \ldots i_{d-1} j \, i_{d+1} \ldots i_D} = \sum_{i_d=1}^{I_d} x_{i_1 i_2 \ldots i_D} u_{j i_d}. \tag{2.11}$$

The core tensor $\mathcal{S}_k$ is in a sense analogous to the diagonal singular value matrix in SVD since it contains coefficients that represent the significance of each mode combination in reconstructing the original tensor, yet there is no diagonality constraint in the core tensor of BTD. BTD achieves a more general low-rank structure compared to CPD while preserving uniqueness under relatively mild conditions, and can better capture the rather complex interactions between the factors where the trilinear model (for third order tensors) of CPD remains insufficient [81].

Other approaches used for analyzing neuroimaging data include K-means clustering, which can identify groups of voxels that exhibit similar responses [82]. Particularly for discovering networks of brain connectivity, graph-based approaches have gained significant attention [83]. In addition, supervised machine learning algorithms such as support vector machines, random forests or deep neural networks have served various purposes in neuroscience research including classification and prediction of brain disorders [84], image reconstruction [85] and development of brain-computer interfaces [86].

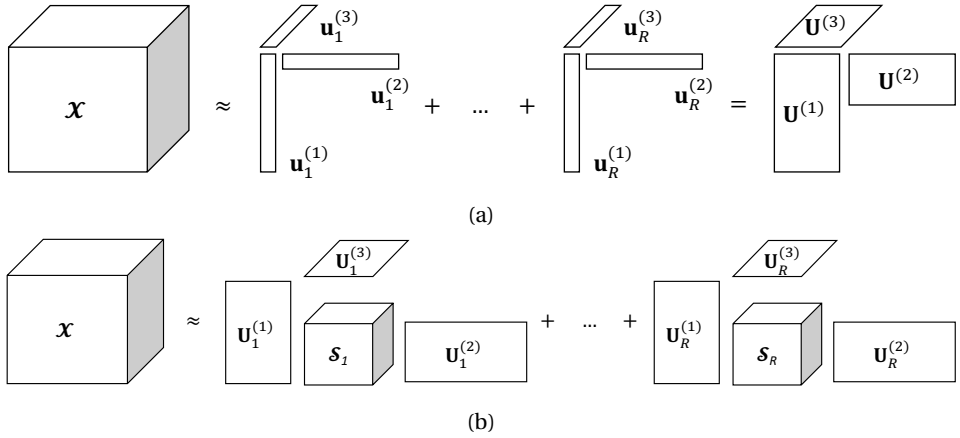

Figure 2.6: CPD (a) and BTD (b) of a third-order tensor.

2.3. fUS ANALYSIS IN THE LITERATURE

Having entered the field of neuroimaging relatively recently, there is ongoing research to elucidate what is measured by functional ultrasound (fUS). The relationship between simultaneously recorded fUS data and neuronal firing has been investigated by [87], revealing a clear link between the two especially at lower (< 0.3 Hz) frequencies, and this link can be well described by a hemodynamic response function. Similarly, [66] delineates a significant connection between the fUS signal and local neuronal activation. However, the authors point out a discrepancy of the modeled response (HRF convolved with neuronal activation) compared to the measured fUS data particularly under very strong stimulation, indicating stimulus-dependent behaviour. Due to the fact that the post-stimulation undershoot response is not present in fluctuations in cerebral blood volume, the HRF model given in Eq. 2.4 is adapted to fUS data as a single gamma function [88, 66].

The sensitivity of fUS in measuring subtle variations of microvascular blood dynamics has led to a variety of studies in the past decade, ranging from detection of responses to sensory stimuli to complex brain states and behavior [87]. Due to the intracranial imaging, most of the fUS studies are conducted on animal subjects including small rodents [19, 24, 34, 89], birds [90] and non-human primates [91, 92]. Nevertheless, fUS has also been used to monitor the brain activity of infants [93] and for intra-operative imaging of awake patients under somatosensory stimuli [44, 94].

2.3.1. LIMITATIONS OF CURRENT METHODS

Majority of the fUS studies so far has employed PCC for examining brain responses, whether to display evoked activity or functional connectivity – see a summary in Fig. 2.7. As such, reported activation values are restricted by the modelled time course (lagged or HRF-convolved stimuli) that is used to compute voxel correlations. Despite that it is possible to find an ideal lag (or HRF) per voxel, commonly one value that provides the best

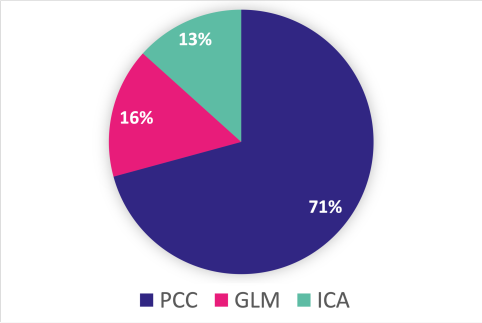

Figure 2.7: Frequency of ICA, GLM and PCC analyses employed in fUS studies. The values presented here are obtained by searching for the keywords "independent component analysis", "general linear model" and "correlation" amongst publications that include "Functional Ultrasound" in their title in Google Scholar.

overall activation in the whole brain is chosen [44]. Indeed, [28] reports that large differences between estimated and true HRFs can cause false negatives. Moreover, PCC is susceptible to outlier samples and trial responses, which can introduce erroneous correlations or mask existing ones [95].

Although spatiotemporal variability throughout an experiment or across different scans of the same subject [96] have been observed before with fUS, the quantification or modelling of this variability have been missing. Examples of such behaviour from the literature are provided in Fig. 2.8. Misinterpretation of the brain's variability can lead to deceptive conclusions. For instance, [97] points out that the consistently lower amplitudes found in the task-averaged response of schizophrenic patients (compared to healthy controls) [98] might be the result of latency variations across tasks [99].

2.4. CONCLUSIONS

In this chapter, we provided an introduction to the commonly used techniques for the analysis of functional neuroimaging data. We described the LTI model, which is at the center of hypothesis-driven techniques owing to its simple yet intuitive portrayal of hemodynamic activity. Moving forward, we delved into data-driven techniques, emphasizing their enhanced flexibility. We underlined that there is still so much to unfold from neuroimaging data, but even more so specifically for fUS. In the sections to follow, we will first describe the spatiotemporal variability of the fUS signal and subsequently propose new approaches for decoding the brain responses with fUS while accounting for this variability.

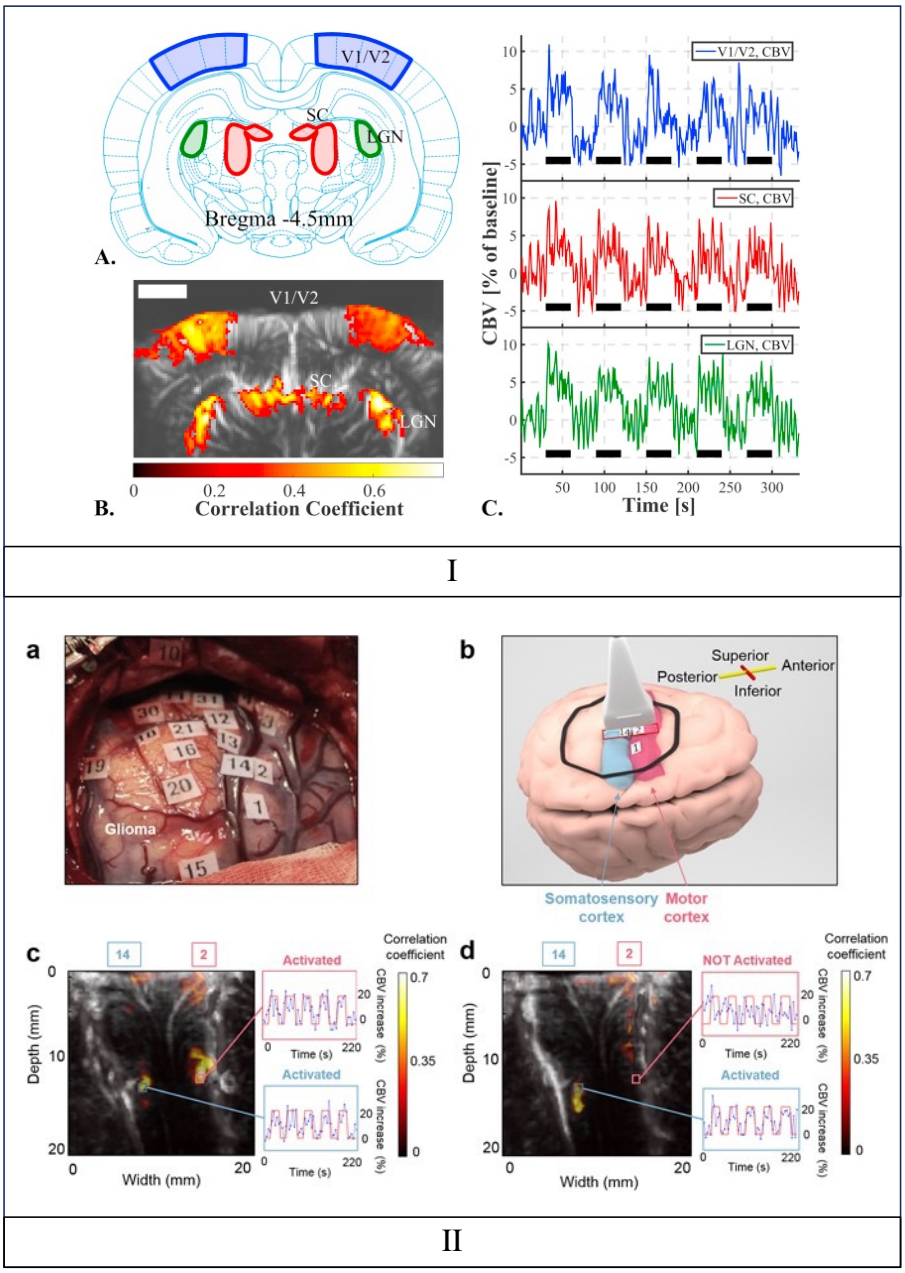

Figure 2.8: Examples from the literature demonstrating task-evoked fUS variability. At the top, the cerebral blood volume (CBV) change in percentage induced by visual stimuli along three regions of interest of the rodent brain are shared [34]. In Figure (I.C), the visual cortex (V1/V2) can be seen to exhibit a much more variable response compared to the other regions. At the bottom, the CBV change (%) in the motor cortex (blue colored) and somatosensory cortex (red colored) of the human brain are provided with respect to a motor (II.c) and sensory (II.d) task [94]. Although both of the investigated regions are activated by the motor task (Figure (II.c)), the somatosensory cortex is observed to maintain a more stable response magnitude across the task repetitions. Images taken from [34] (I) and [94] (II).

3

SPATIOTEMPORAL CHARACTERISTICS OF THE fUS RESPONSE

"The important thing is to never stop questioning."

— Albert Einstein

Part of this chapter is published as: (1) A. Erol, P. Kruizinga and B. Hunyadi (2024). Analyzing Trial-to-Trial Variability in the Mouse Visual Pathway using Functional Ultrasound. In 2024 IEEE 21st International Symposium on Biomedical Imaging (ISBI). IEEE. (2) S. E. Kotti, A. Erol and B. Hunyadi (2023, June). Modeling Nonlinear Evoked Hemodynamic Responses in Functional Ultrasound. In 2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW) (pp. 1-5). IEEE.

3.1. INTRODUCTION

THIS chapter investigates the functional ultrasound (fUS) signal characteristics in two parts. We first examine the validity of the linear convolutional model under different stimulus conditions. We quantify the difference between a linear and a nonlinear time-invariant model of the hemodynamic response in terms of data fitting and prediction performance with respect to changing stimulus duration. This way, we determine a range of durations in which the linear convolutional model can be considered plausible. In the second part, we focus on the time-varying properties of the fUS signal by keeping a fixed stimulus duration throughout the experiments and following how the responses of regions of interest vary across the stimulus repetitions. We track this variation by defining an expected voxel-specific response for one stimulus repetition, and following how much this model fits to the actual measured responses of each stimulus repetition. Our results confirm the time-varying aspect of brain responses found in prior studies performed with other functional neuroimaging modalities.

3.2. LINEARITY OF THE FUS RESPONSE

fUS is a high-sensitivity neuroimaging technique that images cerebral blood volume changes, which reflect neuronal activity in the corresponding brain area through the phenomenon of neurovascular coupling (NVC). The majority of studies addressing the NVC assume a linear time-invariant (LTI) model for the hemodynamic response, which can be described by the convolution of an impulse response known as the hemodynamic response function (HRF) and a binarized representation of the stimulus signal as input. Estimating the HRF, and thus uncovering the dynamics of the NVC, is a prevalent research topic in hemodynamics-based imaging modalities for decades. Although the hemodynamic response is known to exhibit nonlinear characteristics, linear modeling has maintained its popularity due to its good performance and simplicity. However, under certain circumstances, deviations from the behavior of an LTI system can become too severe to neglect, examples of which include experiments with closely repeated stimuli [100], or under changing duration [101] or strength [102] of stimuli.

Volterra series are high order extensions of the Taylor series. A Volterra series of order one simply describes a one-dimensional convolution operation; higher order Volterra kernels, e.g., second-order kernels that describe pairwise interactions of the input signal at different time instances, were found to sufficiently capture nonlinearities of the functional magnetic resonance imaging (fMRI) signal caused by interactions between successive stimuli in an auditory experiment [100].

The contribution of this work is twofold. On one hand, fUS measurements on mice during a visual experiment with varying stimulus durations were used to show and quantify under which stimulus conditions the fUS data can be better approximated using a nonlinear convolution, as opposed to a linear one. On the other hand, the applicability of Volterra series to characterize the nonlinear aspects of the HRF was shown by estimating the kernels using a series of different stimulus durations and predicting individual responses.

METHODS

DATA ACQUISITION AND PREPROCESSING

For our HRF modeling on fUS data, two experiments (on two different mice and brain slices) were conducted at the Center for Ultrasound and Brain imaging at Erasmus University Medical Center (CUBE). During data acquisition, the subject mouse was presented with randomly generated high-contrast images (a rectangular patch of white "speckles" against a black background [35, 34, 103]), which succeeded each other with 25 frames per second, on two stimulation screens simultaneously. In the first experiment (Exp. 1), the stimulus was shown with 4 different durations (1, 4, 10 and 20 s), 10 repetitions each, in random order. Each repetition was followed by a rest period (completely black screens) of random duration between 20 and 25 s. In the second experiment (Exp. 2), the durations were 0.25, 0.5, 1, 2, 3, ..., 10 s and there were 11 repetitions. For Exp. 1, we report results on all durations; for Exp. 2, we report results on 1, 4, 8 and 10 s.

The time series of a pixel corresponds to its power variation across the Power-Doppler Image (PDI) stream. The final frame rate was 4 Hz in Exp. 1 and 7.44 Hz in Exp. 2. The region-of-interest (ROI) per experiment was selected as an area of 16 pixels in the left superior colliculus, a region known to be involved in the processing of visual stimuli [104]. Pixels with a Pearson Correlation Coefficient (PCC) larger than 0.3 with the stimulus are indicated in Fig. 3.1 for Exp. 1.

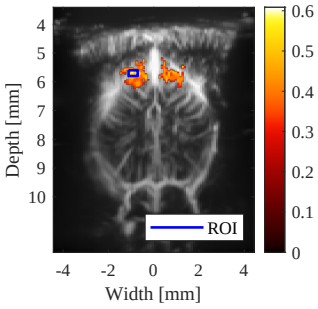

Figure 3.1: ROI in experiment 1.

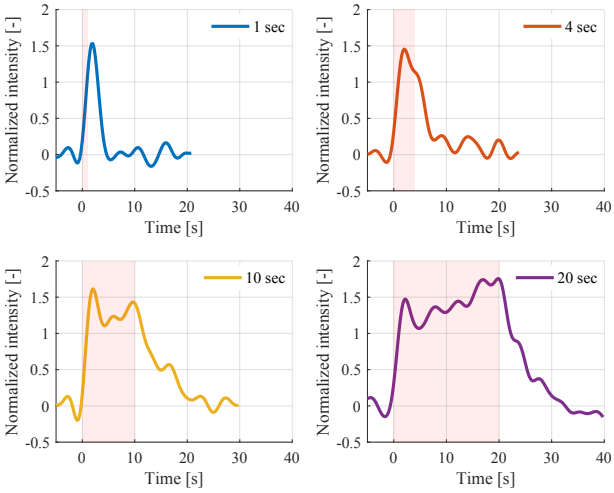

Figure 3.2: Average filtered ROI response for all stimulus durations. Each response is plotted from 5 seconds prior to stimulus onset to 20 seconds after the stimulus ends, i.e. till the end of the rest period. The stimulus is displayed in red color in the background of each plot.

The next step of data preprocessing included standardizing each ROI pixel time series and averaging them over the ROI to obtain a single time series. This signal was subsequently subjected to low-pass filtering using a fifth-order Butterworth filter with 0.3 Hz cutoff frequency, in order to only preserve the fUS signal content that associates with

the underlying neuronal activity [87]. We applied baseline correction to each stimulus repetition separately.

In Fig. 3.2, the resulting signals per duration are shown for Exp. 1, after averaging over the multiple stimulus repetitions. The response to 1 sec stimulus shows a single peak before it returns to the baseline level. The observed signal takes more intricate shapes for the stimuli of 10 and 20 sec: the responses drop after the initial peak and then increase again. This raised the question of whether such responses can be modeled as the linear convolution of an HRF with a block-type signal. A similarly delayed component in the fUS response to strong inputs (high concentration odors) has been observed in [66], which was modelled using a second linear HRF term that peaks at a late point. However, the authors assume that the neuronal Ca^{2+} signal is the driver of that LTI system, which is not block type.

VOLTERRA SERIES

In this part, we evaluate the LTI assumption and further explore Volterra series [100] as an alternative to perform nonlinear system identification given fUS measurements. The Volterra series is an extension of the Taylor series to cover dynamic systems and can represent any analytic time-invariant system [105]. In our case, the second-order Volterra series approximation is given by

$$y(t) \approx h^{(0)} + \sum_{t_1=0}^{L} h^{(1)}(t_1)s(t-t_1) + \sum_{t_1=0}^{L}\sum_{t_2=0}^{L} h^{(2)}(t_1,t_2)s(t-t_1)s(t-t_2), \tag{3.1}$$

where $y(t)$ for $t = 1,\dots,N$ is the hemodynamic response captured in the fUS signal, $h^{(m)}(\cdot)$, $m = 0,1,2$, is the m-th order Volterra kernel and $s(t)$ is the block-type stimulus signal. The high-order Volterra kernels allow for modelling of the output response as a nonlinear function of the input samples, as shown in Fig. 3.3.

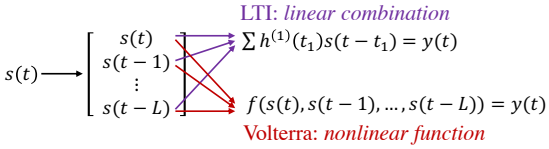

Figure 3.3: Linear (LTI) vs. nonlinear (Volterra series) convolution. Image adapted from [100].

In order to estimate the coefficients of this expansion using linear estimation methods and ensure that the resulting kernel shapes will not be arbitrary, we expand the kernels in terms of B temporal basis functions $g_i(t)$, $i = 1,\dots,B$, following [100]. The kernels are expanded as

$$h^{(0)} = \beta^{(0)}$$
$$h^{(1)}(t_1) = \sum_{i=1}^{B} \beta_i^{(1)} g_i(t_1) \tag{3.2}$$

$$h^{(2)}(t_1,t_2)=\sum_{i=1}^{B}\sum_{j=1}^{B}\beta_{ij}^{(2)}g_i(t_1)g_j(t_2).$$

Using the one-dimensional convolution of the input $s(t)$ and the basis functions $g_i(t)$, we can define a new set of response variables $x_i(t)$ as $x_i(t)=\sum_{t_1=0}^{L}g_i(t_1)s(t-t_1)$. Substituting this into 3.1 gives

$$y(t)\approx\beta^{(0)}+\sum_{i=1}^{B}\beta_i^{(1)}x_i(t)+\sum_{i=1}^{B}\sum_{j=1}^{B}\beta_{ij}^{(2)}x_i(t)x_j(t). \tag{3.3}$$

Stacking all $y(t)$ into the observation vector $\mathbf{y}\in\mathbb{R}^{N\times 1}$ and the B' explanatory variables 1, $x_i(t)$ and $x_i(t)x_j(t)$ as columns of a design matrix $\mathbf{X}\in\mathbb{R}^{N\times B'}$, we obtain the model

$$\mathbf{y}=\mathbf{X}\boldsymbol{\beta}+\boldsymbol{\epsilon}, \tag{3.4}$$

where the error vector $\boldsymbol{\epsilon}\in\mathbb{R}^{N\times 1}$ includes possible measurement noise and uncaptured higher order dynamics. The vector $\boldsymbol{\beta}\in\mathbb{R}^{B'\times 1}$ includes all unknown coefficients $\beta^{(0)}$, $\beta^{(1)}$, and $\beta^{(2)}$ from which the kernel coefficients $h^{(m)}(\cdot)$, $m=0,1,2$, can be calculated using 3.2. Notice that, Eq. 3.4 has the same form as Eq. 2.3 due to the basis function representation, however, $\mathbf{X}$ here also includes higher-order interactions between input samples.

The basis functions in this work were chosen to follow the HRF model described by a single gamma function [66, 88], that is

$$g(t,\boldsymbol{\theta})=\theta_1(\Gamma(\theta_2))^{-1}\theta_3^{\theta_2}t^{\theta_2-1}\mathrm{e}^{-\theta_3 t}, \tag{3.5}$$

for $t\geq 0$, where $\boldsymbol{\theta}=[\theta_1,\ \theta_2,\ \theta_3]^{\mathrm{T}}$ and θ_1, θ_2, $\theta_3>0$ control the response height, the delay and the dispersion of the function [106]. The parameters of the basis functions were fixed so that the functions peak during different possible components of the hemodynamic response, i.e., early, intermediate and late components, and were not optimized for. Fig. 3.4 shows the 8 basis functions that were used; their peaks were spaced at 2 seconds.

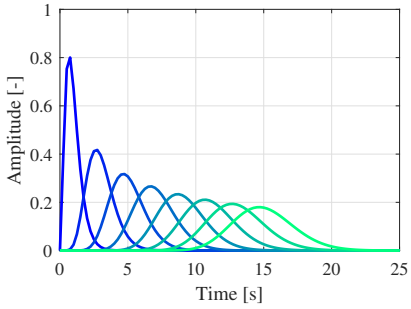

Figure 3.4: Basis functions.

Finally, the following optimization problem is solved to estimate $\hat{\boldsymbol{\beta}}$:

$$\begin{aligned}&\underset{\boldsymbol{\beta}}{\text{minimize}} && \|\mathbf{X}\boldsymbol{\beta}-\mathbf{y}\|_2+\lambda\|\boldsymbol{\beta}\|_1\\ &\text{subject to} && \textstyle\sum_{i=1}^{B}\beta_i^{(1)}g_i(t)\geq 0 \quad \forall t.\end{aligned} \tag{3.6}$$

The ℓ_1-norm regularization is used to promote sparsity in the solution, so that interpretable shapes for the HRF can be obtained, given the number of columns in the design matrix. The parameter λ controlling the regularization was set empirically to 10^{-4} (for unit energy $\mathbf{X}$ and $\mathbf{y}$). The constraint requires that the first order kernel $h^{(1)}(t)$ be positive for all t, as the linear fUS HRF is assumed to be always positive [66, 106].

RESULTS

TRAINING PER STIMULUS DURATION

In order to evaluate and compare the potential of a linear and a nonlinear HRF to describe the obtained fUS measurements, we solve Eq. 3.6 assuming a maximum Volterra order $M = 1$ (linear) and $M = 2$ (nonlinear) per each stimulus duration separately. Given the estimated coefficients $\hat{\boldsymbol{\beta}}$, we then reconstructed the observation vector $\mathbf{y}$ as $\hat{\mathbf{y}} = \mathbf{X}\hat{\boldsymbol{\beta}}$.

The mean squared error (MSE) achieved by the reconstruction for $M = 1$ and $M = 2$, is given in Tables 3.1 and 3.2, for the two experiments respectively. The difference in the reconstruction achieved by the two in the case of stimulus of 1 sec is rather low. The MSE improvement is much larger for longer stimuli. Focusing on the 4 sec stimulus, in Exp. 2 the MSE improvement achieved when incorporating second-order kernels is comparable to that achieved for 1 sec stimulus. In Exp. 1 the improvement factor is much larger. This difference does not come as a surprise, as it is known that the hemodynamic response can differ per subject. These results show that the responses to longer stimuli can be reconstructed more faithfully when a maximum Volterra order $M = 2$, thus a nonlinear HRF, is used. The above observation is justified by the fact that a binary stimulus in combination with an LTI hemodynamic system cannot account for responses that depart from the form of the 1 sec response in Fig. 3.2.

Table 3.1: Reconstruction MSE - Exp. 1.

Stimulus duration [sec]	**MSE** [$\times 10^{-5}$]		Improvement factor [$\times 10^3$]
	$M = 1$	$M = 2$	
1	9.014	0.0931	0.097
4	64.672	0.005	12.992
10	31.868	0.007	4.569
20	19.239	0.003	7.746

Table 3.2: Reconstruction MSE - Exp. 2.

Stimulus duration [sec]	**MSE** [$\times 10^{-6}$]		Improvement factor [$\times 10^3$]
	$M = 1$	$M = 2$	
1	1.845	0.293	0.006
4	31.183	0.081	0.385
8	290.847	0.005	62.371
10	436.196	0.004	98.407

TRAINING AND TESTING ACROSS STIMULUS DURATIONS

After confirming that the Volterra series can be used to model hemodynamic responses to long durations of this visual stimulus, the next step was to investigate whether kernels trained on multiple durations can predict the responses to individual stimuli.

Using the data of Exp. 1, the temporally averaged responses to the stimulus durations of 1, 4, 10, and 20 seconds (one time series per duration) were concatenated, as well as the stimuli themselves. It should be noted that the energy of each response was matched to the energy of the corresponding stimulus, to alleviate differences between responses to different durations. During the training phase, a single set of Volterra kernels was obtained by solving Eq. 3.6 using the concatenated signals (70% of the repetitions were averaged) for $M = 1$ and $M = 2$ separately. These kernels were then used on the testing set (30% of the repetitions) to predict individual responses to stimuli of 1, 4, 10 and 20 seconds. The performance was evaluated against the average of the repetitions in the testing set, and also against the individual repetitions. The prediction results can be found in Fig. 3.5. In Tables 3.3 and 3.4 the PCC between the true response and the predicted response is given, for the average data and each repetition separately.

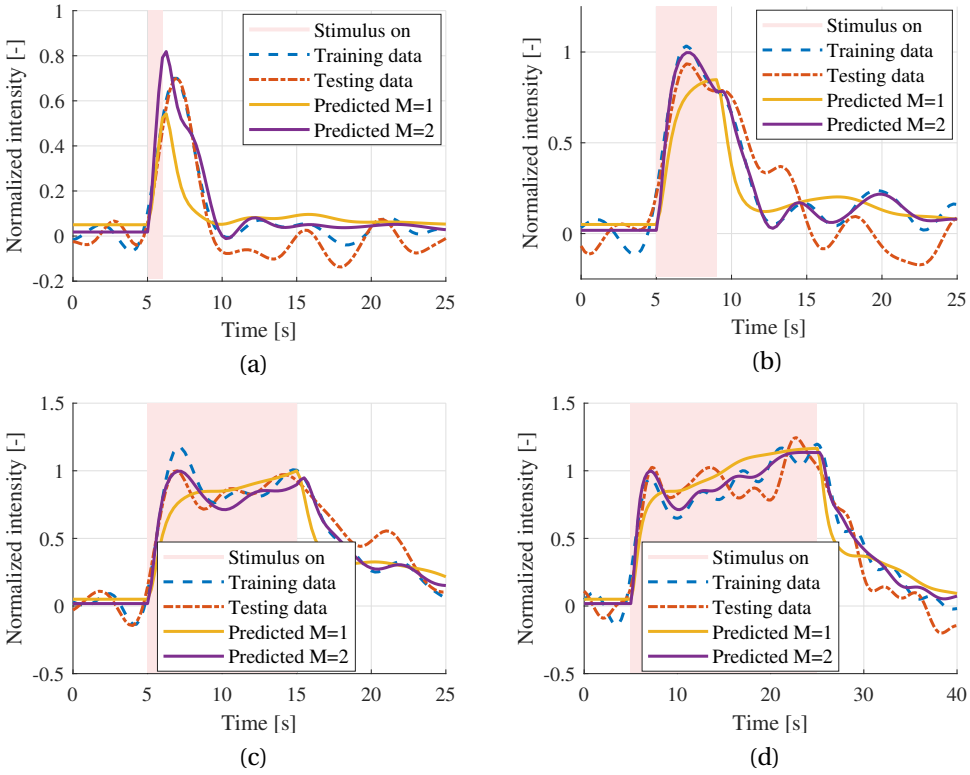

Figure 3.5: Training, testing and predicted signals for maximum Volterra order $M = 1$ and $M = 2$ in Exp. 1, for stimulus duration 1 sec (a), 4 sec (b), 10 sec (c) and 20 sec (d), when all durations are used to estimate the kernels. The testing data shown is the average.

Notable is the underestimation of the first peak in the responses for $M = 1$ in Fig. 3.5, compared to $M = 2$, and the noncapture of the dip after the peak, although the overall duration of the response is correctly predicted. Contrary to [66], where HRF amplitude can vary based on the stimulus, we assume a fixed amplitude for the kernels over all durations. This leads to $M = 2$ being able to capture amplitude differences in the responses better than $M = 1$. When it comes to the PCC results, there is improvement in all durations (averaged or individual data) when kernels are trained with $M = 2$. Training with $M = 1$ using all durations provides rather insufficient kernels for the 1 sec stimulus.

Table 3.3: PCC between the actual and the predicted responses, $M = 1$.

Stimulus duration [sec]	Testing data average	Testing data 1	Testing data 2	Testing data 3
1	0.773	0.740	0.681	0.830
4	0.826	0.733	0.735	0.858
10	0.876	0.859	0.662	0.783
20	0.914	0.868	0.877	0.839

Table 3.4: PCC between the actual and the predicted responses, $M = 2$.

Stimulus duration [sec]	Testing data average	Testing data 1	Testing data 2	Testing data 3
1	0.928	0.898	0.890	0.875
4	0.880	0.773	0.825	0.870
10	0.944	0.901	0.675	0.920
20	0.961	0.897	0.918	0.905

The estimated kernels for $M = 2$ are shown in Fig. 3.6. It is clear that the first-order kernel is mostly concentrated in the first 5 seconds, but also shows fluctuations in the tail. This is a result of training on all durations and using basis functions that extend in time. This will, in turn, affect the prediction for the shorter durations, which are overfitted in the tail (Fig. 3.5(a), (b)). The second-order kernel is also mostly concentrated in the first 5 seconds. The negative values in this area suggest that if the stimulus has been on in the last few seconds, the hemodynamic response will be suppressed. The estimated kernels in Exp. 2 are not included, but overall agree that the larger part of the hemodynamic response is contained within the first 5 seconds.

DISCUSSION

The aim of this part was to complement ongoing modeling efforts in the fUS domain. We drew attention to the fact the nonlinearities in the responses due to stimulus interactions over time can be significant and should be considered when designing experimental paradigms and modeling the hemodynamic system. The results show that a Volterra series approximation can be used to model the nonlinear character of the HRF. When trained using the responses to stimuli of different durations, it is possible to predict individual responses to stimuli with good accuracy. While the difference between the model

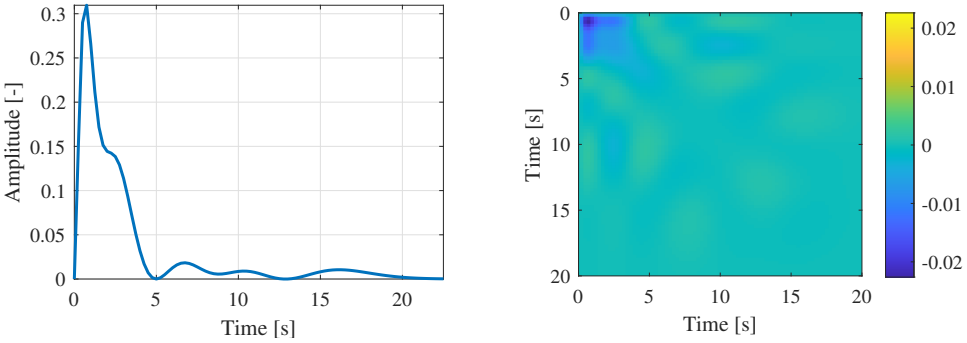

Figure 3.6: Estimated kernels $h^{(1)}(\cdot)$ (a) and $h^{(2)}(\cdot)$ (b), for maximum Volterra order $M = 2$ in Exp. 1.

fit achieved by first and second-order kernels is not significant for short stimuli, the intricacies of the responses to longer stimuli become too strong to disregard, and cannot be justified by the combination of an LTI system and a binary input signal.

Last but not least, we would like to discuss the significance of the oscillations in the fUS response prior to stimulus presentation (Fig. 3.5). While the low-pass filtering employed in this study might make these oscillations appear more pronounced, we could still observe them without any additional filtering on the power-Doppler stream. Brain's pre-stimulus activity is commonly studied by investigating neuronal oscillations at high frequencies (~ 10 Hz), which are argued to occur spontaneously [107] and have influence on the task response that follows in cognitive paradigms [108, 109]. For hemodynamics-based modalities such as fMRI and fUS with lower temporal resolution, pre-stimulus oscillations are very subtle and are generally not the focus of analysis. Although the literature on it is quite limited, pre-stimulus hemodynamic activity is associated with developing anticipation, presenting itself as an overshoot [110, 111]. Indeed, pre-stimulus activity is often used solely to define the baseline (by selecting a window of a few seconds prior to stimulus presentation and averaging the activity within this window — thus flattening any oscillation) for quantifying the magnitude of post-stimulus peak responses in both fMRI and fUS [35, 112]. Considering the challenges in accurately measuring and interpreting the oscillatory effects in hemodynamics-based studies, we proceed under the assumption of a causal system throughout this thesis.

CONTRIBUTIONS

Part of this section is published as S. E. Kotti, A. Erol and B. Hunyadi (2023, June). Modeling Nonlinear Evoked Hemodynamic Responses in Functional Ultrasound. In 2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW) (pp. 1-5). IEEE. The authors confirm contribution to this work as follows. SE Kotti: Methodology, Implementation, Manuscript Preparation. A Erol: Conceptualization, Experimental Design, Manuscript Preparation. B Hunyadi: Methodology, Supervision, Manuscript Preparation.

3.3. TRIAL-TO-TRIAL VARIABILITY OF THE FUS RESPONSE

fUS, similar to fMRI, measures hemodynamics as a proxy for neural activity [17]. Both modalities frequently employ multiple repetitions (*trials*) of stimuli in order to improve the signal-to-noise ratio (SNR) for capturing evoked activity and enhance the statistical power of the study [113]. The majority of fUS and fMRI analyses use the time courses of stimuli to extract spatial activation maps, either via linear models (known as the general linear model - GLM) or by computing the per-voxel PCC. For both approaches, the stimulus time course, represented as a boxcar function indicating when the stimulus is on or off, is convolved with a hemodynamic response function (HRF).

An inherent assumption of using the convolutional model with a binarized input is that the brain response is time-invariant. However, we know this assumption does not hold due to various neurobiological mechanisms such as neuronal adaptation, stimulus expectation, time-varying functional connectivity [114], variations in arousal [115], and uninstructed movements [116]. While the extent to which fUS detects this variability based on the measured hemodynamics remains an open question, not taking it into account can lead to an imprecise portrayal of the brain's actual response. In more extreme scenarios, this imprecision could potentially yield erroneous conclusions in neurocognitive research [97].

In the literature, there are several approaches proposed for tackling trial variability. For example, a common procedure is *selective* averaging of trials by rejecting outlier trials. To determine which trials should be considered as unreliable, various methods have been employed such as amplitude-based thresholding (attributed to motion or eye blink artifacts) [24, 117] or visual inspection [118]. Nevertheless, experiments have shown that repeating of stimuli not only induces changes in the magnitude, but also in the peak latency or duration [119] of individual trial responses. As such, only tracking the changes in magnitude may not be adequate to represent the full extent of trial variability.

The GLM framework has also been used to model trial variability, for instance, by defining a separate regressor for each trial, or picking a trial of interest and grouping every other trial under another regressor [120]. Allowing for trial variability was shown to improve classification accuracy with fMRI [121]. However, it's important to note that these methods employ a fixed HRF in their designs, despite the known variations in HRF across individuals and brain regions [28]. For exploring trial variability with fUS, we define a data matrix for each voxel whose columns correspond to the distinct trial responses of the voxel. Likewise, we characterize the model time course for a single trial response, i.e. we convolve a voxel-specific HRF with one stimulus trial. In our analysis, we focus on two crucial regions of interest (ROIs) along the mouse brain's visual-processing pathway: the lateral geniculate nucleus (LGN) and visual cortex (VIS). Finally, we estimate trial activation coefficients for each ROI using linear regression.

METHODS

For each voxel, we segment the observed response into trials, and place these segments to columns of a matrix $\mathbf{Y}$. We define our data model as:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \tag{3.7}$$

where $\mathbf{Y} \in \mathbb{R}^{N\times K}$ is the data matrix, $\boldsymbol{\beta} \in \mathbb{R}^{R\times K}$ is the matrix of activation coefficients, $\mathbf{X} \in \mathbb{R}^{N\times R}$ contains the design variable(s), $\boldsymbol{\epsilon} \in \mathbb{R}^{N\times K}$ represents noise, K is the number of stimulus repetitions, N is the number of time samples included within one trial and R is the number of different types of stimuli. The proposed method is illustrated over example simulated data in Fig. 3.7. Note that we employed a single type of visual stimulus in our fUS experiments, hence for consistency the illustration is given for $R = 1$. It is possible to incorporate different stimuli by including their responses together in the trial segments. Then, a separate regressor can be defined with the onset of each event, as a new column in $\mathbf{X}$.

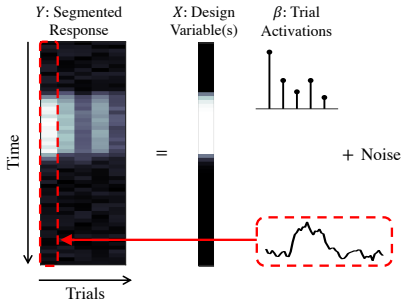

Figure 3.7: Depiction of the proposed method for $R = 1$. Each column of $\mathbf{Y}$ stands for a trial response. The design variable is given by the convolution of a stimulus trial with an HRF.

For constructing the design variable, we estimated voxel-specific HRFs by shifting the HRF peak latency at various values using a single gamma function as in Eq. 3.5: $h(t,\boldsymbol{\theta}) = \Gamma(\theta_1)^{-1}\theta_2^{\theta_1} t^{\theta_1-1}\mathrm{e}^{-\theta_2 t}$. We determined the HRF of a voxel as the one that provides the highest PCC with the stimulus across slices. Finally, we estimate $\hat{\boldsymbol{\beta}} = \mathbf{X}^{\dagger}\mathbf{Y}$ where $(.)^{\dagger}$ is the pseudo-inverse. We aimed at answering the following research questions:

Q1. Do the trial activations stay constant, or do they vary during the experiment? If they do, is this variation random or could we track a trend?

Q2. Do the variations occur simultaneously at the ROIs?

Q3. Can we find a reproducible difference in variability between LGN and VIS across slice recordings?

To unveil the trend of trial activations, we fit a linear model to the estimated $\boldsymbol{\beta}$ coefficients of each ROI (Q1). Subsequently, we express the trial activations of VIS as a function of those of LGN to see if they vary in a correlated manner (Q2). Finally, we compute the coefficient of variation (CoV) of the trial activations of the ROIs (Q3).

EXPERIMENTAL SETUP

fUS records the hemodynamic activity of the brain in 2D slices using ultrafast Doppler imaging. For fUS acquisition, 20 tilted plane waves ($\pm 8°$) were transmitted from the ultrasonic transducer, placed on the cranial window of a 7-months old male mouse (C57BL/6J). After Fourier-domain beamforming and angular compounding of the echo waves, singular value decomposition-based clutter filtering [20] was utilized to remove

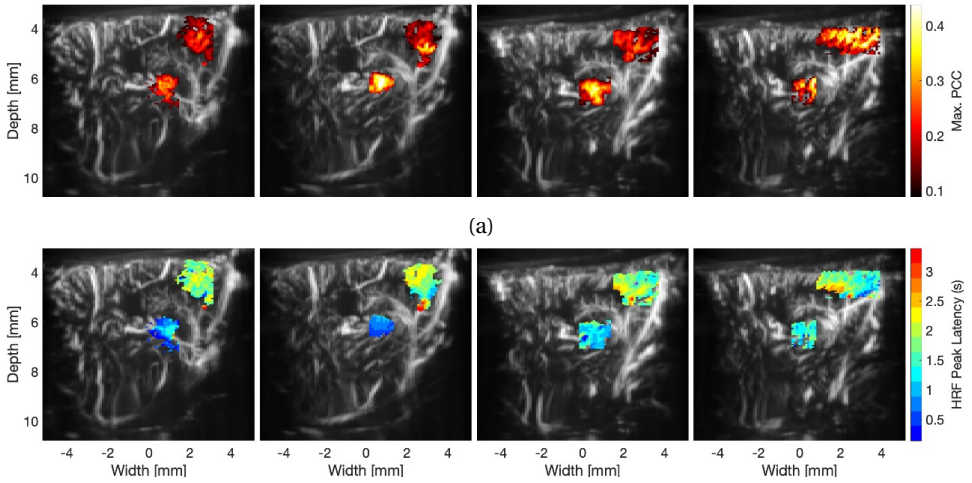

Figure 3.8: Thresholded correlation images (a) obtained using the optimal HRF peak latency (b), overlaid against the mean PDI, displaying LGN (bottom region) and VIS (top region). The slices were imaged at Lateral +2.55, -2,55, +2.15 and -2.15 mm from left to right respectively.

tissue components. We worked on PDIs sampled at a rate of 4 Hz. The mouse brain was imaged sagittaly at 4 locations (Lateral $\pm 2.15, \pm 2.55$ mm), all of which captured LGN and VIS (Fig. 3.8). The visual stimulus consisted of randomly generated high contrast images (white speckles against a black background). The stimulus was repeated 20 times throughout the experiment and its duration was kept at 4 s, while the rest periods in-between stimuli were randomized from 10 to 15 s. We applied spatial smoothing using a Gaussian kernel with a standard deviation of 0.5 voxels in size, and standardized the voxel time-series to zero-mean and unit variance prior to further processing.

To determine significantly active voxels for each ROI, we selected a P-value threshold of 0.0001 and converted this value to z-score using a two-tailed test (z-score: 3.71). Next, we applied Fisher's transform to this z-score and arrived at a PCC threshold of $c = 0.1$ [96]. We calculated the PCC of each voxel by correlating the voxel time-series with the HRF-convolved stimulus time course. The thresholded correlation images are displayed in Fig. 3.8. We segmented the voxel time-series into trials that start 3 s before a stimulus onset and ends 12 s after. Note that these values are determined in accordance with our stimulus design, but can be re-adjusted based on the paradigm. All segments were baseline corrected by subtracting the mean amplitude of the pre-stimulus response ([-3,0] s with respect to the stimulus onset time).

RESULTS

Using the optimal HRFs, we estimated trial activations per voxel. We averaged the trial activations across the voxels of a ROI at each slice (Fig. 3.9(a)). To start with, we observe a declining trend of trial activations as the trials progress. As discussed above, many factors including neuronal adaptation or habituation can cause such descent.

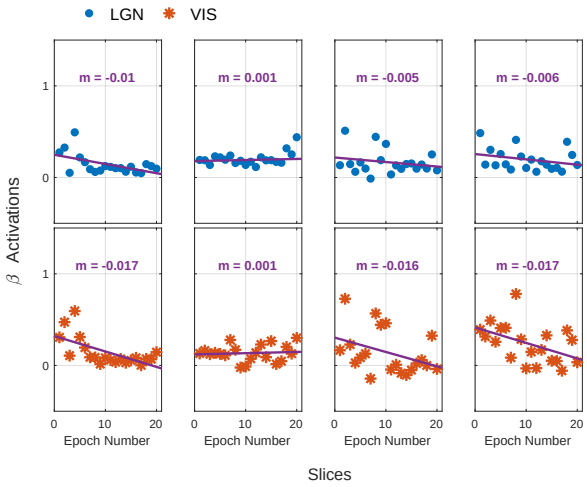

(a)

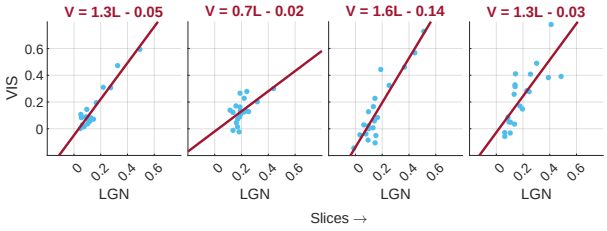

(b)

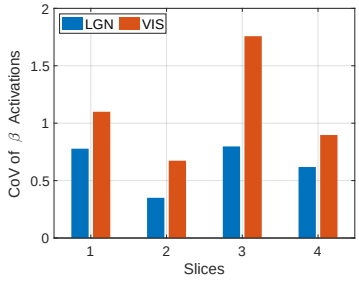

(c)

Figure 3.9: (a) Variation of β across trials for each region and brain slice. We plotted the best-fitting line (in least-squares sense) to β values in each case to highlight the trend of activations across trials. The slope m of the predicted line is indicated on top of each plot. (b) Trial activations of LGN vs. VIS. The equation written on each plot describes the line of best fit in least-squares sense (shown in red color) between VIS and LGN, denoted by V and L respectively. (c) Coefficient of variation (CoV) of β across regions and slices.

For studying the correlation of the timing of these variations between LGN and VIS, we plotted the trial activations of VIS as a function of those of LGN. The link between the two regions can be clearly observed in Fig. 3.9(b), as the best-fitting line shares similar parameters (slope and intersection) across slices. Note that this correlation is expected, at least up to some degree, as VIS is anatomically connected to, and receives direct input from the LGN.

Another point worth discussing is the offset in the line equations relating LGN to VIS (Fig. 3.9(b)). This offset is caused by the occasional negative activations of VIS, albeit small both in absolute value and quantity. This negativity comes as a result of VIS amplitudes dropping below the baseline in certain trials. We can consider the possibility that the mouse was not responsive in those trials, and the effects of noise (physiological and/or instrumentation related) became too dominant. However, in those same trials, negative activations were almost never found in LGN. Indeed, stimulus-evoked negative activations in VIS were reported before with fMRI as well [122]. For both Fig. 3.9(a) and (b), the results of the second slice are slightly different from the rest. Although this difference might be due to the particular slice that was imaged or a change in the mouse's attention, it is worth to mention that this slice has the least amount of active voxels, possibly affecting it's generalizability compared to the rest.

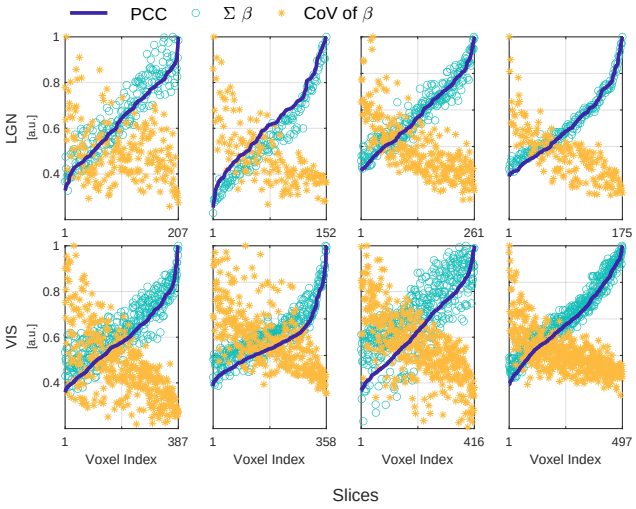

Figure 3.10: Change of PCC values, summation of non-negative trial activations and CoV of trial activations across voxels. Note that the voxels were sorted in ascending PCC and all measures were normalized to have a maximum amplitude of 1 for easier interpretability of the plots.

In order to assess the overall variability of the two regions, we calculated the CoV of β's, defined as the ratio of the standard deviation to the mean (Fig. 3.9(c)). We can observe that VIS exhibits a variability that is twice as high as that of LGN for all slices, in accordance with neuronal findings [27]. In fact, the transformation between LGN and VIS is argued to be the point at which the large response variability of VIS originates [123]. As a last note, we observed that the PCC values calculated conventionally (i.e., via correlating the whole HRF-convolved stimulus time course with the voxel time-series,

Fig. 3.8(a)) share a very similar trend with the summation of trial activation coefficients across voxels. On the contrary, the CoV of trial activations follows an opposing trend (Fig. 3.10), indicating that PCC values might drop as a result of higher variability.

DISCUSSION

In this part, our objective was to investigate trial variability using fUS. By segmenting the observed data into individual trial responses, we computed activation coefficients for each trial within specific regions of interest. We noted a general decline in activations throughout the experiment in both regions, potentially influenced by intricate brain mechanisms such as neuronal adaptation. The consistently higher variability in VIS compared to LGN observed in all of our recordings aligns with established neuronal findings. For example, neurons in the primary visual cortex were shown to encode both visual information and motor activity that arise from facial movements [124]. While we referred to the anatomical and functional connection between LGN and VIS to explain the similarity in trials where high variability is observed in each region, it is important to emphasize that spontaneous brain activity can also contribute to coherent trial variability across different brain areas [125].

3.4. CONCLUSIONS

This chapter discusses the implications of the conventional LTI model over fUS responses. We first showed that the linearity assumption is plausible under short stimulus durations by comparing the data fitting and prediction performance achieved by linear and nonlinear convolution using Volterra series, assuming the input of the Volterra series is the binarized stimulus signal, $s(t)$. Subsequently, we demonstrated that the fUS response to repeated stimuli differs, indicating time-variance. Furthermore, we showed that the amount of variability is region-dependent. Overall, we conclude that using the LTI model with a fixed input, that is the binarized stimulus, is an over-simplification of the brain's response, and does not truly represent the spatiotemporal variability. In the following sections, we continue to utilize the LTI model, albeit with a flexible characterization of the input sources for a better depiction of brain's complex activity.

4

DECONVOLUTION OF THE FUS RESPONSE USING BLOCK-TERM DECOMPOSITION

"It is imperfection –not perfection– that is the end result of the program written into that formidably complex engine that is the brain."

— Rita Levi-Montalcini

Part of this chapter is published as: A. Erol, C. Soloukey, B. Generowicz, N. Van Dorp, S. Koekkoek, P. Kruizinga, and B. Hunyadi (2023). Deconvolution of the Functional Ultrasound Response in the Mouse Visual Pathway Using Block-Term Decomposition. Neuroinformatics, 21(2), 247-265.

Functional ultrasound (fUS) indirectly measures brain activity by detecting changes in cerebral blood volume following neural activation. In the previous chapter, we presented the existing approaches for modelling such functional neuroimaging data, which are commonly based on the convolution between an impulse response, known as the hemodynamic response function (HRF), and a binarized representation of the input signal based on the stimulus onsets, the so-called experimental paradigm (EP). We discussed how the binarized EP form may not characterize the whole complexity of the activity-inducing signals that evoke hemodynamic changes due to the non-LTI aspects of brain function.

In this chapter, we propose an adaptable framework that can capture such complex dynamics of the brain function by modelling the multivariate fUS time-series as convolutive mixtures. We apply block-term decomposition on a set of lagged fUS autocorrelation matrices, revealing both the region-specific HRFs and the source signals that induce the hemodynamic responses. We test our approach on two mouse-based fUS experiments. In the first experiment, we present a single type of visual stimulus to the mouse, and deconvolve the fUS signal measured within the mouse brain's lateral geniculate nucleus, superior colliculus and visual cortex. We show that the proposed method is able to recover back the time instants at which the stimulus was displayed, and we validate the estimated region-specific HRFs based on prior studies. In the second experiment, we alter the location of the visual stimulus displayed to the mouse, and aim at differentiating the various stimulus locations over time by identifying them as separate sources.

4.1. INTRODUCTION

Hemodynamic activity is commonly modelled as the convolution between a function representing the impulse response of the neurovascular system, known as the hemodynamic response function (HRF), and a binarized representation of the input signal based on the stimulus onsets, the so-called experimental paradigm (EP) [126]. However, not all stimuli can be predefined, i.e. under the experimenters' control [53]. For example, brain reaction to mental imagery is shown to be almost as strong as the activity evoked by real perception in certain brain regions under a variety of experimental designs, such as visual [127] or auditory [51]. Moreover, spontaneous brain activity [125], changes in alertness [115], and uninstructed movements [116] can cause coherent fluctuations throughout the brain, and contribute to trial-by-trial variability. Therefore, the input signals that represent such tasks or events that evoke brain activity should be generalized beyond merely the preset paradigms. This issue has been addressed by [53, 128, 129], where the authors have defined the term *activity-inducing* signal, which, as the name suggests, comprises any input signal that induces hemodynamic activity. We will refer to activity-inducing signals as *source signals* in the rest of this chapter, which steers the reader to broader terminology not only used in biomedical signal processing, but also in acoustics and telecommunications [130], and emphasizes that recorded output data are *sourced* by such signals.

An accurate estimation of the HRF is crucial to correctly interpret both the hemodynamic activity itself and the underlying source signals. Furthermore, the HRF has shown

potential as a biomarker for healthy aging [60] or pathological brain functioning; examples of which include obsessive-compulsive disorder [61], mild traumatic brain injury [62], Alzheimer's disease [63], epilepsy [64] and severe psychosocial stress [65]. While HRFs can as well be defined in nonlinear frameworks with the help of Volterra kernels [131], linear models have particularly gained popularity due to the combination of their remarkable performance and simplicity. Several approaches have been proposed in the literature which employ linear modelling for estimating the HRF. The strictest approach assumes a constant a priori shape of the HRF, i.e. a mathematical function with fixed parameters, and is only concerned with finding its scaling (the activation level). The shape used in this approach is usually given by the canonical HRF model [132]. As such, this approach does not incorporate HRF variability, yet the HRF is known to change significantly across subjects, brain regions and triggering events [28, 133, 134]. A second approach is to estimate the parameters of the chosen shape function, which leads to a more flexible solution [66]. Alternatively, HRF estimation can be reformulated as a regression problem by expressing the HRF as a linear combination of several basis functions (which are often chosen to be the canonical HRF and its derivatives). This approach is known as the general linear model (GLM) [135]. Finally, it is also possible to apply no shape constraints on the HRF, and predict the value of the HRF distinctly at each time point. This approach suffers from high computational complexity and variance of the estimated HRFs, which might be of arbitrary or physiologically meaningless forms [136].

Note that the majority of studies which tackle HRF estimation presume that the source signal is known and equal to the EP, leaving only one unknown in the convolution: the HRF [50]. However, as mentioned earlier, a functional brain response can be triggered by more sources than the EP alone. These sources can be extrinsic, i.e., related to environmental events, such as unintended background stimulation or noise artefacts. They might also be intrinsic sources, such as mental imagery. Under such complex and multi-causal circumstances, recovering the rather 'hidden' source signal(s) can be of interest. Moreover, even the EP itself can be much more complex than what a simple binary pattern allows for. Indeed, the hemodynamic response to, for instance, a visual stimulus, can vary greatly depending on its parameters, such as its contrast [34], demanding a continuous variable to represent the "on" times of the stimulus. In contrast to the aforementioned methods, where the goal was to estimate HRFs from a known source signal, there have also been attempts to predict the sources by assuming a known and fixed HRF [52] [53]. However, these methods fall short of depicting the HRF variability.

To sum up, neither the sources nor the HRF are straightforward to model, and as such, when either is assumed to be fixed, it can easily lead to misspecification of the other. Therefore, we consider the problem of jointly estimating the source signals and HRFs from multivariate fUS time-series. This problem has been addressed by [54], [55], [67] and [56]. In [54], it is assumed that the source signal (here considered as neural activity) lies in a high frequency band compared to the HRF, and can thus be recovered using homomorphic filtering. On the other hand, [55] first estimates a spike-like source signal by thresholding the fMRI data and selecting the time points where the response begins, and subsequently fits a GLM using the estimated source signal to find the HRF. Both of the aforementioned methods are univariate: although they analyze multiple regions and/or subjects, the analysis is performed separately on each time series, thereby

ignoring any mutual information shared amongst biologically relevant ROIs.

Recently, a multivariate deconvolution of fMRI time series has been proposed in [56]. The authors proposed an fMRI signal model, where neural activation is represented as a low-rank matrix - constructed by a certain (low) number of temporal activation patterns and corresponding spatial maps encoding functional networks - and the neural activation is linked with the observed fMRI signals via region-specific HRFs. These regions are determined *a priori* by anatomical parcellation, and the HRFs characterizing them are modelled via a single dilation parameter– which may not be enough to capture all possible variations of the HRF, as the width and latency of the HRF are united. The main advantage of this approach is that it allows whole-brain estimation of HRF and neural activation. However, due to the concordance problem associated with brain atlases and inter-subject variations, accurately parcellating the brain poses its own challenges [57]. Moreover, the estimated HRFs are region-specific, but not source-specific. Yet, the length and intensity of stimuli appear to have a significant effect on HRF shape even within the same region, as observed in recent fast fMRI studies [137].

In order to account for the possible variations of the HRF for both different sources and regions, we model the fUS signal in the framework of convolutive mixtures, where multiple input signals (sources) are related to multiple observations (measurements from a brain region) via convolutive mixing filters. In the context of fUS, the convolutive mixing filters stand for the HRFs, which are unique for each possible combination of sources and regions, allowing variability across different brain areas and triggering events. In order to improve identifiability, we make certain assumptions, namely that the shape of the HRFs can be parametrized and that the source signals are uncorrelated. Considering the flexibility of tensor-based formulations for the purpose of representing such structures and constraints that exist in different modes or factors of data [138], we solve the deconvolution by applying block-term decomposition (BTD) on the tensor of lagged measurement autocorrelation matrices.

We evaluate our method on a simulational study and two fUS experiments recorded from mice during visual stimulation. In the first dataset, we track the visual information pathway by investigating the peak latency of the HRF in key anatomical structures involved within the mouse brain's colliculo-cortical, image-forming visual pathway: the lateral geniculate nucleus (LGN), the superior colliculus (SC) and the primary visual cortex (V1). We show that the ordering of the peak latencies agrees with prior works [34], confirming with fUS that visual information first travels through the subcortical targets SC and LGN, before being relayed to V1. In the second experiment we repeatedly display the visual stimuli at 5 distinct locations. We show that our technique is able to extract 5 underlying sources, and the time course of each of these sources have a one-to-one correspondence with the timing of the 5 distinct stimulus locations.

The rest of this chapter is organized as follows. First, we describe our data model and the proposed tensor-based solution for deconvolution. Next, we describe the experimental setup and data acquisition steps used for fUS imaging of a mouse subject. This is followed by the deconvolution results, which are presented in two-folds: *(i)* Numerical simulations, and *(ii)* Results on real fUS data. Next, under discussion, we review the highlights of our modelling and results, and elaborate on the neuroscientific relevance of our findings. Finally, we conclude our work and suggest several future extensions.

4.2. SIGNAL MODEL

Naturally, fUS images contain far more pixels than the number of anatomical or functional regions. We therefore expect certain groups of pixels to show similar signal fluctuations along time. We consider the fUS data as parcellated in space into several regions, represented as an $M \times N$ matrix, where each of the M rows contain the average pixel time-series within a region-of-interest (ROI), and N is the number of time samples.

Assuming a single source signal, a single ROI time-series $y(t)$ can be written as the convolution between the HRF $h(t)$ with length $L+1$ and the input source signal $s(t)$ as:

$$y(t) = \sum_{l=0}^{L} h(l)s(t-l) \tag{4.1}$$

However, a single ROI time-series may be affected by a number of (R) different source signals. Each source signal $s_r(t)$ may elicit a different HRF, $h_r(t)$. Therefore, the observed time-series is the summation of the effect of all underlying sources:

$$y(t) = \sum_{r=1}^{R} \sum_{l=0}^{L} h_r(l)s_r(t-l). \tag{4.2}$$

Finally, extending our model to multiple ROIs, where each ROI may have a different HRF, we arrive to the multivariate convolutive mixture formulation:

$$y_m(t) = \sum_{r=1}^{R} \sum_{l=0}^{L} h_{mr}(l)s_r(t-l) \tag{4.3}$$

where $h_{mr}(l)$ is the convolutive mixing filter, belonging to the ROI m and source r [139]. Note that in this work we consider that the ROIs are known (for instance, via anatomical labelling [55]), or can be predicted from the data as a pre-processing step before deconvolution. We employ the latter approach in this work, and apply independent component analysis (ICA) on the fUS data, for which more details will follow in Section 4.4.

In the context of fUS, the sources that lead to the time-series can be task-related (T), such as the EP, or artifact-related (A). The task-related sources are convolved with an HRF, whereas the artifact-related sources are directly additive on the measured time-series [140]. Artifact sources in general are used to represent fUS signal variation of non-neural origin. Under this definition, we consider physiological noise, e.g. movement of the subject causing signal fluctuations in the entire field-of-view. Moreover, a recent study [87] found out that only the low-frequency content of the fUS signal reflects neural activity. Artifact sources can as well incorporate instrumentation noise, such as thermal or electronic noise (commonly modeled as additive, [20, 141]) introduced by the ultrasound acquisition system, which can be spatially varying, becoming more prominent at deeper areas of the brain. As such, the strength of the effect that an artifact source exerts on a region should depend on the artifact type and the brain region. Hence, each $h_{mr}(l)$ with $r \in A$ should correspond to a scaled (by a_{mr}) unit impulse function (ensuring addi-

tivity). Finally, we rewrite Eq. 4.3 as:

$$\begin{aligned} y_m(t) &= \sum_{r\in T}\sum_{l=0}^{L} h_{mr}(l)s_r(t-l) + \sum_{r\in A}\sum_{l=0}^{L} a_{mr}\delta(l)s_r(t-l) \\ &= \sum_{r\in T}\sum_{l=0}^{L} h_{mr}(l)s_r(t-l) + \sum_{r\in A} a_{mr}s_r(t). \end{aligned} \tag{4.4}$$

We aim at solving this deconvolution problem to recover the sources and HRFs of interest separately at each ROI m.

4.3. PROPOSED METHOD

In this section, we will present the steps of the proposed tensor-based deconvolution method. We will first introduce how deconvolution of the observations modeled as in Eq. 4.4 can be expressed as a BTD. Due to the fact that this problem is highly non-convex, we will subsequently explain our approach to identifying a final solution for the decomposition. Finally, we will describe source signal estimation using the estimated HRFs.

4.3.1. FORMULATING THE BLOCK-TERM DECOMPOSITION

We start by expressing the convolutive mixtures formulation in matrix form. First, the output time-series and source signals are re-organized into block-Hankel matrices $\mathbf{Y}$ and $\mathbf{S}$ of size $ML'\times(N-L')$ and $R(L+L')\times(N-L')$ respectively. More specifically, the columns of $\mathbf{Y}$ and $\mathbf{S}$ contain (lagged versions of) the output and source signals, denoted by $\mathbf{y}(n)$ and $\mathbf{s}(n)$, $n=1,\dots,N-L'$ respectively. Here, the parameter L' controls the size of the tensor of lagged output autocorrelations to be decomposed, regarding which more details will follow later within this section. The column vectors $\mathbf{y}(n)$ and $\mathbf{s}(n)$ are constructed as follows [142]:

$$\begin{aligned} \mathbf{y}(n) = [&y_1(n),\dots,y_1(n+L'-1),\\ &\dots,y_M(n),\dots,y_M(n+L'-1)]^{\mathrm{T}} \text{ and}\\ \mathbf{s}(n) = [&s_1(n-L),\dots,s_1(n+L'-1),\\ &\dots,s_R(n-L),\dots,s_R(n+L'-1)]^{\mathrm{T}}. \end{aligned} \tag{4.5}$$

This way, Eq. 4.3 can be written as $\mathbf{Y}=\mathbf{HS}$, where $\mathbf{H}$ is the mixing matrix containing the convolutive mixing filters in the form of Toeplitz matrices:

$$\mathbf{H} = [\mathbf{H}_1 \quad \dots \quad \mathbf{H}_R] = \begin{bmatrix} \mathbf{H}_{11} & \dots & \mathbf{H}_{1R} \\ \vdots & \ddots & \vdots \\ \mathbf{H}_{M1} & \dots & \mathbf{H}_{MR} \end{bmatrix} \tag{4.6}$$

whose any block-entry $\mathbf{H}_{mr}$ is the Toeplitz matrix of $h_{mr}(l)$:

$$\mathbf{H}_{mr} = \begin{bmatrix} h_{mr}(L) & \dots & h_{mr}(0) & \dots & 0 \\ & \ddots & \ddots & \ddots & \\ 0 & \dots & h_{mr}(L) & \dots & h_{mr}(0) \end{bmatrix}. \tag{4.7}$$

Next, the autocorrelation $\mathbf{R}_\mathbf{y}(\tau)$ for a time lag τ is expressed as:

$$\begin{aligned}\mathbf{R}_\mathbf{y}(\tau) &= \mathrm{E}\{\mathbf{y}(n)\mathbf{y}(n+\tau)^\mathrm{T}\} = \mathrm{E}\{\mathbf{H}\mathbf{s}(n)\mathbf{s}(n+\tau)^\mathrm{T}\mathbf{H}^\mathrm{T}\} \\ &= \mathbf{H}\mathbf{R}_\mathbf{s}(\tau)\mathbf{H}^\mathrm{T}, \quad \forall\tau. \end{aligned} \tag{4.8}$$

Notice that L' determines the number of variables (per region) for computing the autocorrelations $\mathbf{R}_\mathbf{y}(\tau)$ of size $ML' \times ML'$, whereas $N-L'$, the number of columns of $\mathbf{Y}$, corresponds to the number of observations. Assuming that the sources are uncorrelated, the matrices $\mathbf{R}_\mathbf{s}(\tau)$ are block-diagonal, i.e. non-block-diagonal terms representing the correlations between different sources are 0. Therefore, the output autocorrelation matrix $\mathbf{R}_\mathbf{y}(\tau)$ is written as the block-diagonal matrix $\mathbf{R}_\mathbf{s}(\tau)$ multiplied by the mixing matrix $\mathbf{H}$ from the left and by $\mathbf{H}^\mathrm{T}$ from the right. Then, stacking the set of output autocorrelation matrices $\mathbf{R}_\mathbf{y}(\tau)$ for K different values of τ will give rise to a tensor $\mathscr{T}$ of size $ML' \times ML' \times K$ that admits a so-called block-term decomposition (BTD). Eventually, $\mathscr{T}$ can be written as a sum of low-multilinear rank tensors, in this specific case a rank of $(L+L', L+L', \cdot)$ [143]. The decomposition for $R=2$ is illustrated in Fig. 4.1.

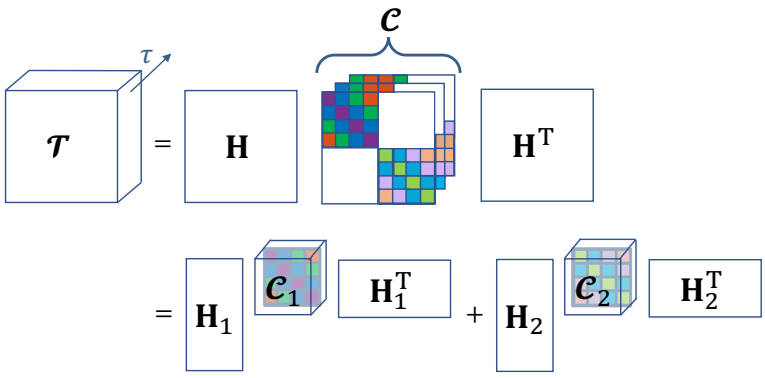

Figure 4.1: Demonstration of BTD for $R=2$. The tensor $\mathscr{T}$ of stacked output autocorrelations $\mathbf{R}_\mathbf{y}(\tau)$, $\forall\tau$ is first expressed in terms of $\mathbf{H}$ and a core tensor $\mathscr{C}$, which stores the stacked source autocorrelations $\mathbf{R}_\mathbf{s}(\tau)$, $\forall\tau$. Each $\mathbf{R}_\mathbf{s}(\tau)$ corresponds to a frontal slice of $\mathscr{C}$ and exhibits a block-diagonal structure with inner Toeplitz-blocks. Note that each slice comes as a lagged version of the preceeding slice. $\mathscr{T}$ is decomposed into R terms, each of which contains a core tensor ($\mathscr{C}_1$ or $\mathscr{C}_2$, representing the autocorrelation of the corresponding source) and a block column of $\mathbf{H}$ ($\mathbf{H}_1$ or $\mathbf{H}_2$).

Due to the Hankel-block structure of $\mathbf{Y}$ and $\mathbf{S}$, $\mathbf{R}_\mathbf{y}(\tau)$ and $\mathbf{R}_\mathbf{s}(\tau)$ are Toeplitz-block matrices. Furthermore, as each frontal slice of $\mathscr{T}$ is constructed using a different time-lag τ in $\mathbf{R}_\mathbf{y}(\tau)$, the slices come as shifted versions of one another (the same shift-structure is valid for $\mathscr{C}$, constructed with $\mathbf{R}_\mathbf{s}(\tau)$'s, as shown in Fig. 4.1). As such, the construction of each whole core tensor ($\mathscr{C}_r$ for the rth source) is based on a single vector, $\mathbf{z}_r$. We will denote the aforementioned transformations (first the formation of a Toeplitz matrix out of $\mathbf{z}_r$, and later shifting this Toeplitz matrix at various lags such that they are placed at different slices of a tensor) using the operator λ, such that $\mathscr{C}_r = \lambda(\mathbf{z}_r)$. Note that the

number of time-lags to be included is a hyperparameter of the algorithm, and we take it as equal to the filter length in this work.

Considering our signal model, where we have defined two types of sources, we can rewrite the block-columns of $\mathbf{H} = [\mathbf{H}_1\ \mathbf{H}_2]$ (Eq. 4.6) simply as $\mathbf{H} = [\mathbf{H}_T\ \mathbf{H}_A]$. Here, $\mathbf{H}_T$ relates to the task-sources, i.e. includes the region-specific HRFs, whereas $\mathbf{H}_A$ includes the region-specific scalings of the artifact sources.

In addition, we impose a shape constraint to the HRFs such that they are physiologically interpretable. To this end, we adapted the canonical HRF model used predominantly in fMRI studies [132] for depicting CBV changes by removing the second gamma function leading to the undershoot response (as similarly applied by [66]). Thus, we expresses the HRF in terms of a single gamma function defined on a parameter set $\boldsymbol{\theta}$:

$$h(t,\boldsymbol{\theta}) = \theta_1(\Gamma(\theta_2)^{-1}\theta_3^{\theta_2} t^{\theta_2-1}\mathrm{e}^{-\theta_3 t}) \tag{4.9}$$

where θ_1, θ_2 and θ_3 control the response height, delay and dispersion of the HRF respectively. By adjusting these parameters, it is possible to model a wide range of HRF shapes, of which some examples are shown in Fig. 4.2. In order to preserve this variety in our solutions, we do not apply any bounds or constraints on the HRF parameters except for non-negativity, i.e. $\boldsymbol{\theta} > \mathbf{0}$.

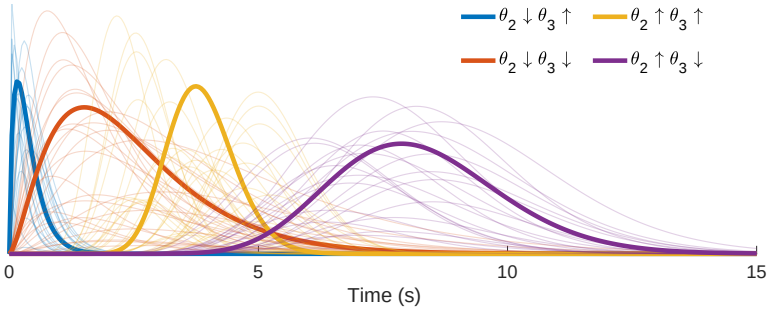

Figure 4.2: Example HRF shapes constructed using Eq. 4.9. In order to illustrate the effect of θ_2 and θ_3 individually, we generated HRF shapes using different combinations of "low range" and "high range" values of each parameter. The low range of θ_2 and θ_3 (indicated by the downward arrow in the legends) were both selected randomly in the interval $[1,3]$, whereas their high range (indicated by the upward arrow in the legends) were selected randomly amongst the intervals $[15,25]$ and $[4,6]$ respectively.

Finally, the BTD is computed by minimizing the following cost function:

$$\begin{aligned} J(\mathscr{C},\boldsymbol{\theta},\mathbf{a}) = \| \mathscr{T} &- \sum_{r\in T} \mathscr{C}_r \times_1 \mathbf{H}_r(\boldsymbol{\theta}_r) \times_2 \mathbf{H}_r(\boldsymbol{\theta}_r) \\ &- \sum_{r\in A} \mathscr{C}_r \times_1 \mathbf{H}_r(\mathbf{a}_r) \times_2 \mathbf{H}_r(\mathbf{a}_r) \|_F^2 \\ \text{s.t. } \mathscr{C}_r &= \lambda(\mathbf{z}_r), \\ \boldsymbol{\theta}_r &> \mathbf{0} \end{aligned} \tag{10}$$

where $\mathbf{H}_r$ shows the rth block column of $\mathbf{H}$ and $\boldsymbol{\theta}_r$ shows the dependency of $\mathbf{H}_r$ (i.e., regional HRFs assigned to the rth source) on its own set of parameters. The operator $||\cdot||_F$ is the Frobenius norm.

We optimize Eq. 10 using the structured data fusion (SDF) framework, more specifically using the quasi-Newton algorithm `sdf_minf`, offered by Tensorlab [144]. In order to overcome the sensitivity of this algorithm to initial point selection [145], we run the BTD several times with random initializations, and use a clustering-based approach to determine the best solution from these runs. In the next section, we will elaborate on our selection procedure.

4.3.2. IDENTIFYING A STABLE SOLUTION FOR BTD

For many matrix and tensor-based factorizations, such as the BTD described above, the objective functions are non-convex. As such, the algorithm selected for solving the non-convex optimization might converge to local optimas of the problem [146]. In order to identify a stable solution, it is common practice to run the optimization multiple times, with a different initialization at each run. Finally, a choice needs to be made amongst different solutions of the decomposition. Unfortunately, the solution with the lowest cost value is not always guaranteed to yield the most meaningful result, especially under strong levels of noise. For instance, when fitting a BTD model to the data, the solution where task components only reflect true sources of interest may produce a higher cost value compared to a solution where the extracted components overfit to noise or artifacts. However, the former solution will be more suitable and interpretable from a biological perspective [147].

For our problem, each BTD repetition produces M HRFs, characterized by their parameters $\boldsymbol{\theta}_m, m = 1,2,\ldots,M$. We follow a similar approach as in the *Icasso* software developed for instantaneous independent component analysis [148], where multiple solutions of the same problem are clustered, revealing that reliable estimates reside in tight clusters, whereas unreliable ones do not belong to any such cluster. Likewise, we use the peak latencies of the estimated HRFs as our features and cluster the BTD solutions. The steps of our clustering approach are as follows:

1. Run BTD 20 times with random initializations, and from each run, store: *(i)* Final value of the cost (i.e., objective) function, *(ii)* M HRFs.

2. Eliminate the P outlier BTD repetitions having significantly higher cost values (We use Matlab's `imbinarize` for the elimination which chooses an optimal threshold value based on Otsu's method [149], as we expect the best solution to be amongst the low-cost solutions)

3. Form a matrix with M columns (standing for the peak latencies of M HRFs, these are the features) and $20-P$ rows (standing for the retained BTD repetitions, these are the observations)

4. Apply agglomerative hierarchical clustering to the columns of the matrix in Step 3

5. Compute the following intracluster distance metric for each cluster as:

$$d_C = \frac{\max_{c_1,c_2 \in C} d(c_1, c_2)}{n_C} \tag{4.11}$$

   where the numerator gives the Euclidean distance between the two most remote observations inside the cluster C (known as the complete diameter distance [150]), and the denominator, n_C, is the number of observations included in C

6. Determine the most stable cluster as the one having the minimum intracluster distance

7. Compute the mean of the estimated HRFs belonging to the cluster of Step 6

To sum up, the clustering approach described above assumes that the best possible solution will be low-cost (Step 2), have low intracluster distance (numerator of Eq. 4.11) and frequently-occurring (denominator of Eq. 4.11). Note that the average run-time for the BTD of, for instance, a $192 \times 192 \times 32$ tensor (representing the lagged autocorrelation tensor of size $ML' \times ML' \times K$ from the first fUS experiment with the following parameters: $M = 3$ regions, $N = 1430$ time points, $R = 2$ sources, $L = K =$ (fUS sampling rate = 4 Hz) $*$ (8 seconds) = 32 samples, $L' = 64$ samples) is ~ 30 seconds. Eventually, repeating the BTD 20 times to reach to a final solution leads to a total run-time of 10 minutes.

After computing the final HRF predictions, the last step is to estimate the sources.

4.3.3. ESTIMATION OF THE SOURCE SIGNALS

The final HRF estimates are reorganized in a Toeplitz-block matrix as shown in Equations 4.6 and 4.7. This gives rise to $\hat{\mathbf{H}}_r$'s ($r = 1, 2, ..., R$), i.e., the block columns of $\hat{\mathbf{H}}$ which contain the estimated convolutive mixing filters that are associated to source r at different regions. After column-wise concatenation, as shown in Eq. 4.6, we can obtain $\hat{\mathbf{H}}$ and estimate the source matrix $\hat{\mathbf{S}}$ by:

$$\hat{\mathbf{S}} = \hat{\mathbf{H}}^{\dagger}\mathbf{Y} \tag{4.12}$$

where $(.)^{\dagger}$ shows the Moore-Penrose pseudo-inverse.

In order to obtain the pseudo-inverse of $\hat{\mathbf{H}}$, we used truncated singular value decomposition (SVD). Truncated SVD is a method for calculating the pseudo-inverse of a rank-deficient matrix, used particularly for extraction of signals from noisy environments [151]. Stabilization of the pseudo-inverse in presence of noise can be achieved by choosing the optimal number of singular values of $\hat{\mathbf{H}}$ to be discarded, which can be viewed as a regularization problem [152]. In this work, we determine this number heuristically. Last but not least, while ideally $\hat{\mathbf{S}}$ would turn out as a block-Hankel matrix, in practice this may not hold. As such, we take the mean of the anti-diagonal terms in each of the R blocks of $\hat{\mathbf{S}}$ to reconstruct back a vector that represents the corresponding source.

4.4. EXPERIMENTAL SETUP AND DATA ACQUISITION

We used two mice (C57BL/6J in the single-stimulus, and B6CBAF1/JRj in the multiple-stimuli experiment; both 7 months-old and male) for the *in vivo* fUS experiments. The

experimental setup depicted in Fig. 4.3. The mice were housed with food and water *ad libitum*, and maintained under standard conditions (12/12 h light-darkness cycle, 22°C). Preparation of each mice involved surgical pedestal placement and craniotomy. First, an in-house developed titanium pedestal (8 mm in width) was placed on the exposed skull using an initial layer of bonding agent (OptiBond™) and dental cement (Charisma®). Subsequently, a surgical craniotomy was performed to expose the cortex from Bregma -1 mm to -7 mm. After skull bone removal and subsequent habituation, the surgically prepared, awake mouse was head-fixed and placed on a movable wheel in front of two stimulation screens (Dell 23,8" S2417DG, 1280 x 720 pixels, 60 Hz) in landscape orientation, positioned at a 45° angle with respect to the antero-posterior axis of the mouse, as well as 20 cm away from the mouse's eye, similar to [24]. All experimental procedures were approved *a priori* by an independent animal ethical committee (DEC-Consult, Soest, the Netherlands), and were performed in accordance with the ethical guidelines as required by Dutch law and legislation on animal experimentation, as well as the relevant institutional regulations of Erasmus University Medical Center.

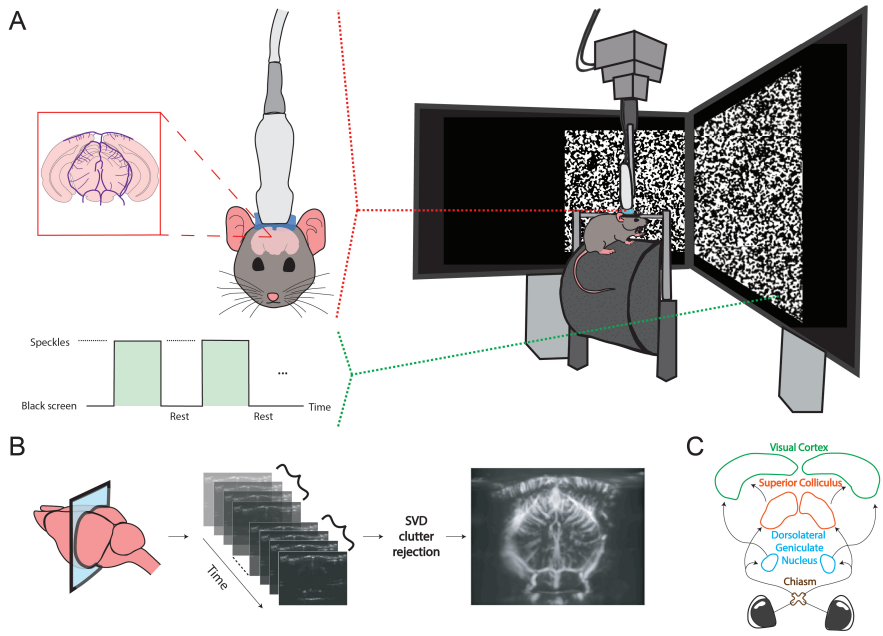

Figure 4.3: The setup and flowchart for fUS imaging of the ROIs. The setup is shown in A, with the awake, head-fixed mouse walking on a movable wheel. During an experiment, either a visual stimulus (here the speckles) or an entirely black screen (rest) is displayed across both monitors. In B, the process of forming a PDI is demonstrated for a coronal slice. First, backscattered ultrasonic waves obtained at different imaging angles are beamformed, resulting in compound images. Next, the compound images are progressed to SVD clutter filtering in batches to remove the tissue motion from the vascular signal. From each batch, a PDI is constructed by computing the power per-pixel. In C, the ROIs that we will focus on in the rest of this work are shown. The pointed arrows represent the signal flow for processing of visual information.

In the first experiment, the visual stimulus consisted of a rectangular patch of randomly generated, high-contrast images - white "speckles" against a black background - which succeeded each other with 25 fps, inspired by [35, 34, 103]. The rectangular patch spanned across both stimulation screens such that it was centralized in front of the mouse, whereas the screens were entirely black during rest. The visual stimulus was presented to the mouse in 20 blocks of 4 seconds in duration. Each repetition of the stimulus was followed by a random rest period between 10 to 15 seconds.

In the second experiment, the visual stimulus was a sinusodial grating presented randomly at one of 5 predefined locations, determined by dividing the mouse horizontal field-of-view (140° [153]) as projected on the screens into 5 equal parts. These locations are displayed in Fig. 4.4. The grating was drifted at 4 degrees per cycle, and at a temporal frequency (cycles per second) of 8.3 Hz. The stimulus and rest duration were fixed at 10 and 15 seconds respectively.

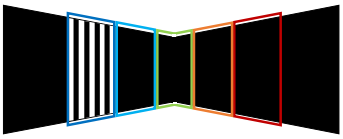

Figure 4.4: Five locations of the grating stimulus displayed on the screens.

Before experimental acquisition, a high-resolution anatomical registration scan was made of the exposed brain's microvasculature so as to locate the most ideal imaging location for capturing the ROIs aided by the Allen Mouse Brain Atlas [154]. For data acquisition, 20 tilted plane waves were continuously transmitted from an ultrasonic transducer (Vermon L22−14v, 15 MHz) at 800 Hz, which was coupled to the the mouse's cranial window with ultrasound transmission gel (Aquasonic). A compound image was obtained by Fourier-domain beamforming and angular compounding, and non-overlapping ensembles were formed by concatenating 200 consecutive compound images. Next, we applied SVD-filtering to denoise the images and separate the blood signal from stationary and slow-changing ultrasound signals arising from other brain tissue. More specifically, SVD-filtering was performed on each ensemble by setting the first (i.e., largest) 30% and the last (i.e., smallest) 1% of the singular values to 0, the former rejecting tissue clutter [20] whereas the latter removing noise [155]. Afterwards, the vascular signal of interest was reconstructed back from the remaining singular components [20]. Images were upsampled in the spatial frequency domain to an isotropic resolution of 25μm, matching to that of the Allen Reference Atlas. Lastly, a Power-Doppler Image (PDI) was obtained by computing for every pixel the power of the SVD-filtered signal over the frames within the ensemble, providing a final sampling rate of 4 Hz for the PDIs. The time-series of a pixel (Eq. 4.4) corresponds to the variation of its power across the PDI stream.

For the selection of ROIs, the experimental data was first parcellated using spatial ICA with 10 components [156]. The components of interest were thresholded to reveal a spatial binary mask standing for an anatomical ROI. To obtain a representative time-series for each ROI, we averaged the time-series of pixels which are captured within the boundaries of the corresponding spatial mask. Finally, the ROI time-series were normalized to zero-mean and unit-variance before proceeding with the BTD.

4.5. RESULTS

To demonstrate the power of our method, the following sections discuss a simulation study and the results of the *in vivo* mouse experiments respectively.

4.5.1. NUMERICAL SIMULATIONS

We simulated three ROI time-series at a sampling rate of 2 Hz, where each time-series was characterized with a unique HRF, i.e., with a different parameter set $\boldsymbol{\theta}$ (Eq. 4.9). We assumed that there are two common source signals that build up to the ROI time-series. The first source signal is a binary vector representing the EP. The EP involves 20 repetitions of a 4-sec stimulus (where the vector takes the value 1) interleaved with 10 – 15 sec of random non-stimulus intervals (where the vector becomes 0). This is the same paradigm that will be used later for deconvolution of the first *in vivo* fUS experiment (Section 4.5.2). The EP is assumed to drive the hemodynamic activity in all ROIs, but the measured fUS signals are linked to the EP through possibly different HRFs. The second source signal stands for the artifact component and is generated as a Gaussian process with changing mean, in accordance with the system noise and artifacts modeled in [157].

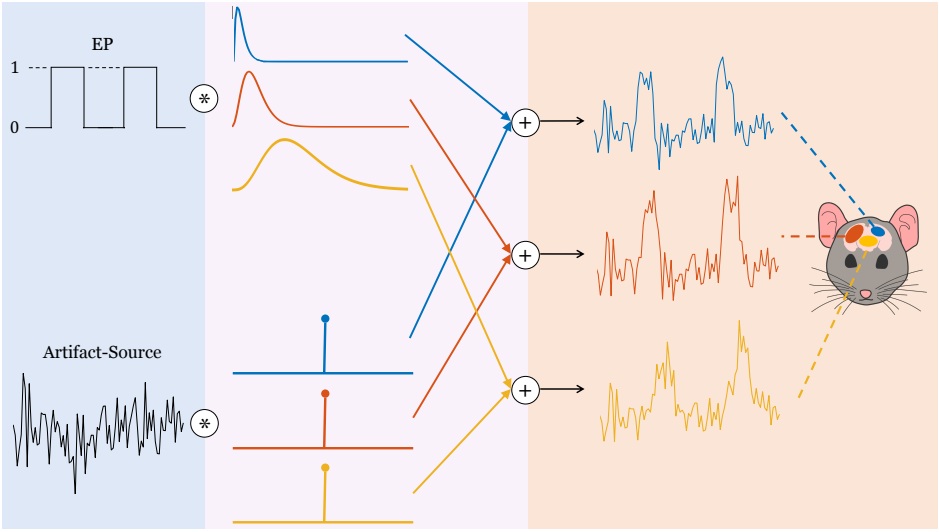

Figure 4.5: Illustration of the simulator. Both of the simulated sources are shown in the left section, one task-related (EP) and one artifact-related. In the middle section, convolutive mixing filters are depicted. The filters which are convolved with the EP are the HRFs, whereas the filters which are convolved with the artifact source only differ by their scaling and modeled as impulses, such that their convolution with the artifact source lead to a direct summation over the time-series. In the last section, the convolved results are added together to deliver the time-series at each ROI.

Each ROI time-series is obtained by convolving the corresponding HRF and the EP, and subsequently adding on the noise source, whose variance depends on the region. The noise variances are adjusted in order to assess the performance under various signal-to-noise ratios (SNRs). The data generation steps are illustrated in Fig. 4.5.

We normalized the time-series to zero-mean and unit-variance before proceeding with the BTD. While solving the BTD, we assumed that there was one task-source of interest and one artifact source. We performed a Monte Carlo simulation of 100 iterations for different SNR values. In each iteration, the HRF parameters were generated randomly such that the HRF peak amplitude was random in the range $[0, 1]$, whereas the peak la-

tency (PL, also referred as the rise time or time-to-peak) and width (measured as full-width at half-maximum; FWHM) of the simulated HRFs varied between [0.25, 4.5] and [0.5, 4.5] seconds respectively. These ranges generously cover the CBV-based HRF peak latencies (reported as 2.1 ± 0.3 s in [87], and between 0.9 and 2 seconds in [66, 158, 159]) and FWHMs (reported as 2.9 ± 0.6 s in [87]) observed in previous mouse studies.

Finally, we defined the following performance metrics at each Monte Carlo iteration:

- The Pearson correlation coefficient of the EP and the estimated source signal, and
- The absolute PL difference (in seconds) between the true and estimated HRFs, averaged for $M = 3$ ROIs.

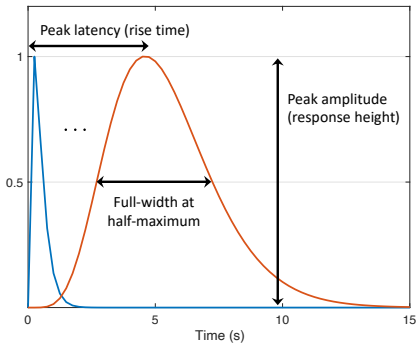

(a) Range of simulated HRFs. The first HRF (blue) has a peak latency and width of 0.25 and 0.5 seconds, whereas the second HRF (orange) has both its peak latency and width as 4.5 seconds respectively.

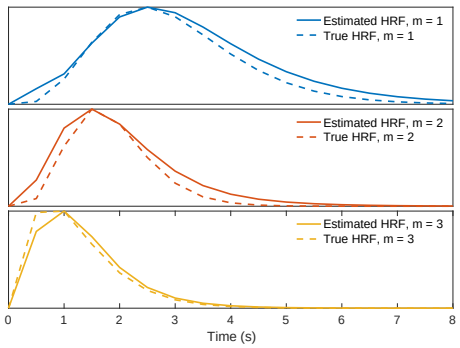

(b) Visualization of the simulated HRFs and their corresponding estimates under 0 dB SNR (from one Monte-Carlo iteration).

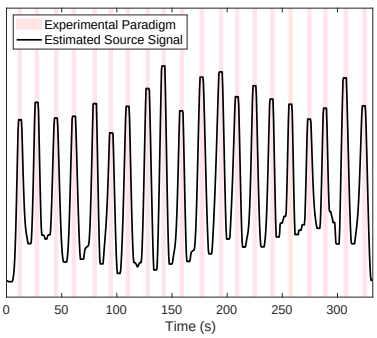

(c) Visualization of the estimated source signal versus the true EP under 0 dB SNR (from one Monte-Carlo iteration).

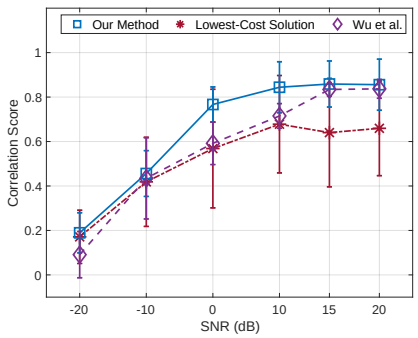

(d) Source estimation performance with respect to SNR. The markers and errorbars denote the mean and standard deviation of the correlation between the true EP and the estimated source signal.

Figure 4.6: Simulation results.

Simulation results are provided in Fig. 4.6. Under 0 dB SNR, the estimated HRFs have an error of 0.3 ± 0.4 (median $\pm$ standard deviation) seconds in the peak latencies across the Monte-Carlo iterations. In addition, we compared our EP estimation performance to: *(i)* When the BTD solution providing the lowest cost value is picked, as opposed to the selection procedure proposed in Section 3.2, and *(ii)* The method proposed by [55]. Note that as the method by [55] is univariate, we computed their average source signal estimate from all the regions for calculating the correlation score with respect to the simulated EP). The results (Fig. 4.6(d)) highlight that the clustering approach proposed for converging to a final solution with BTD yields to a significant increase in the correlation values compared to the lowest-cost solution. The method by [55] performs close to ours at high SNR values (15 to 20 dB), yet, their performance significantly deteriorates as the noise power is increased. In the context of real neuroimaging data, this difference could cause a misinterpretation of the underlying source signals and neurovascular dynamics.

Lastly, we generated an HRF that is not characterizable by the gamma-model given in Eq. 4.9 to see how the proposed method will handle an outlier. In this particular case, we assumed that the HRF shape exhibits a plateau, corresponding to a sustained peak response along time, as shown in the bottom plot of Fig. 4.7. We observed that the proposed method tries to approximate the HRF in the best way possible, while the estimated source still achieves a correlation of 0.72 with the EP under 0 dB SNR. This is slightly less than the mean correlation score (0.77) at this SNR (Fig. 4.6(d)). For the same outlier scenario, the source estimate by [55] has a correlation of 0.62 with the EP.

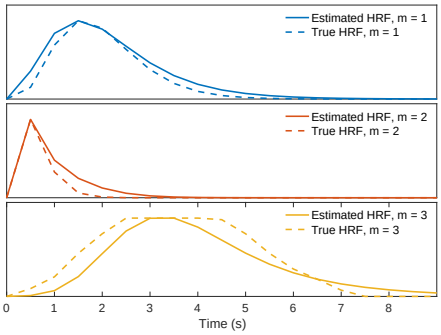

Figure 4.7: Estimated HRFs in case of an outlier HRF shape (bottom plot), under 0 dB SNR.

4.5.2. EXPERIMENTAL DATA

SINGLE STIMULUS

In this experiment, we imaged the mouse brain at two slices (one coronal at Bregma -3.80 mm, and one sagittal at Bregma -2.15 mm [160]) to capture the ROIs that we wished to analyze: SC, LGN and the primary visual cortex (V1). We first applied ICA to select groups of pixels that involve these ROIs as demonstrated in Fig. 4.8(a) and Fig. 4.8(b), showing SC in the former (coronal slice); LGN and V1 in the latter plot (sagittal slice). Note that we needed to image two different slices in order to achieve a good capturing of all the ROIs. Although two of them are presented here, we actually imaged more slices with the same experimental paradigm so as to select the best location for the ROIs. We investigated the time-series of the different slices and established that the responses within the same ROI were reproducible across runs, concluding that jointly decomposing the ROI time-series acquired from different slices is a valid approach.

The raw, normalized fUS time-series belonging to each ROI are displayed in Fig. 4.8(c). By deconvolving this multivariate time-series data, we estimated the region-specific HRFs and the underlying source signal of interest.

We compared our deconvolution results to those by [55]. According to the HRFs estimated by our method (Fig. 4.8(d), top plot), LGN and V1 have a clear lag in time compared to SC. On the other hand, the HRFs estimated by [55] ((Fig. 4.8(d), bottom plot) are observed to precede each other more closely. The ordering of the HRFs is the same in both methods. The difference between the predicted HRFs by two methods likely arises from the fact that [55] estimates a different source signal for each ROI. In other words, as their input signals are assumed different, the impulse responses of each ROI can as well be different compared to when a single common source is assumed, as by our method.

More specifically, [55] offers three source signals for three ROIs, and the source signal in SC is highly aligned with the EP with a correlation coefficient of 0.57, whereas the correlation drops to 0.4 in V1 and 0.16 in LGN. The estimated source signal by our method is assumed to evoke the responses of all ROIs, and has a correlation coefficient of 0.5 with the EP. The low correlation of the source signal in LGN estimated by [55] makes it challenging to decipher its HRF, such as for understanding how swiftly it reacts to the visual stimulus compared to SC or V1. Although a univariate approach might be advantageous in certain situations, when dealing with ROIs that are known to be alerted by the same external or internal stimulus, the presumption that they share a common input signal not only intuitively makes sense, but also makes the interpretation of the HRFs easier. On the other hand, our method aims at finding the best-fitting transfer function between estimated stimulus events and the measurements. This way, the estimated HRFs provide an insight into how fast (by the peak latency) or how long (by the width) a region reacts to a common triggering event. The HRFs estimated by our method point to a peak latency of 1 s in SC, 1.75 s in LGN and 2 s in V1. Similarly, the FWHMs are found as 1.25 s in SC, 1.75 s in LGN and 1.75 s in V1. These results manifest that SC gives the fastest reaction to the visual stimulus amongst the ROIs, followed by the LGN. In addition, the HRF in SC is observed to be steeper than in LGN and V1.

Fig. 4.8(e) demonstrates the source signal estimated by the proposed method. Unlike the simulations, we see that the source signal exhibits a substantial variation in amplitude across time. In order to interpret this behavior of the estimated source signal, we further investigated the raw fUS signals (Fig. 4.8(c)). When the responses given to consecutive repetitions of the stimulus are compared within each region, it can be observed that SC reacts most consistently to the stimulus, while the reproducibility of the evoked responses in LGN and V1 (particularly in V1) are much lower, especially in the second half of the repetitions. To better quantify and compare the region-specific differences in response-variability, we computed the Fano factor (FF) as the ratio of the variance to mean peak amplitude of each region's post-stimulus response [161], defined in a window $[0, 10]$ seconds after a stimulus has been shown. We found an FF value of $0.23, 0.42$ and 0.8 respectively for SC, LGN and V1. These findings indicate that the consistency of the hemodynamic response strength is halved from SC to LGN, and again from LGN to V1.

There are cases where there is a very subtle reaction (as detected by fUS) to the stimulus in V1, such as in repetitions $10, 12, 15, 16$ and 20. These repetitions coincide with the points in Fig. 4.8(e) wherein the most considerable drops in the estimated source

signal were observed. As such, the variability of responses can explain the unexpected amplitude shifts of the estimated source signal.

Due to its changing amplitude, binarizing the estimated source signal using a single global threshold would not work well (Fig. 4.8(e)) for the sake of discovering the exact on- and off- times of the stimulus as found by our method. However, it is still possible to observe local peaks of the estimated source signal occurring around the times that the stimulus was shown. While applying a global threshold can uncover 13 out of 20 repetitions, with a detection of local peaks, this number increases to 19 out of 20 repetitions. After detecting the peaks, we located the time points where for the first time a significant rise (and drop) was observed before (and after) the peak, leading to the starting (and ending) times of the estimated repetitions. Hence, we obtained an estimation of the EP by constructing a binary vector of all 0's with the exception of the time periods in between the predicted starting and ending points.

In Fig. 4.8(f), we compared our EP estimation (averaged across repetitions) with the true EP. We can appreciate that our EP estimation is a slightly shifted (< 0.5 seconds) version of the true EP. In this figure, we also displayed the repetition-averaged responses in SC, LGN and V1; which as well support the HRFs found by our method- with the SC response preceding LGN and V1 by a relatively large separation in time.

Note that the observed trial-by-trial variability in temporal profile across the measured HRs underlines the importance of estimating the source signal. The conventional definition of the EP strictly assumes that the input of the convolution leading to the neuroimaging data (Eq. 4.1) is the same ($= 1$) at each repetition of the stimulus. This would mean that the exact same input, shown at different times, outputs different responses, which would evidence a dynamic system [49, 162]. However, estimating the source signal allows for a flexible characterization of the input, and thus LTI modelling *can* remain plausible. Although extensive analysis of the repetition-dependent behavior of the vascular signal is beyond the scope, we will mention its possible foundations in Discussions.

Lastly, we explored how the estimated source signal and HRFs can be used to generate different correlation images of the brain (Fig. 4.9). Although the active regions do not change between using an optimally-delayed version of the EP ((a) or (d)) or an HRF ((b) or (e)), the maximum correlation value increases slightly when the HRF is utilized. For (a) and (d), the EP was delayed by the value, between 0 to 10 seconds, that provided the highest overall correlation, measured by the mean of non-negative correlations. The optimal delay for the coronal slice (containing the SC) was 0.75 seconds (which is 0.25 s less than the peak latency of the estimated HRF in SC), whereas the one in the sagittal slice was found as 1.75 seconds (which is equal to the peak latency of the estimated HRF in LGN).

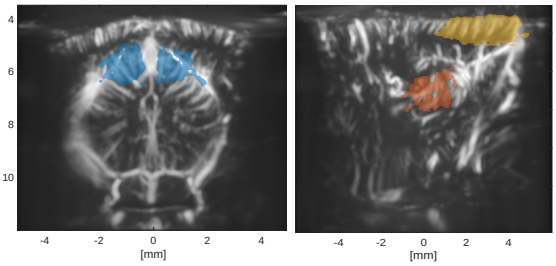

(a) ICA spatial map of SC in the coronal slice.

(b) ICA spatial maps of V1 (yellow) and LGN (orange) in the sagittal slice.

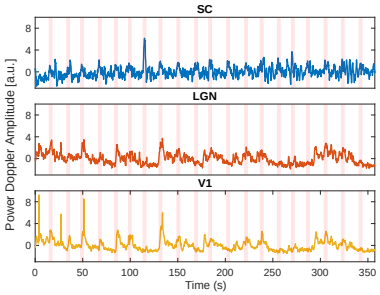

(c) Normalized fUS responses of the ROIs shown in (a) and (b) (EP is displayed in the background).

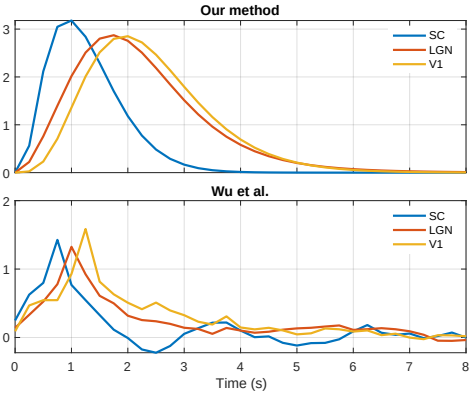

(d) HRFs estimated by our method and by [55].

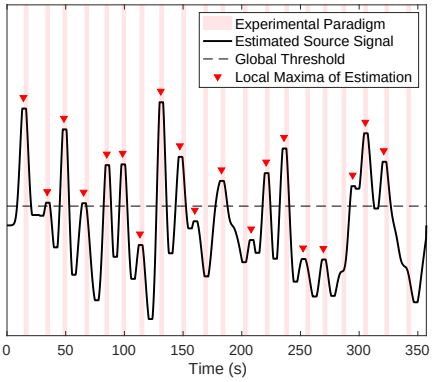

(e) Estimated source signal.

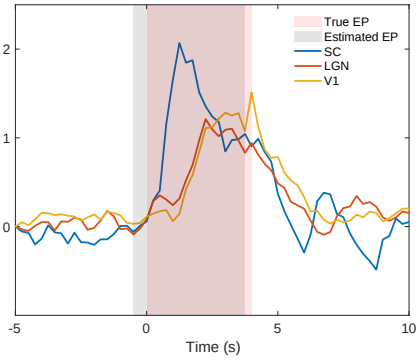

(f) True EP, estimated EP and normalized ROI responses, averaged across stimulus repetitions.

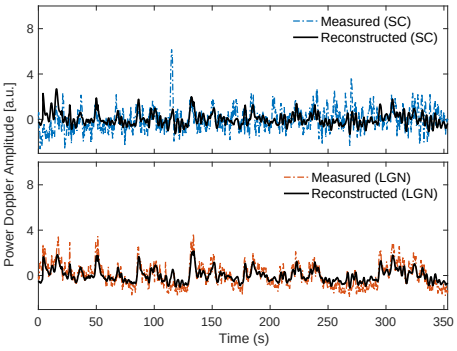

(g) Reconstructed ROI time-series obtained by convolving the estimated HRF (at the corresponding ROI) and the estimated source signal, shown for SC and LGN.

Figure 4.8: Deconvolution results of the first fUS experiment.

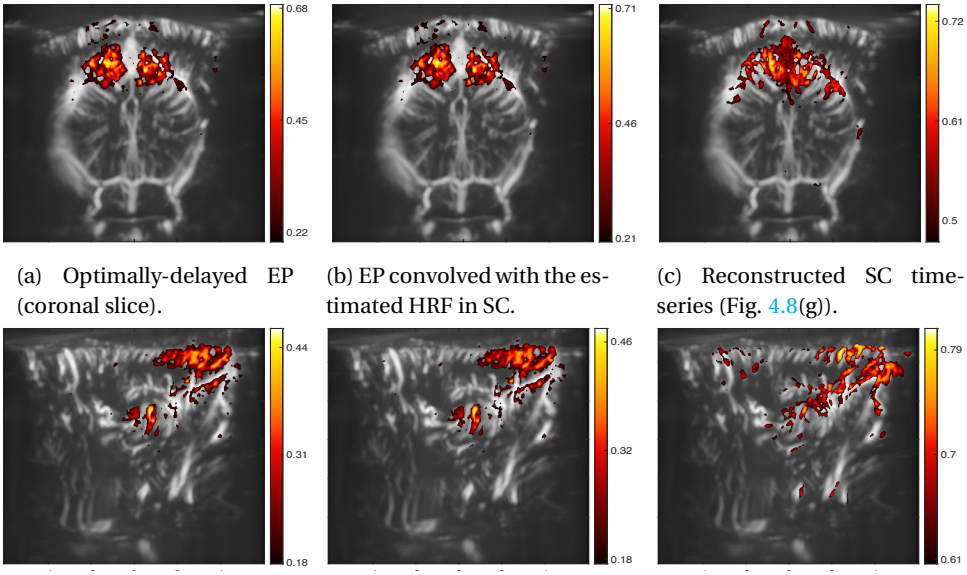

(a) Optimally-delayed EP (coronal slice).

(b) EP convolved with the estimated HRF in SC.

(c) Reconstructed SC time-series (Fig. 4.8(g)).

(d) Optimally-delayed EP (sagittal slice).

(e) EP convolved with the estimated HRF in LGN.

(f) Reconstructed LGN time-series (Fig. 4.8(g)).

Figure 4.9: Correlation images obtained by correlating the pixel time-series with a different signal at each plot. All of the correlation images are thresholded such that only the pixels with a significant correlation value (with a z-score ≥ 2) are displayed. The colorbars denote the value of the Pearson correlation coefficient, and they are scaled up to the maximum value achieved with the corresponding approaches. Note that the approach used in (c) and (f) leads to a similar result as using SC and LGN respectively as a seed region, since the result of the convolution matches well to the average ROI responses, as shown in Fig. 4.8(g).

MULTIPLE STIMULI

In this experiment, we imaged the mouse brain at Bregma -3.52 mm, and defined 5 different stimulus conditions based on the location where the grating stimulus was presented: Leftmost (LM), Slight-Left (SL), Front (F), Slight-Right (SR), and Rightmost (RM). In this experiment our aim was to demonstrate that the proposed method is able to separate multiple underlying sources. More specifically, we assume that the 5 different stimulus locations will evoke hemodynamic activity at different areas within the brain. Therefore, we expect that we can recover the timings of the different stimulus locations as different sources in our model (Eq. 4.12).

We again started by selecting the ROIs by applying ICA. As it can be seen in Fig. 4.10(a), ICA is able to extract all the regions exhibiting a significant correlation with the various stimulus conditions. Next, we ran BTD on the displayed regions while assuming five task-sources and one artifact-source. The estimated source signals are shown in Fig.4.10(b), which reveals that each estimated source signal tends to fit to the timings of one particular stimulus condition, while suppressing the rest of the stimulus conditions. In order to visualize this more clearly, we plotted the estimated sources against the stimulus condition they designate.

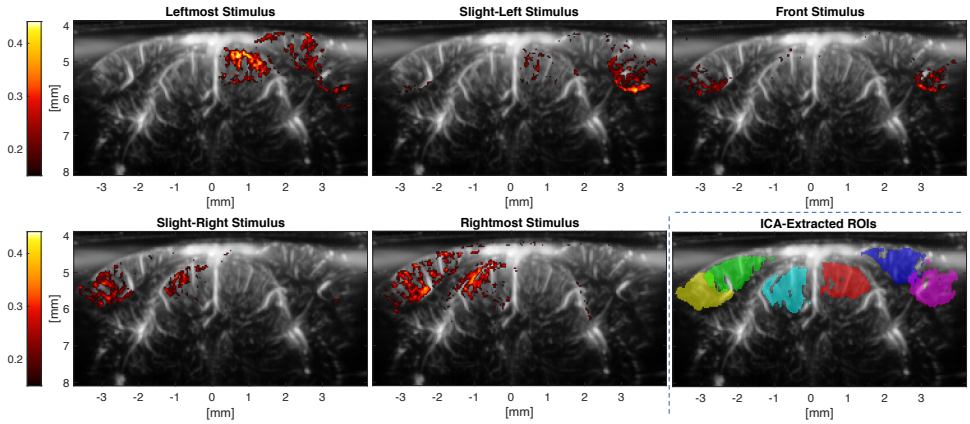

(a) Correlation images for each stimulus location and the ROIs extracted by ICA. The correlation images are obtained by correlating the pixel time-series with an optimally-delayed version of the respective stimulus.

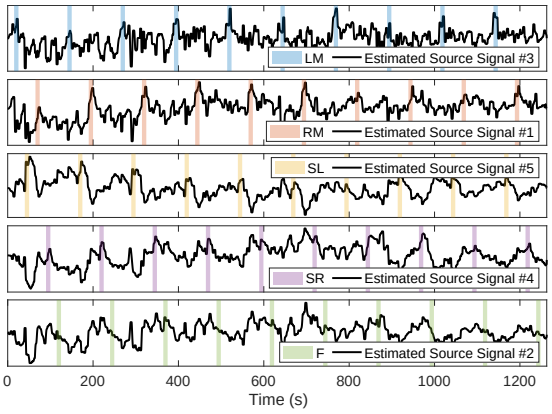

(b) Estimated source signals for 5 task-sources.

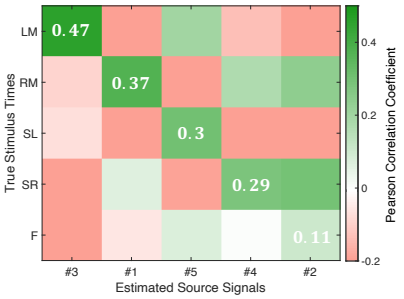

(c) Correlation matrix indicating the Pearson correlation coefficient between each estimated source signal and the onsets of various stimulus locations.

Figure 4.10: Results of the multiple-stimuli experiment.

For a more precise evaluation of the proposed method's performance in this scenario, we computed the matrix given in Fig. 4.10(c), where each column shows the correlation values of one estimated source signal with all the stimulus conditions. Note that an ideal version of this matrix would correspond to an identity matrix (i.e., it would have 1's at the diagonal and 0's at the non-diagonal entries), meaning that each source is matched with the timings of exactly one stimulus location. In other words, as important it is to detect the timings of a stimulus condition via an estimated source (which would be assured by 1's at the diagonal), a perfect separation of sources would also imply that no two or more stimulus conditions exist simultaneously within one estimated source signal (which would be assured by 0's at the non-diagonal terms). For example,

our source estimation results show that the stimuli on the right screen (RM and SR) are more prone to be confused for one another, which is rather understandable, as the correlation images (Fig. 4.10(a)) suggest that the same regions can react similarly to both conditions, making it challenging to discriminate between the two (the same is valid for the stimuli on the left screen as well).

Furthermore, the conditions that were identified most accurately are RM and LM, which are also the ones that provide the best overall correlation images. Similarly, the condition that was identified most poorly, the front location, is also the one that has the lowest correlation with the stimulus. These observations confirm that the more the brain reacts to a stimulus, the better that stimulus can be recovered with the proposed deconvolution approach, and the more regions differ in their reaction to two stimulus conditions, the easier our method can distinguish between them. Last but not least, we compared our source estimation results with [55] in Table 4.1. Our method achieves a better score in both identifying the stimulus conditions, and their separation from each other.

Table 4.1: Results of the multiple-stimuli experiment. We first identified the source signal which achieves the highest correlation with a stimulus condition (ideally, 1). Next, we evaluated how much this source signal correlates with the other stimulus conditions (ideally, 0) by summing up these correlation values. Note that at this step, we considered negative correlation values as 0. We compared our results to those by [55], and the best result achieved by either method is marked as bold for each stimulus condition.

	Highest Correlation		False Correlations	
Stimulus Location	[55]	Ours	[55]	Ours
LM	**0.49**	0.47	0.18	**0**
RM	0.34	**0.37**	0.07	**0.06**
SL	**0.31**	0.3	**0.19**	0.26
SR	0.21	**0.29**	0.33	**0.17**
F	**0.12**	0.11	**0.4**	0.5
Mean Score	0.29	**0.31**	0.23	**0.2**

The proposed method estimates HRFs that are source-specific as well as region-specific. We looked at the HRFs of the two sources that were extracted most accurately, i.e. first (corresponding to RM) and third (corresponding to LM) source signals, and we examined the HRFs estimated at the contralateral regions, i.e. left-SC and V1 for the first source; and right-SC and V1 for the third source. We found a peak latency of 0.8 s in both left and right SC, and 1.1 and 1.4 s in left and right V1 respectively. Although SC in both sides is again found to respond earlier than V1, the difference between the HRF peaks is not as high as in the first experiment, which is also the case for the repetition-averaged fUS responses of the ROIs. This could be the result of displaying different types of visual stimulus in the experiments (sinusodial grating versus speckles), as the HRF is shown to exhibit stimulus-dependent behavior [163, 66].

4.6. DISCUSSION

In this chapter, we considered the problem of deconvolving multivariate fUS time-series by assuming that the HRFs are parametric and source signals are uncorrelated. We formulated this problem as a block-term decomposition, which jointly estimates the underlying source signals and region-specific HRFs. In other words, the proposed method for deconvolution of the hemodynamic response has the advantage of not requiring the source signal(s) nor the HRFs to be specified. As such, it can take into account regional differences of the HRF, as well as recover numerous sources besides the EP, that are unrelated to the intended task and/or outside of the experimenters' control. We investigated the fUS response in several regions of the mouse brain, namely the SC, LGN and the visual cortex, which together compose significant pathways between the eye and brain.

We tested our method in two *in vivo* fUS experiments. In the first experiment, we assumed one common task-source for the ROIs (SC, LGN and V1) and inspected the estimated source signal and HRFs. In the estimated source signal, we observed trial-dependent amplitude variations along time. To better depict this behavior, we investigated the hemodynamic responses in the selected ROIs across repetitions. We noticed that the variability in the visual system increases from the subcortical to the cortical level, confirming our observations in Section 3.3 and prior electrophysiological studies such as [164]. The literature points to eye movements as a source of the high response variability in V1, which can be found in head-fixated awake mice following attempted head rotation [165] and lead to extraordinary alterations stimulus-evoked responses [166]. In addition, our setup allowed mice to run, which can strongly modulate the activity in V1. On this regard, [167] has compared the neuronal response in LGN and V1 in both running and stationary mice, noting that LGN was not particularly affected by running but V1's response more than doubled during motion. On the other hand, the transient responses of both LGN and V1 have also been attributed to lowered levels of arousal [168], which can be considered as a proxy for running. To sum up, prior studies evidence that motion can be an important source of brain activity, even at regions that were previously associated only with visual information processing. Hence, understanding the effects of motion in estimated sources presents itself as an important future aspect of our work.

The SC, LGN and V1 are known to be modulated by the same visual stimuli [169, 34, 163], as evident in the stimulus-correlation maps of our experiment (Fig. 4.9), which led us to assume shared sources for these ROIs in this work. Nevertheless, the true neural activity within each ROI is actually unknown, and is likely to vary across the ROIs due to differences in their regulatory mechanisms, cell type compositions, and neural connections [170], which would partly explain the variations in their hemodynamic responses. While incorporating more sources in our model could account for the regional differences in neural responses, the estimated sources can only be verified with simultaneous measurements of neural activity. However, neurovascular coupling also shows irregularities across the brain as a result of distinct functional (i.e., differences in metabolic demands) and structural (i.e., different vascular architectures) organizations [171], which underscores the importance of allowing the HRF to spatially vary [172].

The HRFs estimated in this work correspond to the transfer functions linking stimuli (shared source(s)) to hemodynamic activity. Decoding this link can enhance our

understanding of cognitive states [173] and sensory processing and perception using hemodynamics-based neuroimaging modalities, such as in discovering how sensory information flows within the brain [163, 174]. Based on the estimated HRFs of our first experiment, we found that SC has the fastest reaction to stimuli, followed respectively by LGN and V1. As V1 does not receive direct input from the retina, but via LGN and SC, its delayed HRF is consistent with the underlying subcortical-cortical connections of the visual processing pathway, as has also been reported by [96, 175, 163]. What's more, the SC's particular aptness to swiftly respond to the visual stimulus aligns with its biological function to indicate potential threats (such as flashing, moving or looming spots [31, 176, 177]).

In the second experiment, we explored the source-detection capabilities of our method further in a more complex setting, with five different stimulus conditions determined by five different locations that the stimulus was presented. The proposed method was able to successfully detect the timings of these locations and differentiate them from one another, particularly those that elicited a stronger and a more discriminative response in the mouse brain.

For both experiments, we considered the number of sources to be known, which is usually not the case in practice. In order to investigate the sensitivity of our approach to this choice, we tried selecting less or more sources in the second fUS experiment. Selecting less number of sources (3 sources) resulted in similar stimuli being grouped together under one source signal. Particularly, one of the estimated sources revealed the timings of both the leftmost and slight-left stimuli, whereas another estimated source reflected both the rightmost and slight-right stimuli. On the other hand, when the number of sources was overestimated, we observed that the estimated sources were more interleaved with one another, resulting in slightly lower correlation values with the true stimulus times (at least while using the same set of parameters). Overall, we conclude that a sub-optimal choice for the number of sources still leads to reasonable results. Nevertheless, a possible future extension of this work would be to explore methods for estimating the number of sources, such as information criterion-based approaches [178].

Compared to our previous BTD-based deconvolution [179], we have made several improvements in this work. To start with, the current method exploits all the structures in the decomposition scheme. In particular, previously the core tensor representing the lagged source autocorrelations was structured to be having Toeplitz slices, yet, these slices were not enforced to be lagged versions of each other. Incorporating such theoretically-supported structures significantly reduced the computation time of BTD by lowering the number of parameters to be estimated. In addition, we increased the robustness of our algorithm by applying a clustering-based selection of the final HRF estimates amongst multiple randomly-initialized BTD runs. Nevertheless, the formulated optimization problem is highly non-convex with many local minima, and the simulations show that there is still room for improvement. For instance, different clustering criteria can be applied for choosing the "best" solution [148]. Furthermore, there are several hyperparameters - namely the HRF filter length, number of time lags in the autocorrelation tensor, number of BTD repetitions - which affect the algorithm's performance as well as convergence speed. For instance, although the run-time of one BTD is not very long, repeating it 20 times to converge to a better solution, as done in our current setting,

considerably increases the overall time. Instead, it is possible to seek for different initialization strategies and run fewer BTDs. Note that the computational complexity of BTD is linear in the number of sources, in the order of the tensor and in the product of the sizes of each mode [145]. The autocorrelation tensor in our solution is of size $ML' \times ML' \times K$, meaning that the computational complexity increases quadratically with the number of ROIs (M) and the number of variables used for calculating the autocorrelations (L'). For a higher number of M, the complexity level can be preserved by lowering L', which could result in a loss of accuracy. In that case, the trade-off between complexity and accuracy should be thoroughly analyzed. All in all, our purpose in this study was to show that the convolutive mixtures modelling and the BTD formulation together offer a deconvolution framework that is able to accommodate and unveil many unknowns regarding hemodynamic activity, yet there is more to discover by conducting more real-life experiments for fully acknowledging the potential of the proposed method.

4.7. CONCLUSIONS

In this chapter, we deconvolved the fUS-based hemodynamic response in several regions of interest along the mouse visual pathway. We started with a multivariate model of fUS time-series using convolutive mixtures, which allowed us to define region-specific HRFs and multiple underlying source signals. By assuming that the source signals are uncorrelated, we formulated the blind deconvolution problem into a block-term decomposition of the lagged autocorrelation tensor of fUS measurements. The HRFs estimated in SC, LGN and V1 are consistent with the literature and align with the commonly accepted neuroanatomical functions and interconnections of said areas. In the meantime, the estimated source signals, whether a single task or multiple tasks were employed throughout the experiment, can be identified successfully in terms of the timings they were presented. Overall, our results show that convolutive mixtures with the accompanying tensor-based solution provides a flexible framework for deconvolution by revealing an elaborate characterization of hemodynamic responses in functional neuroimaging data.

INFORMATION SHARING STATEMENT

The data and MATLAB scripts that support the findings of this chapter are publicly available at https://github.com/ayerol/btd_deconv.

5

Evoked Component Analysis (ECA): Decomposing the fUS Signal with GLM-Regularization

"Simplicity is a prerequiste for reliability."
— Edsger W. Dijkstra

Part of this chapter is published as: A. Erol, B. Generowicz, P. Kruizinga and B. Hunyadi (2024). Evoked Component Analysis (ECA): Decomposing the Functional Ultrasound Signal with GLM-Regularization. IEEE Transactions on Biomedical Engineering. doi: 10.1109/TBME.2024.3395154.

Analysis of functional neuroimaging data aims to unveil spatial and temporal patterns of interest. Existing analysis methods fall into two categories: fully data-driven approaches and those reliant on prior information, e.g. the stimulus time course. While using the stimulus signal directly can help identify the activated brain areas, we showed that the relationship between stimuli and the brain's response exhibits nonlinear and time-varying characteristics (Chapter 3). As such, relying completely on the stimulus signal to describe the brain's temporal response leads to a restricted interpretation of the brain function. In the previous chapter, we proposed a tensor-based technique to jointly estimate the sources that induce brain activity and hemodynamic response functions. This technique estimates both unknowns of the brain function while assuming a convolutive mixing of underlying (uncorrelated) sources, while making no other presumptions on the data. However, in case of stimulus-evoked activity, using the known stimulus time courses can significantly help with the interpretability and scalability of the results.

In this chapter, we present a new technique called Evoked Component Analysis (ECA), which leverages prior information up to a defined extent. This is achieved by including the general linear model (GLM) design matrix as a regulatory term and estimating the factor matrices in both space and time through an alternating minimization approach. We apply ECA to 2D and swept-3D functional ultrasound (fUS) experiments conducted with mice. When decomposing 2D fUS data, we employ GLM regularization at various intensities to emphasize the role of prior information. Furthermore, we show that incorporating multiple hemodynamic response functions within the design matrix can provide valuable insights into region-specific characteristics of evoked activity. Finally, we use ECA to analyze swept-3D fUS data recorded from five mice engaged in two distinct visual tasks. Swept-3D fUS images the 3D brain sequentially using a moving probe, resulting in different slice acquisition time instants. We show that ECA can estimate factor matrices with a fine resolution at each slice acquisition time instant and yield higher t-statistics compared to GLM and correlation analysis for all subjects.

5.1. INTRODUCTION

Functional ultrasound (fUS) is a neuroimaging modality that uses plane-wave irradiation to generate 2D images of changing local blood dynamics [19]. Similar to functional magnetic resonance imaging (fMRI), fUS records hemodynamic changes induced by the local alterations in neuronal activity. In other words, the fUS signal reflects neuronal activity indirectly through the impulse response of the underlying neurovascular system, known as the hemoydnamic response function (HRF) [66].

For studying the relation between brain's hemodynamic response and external stimuli, two methods come forward in the literature. First one of these methods is correlation analysis, where temporal synchrony between the stimulus signal and voxel time series is explored by computing the Pearson Correlation Coefficient (PCC) per voxel. The second method is the general linear model (GLM). GLM can be considered as an extension of correlation analysis which allows to define multiple design variables for modelling the observed neuroimaging data in a multivariate regression problem [180]. Typically, a design variable is calculated as the convolution between a stimulus signal and an HRF, reflecting the expected response pattern when a brain voxel is activated by the stimulus.

GLM can incorporate multiple stimuli and HRF shapes [28] as additional columns in the design matrix. While the former enables the extraction of distinct spatial activation maps per stimulus, the latter models HRF variations across voxels.

Both of these methods are quite intuitive and can provide valuable insight for identification of activated brain areas. However, they both rely on an overly-simplified representation of the stimulus signal. For example, it is known that the brain can habituate to repeated stimuli, resulting in weakened responses over time – a phenomenon known as Repetition Suppression (RS) [181]. These epoch-dependent variations would cause lower correlation values and lower GLM coefficients to be estimated in responsive brain areas, as the overall match between the corresponding voxel time series and the stimulus signal will be reduced [182]. To understand why low activations are observed would require a subsequent step of choosing ROIs (either by anatomical labelling or by thresholding the spatial activation maps obtained by either method) and investigating their time courses thoroughly.

On the other hand, there are also fully data-driven approaches, such as independent component analysis (ICA), which do not make use of any prior information of the stimulus, and returns spatio-temporal signatures for each independent component [183]. Using a data-driven approach is crucial when no prior knowledge is available on the expected activation time courses, such as in case of resting-state experiments [69], or in presurgical electroencephalogram (EEG)-fMRI studies in epilepsy, when no epileptic spikes are observed in the EEG [184]. However, interpreting the large number of resulting independent components requires time-consuming manual classification [185] or subsequent machine learning in order to identify artifact sources and components of interest [186, 187].

Overall, model-based approaches offer an undeniable ease in detection and interpretation of evoked activities of interest, but fall short in capturing sources of variability, whereas the opposite holds for data-driven methods. In this work, we propose a novel method called as Evoked Component Analysis (ECA) to balance between the two extremities by utilizing prior information only as a guiding factor. We achieve this by incorporating the GLM design matrix as a regulatory term in a low-rank decomposition framework instead of enforcing it. This flexibility allows for extracting time courses that follow the stimulus onsets but still get nourished from the measured data. An illustration of ECA is provided in Fig. 5.1.

We apply ECA on mouse-based 2D and 3D fUS experiments. In 2D fUS data, we investigate the impact of algorithmic parameters, namely the rank and regularization coefficient. For 3D ultrasound acquisition, we employ a new technique known as swept-3D fUS [188, 112], where a full volumetric scan of the brain is obtained by moving the ultrasound probe across the exposed brain. In other words, swept-3D fUS does not capture the whole brain at once but instead images it sequentially – resulting in an incomplete data array. We demonstrate how ECA can also be used to decompose such 3D neuroimaging data at the full resolution without relying on slice timing correction techniques. We explore the effect of GLM regularization thoroughly at different intensities and assess the performance of ECA at higher ranks, various stimulus conditions and multiple subjects.

The rest of this chapter is organized as follows. We first explain our problem formula-

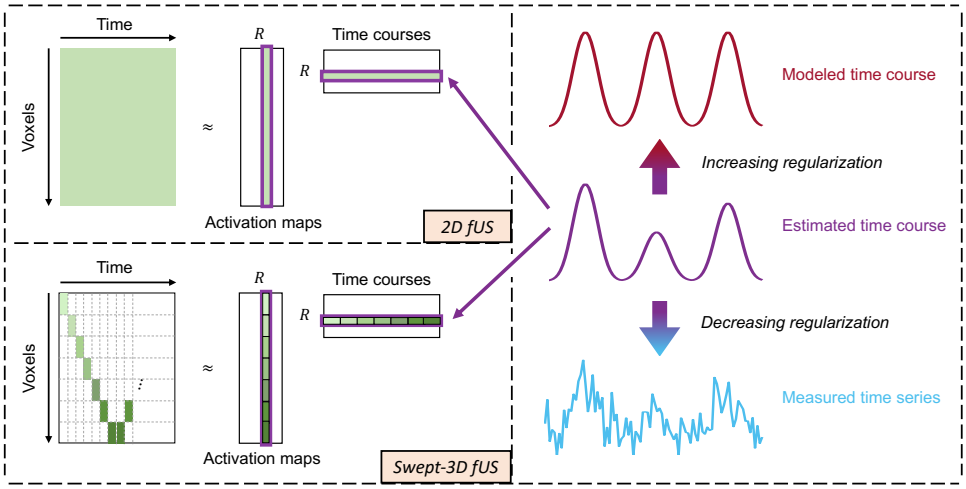

Figure 5.1: Illustration of ECA for decomposing neuroimaging data into its temporal and spatial factor matrices of rank R. Modeled time courses are the design variables, obtained by convolving the known stimulus signal(s) with an HRF. By regularizing the temporal factor matrix to the design matrix, ECA achieves a non-restricting estimation of the brain's evoked response. In case of swept-3D fUS, where slice timings are different as a result of probe motion, ECA can recover the evoked response at the whole resolution without necessitating interpolation to a common time axis.

tion, accompanied by the proposed solution. We then describe the fUS data acquisition pipeline for both standard 2D-fUS and swept-3D fUS. Next, we present our results on *in-vivo* 2D and swept-3D fUS experiments. We compare our findings with correlation and GLM analysis. Finally, we conclude this chapter with discussion and possible future extensions.

5.2. PROBLEM FORMULATION

The objective of this work is to estimate evoked brain activity both in space and time (i.e. its anatomical location and the activity time course), considering that the stimuli inducing this activity are known. However, stimuli might not always translate to brain activity as expected. Indeed, it is known that the relationship between stimuli and the brain's response is quite complex and can exhibit nonlinear and time-varying traits [189, 26]. Hence, we aim to devise a novel decomposition technique that exploits prior information while not completely relying on it.

Matrix decompositions are used for a variety of purposes in neuroimaging, including denoising [190], compression [191] and demixing [192]. In general, decomposition of neuroimaging data reveals underlying patterns of interest that may not be immediately apparent in its raw form [193]. The goal of matrix decompositions is to express the input data matrix $\mathbf{Y}$ as a product of factor matrices $\mathbf{U}$ and $\mathbf{V}$:

$$\mathbf{Y} \approx \mathbf{U}\mathbf{V}^{\mathrm{T}}. \tag{5.1}$$

The column size of $\mathbf{U}$ and $\mathbf{V}$, denoted by R, corresponds to the rank, i.e. number

of components used to approximate $\mathbf{Y}$. In case of neuroimaging data, $\mathbf{Y} \in \mathbb{R}^{N \times T}$ is a space-time matrix composed of time series of all voxels, where N is the total number of voxels and T is the number of time samples. Columns of $\mathbf{U} \in \mathbb{R}^{N \times R}$ and $\mathbf{V} \in \mathbb{R}^{T \times R}$ (also referred to as signatures) store the estimated spatial activation maps and time courses, respectively.

Different decomposition techniques make different assumptions on the factor matrices. For instance, SVD achieves an exact decomposition of $\mathbf{Y}$ by extracting orthogonal factor matrices [71], while ICA attempts to find a set of maximally independent components [74]. While both SVD and ICA are fully data-driven, GLM specifies the factor matrix in time by defining design variables using the known time traces of triggering events (stimuli, motion, etc.) and HRFs. This way, a spatial activation map is estimated for each design variable. By definition, GLM assumes the same response to different repetitions of the same stimulus, although in reality, brain response exhibits dynamic characteristics [194].

5.2.1. 2D fUS

We address the problem of finding evoked activity using a decomposition framework that exploits the GLM design matrix as a regulatory term. The problem that we aim to solve can be expressed as follows:

$$(\hat{\mathbf{U}}, \hat{\mathbf{V}}) = \underset{(\mathbf{U},\mathbf{V})}{\operatorname{argmin}} \|\mathbf{Y} - \mathbf{U}\mathbf{V}^{\mathrm{T}}\|_{\mathrm{F}}^2 + \lambda \|\mathbf{V}_{:,j\in\Psi} - \mathbf{X}\|_{\mathrm{F}}^2 + \eta \|\mathbf{D}\mathbf{V}_{:,j\in\Psi}\|_1, \tag{5.2}$$

where $\mathbf{X} \in \mathbb{R}^{T \times K}$ is the GLM design matrix, K is the number of design variables, $j \in \Psi$ denotes the columns of $\mathbf{V}$ that are associated with the design matrix, $(\cdot)_{\mathrm{F}}$ is the Frobenius norm and $(\cdot)_1$ is the l_1-norm (for matrices, the l_1-norm is calculated as the maximum absolute column sum). The regularization coefficient λ adjusts the influence of design variables on these time courses. The remaining time courses (i.e. rest of $\mathbf{V}$'s columns) are kept unregularized to model the content that can not be pre-defined in a design matrix, such as background hemodynamic activity. In this work, we obtain the design variables by convolving each stimulus signal with an HRF, while in standard GLM, it is quite common to accommodate nuisance regressors in the design matrix as well [180]. These regressors, called as covariates, are used to suppress confounding effects, but are generally of no particular interest by themselves. Our decomposition framework allows for modelling of all such components in the non-regularized column(s) of $\mathbf{V}$, therefore we keep the design matrix simple by including only the stimuli of interest. Lastly, we impose sparsity on the derivatives of the same columns of $\mathbf{V}$ ($j \in \Psi$) using the l_1-norm to avoid sudden temporal changes in the components of interest, which can be introduced by motion. For this purpose, we define the first order difference operator $\mathbf{D}$ [53] and a constant η to control the desired smoothness of the functional response.

5.2.2. Swept-3D fUS

Swept-3D fUS does not image the whole brain at once, instead, captures it sequentially slice-by-slice at each sweep, leading to a slice timing offset (STO) problem. Correction of slice timings is a well-known challenge for fMRI as well, and is often tackled using interpolation techniques to re-align data from all slices to a common reference time axis

[195]. However, the true interpolating function used for transferring neuroimaging data to a new time axis is actually unknown, and it is shown that its selection can have a major impact on further processing [196]. We instead propose to treat the data as it is, considering that one slice is measured at a time. This leads to a data array with missing values as shown in Fig. 5.2(b) [197]. Namely, an $(N_z \times N_x)$ image of a brain slice is acquired at each time instant, where N_z and N_x is the number of voxels in depth and width respectively. Within a sweep, N_y of such slices are acquired with slight shifts in time, creating a 3D image of the brain. By vectorizing the depth-width dimension, the swept-3D fUS data can be represented as an incomplete space-time matrix $\mathbf{Y}$.

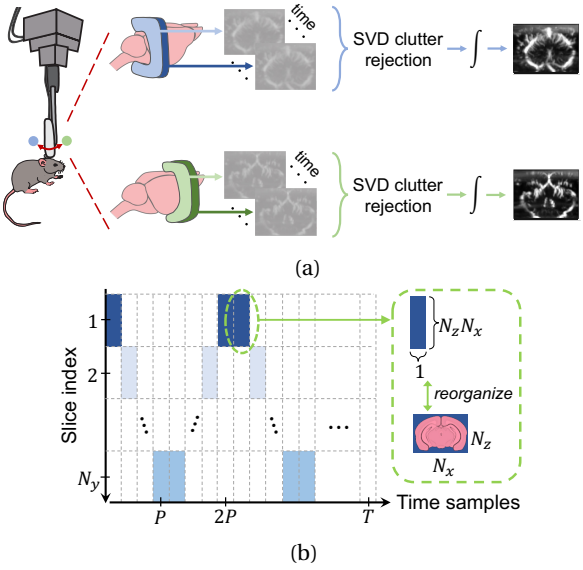

Figure 5.2: Illustration for swept-3D fUS. (a) The ultrasound probe is continuously moved back and forth along the mouse brain. Beamformed echo frames are clutter-filtered with SVD in ensembles and later integrated over the ensemble dimension to create PDIs. (b) Due to the movement of the probe, swept-3D fUS creates an incomplete data matrix. Each colored block stands for an imaged brain slice in vectorized format, whereas the remaining (i.e. blank) blocks refer to points of no acquisition. P denotes the time for one sweep.

Due to missing slice information, the optimization problem given in Eq. 5.2 should be solved only over the observed entries of $\mathbf{Y}$ for swept-3D fUS. This way, ECA also recovers the evoked activity of interest at each acquisition point. The minimization problem for swept-3D fUS data can be expressed as follows:

$$(\hat{\mathbf{U}},\hat{\mathbf{V}}) = \underset{(\mathbf{U},\mathbf{V})}{\operatorname{argmin}} \sum_{(i,j)\in\Omega} [\mathbf{Y}_{ij} - (\mathbf{U}\mathbf{V}^{\mathrm{T}})_{ij}]^2 + \lambda \|\mathbf{V}_{:,j\in\Psi} - \mathbf{X}\|_{\mathrm{F}}^2 + \eta \|\mathbf{D}\mathbf{V}_{:,j\in\Psi}\|_1. \tag{5.3}$$

5.3. PROPOSED METHOD

Although the problems stated in Eq. 5.2 and Eq. 5.3 are not jointly convex in $\mathbf{U}$ and $\mathbf{V}$, they can be reformulated as such by alternating the optimization between the variables. Steps of the proposed two-way alternating minimization (AM) approach are elaborated in Algorithm 1, where ϵ is the error threshold for determining the point of convergence. Note that these steps are written according to swept-3D fUS, where only a subset of $\mathbf{Y}$'s indices are known. For the 2D fUS case, the index set $(i, j) \in \Omega$ refers to the whole matrix. We solve the presented AM scheme using the CVX package in MATLAB.

Algorithm 1: Steps of the proposed AM algorithm.

1: Inputs: $\mathbf{Y}, \mathbf{X}$

2: Initialize: $\mathbf{V}^{(0)}_{:,j\in\Psi} \leftarrow \mathbf{X}, \mathbf{V}^{(0)}_{:,j\notin\Psi} \leftarrow \mathcal{N}(0,1), k = 0, \lambda, \eta, \epsilon$

3: Calculate $\mathbf{U}^{(0)}$, $\mathbf{V}^{(1)}$ by Eq. 5.4 and 5.5 respectively.

4: **repeat**

$$k \leftarrow k+1$$

$$\mathbf{U}^{(k)} \leftarrow \underset{\mathbf{U}}{\operatorname{argmin}} \sum_{(i,j)\in\Omega} [\mathbf{Y}_{ij} - (\mathbf{U}\mathbf{V}^{(k)^{\mathrm{T}}})_{ij}]^2 \tag{5.4}$$

$$\mathbf{V}^{(k+1)} \leftarrow \underset{\mathbf{V}}{\operatorname{argmin}} \sum_{(i,j)\in\Omega} [\mathbf{Y}_{ij} - (\mathbf{U}^{(k)}\mathbf{V}^{\mathrm{T}})_{ij}]^2 + \lambda\|\mathbf{V}_{:,j\in\Psi} - \mathbf{X}\|_{\mathrm{F}}^2 + \eta\|\mathbf{D}\mathbf{V}_{:,j\in\Psi}\|_1 \tag{5.5}$$

until $\dfrac{\|\mathbf{U}^{(k)} - \mathbf{U}^{(k-1)}\|_{\mathrm{F}}^2}{\|\mathbf{U}^{(k)}\|_{\mathrm{F}}^2} + \dfrac{\|\mathbf{V}^{(k+1)} - \mathbf{V}^{(k)}\|_{\mathrm{F}}^2}{\|\mathbf{V}^{(k+1)}\|_{\mathrm{F}}^2} > \epsilon$;

5: Outputs: $\mathbf{U}, \mathbf{V}$

5.4. FUS DATA ACQUISITION

fUS imaging uses angled plane waves sent to the brain through a cranial window. In the regular (i.e., 2D) case, a linear array transducer is used, such that the backscattered signals, that are later beamformed and coherently compounded, constitute a 2D image of a given brain slice. Next, ensembles of adjacent compound frames are formed and SVD-filtered to reject undesired tissue artifacts. Finally, the SVD-filtered frames are integrated over the ensemble to create power-Doppler images (PDIs), whose pixel amplitude varies in proportion to the changes in local cerebral blood volume [17].

In our 2D fUS experiment, we displayed visual stimulus to a mouse (7-months old, C57BL/6J male) in 20 blocks of 4 seconds in duration. Each stimulus epoch was followed by a random rest period of [10,15] seconds. The mouse brain was imaged sagitally at Bregma -2.15 mm and PDIs were sampled at 4 Hz. Details of our imaging pipeline are shared in the Appendix.

Although initially developed for acquisition of a single brain slice, various 3D-volumetric extensions of fUS imaging have been employed since then. One such extension is obtained by concatenating 2D fUS images of different slices by repeating the

same experiment at each slice [24]. Alternatively, a 2D matrix array can be used for ultrasound transmission [25], which requires expensive hardware and acquired volumes suffer from lower sensitivity. Recently, another solution known as swept-3D fUS (Fig. 5.2(a)) has been proposed, which uses a continuously-moving ultrasound probe. The probe is moved continuously back-and-forth over the craniotomy during the experiment, resulting in a 3D volume after every half-cycle of the probe's movement (called as a *sweep*). Eventually, a full PDI volume of the mouse brain (roughly from Bregma -4 mm to +2 mm) is completed in ~ 1 second. Further information on our acquisition setup can be found in the Appendix. For a more comprehensive understanding of swept-3D fUS as an imaging technique, including how the sweeping speed affects the resulting PDIs, we refer the reader to [188, 112].

For swept-3D fUS experiments, we used LED stimuli flickering at 3 Hz during the on-periods [34], which lasted 5 seconds and were followed by a random rest period of [10,16] seconds. The LED brightness level was altered randomly between 3 values (25%, 50% and 100%). Furthermore, the LED was presented to a single eye at each on-period, possibly alerting different regions of the brain depending on which side it was shown to. The experiment was repeated on five adult C57BL6/J mice (12-14 weeks of age).

A film of transparent plastic (TPX) (CS Hyde Company, IL, USA) was used to cover the cranial windows of mice. Acoustic contact between the transducer and the TPX film was ensured by a small layer of milliQ water topped with a layer of ultrasound transmission gel (Aquasonic 100, Parker Laboratories, NJ, USA). All experimental procedures were approved *a priori* by the national authority (Centrale Commissie Dierproeven, The Hague, The Netherlands; license no. AVD1010020197846) as required by Dutch law, and all experiments were performed according to institutional, national, and European Union guidelines and legislation.

Prior to applying ECA, we standardized each PDI voxel time series to zero-mean and unit variance and applied spatial smoothing with a Gaussian kernel with standard deviation of a half-voxel in size. Lastly, we registered the corresponding slices of Allen Brain Atlas on the mean PDIs to locate the regions-of-interest (ROIs) [154].

5.5. RESULTS

We first present our results on 2D fUS data and thoroughly examine the impact of hyperparameter selection in ECA. Next, we apply ECA on swept-3D fUS data and use t-statistics for comparing our findings to correlation and GLM analyses. Note that while ECA allows the observation matrix to encompass missing entries, accounting for slice timing differences illustrated in Fig. 5.2(b), slice measurements should be carried to the same reference time axis for the standard GLM formulation. To that end, we employ sinc interpolation as offered by statistical parametric mapping (SPM) [59] for slice timing correction while applying GLM.

5.5.1. RESULTS ON 2D FUS

We first performed a conventional correlation analysis (Fig. 5.3(a)) on the 2D fUS data, which reveals activations in three ROIs: visual cortex (VIS), lateral geniculate nucleus (LGN) and primary motor area (M1). The PCC image suggests a clear response to the

stimulus in LGN and VIS, and a lower level of activation in M1.

REGULARIZATION INTENSITY

For this part, we assumed one regularized and one non-regularized component. We start with exploring the effect of regularization intensity (λ) to first provide an intuition on the main working principle of ECA. An in-depth analysis for rank selection will follow in the later section.

We applied ECA at 4 different values of λ. The estimated spatial maps and corresponding time courses are shown in Fig. 5.3(b). When $\lambda = 5000$, the method attempts at discovering voxels whose response aligns almost perfectly with the stimulus signal, which is only observed in parts of LGN and VIS with not as high activation levels as in less strict cases of regularization. As λ is decreased, epoch-dependent variations appear in the estimated time course, and activation levels rise in LGN, VIS and M1. LGN and VIS appear brightest at $\lambda = 1000$, meaning that the time course obtained at $\lambda = 1000$ is the most representative for the response of these regions.

When λ is reduced further to 100, activation levels in LGN and VIS drop, while M1 reaches its highest activation. As a result of capturing of the common fluctuations in these regions instead of the stimulus-based design variable, the contrast between LGN/VIS and M1 is almost vanished at $\lambda = 100$. The changing temporal patterns can also be visualized epoch-by-epoch to highlight epoch-specific variations of the brain's functional response (Fig. 5.7). Our results indicate that the response of M1 to the stimulus is not consistent across all epochs, hence activations in this area become more prominent at low regularizations. Notice that even when $\lambda = 0$, the estimated time course can still reflect a link to the stimulus onsets (Fig. 5.8). This link stems from the algorithm's initialization approach, wherein $\mathbf{V}$'s column(s) of interest are initialized directly as the design variable(s). It should as well be noted that the stopping criterion is fixed at $\epsilon = 0.1$ for all results presented in this work, with an example convergence plot provided in Fig. 5.9.

On the other hand, the estimated time course of the non-regularized component (Fig. 5.3(c)) manifests great baseline shifts, and is unrelated from the stimulus signal. Likewise, the associated spatial map exhibits much higher activity in large brain vessels and outside the brain instead of functional brain areas, outlining the brain's vasculature. The non-regularized component was observed to be less affected by varying λ (Fig. 5.10).

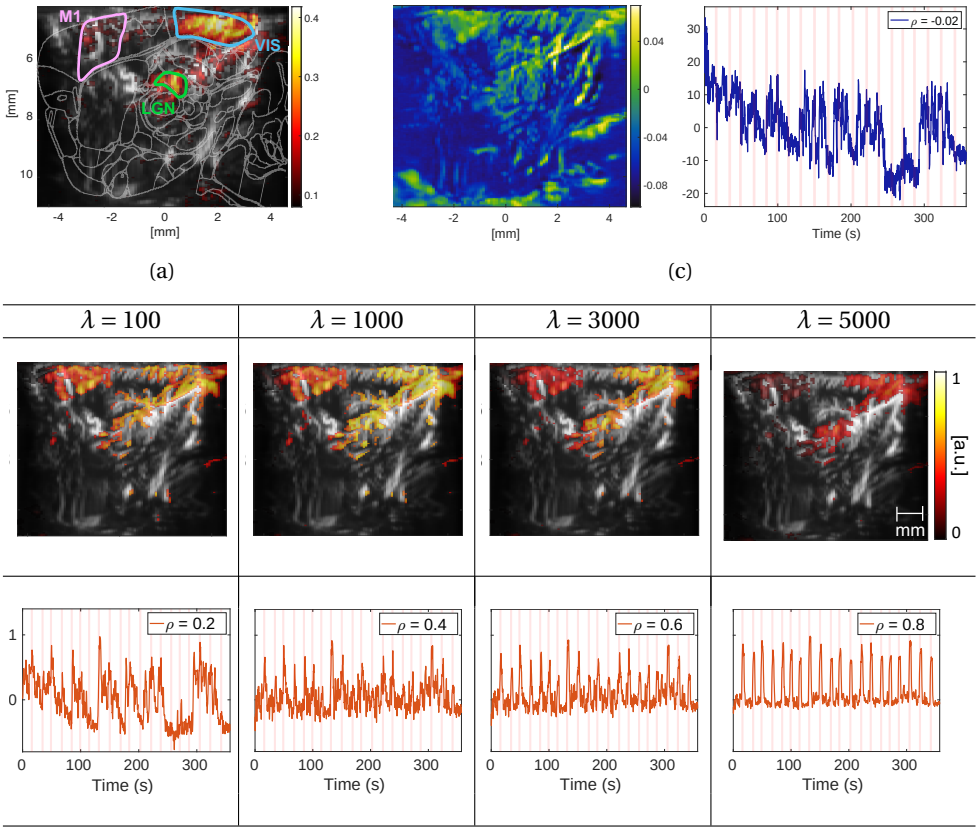

Figure 5.3: Analysis of 2D fUS data. (a) PCC image overlaid against the mean PDI. The ROIs are highlighted over the registered atlas (VIS: visual cortex, M1: primary motor area, LGN: lateral geniculate nucleus). (b) ECA of 2D fUS data with different degrees of regularization. The PCC (ρ) between each estimated time course and the stimulus signal is indicated in parenthesis at the bottom plots. The estimated spatial maps (top plots) display voxels with z-score ≥1.5 overlaid against the mean PDI. Due to the ambiguity in sign and scaling of the estimated factor matrices, we normalized each time course to have a maximum amplitude of 1 and applied a reversed scaling on the corresponding spatial map for a fair comparison of voxel activation levels across various regularizations. Afterward, we used the same color range for all spatial maps. (c) Estimated spatial activation map and time course of the non-regularized component at $\lambda = 1000$. Notice that, z-score thresholding was not applied to the non-regularized component to showcase the vascular structure with full contrast without overlaying of the mean-PDI.

RANK SELECTION

In this section, we will discuss the effect of rank on the estimated factors. When a single regularized component is assumed ($R = 1$), we observed that it is severely affected by noise compared to the results in Fig. 5.3. While adding at least one non-regularized component significantly helps with noise suppression in task-relevant component(s) of interest, introducing more non-regularized terms cause no remarkable change (Fig. 5.11). Hence, we kept the number of non-regularized components as 1 in the results to follow.

Subsequently, we applied ECA with two task-relevant components (both regularized at η=10 and λ=200, selected heuristically) and one non-regularized component. Notice that, we reduced the value of λ here compared to its ideal value found in rank-2 analysis. We observed that keeping λ in the order of thousands in the high-rank case outputs time courses that are very close to what was hypothesised in the design matrix (Fig. 5.13). Increasing the rank allows for a more detailed representation of the data through the **UV** factorization, thus part of the cost function corresponding to the data reconstruction error is reduced. Hence, for preserving the balance between data-fitting (first term in Eq. 5.2) and model-fitting (second term in Eq. 5.2), λ was decreased as well [198].

For the newly introduced task-relevant component, we utilized a different HRF shape in the design matrix. The results are provided in Fig. 5.4. To start with, we can observe the voxels that were mostly captured together in the lower-rank case being divided into two groups. Particularly, M1 and the vessels surrounding VIS and LGN are revealed to exhibit a more dynamic response to the stimulus, resulting in their separation from VIS and LGN. These results are in accordance with our observations in rank-2 analysis, where the highest activation of M1 was detected in the low regularization case, with again a PCC of 0.2 shared with the stimulus signal. The non-regularized component (Fig. 5.12) similarly highlights the brain's vasculature as in the rank-2 case. However, the activations seem to have receded from the pronounced influence of M1 and blood vessels surrounding VIS, since now their stimulus-evoked content is captured by the second HRF. Moreover, almost the entirety of the response in VIS is attributed to the regularized components, resulting in activation values near 0 for VIS in the non-regularized component.

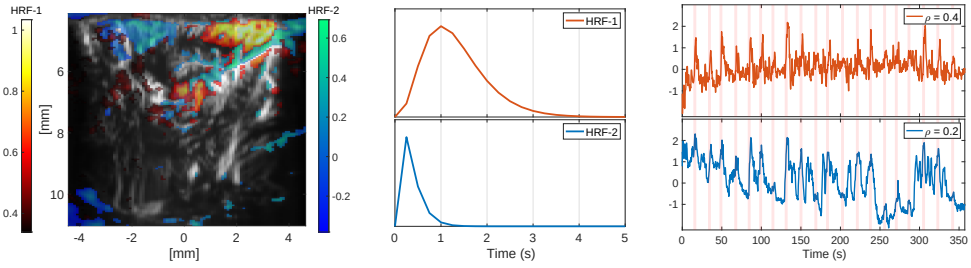

(a) Estimated spatial maps of the regularized components.

(b) HRFs used in the design matrix.

(c) Estimated time courses of the regularized components.

Figure 5.4: Decomposition of 2D fUS data at $R = 3$ with two regularized components. Only the voxels that are found significantly active (z-score ≥1.5) for eact spatial component are displayed in (a). The newly introduced regularized component is shown in blue color.

5.5.2. RESULTS ON SWEPT-3D FUS

Our goal in this experiment was to observe if ECA can capture activity evoked by different tasks, associated with the position of the LED stimulus (either left or right side). For this purpose, we constructed two design variables based on the stimulus onsets of each task. We assumed a total rank of 3, containing two task-regularized components and one non-regularized component to model spatially structured fluctuations due to background activity or noise. We set $\lambda = 1000$ and $\eta = 10$. Our results from one subject are displayed in Fig. 5.5 (time courses and corresponding spatial maps of the GLM-regularized components). Although we did not discover any significant brightness level-dependent changes in the magnitude of estimated time courses, we observed an overall decreasing trend in their magnitudes across epochs (Fig. 5.5(a)). Compared to the 2D fUS experiment with on-screen stimuli, we found the LED stimulus to elicit a more consistent brain response, with less epoch-specific highs or lows.

For assessing the statistical significance of voxel activation levels estimated with ECA as opposed to standard analyses, we checked the mean of top 20 t-statistics [196] achieved at the ROIs, namely superior colliculus (SC), LGN and VIS, and a control region. We selected the control region outside of the functional areas to ensure that the higher statistics obtained in the ROIs by either method is indeed a result of the task-evoked activity in the ROIs. The t-statistic evaluates the accuracy of a model fit by computing the standard error (SE) based on the variance of the residual between the modeled time series and measured data [199].

Specifically, for each voxel in the brain, the null hypothesis asserts that there is no effect of the predictor variables (i.e., columns of the temporal factor matrix) on the response variable (i.e., voxel time series) by assuming that the true coefficient value for the predictor variable is zero, generating a high residual. For GLM (with sinc interpolation), this residual corresponds to the difference between the modeled design variables and voxel time series. For Pearson correlation, the t-statistic is given by $t = \frac{\rho\sqrt{n-2}}{1-\rho^2}$, where n denotes the sample size [200]. For ECA, the predictor variables correspond to the estimated time courses, which are regularized by, but not solely based on stimulus information. As such, ECA time courses constitute a much better fit for the actual data, resulting in very low SE values and much higher t-statistics for all subjects, as shown in Fig. 5.6.

To further support our conclusions, we provide a visual comparison between ECA and the PCC maps obtained from the same subject (Fig. 5.14) by computing the difference in voxel activations (normalized to unit energy) as estimated by both methods in Fig. 5.15. These results confirm that the activations in regions of interest found by ECA are stronger than those found by PCC. In addition, we share the results of ECA on swept-3D fUS data at lower λ in Fig. 5.16, which aligns with our prior observations in the 2D fUS case as the estimated factor matrices in both space and time absorb more of the task-irrelevant content. Consequently, the contrast of the regions of interest in the spatial activation maps is somewhat diminished compared to Fig. 5.5. However, even at lower or higher values of λ, ECA retains higher t-statistics than conventional methods in regions of interest in contrast to the control region, indicated by the ratio between activations in ROIs and the control region as provided in Fig. 5.17, preserving its favorable stance for analyzing evoked components within the brain.

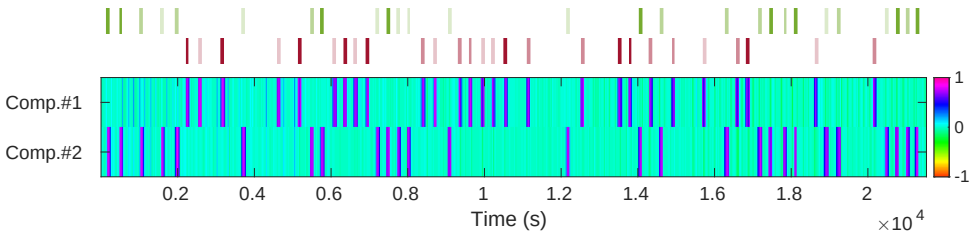

(a) Estimated time courses of the two regularized components. The red and green colored bars on top indicate the left and right-side LED stimuli respectively, with different shades referring to the LED brightness level. The colorbar denotes the amplitude changes of the time courses.

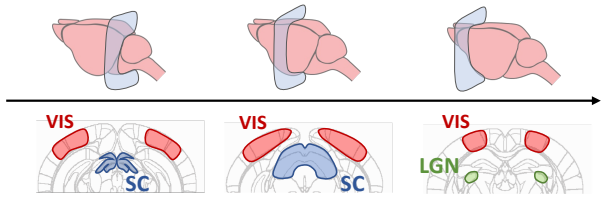

(b) Illustration of the probe's movement and evolution of ROIs within the mouse visual pathway.

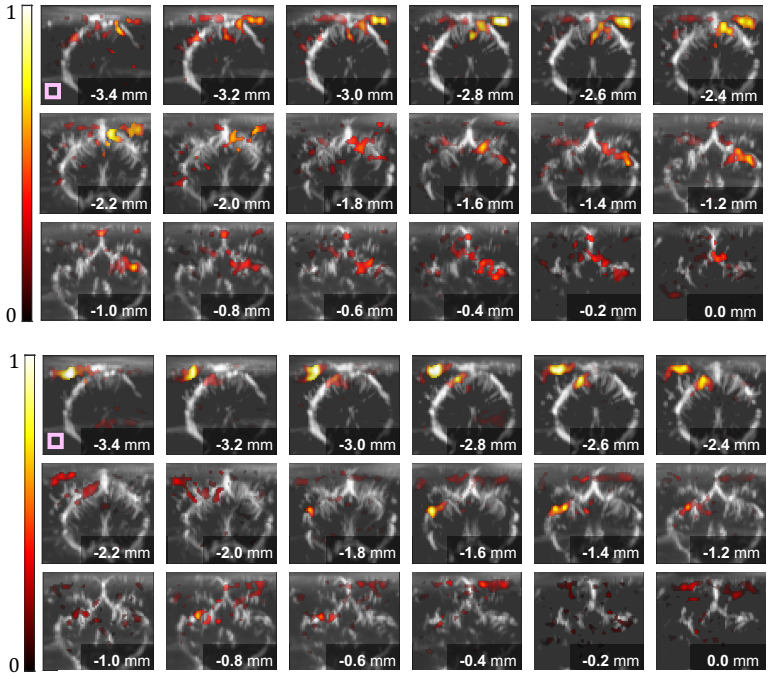

(c) Estimated spatial maps corresponding to left-side (top) and right-side (bottom) stimuli.

Figure 5.5: Estimated time courses (a) and associated spatial maps (c) for one subject (N=1). The ROIs shown over the mouse brain atlas (b) are found to be activated in (c), where only voxels with z-score ≥ 1.5 are displayed against the mean PDI. The selected control region is shown only in the first slice, but was in fact defined over the 3D brain by repeating the same mask at every slice.

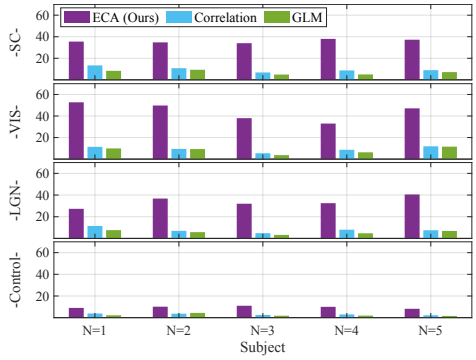

Figure 5.6: Mean of top 20 t-statistics per region for all subjects.

5.6. DISCUSSION

In this study, we introduced ECA as a novel decomposition technique for analyzing neuroimaging data. Existing decompositions in the literature are typically either data-driven or completely reliant on prior information. This prior information is commonly given by the stimulus time course that entails when a stimulus is shown, represented as a boxcar function. However, the brain response may not always vary in sync with stimuli due to various factors including neuronal adaptation and stimulus expectation [201]. In our approach, we use the stimulus time course only as a guiding term, striking a balance between expectation (design variables defined based on stimuli) and reality (measured data). This balance is controlled through regularization of the temporal factor matrix with respect to the design variables.

We employed ECA on 2D fUS data and investigated the impact of regularization. Assuming a rank of 2, we observed that reducing the influence of prior information leads to a higher amount of deviation from the design variable in the estimated time course of interest. For instance, we discovered that responses to certain epochs were notably lower in magnitude compared to others. This dynamic behavior was more pronounced in M1 compared to VIS or LGN. As such, when using stricter regularizations, VIS and LGN appeared more prominently in the estimated spatial maps, whereas reducing regularization led to higher activation levels in M1. It is worth noting that the regularized component was primarily responsible for modeling the stimulus-evoked response, while the non-regularized component effectively captured shared fluctuations unrelated to the stimulus. These effects highlight the brain's vasculature with accumulated activity on the vessels. Next, we incorporated another regularized component using a different HRF, which provided a deeper insight into evoked activity by distinguishing voxels that react differently to the stimulus, i.e. with lower epoch-consistency. From this perspective, we can conclude that the choice of rank depends on the application. In order to observe evoked activity in a collective but summarized manner, meaning that small differences in various voxel responses might be represented together in a single source, choosing $R = 2$ with one regularized column would be reasonable. On the other hand, for further categorization of the task-induced content, the number of regularized com-

ponents should be increased. Similarly, to achieve a more intricate examination of the background hemodynamic activity, the number of non-regularized columns should be elevated.

Subsequently, we applied our method on swept-3D fUS data that is subject to a similar slice timing offset problem as in fMRI. Instead of resolving the problem by interpolation (as commonly done in fMRI), we proposed to treat the data as it was acquired and implicitly recover the evoked activity at each acquisition point. We showed that ECA more precisely describes temporal and spatial brain responses than standard approaches. To elaborate, we experimented on five mice and compared the resulting t-statistics to those obtained via GLM (with sinc interpolation, [59]) and correlation analysis. While applying ECA, we assumed a rank of 3, where the first component was regularized according to stimuli presented to the left eye, the second component towards stimuli shown to the right eye, and the last component was not subject to any regularization. We observed that the estimated spatial maps showed clear differences in the activated areas associated with left and right side LED stimulus. Our approach yielded significantly higher t-statistics in the visual processing pathway of the mouse brain. Note that low-rank completion has been proposed for recovery of data points corrupted with artifacts and slice-timing correction before for fMRI [202]. This approach is based on the assumption that the temporal signal of a voxel at any time point can be expressed as a linear combination of its previous samples, and that these linear weights are shared within all voxels of functionally connected regions. To the best of our knowledge, ours is the first study to achieve recovery of evoked activity at the full scale, i.e. for all slice acquisition points using only the known stimulus onsets. This perspective can be beneficial for future work on task-based swept-3D fUS and fMRI studies, since it does not require selecting an appropriate interpolating function, which is shown to have a significant effect in processing of the data [203].

Notice that ECA necessitates the use of a design matrix as input for regularization, which in turn requires the specification of an HRF shape (or shapes). Indeed, similar specifications are unavoidable for GLM and correlation-based analyses as well. To elaborate, standard GLM also requires a design matrix with user-defined HRF shapes, commonly picked as only the canonical form or canonical form with its derivative(s) [204]. For correlation analysis, the stimulus time course (or if known, neuronal activity) is either subjected to a certain amount of delay or convolved with an HRF [45], which is again a decision left to the user. In fact, ECA provides more flexibility than GLM or correlation analyses since it does not absolutely depend on the given prior information, including the pre-specified HRF shape(s). At the same time, it is worth emphasizing that the results of ECA are influenced by the regularization intensity, λ. While we suggest $\lambda = 1000$ as a point of reference, its choice is indeed data and purpose-dependent, such as whether the application requires a more model-based (high λ) analysis or a data-driven (low λ) one. For example, while Fig. 5.17 reveals a lower contrast in regions of interest compared to the control region at lower λ, the better capturing of shared fluctuations also result in a more accurate portrayal of the brain's activity along time for all regions. In other words, each stage of Fig. 5.3 provides us with valuable insights into evoked and inherent activity of the brain, yet the full picture comes together when inspecting the data at various λ values, which might be favorable for exploratory analyses. Although the optimal λ

is indeed purpose-dependent and influences the relative contrast between interest and control regions, note that ECA, regardless of the choice of λ, provides better contrast than alternative methods (Fig. 5.17).

Another important aspect to address is the computational efficiency of the algorithm. Unlike GLM or correlation analysis, which assume known temporal regressors and only predict their spatial counterparts, ECA estimates time courses as well. This flexibility allows for capturing the time-varying characteristics of the brain response, but comes at the expense of computational time. The estimation of spatial activation maps in 2D fUS takes up a negligible amount of time, as Eq. 5.4 becomes equivalent to least-squares. On the contrary, due to missing entries within our swept-3D fUS matrix formulation, again a point-wise minimization is required in space. Ultimately, the complete decomposition of the swept-3D fUS data (of size 71280×21402) requires approximately half an hour when executed using MATLAB 2021a on a high-performance computing system running Linux, equipped with two AMD EPYC 32-Core Processors and 528 GB of memory. To enhance the algorithm's efficiency, additional structures can be imposed in space or time. For instance, assuming that the estimated spatial maps can be reconstructed from low-rank factors in each space dimension (depth, width and height) can help speed up the process [79, 205].

5.7. CONCLUSION

In this work, we aimed at closing the gap between two sets of approaches used in analysing neuroimaging data: those that are data-driven and those that are completely dependent on existing prior information of the stimulus time course. Within ECA, we use prior information only as a regulatory term through the GLM design matrix, which allows for an informed yet flexible characterization of the brain's response. We showed that ECA can model epoch-dependent changes of the underlying hemodynamic response in 2D fUS data. We employed various regularization strengths for identifying how strong each voxel responds to the stimulus. We demonstrated that the non-regularized component models the global fUS signal, unveiling the brain's vascular structure. We showed that we can further distinguish the response of ROIs by extending the design matrix with a new HRF. Finally, we used ECA to extract task-relevant content of interest from swept-3D fUS data at the full resolution, accounting for slice timing differences that occur as a result of the sequential imaging of the 3D brain. We observed that ECA significantly reduces the standard error between the modelled time courses and measured data, resulting in higher t-statistics for all subjects than conventional analyses.

APPENDIX

2D fUS

We used a high-frequency linear array transducer (Vermon L22 – 14v, 15 MHz), coupled to the subject's cranial window with ultrasound transmission gel (Aquasonic) and to a Verasonics acquisition system. For the 2D fUS experiment, 20 tilted plane waves (±15°) were transmitted. After Fourier-domain beamforming and angular compounding, ensembles of 200 compound images were formed. SVD-filtering was performed on each ensemble by setting the first (i.e., largest) 30% and the last (i.e., smallest) 1% of the singular values to 0, the former rejecting tissue clutter [20] whereas the latter removing noise [155]. PDIs were produced at a final rate of 4 Hz.

The visual stimulus consisted of a rectangular patch of randomly generated, high-contrast white speckle images against a black background, succeeding each other with 25 frames per second [103]. The rectangular patch spanned across two stimulation screens placed 20 cm away from the mouse's eye (Dell 23,8" S2417DG, 1280 x 720 pixels, 60 Hz), such that it was centralized in front of the mouse, whereas the screens were kept entirely black during the rest periods [106].

SWEPT-3D fUS

The trajectory of the ultrasound probe was controlled by a motorized stage (Zaber X-LDA025A). Plane waves were transmitted at 8 angles (±9°). The ensemble size for applying SVD filtering and generating PDIs was selected according to spatial and temporal constraints that allow for an assumption of stationarity in both domains. In contrast to 2D fUS, where ultrasonic frames are acquired over time at a fixed location, swept-3D fUS acquires frames which shift continuously both in space and time. The integration process necessary for calculating PDIs inherently assumes that substantial changes do not transpire within an ensemble. On the other hand, the less frames are used for integration to ensure stationarity, the more PDIs suffer from poor signal-to-noise ratio (SNR). Using ECA, we also aim at mitigating the effect of high noise levels through low-rank decomposition using prior information [206]. In the end, we assumed a slice thickness of 0.2 mm for generating PDIs [19]; corresponding to a time window of ~ 0.1 s towards the ends of the brain, and ~ 0.02 s in the middle section. Overall, the resulting data matrix $\mathbf{Y}$ is of size 71280×21402 $\big((N_z = 60) \cdot (N_x = 66) \cdot (N_y = 18) \times (N_y = 18) \cdot (N_t = 1189)\big)$, with a missing value ratio of 94.44%, as only $1/N_y$ of the matrix entries are available (Fig. 5.2(b)).

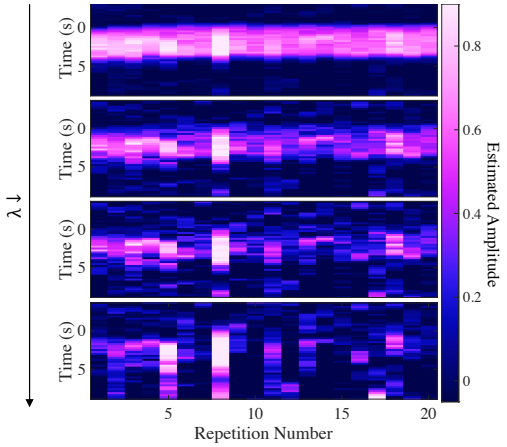

Figure 5.7: ECA results of the 2D fUS experiment with one regularized and one non-regularized component. This figure uses the same time courses as in Fig. 5.4, however, here the time courses are segmented into windows of 15 seconds for an enhanced visual comparison of the results obtained with different λ values. Each segment starts 5 seconds before a stimulus onset (marked by 0 s on the y-axes). The top plot displays the response that is closest to the modeled design variable (high λ). The more λ is decreased, the more epoch-specific variations appear.

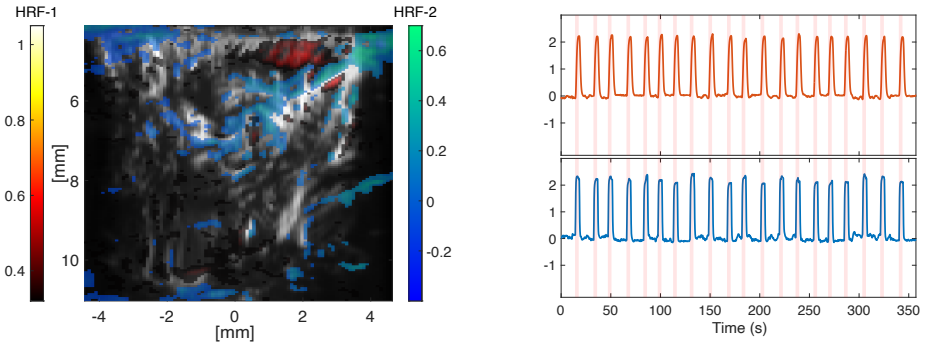

(a) (b)

Figure 5.13: Estimated spatial map (a) and time course (b) of the regularized components at $R = 3$ (2 regularized components) at $\lambda = 1000$.

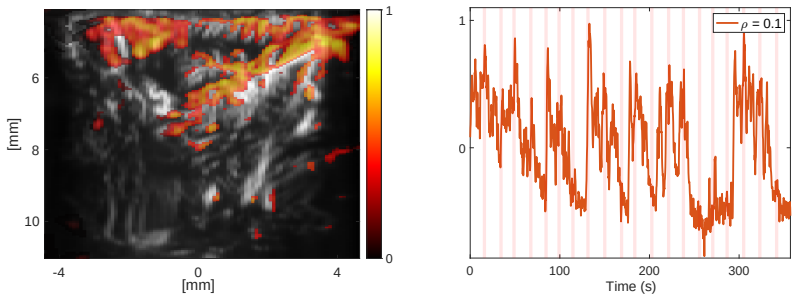

(a) (b)

Figure 5.8: ECA results at $R = 2$ with one component initialized as the design variable and the other component initialized randomly. Neither component was regularized ($\lambda = \eta = 0$). The activation map and time course of the first component are shown in (a) and (b) respectively. Although no regularization was applied, these signatures faintly associate to the stimulus due to the initialization strategy.

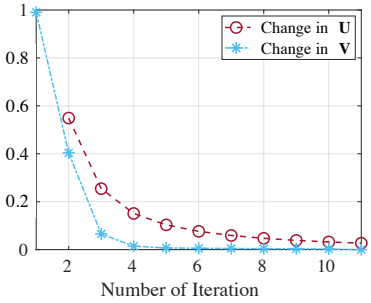

Figure 5.9: Relative changes in the estimated temporal and spatial factor matrices (computed according to line 3, Algorithm 1) when ECA is applied on 2D fUS data at $R = 2, \lambda = 1000$ with a single regularized component. Since the algorithm works iteratively and begins with a completely random initialization of the non-regularized column(s) of $\mathbf{V}$ (line 2, Algorithm 1), the change in $\mathbf{V}$ at the first iteration is drastic.

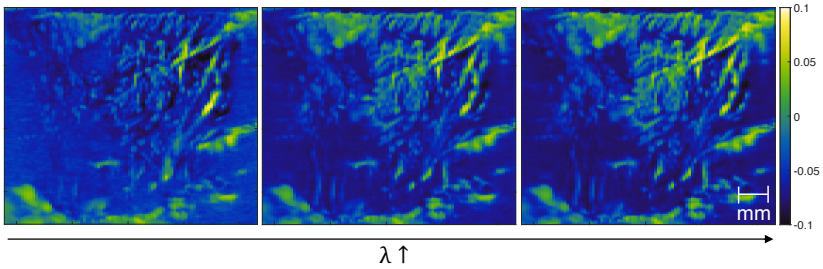

Figure 5.10: Estimated spatial map of the non-regularized component at $R = 2$ for varying λ ($\lambda = 100, 3000$ and 5000 from left to right). The more regularization is weakened, the more the regularized component absorbs from the common activity in the brain (Fig. 5.4(b)). As the weight on the non-regularized component gets lifted, the vascular structure becomes less apparent.

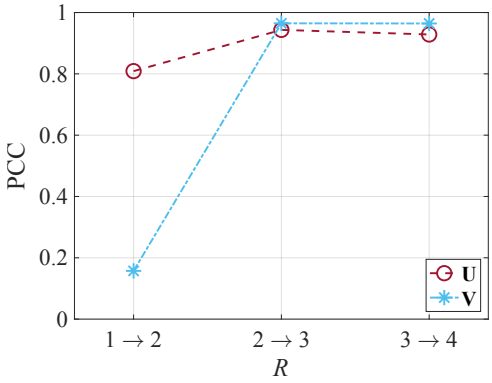

Figure 5.11: PCC between the estimated factor matrices with respect to the rank R. The number of regularized components was always kept at 1, whereas the number of non-regularized components was 0, 1, 2 or 3, meaning the total rank was changed from 1 to 4. Each time R was increased, we computed the PCC between the factor matrices estimated using the previous value of the rank (left side of the arrows) and the ones estimated at its current value (right side of the arrows). The low PCC from R=1 to R=2 indicates that there is a notable change when there is a non-regularized component compared to when there is not.

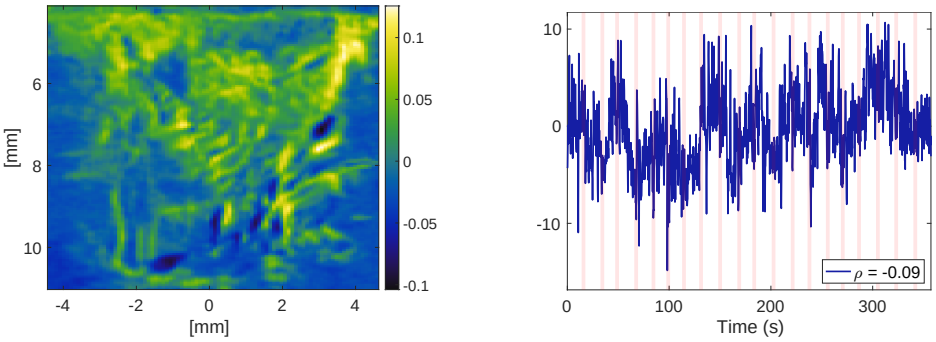

(a) (b)

Figure 5.12: Estimated spatial map (a) and time course (b) of the non-regularized component at $R = 3$ with 2 regularized components at $\lambda = 200$.

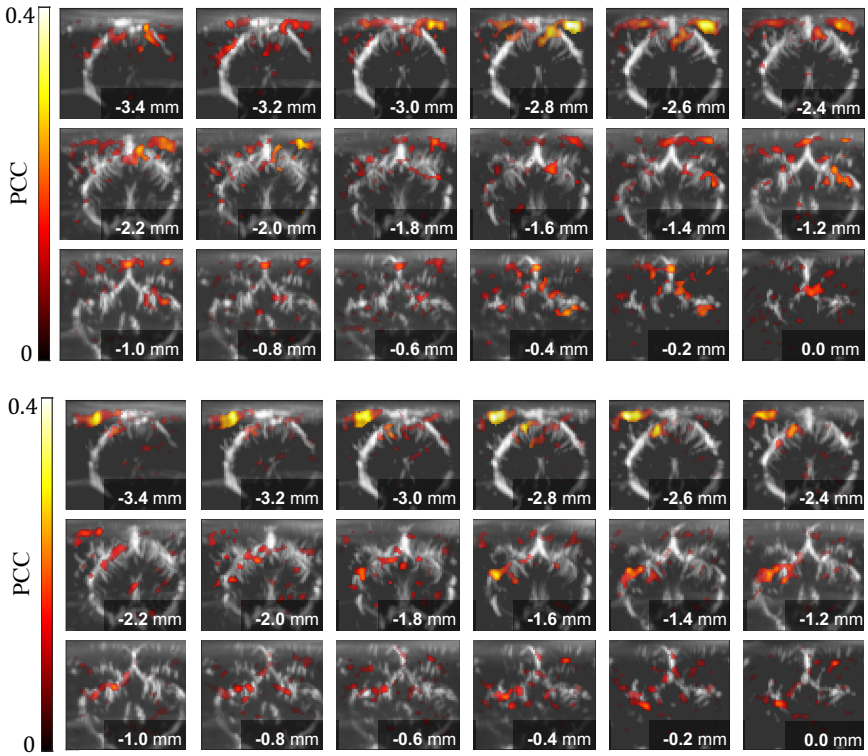

Figure 5.14: Correlation images for left-side (top) and right-side (bottom) stimuli. Only voxels with z-score ≥ 1.5 are displayed.

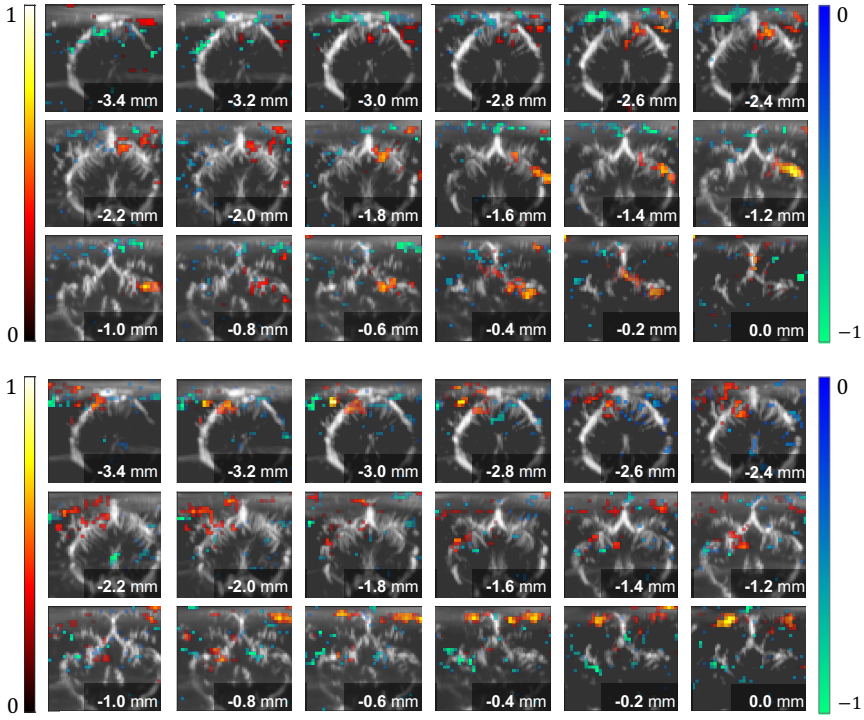

Figure 5.15: The difference in voxel activations calculated by subtracting the activations computed using Pearson correlation coefficient (PCC) from those estimated by ECA. Note that the activations from both methods were normalized to unit energy (i.e., divided by the l_2-norm) prior to subtraction. The top figure shows the activation differences for the left-side stimuli, whereas the bottom figure shows those for the right-side stimuli. Only voxels whose *absolute value* of the z-score ≥ 1.5 are displayed. The voxels who are significantly positive (negative), i.e. when ECA has predicted a higher (lower) activation value than PCC, are shown with the red (blue) colormap. For left-side stimuli, notable differences include the higher activations estimated by PCC in the visual cortex (VIS) of the same side (left), and the higher activations estimated by ECA in the contralateral hemisphere, including right lateral geniculate nucleus (LGN) and right superior colliculus (SC). For right-side stimuli, notable differences are the higher activations estimated by PCC around the left LGN, and the higher activations estimated by ECA in the left VIS and left SC.

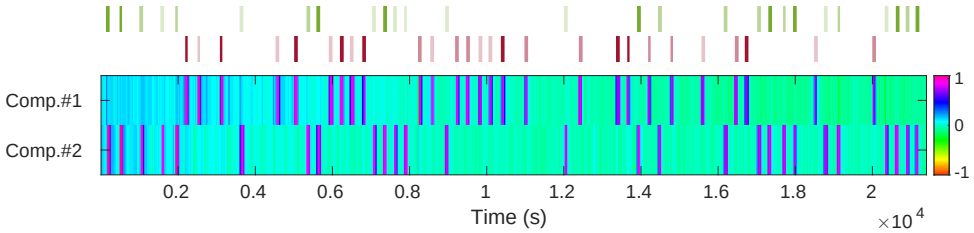

(a) Estimated time courses of the two regularized components. The red and green colored bars on top indicate the left and right-side LED stimuli respectively, with different shades referring to the LED brightness level. The colorbar denotes the amplitude changes of the time courses.

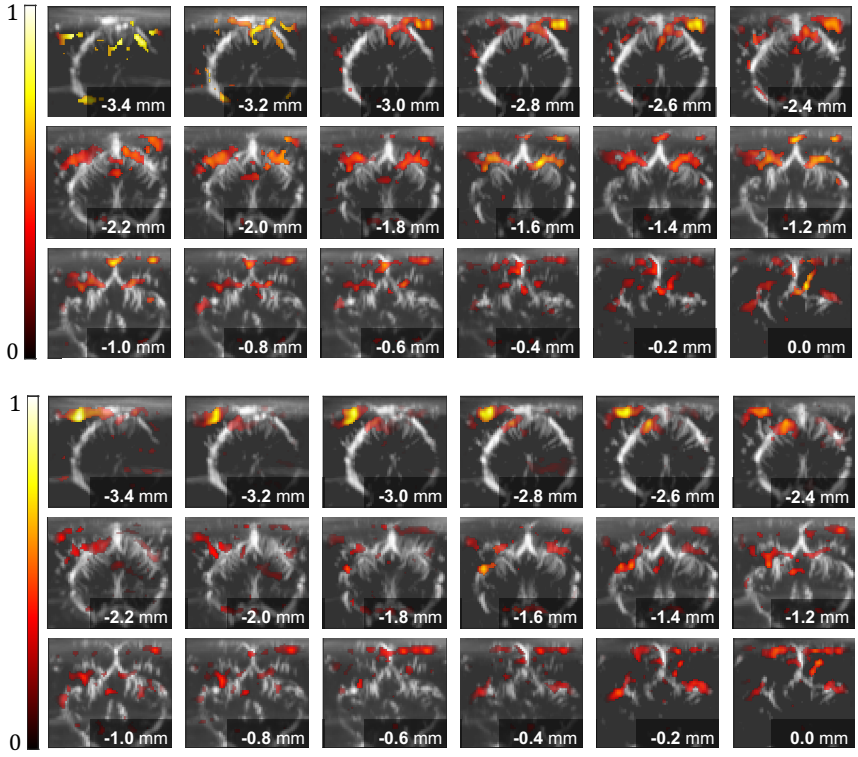

(b) Estimated spatial maps corresponding to left-side (top) and right-side (bottom) stimuli.

Figure 5.16: Estimated time courses (a) and associated spatial maps (b) for one subject (N=1) by ECA at λ=100. Only voxels with z-score ≥ 1.5 are displayed.

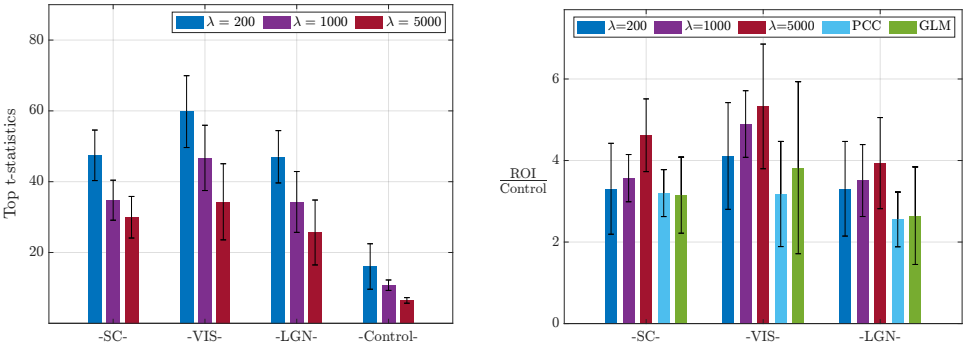

(a) (b)

Figure 5.17: Results of ECA applied on swept-3D fUS data at varying lambda. In (a), the mean of top 20 t-statistics at each region is compared across 3 different values of λ, averaged across subjects. In (b), the mean of top 20 t-statistics in regions of interest are divided by that of the control region for all methods. The errorbars denote the standard deviation across subjects in both figures. While the results in (a) reveal that low λ causes an increase in t-statistics everywhere, i.e. including the control region, the results in (b) demonstrate that ECA still achieves a higher contrast between responsive and unresponsive regions at a wide range of λ.

6

General Discussion

"The brain is a complex system that constantly changes and adapts in response to experience."
— Gerald Edelman

6.1. INTRODUCTION

In the previous chapters we challenged the use of the conventional LTI model for analyzing evoked hemodynamic activity. The LTI model expresses the voxel time series through the convolution of the hemodynamic response function (HRF) and an input signal that evokes the changes in hemodynamics, as explained in Chapter 1 and 2. Subsequently, we demonstrated the nonlinearity and spatiotemporal variability of fUS responses under the assumption of binarized stimuli as input signals in Chapter 3. Accordingly, we avoided complete reliance on binarized stimulus representations for describing temporal dynamics. To elaborate, in Chapter 4, we modelled the fUS time series as convolutive mixtures of underlying sources. We then jointly estimated region-specific HRFs and source signals, allowing us to represent HRF variability across the brain and achieve a flexible representation of the input signal(s), leaving room to incorporate trial-dependent behavior, respectively. This approach is powerful in the sense that it does not require any prior information of stimuli, yet, the high number of parameters to be estimated results in limited scalability. In Chapter 5, we proposed a controlled framework that adjusts the influence of prior information of stimuli through a model-based regularization. This way, it is possible to factorize the evoked fUS activity in the whole brain into its spatial and temporal counterparts while allowing us to observe the spatiotemporal trial variability of the brain at a user-defined extent.

Although the spatiotemporal trial variability in hemodynamic responses is most likely of neuronal origin [207], it is possible that intrinsic vascular fluctuations cause interference [208]. As such, no matter how flexible the sources of hemodynamic responses are estimated, the only way to truly and exactly model the underlying neurovascular behavior is to measure local neuronal activity alongside the fUS signal [87].

From a different perspective, as new technologies evolve into established methodologies, a natural step forward is to combine the different technologies to learn more than we could possibly do from a single modality. Multimodal brain imaging has burgeoned in the last years, with simultaneous electroencephalography (EEG)-functional magnetic resonance imaging (fMRI) taking the lead [209]. The reason that this combination is at the forefront is the complementary nature of the two modalities, i.e. high temporal resolution of EEG and high spatial resolution of fMRI. On the other hand, an EEG-fUS synthesis can potentially provide an even better match by mitigating the drawbacks of fMRI; namely the high costs, lack of portability and low signal-to-noise ratio [210].

In the next section, we will present an overview of processing techniques for multimodal neuroimaging data. Additionally, we will discuss how the methods proposed in this thesis can be modified to incorporate additional information of neuronal activity. The last section of this chapter highlights the pragmatic aspects of working with fUS data.

6.2. MULTIMODAL IMAGING WITH FUS

Joint processing of data from two modalities is generally performed in three different ways. The first one is to use one modality to constrain another. For example, initial

EEG-fMRI integrations have used fMRI to constrain the source locations for solving the inverse problem of EEG [211]. Indeed, EEG measures electrical potentials on the scalp, which are generated by neuronal activity in the brain. Forward models simulate how neuronal activity in different brain regions generate scalp potentials, while inverse models estimate the distribution of neuronal activity in the brain based on scalp measurements, known as source localization.

Alternatively, data from two modalities can be processed separately to cross-check their results. This approach was employed in both multimodal fMRI and multimodal fUS studies up to date. For instance, [87] has reported that the fUS response can be closely predicted by the HRF-filtered firing rate of local neurons. The smoothing effect that HRF exerts on neuronal firing is also displayed through the spectral coherence between the two signals, which is stronger at low (< 0.3 Hz) frequencies. Likewise, [66] has estimated subject and condition-specific HRFs to describe the relationship between neuronal calcium and the fUS signal. A multimodal setup for simultaneous EEG-fUS acquisition has been proposed by [212] in order to investigate the correlation between the θ rhythm (EEG waveforms at $4-7$ Hz, activated during spatial navigation tasks) and the fUS signal. Their results indicate that both the peak value of this correlation and the delay at which the peak is obtained differs across regions.

The last method for processing multimodal data is named fusion, meaning that the data from two modalities are allowed to mutually inform each other during processing. This can either be performed by combining the datasets into one large data array, or via coupled factorizations. The JointICA scheme proposed by [213] concatenates the fMRI activation maps (consisting of intensities of M voxels) and EEG time-series (consisting of V time samples) obtained from P participants, denoted by $\mathbf{Y}^{\mathrm{F}} \in \mathbb{R}^{P\times M}$ and $\mathbf{Y}^{\mathrm{E}} \in \mathbb{R}^{P\times V}$ in a matrix $\mathbf{Y} = [\mathbf{Y}^{\mathrm{F}}\ \mathbf{Y}^{\mathrm{E}}]$. $\mathbf{Y}^{\mathrm{F}}$ and $\mathbf{Y}^{\mathrm{E}}$ are assumed to share the same mixing matrix $\mathbf{A} \in \mathbb{R}^{P\times K}$, where K is the number of neural sources, such that $\mathbf{Y}^{\mathrm{F}} = \mathbf{A}\mathbf{S}^{\mathrm{F}}$ and $\mathbf{Y}^{\mathrm{E}} = \mathbf{A}\mathbf{S}^{\mathrm{E}}$. Hence, the combined data matrix $\mathbf{Y}$ can be decomposed with ICA as $\mathbf{Y} = \mathbf{A}\mathbf{S}$ where $\mathbf{S} = [\mathbf{S}^{\mathrm{F}}\ \mathbf{S}^{\mathrm{E}}]$. The neural sources, each described by a joint EEG-fMRI component based on *where* (fMRI) and *when* (EEG) the activity was occurring, organized in the rows of $\mathbf{S}$, are assumed to be statistically independent from each other. It is possible to incorporate multiple EEG channels by using a coupled matrix-tensor factorization (CMTF) that can jointly factorize the EEG tensor (without vectorizing the channel dimension) and fMRI matrix by assuming a shared participant signature [214]. Another CMTF approach was proposed by [64], where the fMRI temporal signatures are modelled as the convolution between region-specific HRFs and EEG temporal signatures. This way, the two modalities are jointly decomposed with shared temporal signatures while incorporating the convolutive relationship between hemodynamic and neuronal activity.

In what follows, we will go over the methods employed in this thesis, and how the findings of each chapter can be enhanced further by including neuronal activity.

6.2.1. NONLINEARITY

Nonlinearity of hemodynamic responses has been a subject of study in the literature, revealing various causes for it and attempts to model its behavior. For example, [26] reports that the blood-oxygen-level-dependent (BOLD) signal measured with fMRI in response to short stimuli is narrower and much larger in magnitude than would be pre-

dicted by linear models derived from longer duration stimuli. Consequently, when one attempts to predict the response given to a long stimulus by superposing the responses given to short stimuli, the magnitude of resulting curve is an overestimation of the true magnitude. Another cause of nonlinearity was pointed out by [66], revealing a second vascular component in the fUS signal that appears under very strong stimulation, which was modelled by introducing a second HRF.

Volterra series modelling of hemodynamic responses was proposed originally by [100], leveraging the memory capabilities of Volterra systems to capture nonlinearities at high stimulus presentation rates. We have adapted the Volterra series approach to fUS data in Chapter 3 and validated its significance across stimuli of varying durations. We observed that the responses given to long stimuli quickly reach a plateau with subtle oscillations after the stimulus onset. Instead of attempting at cross-predicting responses of different stimulus durations as [26], we inspected how well each stimulus duration can be described by a linear system compared to a nonlinear one described as a second-order Volterra series. Our findings indicated that while a second-order Volterra series better characterizes responses compared to an LTI model for all stimulus durations, the discrepancy between their performances becomes pronounced for longer stimuli.

While it's clear that the fUS signal displays nonlinear characteristics in response to external stimuli, inspecting merely the fUS data does not elucidate whether these nonlinearities stem from neuronal origins. Through a multimodal setup, this would be possible to analyze by simply replacing the input of the Volterra series (which was assumed to be the binarized stimulus time course in Chapter 3) by local neuronal activity. Depending on to what extent an LTI system can describe the fUS measurements compared to a second-order Volterra system in this case, it would be possible to discern whether the nonlinearities primarily arise from neuronal nonlinearities or from the de-coupling between neuronal activity and blood flow [215].

6.2.2. SPATIOTEMPORAL VARIABILITY

In Chapter 3, we demonstrated the temporal variability of a voxel by estimating an activation coefficient for each of the stimulus trials by fitting a convolutional model with a voxel-specific HRF to each trial. Notably, we found the variability observed in the visual cortex (VIS) in mice to be twice as high as that in the lateral geniculate nucleus (LGN), mirroring conclusions drawn in previous studies examining neuronal firing and suggesting a potential neuronal origin of this behavior. Indeed, [27] reports that trial variability gradually increases along the successive stages of the visual-processing pathway, doubling from retinal ganglion cells to the thalamic cells of LGN, and again from the thalamic cells of LGN to the cortical cells of VIS. In addition, the authors point out a significant anti-correlation between the mean firing rate (per trial) and trial variability of a region, as the mean firing rate was the lowest in VIS and highest in retina, following an opposite trend compared to the regions' variabilities. Understanding the exact root of this spatiotemporal variability can provide insights into complex neurobiological mechanisms such as neuronal adaptation, stimulus expectation, time-varying functional connectivity or attentional modulation [114].

When neuronal activity is available, we can calculate the residual between the whole fUS voxel time series and the convolution of neuronal activity with a voxel-optimized

HRF. Analyzing the trial-dependent statistics of this residual — whether it fluctuates randomly or with stimuli by extracting trial-by-trial activation coefficients as proposed in Chapter 3 — would allow us to assess the degree to which hemodynamic trial variability is influenced by neuronal variability. Moreover, by repeating this process over the entire set of voxels and across different stimulus conditions, it would be possible to discover whether the neuronal root of hemodynamic variability exhibits region or context-dependent traits, respectively. For example, [216] has found a consistent divergence between the BOLD response and local-field potentials during perceptual suppression, whereas the two signals were observed to be aligned during conventional stimulus presentation. In addition, if there indeed exists a link between the mean firing rate and variability of a region as suggested by [27], the question then becomes whether the variability of a region can be diminished by choosing an effective stimulus that targets the region's specific role or function, and whether this change in variability will be reflected at the same amount in the hemodynamic and neuronal response of the brain.

6.2.3. CONVOLUTIVE MIXTURES

A convolutive mixture model is suitable to describe hemodynamic responses, for the reasons that can be summarized as follows. To start with, as we mentioned in the previous chapters, the brain is known to operate through networks of interconnected regions that exhibit synchronized activity during specific tasks or resting state. As such, functionally connected regions can be assumed to be modulated via shared sources, which is why (instantaneous) ICA can as well successfully capture both resting state and evoked activity [74]. Secondly, hemodynamic activity is a proxy for underlying neuronal variations, thus, a convolutive rather than instantaneous demixing can reveal the underlying neuronal sources more accurately. In fact, even for modelling neuronal activity, [217] argues that convolutive ICA provides an advantage over instantaneous ICA of EEG data by incorporating delayed interactions between the independent time courses. Although convolutive blind source separation is widely adopted in numerous fields, such as in audio signal processing [218], its applications to neuroscience have been limited, despite the convolutive nature of blood flow responses.

Convolutive source separation is not as straightforward as instantaneous separation at first sight, yet it can be simplified through various approaches. Indeed, it is possible to convert convolutive mixtures to matrix form either in the time domain as described in Chapter 4, or in the frequency domain by applying short-time-Fourier-transform. With the latter approach, an instantaneous ICA is solved at each frequency bin [219]. Alternatively, it is possible to apply independent vector analysis for a simultaneous separation across all frequency bins [220]. For the case where both neuronal activity and fUS data is available, a multimodal convolutive ICA problem can be solved [221], accounting for both neuronal delays and the low-pass filter effect of the HRF reflected in fUS signals.

6.2.4. ECA

With Evoked Component Analysis (ECA), we proposed a new framework that sets a balance between data-driven and model-based approaches by incorporating the design matrix of general linear model (GLM) in a regulatory term. Traditionally, GLM analyzes

brain data by constructing temporal regressors through the convolution of stimuli with an HRF and estimating their influence on each voxel's response as an activation coefficient. However, these regressors, constructed solely based on stimuli, assume a uniform response across each stimulus trial, neglecting the brain's inherent spatiotemporal variability. In Chapter 5, we showed that ECA can capture such complex nuances of brain activity by controlling the influence of prior information of stimuli.

When neuronal activity is available, the regressors within the design matrix can be constructed by convolving the HRF with the neuronal activity of regions of interest instead of stimuli [222]. Each regressor is then used to identify the regions whose fUS measurements reflect associated changes. Alternatively, Yang et al. [223] applied ICA on multichannel EEG data to define regressors. While the stimulus-convolved regressors help untangle the relationship between the fUS signal and external stimuli, those based on neuronal activity can illuminate the connection between the fUS signal and actual neural processes. Within the ECA framework, one can adjust the regularization coefficient in order to reveal the regions that are more (or less) coupled to the neuronal regressors [224].

6.3. PRACTICAL CONSIDERATIONS FOR FUS

Power Doppler has become the standard technique for fUS imaging of the brain, owing to its high sensitivity in detecting subtle changes of blood flow in small cerebral vessels. In this section, we will present several practical considerations when working with Power Doppler Images (PDIs).

fUS data is acquired with ultrafast ultrasound technology. The acquisition rate of beamformed frames is in the order of milliseconds, yet computation of PDIs requires additional filtering and averaging to suppress artifacts and boost the signal-to-noise ratio (Chapter 1). The acquisition rate of PDIs can be adjusted by changing the ensemble size of beamformed frames used for filtering and averaging. In the literature, it is possible to come across different ensemble sizes, such as 120 [44], 200 [19] and 320 [23].

In order to showcase the possible effects of the ensemble size, we used the data from a particular slice examined in Section 3.3. We varied the ensemble size used for averaging from its original value of 200 to 100 and 400. To further insure that our comparisons across different ensemble sizes are not critically affected by the choice of clutter filtering, we will present our results both with singular value decomposition (SVD) -based filtering and high pass filtering (with a cut-off frequency of 60 Hz) of beamformed frames. A brief discussion on this choice will follow shortly.

We started by obtaining Pearson correlation coefficient (PCC) and HRF latency maps as before: We calculated the PCC between the voxel time-series and the HRF-convolved stimulus time course at various HRF latencies, and found the optimal HRF latency per voxel (Fig. 6.1(a)) that generates the highest PCC value (Fig. 6.1(b)). The voxels highlighted in both maps were the ones that were considered to be activated by the stimulus in case of the ensemble size as 200, using a PCC threshold of $c = 0.1$ (corresponding to a P-value of 0.0001, see Section 3.3 for details). For the other ensemble sizes, the same set of voxels were used for each region for consistency.

The HRF latency maps reveal that the layers of LGN exhibiting different HRF laten-

cies disappear at low temporal resolution. Prior research on LGN's layers indeed points to a faster peak response on the first two layers of LGN (magnocellular layers) than layers 3 to 6 (parvocellular layers), which is commonly attributed to the differences in the conduction speeds of the axons they receive their input from [225]. Furthermore, the estimated peak latencies are consistently higher in both regions when using an ensemble size of 100 compared to that of 400. It should be noted that using a smaller ensemble also results in lower levels of activation and Pearson correlation coefficients, and the increase in the noise intensity might influence the accuracy of HRF latency estimations in this case. Last but not least, we visualized the average response in both regions in Fig. 6.1(c-d), which are very much alike for an ensemble size of 200 and 400, whereas there is a visible difference in the ensemble size of 100.

Another point to mention is that as shown in the results of Fig. 6.1(c-d), the filtering step required for rejecting motion artifacts can be handled in various ways. Originally, beamformed frames were high-pass filtered based on the assumption that blood flow is a high-frequency signal whereas artifact sources (such as tissue motion due to respiration or cardiac pulsatility) generate activity at low frequencies. Later, this was replaced by SVD which seeks spatiotempotal coherence in distinguishing signal and noise components [20]. An algorithm based on projection approximation subspace tracking was proposed by [226] that flexibly estimates the number of eigencomponents to remove as new beamformed images arrive. Such an approach might be favorable especially when the threshold between tissue and blood signal varies over time.

Last but not least, PDIs are obtained by averaging the filtered beamformed images over the full Doppler spectrum, where each frequency is weighted equally. However, it is possible to analyze the sub-bands of the Doppler signal, which can elevate the functional signal in certain brain areas as well as their stimulus correlations [227]. When neuronal activity is known, this approach can highlight which sub-bands of the Doppler spectrum relates to neuronal activation.

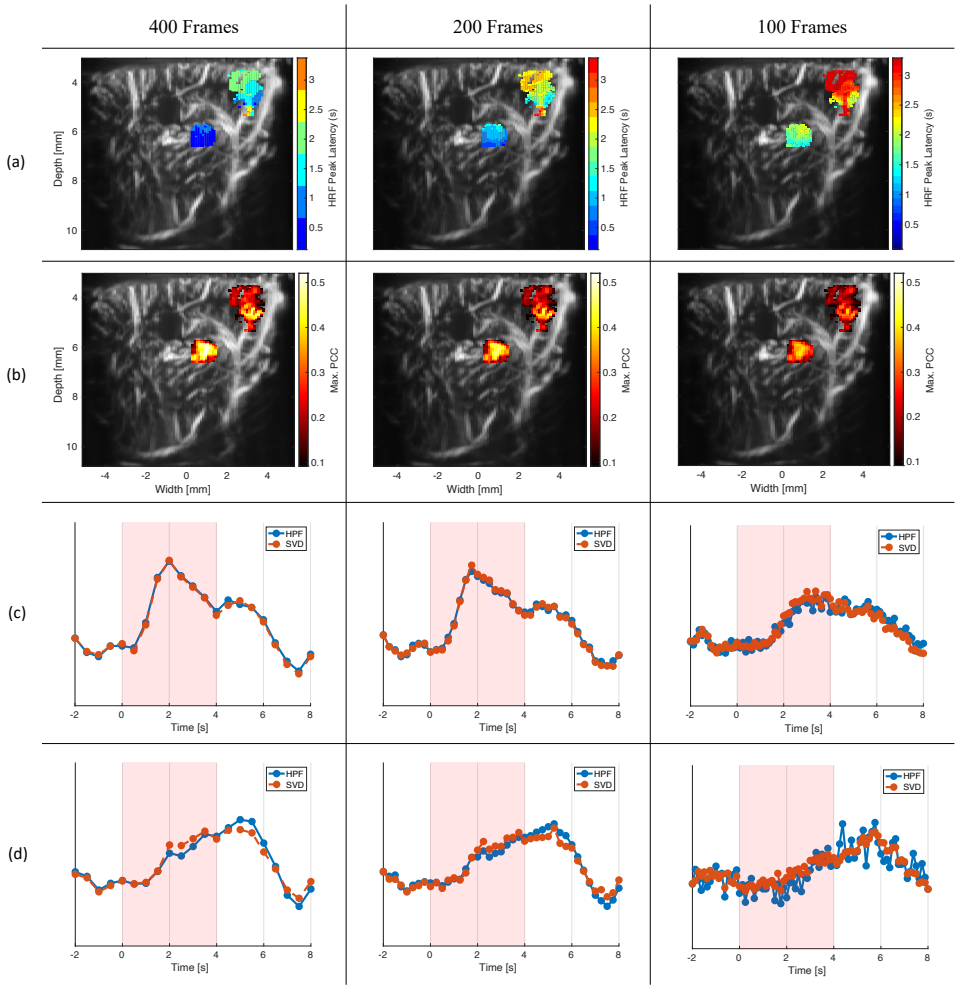

Figure 6.1: (a-b) Spatial maps showing the optimum HRF peak delay per voxel (a) and the highest value of Pearson correlation coefficient (obtained at the optimum delay) per voxel (b). The regions displayed in each spatial map correspond to the visual cortex (top region) and LGN (bottom region). (c-d) Voxel and trial-averaged LGN (c) and VIS (d) responses (only voxels whose Pearson correlation coefficient with the stimulus is above 0.2 were used) using both high-pass filtering (HPF) and SVD-filtering of beamformed frames. The red shade in the background displays the stimulus.

7

Conclusions and Future Research Directions

"Essentially, all models are wrong, but some are useful. However, the approximate nature of the model must always be borne in mind."

— George E.P. Box

7.1. CONCLUSIONS AND SUMMARY OF MAIN RESULTS

UNnderstanding the structure and function of the human brain stands as one of the most captivating, yet challenging, of scientific pursuits. Serving as the center of sensory coordination, emotional regulation and cognitive processing, the brain has attracted scientific interest for many centuries, leading to remarkable developments in imaging technologies.

Functional ultrasound (fUS) is a neuroimaging technique that has emerged recently, offering an unprecedented spatiotemporal resolution for whole-brain imaging using a portable and cost-effective setup. Similar to functional magnetic resonance imaging (fMRI), fUS indirectly measures neuronal activity by detecting changes in local hemodynamics. Variations in neuronal activity cause a delayed response in the blood flow, modelled by a hemodynamic response function representing the underlying neurovascular mechanisms. In this thesis, our goal was to analyze evoked fUS activity considering the spatiotemporal variability of brain responses.

In Chapter 1, we gave a brief introduction to neuroimaging techniques in general, and elaborated on the principles of fUS as a promising new tool for neuroscientific research. Chapter 2 was dedicated to elucidating the brain's physiology and the primary objectives of neuroimaging data analysis. We explained existing techniques for extracting the spatial and temporal information of interest from neuroimaging data. In summary, using a data-driven method allows to capture the complex patterns of brain activity as they are; whereas using a model-based approach enables interpretation of results within a theoretical framework, which may lead to deeper insights.

Chapter 3 focused on our first research question (Q1), namely analyzing the non-LTI characteristics of the fUS response and determining the conditions under which these characteristics become prominent. By comparing the model fit of an LTI system with that of a second-order Volterra series, we showed that the nonlinearities of the fUS signal become stronger under prolonged stimuli. Subsequently, we investigated the time-varying behavior of the fUS responses in two different regions. Our analysis revealed that the trial activation coefficients obtained by fitting a voxel-optimized HRF to each trial response declines as the stimulus repetitions progress, possibly pointing to habituation or neuronal adaptation. Notably, the rate of this decline also altered between the two regions, indicating changing levels of variability across the brain. This led us to conclude that, although an LTI model may be plausible under specific circumstances, the evident time-varying behavior demands a more flexible approach in characterizing input signals beyond simple binary stimulus representations.

In Chapter 4, we proposed a tensor-based solution to address Q2, which can flexibly estimate both the source signals and HRFs, thus taking into account trial variability and the convolutive nature of hemodynamic activity, respectively. This approach models fUS responses as convolutive mixtures of underlying sources, without using prior information of stimuli. We solved the deconvolution problem by applying block-term decomposition on the tensor of lagged autocorrelations of the fUS responses of pre-defined regions, assuming that the sources are uncorrelated. We selected these regions as the ones that were evoked by stimuli; such that that we could later use the known stimulus onsets for validating the estimated sources. On that note, it is important to highlight

that although the proposed method was shown to untangle the timings of multiple stimuli (same visual paradigm repeated at different locations within the mouse field of view) simultaneously, prior selection of regions of interest softens the data-driven prospect.

In Chapter 5, we investigated Q3 and aimed to devise a decomposition tool for the whole brain that can combine the advantages of both worlds: flexibility of data-driven and interpretability of model-based tools. To that end, we defined a matrix factorization which uses the design matrix (defined based on stimuli) as a regulatory term. As such, the influence of prior information can simply be adjusted by controlling the regularization coefficient. By estimating spatial activation maps and their associated time courses at different values of this coefficient, it becomes possible to track both trial variability and which regions respond more (or less) consistently to the stimuli. Moreover, the proposed framework is suited to handle incomplete data, which occurs naturally in swept-3D fUS data due to the different slice acquisition times as a result of probe motion.

Chapter 6 was dedicated to discussing potential multimodal fUS-neuronal activity applications, and how the methods proposed in this thesis can be adapted to further delineate the underlying neurovascular dynamics. In particular, we referred to prior experimental findings that present the (non-)neuronal origin of the nonlinearity and time-variance of hemodynamic responses. We concluded this chapter by sharing several practical considerations related to Power-Doppler imaging, and how the decisions in the acquisition pipeline can affect subsequent processing.

7.2. SUGGESTIONS FOR FUTURE RESEARCH

There are many puzzles yet to solve for understanding brain functioning. In what follows, we will mention future research directions for each chapter individually.

In Chapter 3, we highlighted the nonlinearities of the fUS response, and how they become too strong to reject with prolonged stimuli. While this is an important conclusion to keep in mind when designing an experimental paradigm or choosing an appropriate data model, an interesting research direction would be to explore which other stimulus parameters cause such discrepancies. A similar observation was made by [66] regarding the bimodal behaviour of the fUS response at high stimulus strength. These findings prompt further investigation into the effect of stimulus parameters on the fUS response, such as stimulus frequency, length of rest intervals, or stimulus orientation; examples of which can be found in the fMRI literature respectively at [35, 228, 229]. Again in Chapter 3, we explored the spatiotemporal variability of fUS responses. Although this behavior is expected in general within the brain due to neuronal adaptation, the two regions of interest that we focused on during our analysis, the lateral geniculate nucleus (LGN) and visual cortex (VIS), showed significant differences in their variability across repeated stimuli, aligning with existing neurophysiological studies. Finding such distinctions between regions can provide insights into the way they operate. Indeed, the variability of VIS is shown to be strongly induced by attentional modulations [230]. Therefore, investigating trial-by-trial variability in other brain regions evoked by different types of stimuli presents another intriguing prospect.

In Chapter 4, we presented a novel application for convolutive mixtures in modelling fUS responses and used tensor decomposition to identify source signals and HRFs. Al-

though we used the known stimulus onsets for evaluating our results, the tensor-based deconvolution can also be used in cases with no prior information of the task. For example, the standard spatial independent component analysis (ICA), which assumes instantaneous mixing, is shown to successfully discover resting state networks [69]. Our convolutive model may also be useful for resting state data, as an alternative to the instantaneous ICA that can as well incorporate the convolutive nature of hemodynamic activity.

In Chapter 5, we proposed a novel decopmosition scheme for analyzing fUS responses called Evoked Component Analysis (ECA). In theory, ECA can be adapted to any other hemodynamics-based neuroimaging techniques such as fMRI. It would be interesting to see how the choice of the regularization parameter affects the estimated spatial and temporal signatures in different modalities. Furthermore, the completion of missing values suggested for swept-3D fUS can also be applied for filling in data corrupted with artifacts [202].

In Chapter 6, we delved into potential multimodal imaging applications of fUS, which can combine unique insights obtained from the different modalities. Due to the fact that fUS is a relatively new technique, there are still many questions to answer regarding the exact nature of the relationship between neuronal activity and the fUS signal. To that end, [87] has shared important revelations including the strong correlation between the fUS signal and slow fluctuations (< 0.3 Hz) of local neuronal firing. However, the authors also report a constant HRF shape across subjects, stimulus conditions, and the two regions that were investigated (the hippocampus and visual cortex), while the results of [66] suggest a varying vascular response depending on stimulus conditions and more robust predictions when the HRF is estimated per-subject. Overall, conducting more multimodal experiments is essential to unravel the exact characteristics of the coupling between the fUS signal and neuronal activity.

While fUS has proven efficient for pre-clinical animal studies, its translation to human applications is critical for future clinical development. The first human fUS experiments are performed in human neonates [231] and adults undergoing intra-operative craniotomy surgeries [44, 94]. Indeed, [44] has demonstrated the promise of fUS for tumor delineation based on vasculature characteristics of tumor and healthy tissue. The number of human fUS studies will likely increase in the years to come, with the initial results indicating great potential for the clinical use of fUS and the technology moving toward transcranial applications [17].

As fascinating as it is to investigate brain signals, it is an extremely challenging task due to the vast amount of processes the brain controls simultaneously. All models of the real world are approximations, yet the brain can be one of the most, if not the most, complex mechanisms to model. Moreover, the absence of a definitive “ground truth” complicates matters — we instead rely on experimental data accumulated over time to form such expectations. Hence, a general remark for all neuroscientific endeavours, including this thesis, is that performing more experiments over more subjects and referring to existing literature at every step will always be appreciated to enhance the generalizability of proposed methodologies and conclusions.

ACKNOWLEDGEMENTS

No success in life is achieved alone and this is probably the closest I will get to giving an Oscar speech. So, to whomever is reading this, get ready for a long list of people that I must share my gratitude to.

I would like to start by thanking my supervisor, Borbála Hunyadi. Bori, I will be forever grateful that you returned to my very first e-mail that I sent you, asking if we could meet regarding a PhD position on your webpage, which ended up changing my whole life. You encouraged me to have my own ideas and explore a new field on my own, while not letting me struggle there alone. You always made the time to meet with me or talk to me before I could get lost in my own head. Thank you for all your help, and having the faith in me that I would get to this point. Next, I would like to thank Pieter Kruizinga, first and foremost, for asking me how I am in the most real way, not ever taking 'Fine.' for an answer, every single time he sees me. Thank you Pieter for teaching me your out of the box way of thinking, for your challenging questions -both of which have inspired so much in this thesis-, and for your genuine care and support. I would also like to thank my promotor, Alle-Jan van der Veen, for creating such a welcoming environment for all the students under his wings. Your level of knowledge and humbleness is what I will strive to achieve in the rest of my career.

Needless to say, this thesis wouldn't have happened without my partner and my best friend. Tarik, even when I am upset, deep down I always know that it will be alright because I have you. I am always in awe of your ambition, resilience and passion. I still cannot believe that you managed to teach me how to ride a bike. That was a real challenge. You have a way of comforting me, and making things look easier than they are. You are my rock and my biggest inspiration. Thank you for being in my life.

When I first came to Delft, I did not know a single person in the city, and the friends I made here became my family. Hanie, you were my first friend in the Netherlands, so you were also the first person I shared all my doubts with. Thank you for your endless patience, all the fun memories, the first time you broke the formality and hugged me (on day one), for being my officemate, my go-to person, the delicious meals you've cooked, and of course for getting lost in the canals of Delft with me.

Anurodh, I am finding it very difficult to say normal things to you, but I will force myself to do it anyway. You are very smart, funny and caring of the people around you, and you have been an amazing friend to me. I really enjoy hanging out with you and I will for sure miss our coffee breaks, which are never just ten minutes and always end up with us on the verge of hitting an existential crisis. I plan on annoying you forever with my never-ending Buzzfeed questions to learn all the unnecessary details of your life.

It is almost like I had two PhDs: Before Sofia and After Sofia. Sofia, I cannot count how many times you pulled me out from a spiral of insecurity. I hope you are still writing down the keywords of the stuff you want to talk to me about when you see me again because you don't want to forget them, and because I will want to know all of them.

Thank you for all the laughter, thoughtful gifts, sharing of information (no, not gossip, this was all scientific), and for encouraging me to stand up for myself while still keeping it real by making fun of everything I do and say. You're The Best!

Bas and Sadaf, I have so many unforgettable memories with you that I don't know which to mention. All I know is that I can't wait for the next time we play 'Truth or Truth', tell funny joke stories, share all the drama of our lives and reminisce Corsica. Despite all the crazy-fantastic stuff that happened there, the best part of Corsica for me, undoubtedly, was getting to know the both of you. Sadaf, you are not only extremely smart and inspiring, but you are also so kind and warm, and you have the weirdest-darkest sense of humor that I just love. Whatever you do, you make it memorable for everyone around you because you do it with such passion, including falling asleep on our living room floor. Bas, I don't know if you are aware of this but you have been like a mentor for me ever since Corsica - the most fun and understanding one I could have asked for. I really enjoyed our morning coffees and super-effective work sessions at the lab, and I only wish we had started them earlier. Still, I am very grateful that we had enough time to form the bond we have now, and that I still get to bother you with every small update in my life.

A big thanks to Mario, who whenever made the mistake of asking me how is it going ended up listening to all my problems. Thank you for your friendship and for patiently answering every single one of my questions including ones that felt really obvious, but I would be too shy to ask anybody else.

To my other officemate, Didem, thank you for all your support, our many train rides together, the amazing travel tips, fashion advice, and of course passionate discussions on Turkish TV series characters. I have had a lot of fun with you, and I hope we continue to do that in Eindhoven!

Seline, I am so happy that you started a PhD here — and I must admit that I take a bit of credit for that because I think I was the first PhD student you talked to? Even if I wasn't, don't tell me... I had so much fun with you in the past few years, especially during our morning coffee chats in Rhodes. Thank you for all that we shared, and for making the difficult days feel a bit less so!

While talking about Rhodes, I must thank Ruben and Costas for our many fun and memorable moments. Ruben, I wish you to always keep your positive energy, and to find the best ever pre-processing strategy, I trust you! Costas, I still remember the evening we spent talking non-stop for hours. Thank you for your support and dancing crazy with my cousins at our wedding — who still ask me about you every once in a while...

I would like to thank Alberto for keeping me company during my PhD since the beginning. I will always cherish our first few months of non-stop partying with a big smile on my face. I hope you plan on keeping me in your life so that I can continue to make you feel uncomfortable by asking you inappropriate questions about your personal life.

Cristian, I honestly hate you a bit for not starting your PhD earlier. I could have used your energy, enthusiasm, and your unique perspective on life, which looks simple at first but is actually very wise, so much more. But hey, we will make up for the lost time!

I would like to thank Utku Kumbul for being my friend since I started my PhD, and listening to all my troubles. You made my PhD feel a bit easier. I have had so much fun with you, and I hope that one day we will manage to live/work in the same city again!

Yanbin, just seeing you puts a smile to my face. That's not only because you are one

of the warmest, kindest people I ever met, but because your dark sense of humor doesn't match your bubbly looks at all. I hope you always stay in touch, wherever you end up!

Metin, my first impression of you was that you look very smart and interesting, and you proved me right in all our interactions. I wish our friendship had started back in high school, not just because you are really fun, but because I would have been way more popular. Well, better late than never!

Pim, thank you for guiding me especially in the beginning of my PhD. I still remember how you said you find me a bit similar to yourself, and that's why you are worried. Thanks but as it turns out, you didn't need to worry, I did so much better than you. And I safely claim this since I am quite sure that this is the only part of my thesis you will read.

I also thank Chagajeg for being my companion throughout the pandemic and helping me survive those crazy times. Thank you for your friendship and still having trust in me although one of our meetings literally got interrupted by the police.

I must return back to Corsica and thank my flight-buddy Luuk as well, with whom we also had a very planned trip to Geneva afterwards. Luuk, you are always so kind, and you never fail to entertain me with your adventures. Thank you again for not letting me drown in the sea and all the fun memories!

I would like to thank Hande Pat, my best friend and companion through all the good, bad and ugly for 16 years now. I imagine us in 2085 as two 90 year olds, laughing at our own inside jokes. I thank you and Oguz for all our fun trips, and to many more to come!

Of course, I thank my friends on the other side of the river, who made our moving to Eindhoven much less difficult than I thought it would be. Thank you Emre, Sercan and Candan for your company, may we play many more games together and make endless comparisons between different jobs and countries! Tuğçe and Serkan, thank you for so many fun trips and memories, and I really mean it when I say that if you move to another city we are coming there too!

There are still so many more people that I should thank, but I'm told that my acknowledgements should not be longer than my thesis itself. Giovanni, Changheng, Peiyuan, Yanbo, Zhonggang, Ids, Ellen, Jordi,... I thank you for all the fun memories, and making our office feel like a second home to me. Always feel free to ask for my wisdom.

I always felt indescribably lucky to have been born to my parents, and I still can not believe that life gave me a new family that I find just as close to me. Mirsad, Velida, your warmth and kindness to not only me but everyone around you gives me energy and hope, and seeing you happy makes me happier. I am looking forward to our many holidays in Bosnia and Balikesir together, with Velida's amazing boreks and Mirsad never leaving the water. Ahmed, I don't know if I told you this before but I always wanted a brother and gave up on it years ago. Now I finally have one. Thank you for being in my life, and always having my back against your annoying brother.

Anne, baba, sizin çocuğunuz olmak benim bu hayattaki en büyük ayrıcalığım. Ne için teşekkür etsem, ne kadar teşekkür etsem eksik kalacak. Sizinle birlikte diziler, maçlar seyretmekten; oyunlar oynamaktan, yanınızda hep çocuk olabilmekten, sizden her gün yeni bir şeyler öğrenmekten, sohbetlerimizden, kahvaltılarımızdan, hatta kimi zaman hiç konuşmadan yan yana oturmaktan, kısaca sizinle olan her anımdan tarifsiz bir mutluluk duyuyorum.

Hayatımda başardığım doktoram dahil her şeyin mimarı sizsiniz, sizin sevginiz, sizin desteğiniz... Sizi çok seviyorum.

BIBLIOGRAPHY

[1] Raul Marino Jr and Marco Gonzales-Portillo. “Preconquest Peruvian neurosurgeons: A study of Inca and pre-Columbian trephination and the art of medicine in ancient Peru”. In: *Neurosurgery* 47.4 (2000), pp. 940–950.

[2] Peter J Nestor, Philip Scheltens, and John R Hodges. “Advances in the early detection of Alzheimer’s disease”. In: *Nature medicine* 10.Suppl 7 (2004), S34–S41.

[3] Joel S Perlmutter and Jonathan W Mink. “Deep brain stimulation”. In: *Annu. Rev. Neurosci.* 29 (2006), pp. 229–257.

[4] Luis Fernando Nicolas-Alonso and Jaime Gomez-Gil. “Brain computer interfaces, a review”. In: *sensors* 12.2 (2012), pp. 1211–1279.

[5] Oludare Isaac Abiodun et al. “State-of-the-art in artificial neural network applications: A survey”. In: *Heliyon* 4.11 (2018).

[6] Silvia A Bunge and Itamar Kahn. “Cognition: An overview of neuroimaging techniques”. In: (2009), pp. 1063–1067.

[7] Alex M Aisen et al. “MRI and CT evaluation of primary bone and soft-tissue tumors”. In: *American journal of Roentgenology* 146.4 (1986), pp. 749–756.

[8] Alberto E Pereda. “Electrical synapses and their functional interactions with chemical synapses”. In: *Nature Reviews Neuroscience* 15.4 (2014), pp. 250–263.

[9] Hans Berger. “Über das elektroenkephalogramm des menschen”. In: *Archiv für psychiatrie und nervenkrankheiten* 87.1 (1929), pp. 527–570.

[10] Sylvain Baillet. “Magnetoencephalography for brain electrophysiology and imaging”. In: *Nature neuroscience* 20.3 (2017), pp. 327–339.

[11] Hillman E. M. C. “Optical brain imaging in vivo: techniques and applications from animal to wan”. In: *J. Biomed. Opt.* 12 (2007), pp. 1–28.

[12] Clément Huneau, Habib Benali, and Hugues Chabriat. “Investigating human neurovascular coupling using functional neuroimaging: a critical review of dynamic models”. In: *Frontiers in neuroscience* 9 (2015), p. 467.

[13] Harry T Chugani, Michael E Phelps, and John C Mazziotta. “Positron emission tomography study of human brain functional development”. In: *Annals of neurology* 22.4 (1987), pp. 487–497.

[14] Joana B Balardin et al. “Imaging brain function with functional near-infrared spectroscopy in unconstrained environments”. In: *Frontiers in human neuroscience* 11 (2017), p. 258.

[15] Gary H Glover. “Overview of functional magnetic resonance imaging”. In: *Neurosurgery Clinics* 22.2 (2011), pp. 133–139.

[16] Thomas Deffieux et al. “Functional ultrasound neuroimaging: a review of the preclinical and clinical state of the art”. In: *Current opinion in neurobiology* 50 (2018), pp. 128–135.

[17] Thomas Deffieux, Charlie Demené, and Mickael Tanter. “Functional ultrasound imaging: A new imaging modality for neuroscience”. In: *Neuroscience* 474 (2021), pp. 110–121.

[18] Elodie Tiran et al. “Transcranial functional ultrasound imaging in freely moving awake mice and anesthetized young rats without contrast agent”. In: *Ultrasound in medicine & biology* 43.8 (2017), pp. 1679–1689.

[19] Emilie Macé et al. “Functional ultrasound imaging of the brain”. In: *Nature methods* 8.8 (2011), pp. 662–664.

[20] Charlie Demené et al. “Spatiotemporal clutter filtering of ultrafast ultrasound data highly increases Doppler and fUltrasound sensitivity”. In: *IEEE transactions on medical imaging* 34.11 (2015), pp. 2271–2285.

[21] Thomas L Szabo. *Diagnostic ultrasound imaging: inside out*. Academic press, 2004. Chap. Doppler Modes.

[22] S Campbell et al. “New Doppler technique for assessing uteroplacental blood flow”. In: *The lancet* 321.8326 (1983), pp. 675–677.

[23] Emilie Mace et al. “Functional ultrasound imaging of the brain: theory and basic principles”. In: *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 60.3 (2013), pp. 492–506.

[24] Émilie Macé et al. “Whole-brain functional ultrasound imaging reveals brain modules for visuomotor integration”. In: *Neuron* 100.5 (2018), pp. 1241–1251.

[25] Claire Rabut et al. “4D functional ultrasound imaging of whole-brain activity in rodents”. In: *Nature methods* 16.10 (2019), pp. 994–997.

[26] Alberto L Vazquez and Douglas C Noll. “Nonlinear aspects of the BOLD response in functional MRI”. In: *Neuroimage* 7.2 (1998), pp. 108–118.

[27] Prakash Kara, Pamela Reinagel, and R Clay Reid. “Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons”. In: *Neuron* 27.3 (2000), pp. 635–646.

[28] Daniel A Handwerker, John M Ollinger, and Mark D’Esposito. “Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses”. In: *Neuroimage* 21.4 (2004), pp. 1639–1651.

[29] Suzana Herculano-Houzel. “The human brain in numbers: a linearly scaled-up primate brain”. In: *Frontiers in human neuroscience* (2009), p. 31.

[30] E.M. Ellis et al. “Shared and distinct retinal input to the mouse superior colliculus and dorsal lateral geniculate nucleus”. In: *J. Neurophysiol.,* 116 (2016), pp. 602–610.

[31] S.D. Gale and G.J. Murphy. “Active dendritic properties and local inhibitory input enable selectivity for object motion in mouse superior colliculus neurons”. In: *Journal of Neuroscience,* 36 (2016), pp. 9111–9123.

[32] S.Ito, D.A. Feldheim, and A.M. Litke. “Segregation of visual response properties in the mouse superior colliculus and their modulation during locomotion”. In: *J. Neurosci.,* 12 (2017), pp. 8428–8443.

[33] J. Zhuang et al. “An extended retinotopic map of mouse cortex”. In: *eLife,* 6 (2017), e18372.

[34] Marc Gesnik et al. “3D functional ultrasound imaging of the cerebral visual system in rodents”. In: *Neuroimage* 149 (2017), pp. 267–274.

[35] A. Niranjan et al. “fMRI mapping of the visual system in the mouse brain with interleaved snapshot GE-EPI”. In: *NeuroImage,* 139 (2016), pp. 337–345.

[36] R Matthew Hutchison et al. “Dynamic functional connectivity: promise, issues, and interpretations”. In: *Neuroimage* 80 (2013), pp. 360–378.

[37] Andrew Zalesky et al. “The relationship between regional and inter-regional functional connectivity deficits in schizophrenia”. In: *Human brain mapping* 33.11 (2012), pp. 2535–2549.

[38] Marcus E Raichle et al. “A default mode of brain function”. In: *Proceedings of the national academy of sciences* 98.2 (2001), pp. 676–682.

[39] Qi-Hong Zou et al. “An improved approach to detection of amplitude of low-frequency fluctuation (ALFF) for resting-state fMRI: fractional ALFF”. In: *Journal of neuroscience methods* 172.1 (2008), pp. 137–141.

[40] Anil Man Tuladhar et al. “Default mode network connectivity in stroke patients”. In: *PloS one* 8.6 (2013), e66556.

[41] Niv Tik et al. “Predicting individual variability in task-evoked brain activity in schizophrenia”. In: *Human Brain Mapping* 42.12 (2021), pp. 3983–3992.

[42] Bruno-Félix Osmanski et al. “Functional ultrasound imaging of intrinsic connectivity in the living rat brain with high spatiotemporal resolution”. In: *Nature communications* 5.1 (2014), p. 5023.

[43] Sheldon M Ross. *Stochastic processes*. John Wiley & Sons, 1995.

[44] S. Soloukey et al. “Functional Ultrasound (fUS) During Awake Brain Surgery: The Clinical Potential of Intra-Operative Functional and Vascular Brain Mapping”. In: *Front. Neurosci.* 13 (2020), p. 1384.

[45] Helmut Laufs et al. “EEG-correlated fMRI of human alpha activity”. In: *Neuroimage* 19.4 (2003), pp. 1463–1476.

[46] Vince D Calhoun et al. “fMRI analysis with the general linear model: removal of latency-induced amplitude bias by incorporation of hemodynamic derivative terms”. In: *Neuroimage* 22.1 (2004), pp. 252–257.

[47] Paola Pinti et al. “Current status and issues regarding pre-processing of fNIRS neuroimaging data: an investigation of diverse signal filtering methods within a general linear model framework”. In: *Frontiers in human neuroscience* 12 (2019), p. 505.

[48] Richard NA Henson et al. “Detecting latency differences in event-related BOLD responses: application to words versus nonwords and initial versus repeated face presentations”. In: *Neuroimage* 15.1 (2002), pp. 83–97.

[49] Karl J. Friston et al. “Nonlinear Responses in fMRI: The Balloon Model, Volterra Kernels, and Other Hemodynamics”. In: *NeuroImage* 12 (2000), pp. 466–477.

[50] M. Hütel et al. “Hemodynamic matrix factorization for functional magnetic resonance imaging”. In: *NeuroImage,* 231 (2021), p. 117814.

[51] N. Bunzeck et al. “Scanning silence: mental imagery of complex sounds”. In: *NeuroImage* 26.4 (2005), pp. 1119–1127.

[52] Cesar Caballero Gaudes et al. “Detection and characterization of single-trial fMRI bold responses: Paradigm free mapping”. In: *Human brain mapping* 32.9 (2011), pp. 1400–1418.

[53] Fikret Işık Karahanoğlu et al. “Total activation: fMRI deconvolution through spatio-temporal regularization”. In: *Neuroimage* 73 (2013), pp. 121–134.

[54] Karthik Ramakrishnan Sreenivasan, Martin Havlicek, and Gopikrishna Deshpande. “Nonparametric hemodynamic deconvolution of fMRI using homomorphic filtering”. In: *IEEE transactions on medical imaging* 34.5 (2014), pp. 1155–1163.

[55] Guo-Rong Wu et al. “A blind deconvolution approach to recover effective connectivity brain networks from resting state fMRI data”. In: *Medical image analysis* 17.3 (2013), pp. 365–374.

[56] Hamza Cherkaoui et al. “Multivariate semi-blind deconvolution of fMRI time series”. In: *NeuroImage* 241 (2021), p. 118418.

[57] Jason W Bohland et al. “The brain atlas concordance problem: quantitative comparison of anatomical parcellations”. In: *PloS one* 4.9 (2009), e7200.

[58] Martin A Lindquist and Tor D Wager. “Validity and power in hemodynamic response modeling: A comparison study and a new approach”. In: *Human brain mapping* 28.8 (2007), pp. 764–784.

[59] Wellcome Trust Centre for Neuroimaging. *Statistical Parametric Mapping: The SPM software*. https://www.fil.ion.ucl.ac.uk/spm/software/. Version 12. 2020.

[60] K. L. West et al. “BOLD hemodynamic response function changes significantly with healthy aging”. In: *NeuroImage* 188 (2019), pp. 199–207.

[61] D. Rangaprakash et al. “FMRI hemodynamic response function (HRF) as a novel marker of brain function: applications for understanding obsessive-compulsive disorder pathology and treatment response”. In: *Brain Imaging Behav.,* 15 (June 2021), pp. 1622–1640.

[62] A. R. Mayer et al. “Investigating the properties of the hemodynamic response function after mild traumatic brain injury”. In: *J. Neurotrauma,* 31 (Jan. 2014), pp. 189–197.

[63] D. Asemani, H. Morsheddost, and M. A. Shalchy. “Effects of ageing and Alzheimer disease on haemodynamic response function: a challenge for event-related fMRI”. In: *Healthc. Technol. Lett.,* 4 (June 2017), pp. 109–114.

[64] Simon Van Eyndhoven et al. “Augmenting interictal mapping with neurovascular coupling biomarkers by structured factorization of epileptic EEG and fMRI data”. In: *NeuroImage,* 228 (2021), p. 117652.

[65] I. G. Elbau et al. “The brain’s hemodynamic response function rapidly changes under acute psychosocial stress in association with genetic and endocrine stress response markers”. In: *Proc. Natl. Acad. Sci. USA,* 115 (Oct. 2018), E10206–E10215.

[66] Ali-Kemal Aydin et al. “Transfer functions linking neural calcium to single voxel functional ultrasound signal”. In: *Nature communications* 11.1 (2020), p. 2954.

[67] Hamza Cherkaoui et al. “Sparsity-based Blind Deconvolution of Neural Activation Signal in FMRI”. In: *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2019, pp. 1323–1327.

[68] Mark W Woolrich, Timothy EJ Behrens, and Stephen M Smith. “Constrained linear basis sets for HRF modelling using Variational Bayes”. In: *NeuroImage* 21.4 (2004), pp. 1748–1761.

[69] KA Smitha et al. “Resting state fMRI: A review on methods in resting state connectivity analysis and resting state networks”. In: *The neuroradiology journal* 30.4 (2017), pp. 305–317.

[70] Felix Biessmann et al. “Analysis of multimodal neuroimaging data”. In: *IEEE reviews in biomedical engineering* 4 (2011), pp. 26–58.

[71] Kaiming Li et al. “Review of methods for functional brain connectivity detection using fMRI”. In: *Computerized medical imaging and graphics* 33.2 (2009), pp. 131–139.

[72] Roberto Viviani, Georg Grön, and Manfred Spitzer. “Functional principal component analysis of fMRI data”. In: *Human brain mapping* 24.2 (2005), pp. 109–129.

[73] Michael E Wall, Andreas Rechtsteiner, and Luis M Rocha. “Singular value decomposition and principal component analysis”. In: *A practical approach to microarray data analysis*. Springer, 2003, pp. 91–109.

[74] Vince D Calhoun and Tülay Adali. “Unmixing fMRI with independent component analysis”. In: *IEEE Engineering in Medicine and Biology Magazine* 25.2 (2006), pp. 79–90.

[75] Vince Daniel Calhoun et al. “Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms”. In: *Human brain mapping* 13.1 (2001), pp. 43–53.

[76] Irene Winkler, Stefan Haufe, and Michael Tangermann. “Automatic classification of artifactual ICA-components for artifact removal in EEG signals”. In: *Behavioral and brain functions* 7 (2011), pp. 1–15.

[77] Ludovica Griffanti et al. “ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging”. In: *Neuroimage* 95 (2014), pp. 232–247.

[78] Tülay Adali, Yuri Levin-Schwartz, and Vince D Calhoun. “Multimodal data fusion using source separation: Two effective models based on ICA and IVA and their properties”. In: *Proceedings of the IEEE* 103.9 (2015), pp. 1478–1493.

[79] Aybüke Erol and Borbála Hunyadi. “Tensors for neuroimaging: A review on applications of tensors to unravel the mysteries of the brain”. In: *Tensors for Data Processing* (2022), pp. 427–482.

[80] Andrzej Cichocki et al. “Tensor decompositions for signal processing applications: From two-way to multiway component analysis”. In: *IEEE signal processing magazine* 32.2 (2015), pp. 145–163.

[81] Borbála Hunyadi et al. “Block term decomposition for modelling epileptic seizures”. In: *EURASIP Journal on Advances in Signal Processing* 2014.1 (2014), pp. 1–19.

[82] Cyril Goutte et al. “On clustering fMRI time series”. In: *NeuroImage* 9.3 (1999), pp. 298–310.

[83] Farzad V Farahani, Waldemar Karwowski, and Nichole R Lighthall. “Application of graph theory for identifying connectivity patterns in human brain networks: a systematic review”. In: *frontiers in Neuroscience* 13 (2019), p. 585.

[84] Koji Sakai and Kei Yamada. “Machine learning studies on major brain diseases: 5-year trends of 2014–2018”. In: *Japanese journal of radiology* 37 (2019), pp. 34–72.

[85] Guohua Shen et al. “End-to-end deep image reconstruction from human brain activity”. In: *Frontiers in computational neuroscience* 13 (2019), p. 21.

[86] Noman Naseer and Keum-Shik Hong. “fNIRS-based brain-computer interfaces: a review”. In: *Frontiers in human neuroscience* 9 (2015), p. 3.

[87] Anwar Nunez-Elizalde et al. “Neural correlates of blood flow measured by ultrasound”. In: *Neuron* 110 (Mar. 2022), pp. 1–10.

[88] Julien Claron et al. “Large-scale functional ultrasound imaging of the spinal cord reveals in-depth spatiotemporal responses of spinal nociceptive circuits in both normal and inflammatory states”. In: *Pain* 162.4 (2021), p. 1047.

[89] S. K. E. Koekkoek et al. “High Frequency Functional Ultrasound in Mice”. In: *IEEE Int. Ultrason. Symp.* 2018, pp. 1–4.

[90] R. Rau et al. “3D functional ultrasound imaging of pigeons”. In: *NeuroImage* 183 (2018), pp. 469–477.

[91] K. Blaize et al. “Functional ultrasound imaging of deep visual cortex in awake nonhuman primates”. In: *Proc. Natl. Acad. Sci.* 117.25 (2020), pp. 14453–63.

[92] Sumner L. Norman et al. “Single-trial decoding of movement intentions using functional ultrasound neuroimaging”. In: *Neuron* 109.9 (2021), 1554–1566.e4.

[93] J. Baranger et al. “Bedside functional monitoring of the dynamic brain connectivity in human neonates”. In: *Nat. Commun.* 12.1080 (2021).

[94] M. Imbault et al. “Intraoperative Functional Ultrasound Imaging of Human Brain Activity”. In: *Sci. Rep.* 7.7304 (2017).

[95] Guillaume A Rousselet and Cyril R Pernet. “Improving standards in brain-behavior correlation analyses”. In: *Frontiers in human neuroscience* 6 (2012), p. 119.

[96] C. Brunner et al. “Whole-brain functional ultrasound imaging in awake head-fixed mice”. In: *Nat. Protoc.,* 16 (2021), pp. 3547–3571.

[97] Guang Ouyang, Werner Sommer, and Changsong Zhou. “Reconstructing ERP amplitude effects after compensating for trial-to-trial latency jitter: a solution based on a novel application of residue iteration decomposition”. In: *International Journal of Psychophysiology* 109 (2016), pp. 9–20.

[98] Yang-Whan Jeon and John Polich. “Meta-analysis of P300 and schizophrenia: Patients, paradigms, and practical implications”. In: *Psychophysiology* 40.5 (2003), pp. 684–701.

[99] Alexander Roth et al. “Increased event-related potential latency and amplitude variability in schizophrenia detected through wavelet-based single trial analysis”. In: *International Journal of Psychophysiology* 66.3 (2007), pp. 244–254.

[100] Karl J Friston et al. “Nonlinear event-related responses in fMRI”. In: *Magnetic resonance in medicine* 39.1 (1998), pp. 41–52.

[101] Nikos K. Logothetis. “The Underpinnings of the BOLD Functional Magnetic Resonance Imaging Signal”. In: *The Journal of Neuroscience* 23 (2003), pp. 3963–3971.

[102] G. M. Boynton et al. “Linear Systems Analysis of Functional Magnetic Resonance Imaging in Human V1”. In: *J. Neurosci.* 16.13 (1996), pp. 4207–4221.

[103] S. Ito, D. A. Feldheim, and A. M. Litke. “Segregation of visual response properties in the mouse superior colliculus and their modulation during locomotion”. In: *J. Neurosci.,* 37 (2017), pp. 8428–8443.

[104] Ana F Oliveira and Keisuke Yonehara. “The mouse superior colliculus as a model system for investigating cell type-based mechanisms of visual motor transformation”. In: *Frontiers in Neural Circuits* 12 (2018), p. 59.

[105] Jonathan Wray and Gary G. R. Green. “Calculation of the Volterra kernels of non-linear dynamic systems using an artificial neural network”. In: *Biological Cybernetics* 71 (2004), pp. 187–195.

[106] Aybüke Erol et al. “Deconvolution of the Functional Ultrasound Response in the Mouse Visual Pathway Using Block-Term Decomposition”. In: *Neuroinformatics* 21.2 (2023), pp. 247–265.

[107] Luca Iemi et al. “Spontaneous neural oscillations bias perception by modulating baseline excitability”. In: *Journal of Neuroscience* 37.4 (2017), pp. 807–819.

[108] Hanneke Van Dijk et al. “Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability”. In: *Journal of Neuroscience* 28.8 (2008), pp. 1816–1823.

[109] Katharina Limbach and Paul M Corballis. “Prestimulus alpha power influences response criterion in a detection task”. In: *Psychophysiology* 53.8 (2016), pp. 1154–1164.

[110] Yevgeniy B Sirotin and Aniruddha Das. “Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity”. In: *Nature* 457.7228 (2009), pp. 475–479.

[111] Qingqing Zhang et al. “High-frequency neuronal signal better explains multiphase BOLD response”. In: *Neuroimage* 268 (2023), p. 119887.

[112] Bastian S Generowicz et al. “Swept-3D Ultrasound Imaging of the Mouse Brain Using a Continuously Moving 1D-Array Part II: Functional Imaging”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* (2023).

[113] Wenqiang Yan and Yongcheng Wu. “A time-frequency denoising method for single-channel event-related EEG”. In: *Frontiers in Neuroscience* 16 (2022), p. 991136.

[114] Guang Ouyang and Changsong Zhou. “Characterizing the brain’s dynamical response from scalp-level neural electrical signals: a review of methodology development”. In: *Cognitive Neurodynamics* 14.6 (2020), pp. 731–742.

[115] M Andrea Pisauro, Andrea Benucci, and Matteo Carandini. “Local and global contributions to hemodynamic activity in mouse cortex”. In: *Journal of neurophysiology* 115.6 (2016), pp. 2931–2936.

[116] Simon Musall et al. “Single-trial neural dynamics are dominated by richly varied movements”. In: *Nature neuroscience* 22.10 (2019), pp. 1677–1686.

[117] Farshad Moradi et al. “Consistent and precise localization of brain activity in human primary visual cortex by MEG and fMRI”. In: *Neuroimage* 18.3 (2003), pp. 595–609.

[118] Karen J Mullinger et al. “Post-stimulus fMRI and EEG responses: evidence for a neuronal origin hypothesised to be inhibitory”. In: *Neuroimage* 157 (2017), pp. 388–399.

[119] Maggie W Guy et al. “Peak selection and latency jitter correction in developmental event-related potentials”. In: *Developmental psychobiology* 63.7 (2021), e22193.

[120] Hunar Abdulrahman and Richard N Henson. “Effect of trial-to-trial variability on optimal event-related fMRI design: Implications for Beta-series correlation and multi-voxel pattern analysis”. In: *NeuroImage* 125 (2016), pp. 756–766.

[121] Jeanette A Mumford et al. “Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses”. In: *Neuroimage* 59.3 (2012), pp. 2636–2643.

[122] David Bressler, Nicole Spotswood, and David Whitney. “Negative BOLD fMRI response in the visual cortex carries precise stimulus-specific information”. In: *PLoS One* 2.5 (2007), e410.

[123] Moshe Gur, Alexander Beylin, and D Max Snodderly. “Response variability of neurons in primary visual cortex (V1) of alert monkeys”. In: *Journal of Neuroscience* 17.8 (1997), pp. 2914–2920.

[124] C. Stringer et al. “Spontaneous behaviors drive multidimensional, brainwide activity”. In: *Science* 364.6437 (2019), p. 255.

[125] Michael D Fox et al. “Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses”. In: *Nature neuroscience* 9.1 (2006), pp. 23–25.

[126] Abd-Krim Seghouane and Adnan Shah. “HRF Estimation in fMRI Data With an Unknown Drift Matrix by Iterative Minimization of the Kullback–Leibler Divergence”. In: *IEEE Trans. Med. Imag.,* 31.2 (2012), pp. 192–206.

[127] W. L. Ganis G.and Thopson and S. M. Kosslyn. “Brain areas underlying visual mental imagery and visual perception: an fMRI study”. In: *Cogn. Brain Res.* 20.2 (2004), pp. 226–241.

[128] F. I. Karahanoğlu and D. Van De Ville. “Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks”. In: *Nat. Commun.* 6 (2015), p. 7751.

[129] Eneko Uruñuela et al. *Hemodynamic Deconvolution Demystified: Sparsity-Driven Regularization at Work*. 2021. arXiv: 2107.12026 [q-bio.NC].

[130] Ganesh R Naik and Wenwu Wang. “Blind source separation”. In: *Berlin: Springer* 10 (2014).

[131] K. J. Friston. “Functional integration and inference in the brain”. In: *Prog. Neurobiol.,* 68.2 (2002), pp. 113–143.

[132] Karl J Friston et al. “Event-related fMRI: characterizing differential responses”. In: *Neuroimage* 7.1 (1998), pp. 30–40.

[133] G. K. Aguirre, E. Zarahn, and M. D’esposito. “The variability of human, BOLD hemodynamic responses”. In: *NeuroImage* 8.4 (1998), pp. 360–369.

[134] P. Fransson et al. “Temporal and spatial MRI responses to subsecond visual activation”. In: *Magn. Reson. Imaging* 17.1 (1999), pp. 1–7.

[135] Martin Lindquist et al. “Modeling the hemodynamic response function in fMRI: Efficiency, bias and mis-modeling”. In: *NeuroImage,* 45.1 (2009), S187–S198.

[136] Gary H Glover. “Deconvolution of impulse response in event-related BOLD fMRI1”. In: *Neuroimage* 9.4 (1999), pp. 416–429.

[137] J. E. Chen et al. “Investigating mechanisms of fast BOLD responses: The effects of stimulus intensity and of spatial heterogeneity of hemodynamics”. In: *NeuroImage* 245 (2021), p. 118658.

[138] Laurent Sorber, Marc Van Barel, and Lieven De Lathauwer. “Structured data fusion”. In: *IEEE Journal of Selected Topics in Signal Processing* 9.4 (2015), pp. 586–600.

[139] N. Mitianoudis and M. E. Davies. “Audio source separation of convolutive mixtures”. In: *IEEE Trans. Audio Speech Lang. Process.,* 11 (2003), pp. 489–497.

[140] G. Marrelec et al. “Robust Bayesian estimation of the hemodynamic response function in event-related BOLD fMRI using basic physiological information”. In: *Hum. Brain Mapp.,* 19 (2003), pp. 1–17.

[141] C. Huang et al. “Debiasing-Based Noise Suppression for Ultrafast Ultrasound Microvessel Imaging”. In: *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* 66 (2019), pp. 1281–1291.

[142] H. Bousbia-Salah, A. Belouchrani, and K. Abed-Meriam. “Jacobi-like algorithm for blind signal separation of convolutive mixtures”. In: *Electron. Lett.,* 37 (2001), pp. 1049–1050.

[143] F. Van Eeghem and L. De Lathauwer. “Second-order tensor-based convolutive ICA: Deconvolution versus tensorization”. In: *Proc. Int. Conf. Acoust. Speech, Signal Process.* (2017), pp. 2252–2256.

[144] N. Vervliet et al. “Tensorlab 3.0 [Online]. Available: https://www.tensorlab.net/”. In: (2016).

[145] Laurent Sorber, Marc Van Barel, and Lieven De Lathauwer. “Optimization-Based Algorithms for Tensor Decompositions: Canonical Polyadic Decomposition, Decomposition in Rank-$(L_r,L_r,1)$ Terms, and a New Generalization”. In: *SIAM Journal on Optimization* 23.2 (2013), pp. 695–720.

[146] Dijun Luo, Heng Huang, and Chris Ding. “Are Tensor Decomposition Solutions Unique? On the Global Convergence HOSVD and ParaFac Algorithms”. In: *Proc. PAKDD* (Feb. 2011), pp. 148–159.

[147] Simon Van Eyndhoven et al. “Identifying Stable Components of Matrix /Tensor Factorizations via Low-Rank Approximation of Inter-Factorization Similarity”. In: *2019 27th European Signal Processing Conference (EUSIPCO)*. 2019, pp. 1–5.

[148] J. Himberg, A. Hyvärinen, and F. Esposito. “Validating the independent components of neuroimaging time series via clustering and visualization”. In: *NeuroImage* 22.3 (2004), pp. 1214–1222.

[149] N. Otsu. “A Threshold Selection Method from Gray-Level Histograms”. In: *IEEE Trans. Syst. Man Cybern.* 9.1 (1979), pp. 62–66.

[150] N.Bolshakova and F.Azuaje. “Cluster validation techniques for genome expression data”. In: *Signal Processing,* 83.4 (2003), pp. 825–833.

[151] J. W. Demmel. “Applied Numerical Linear Algebra”. In: *SIAM* (1997).

[152] A. Sano. “Optimally regularized inverse of singular value decomposition and application to signal extrapolation”. In: *Signal Process.* 30.2 (1993), pp. 163–176.

[153] J. H. Marshel et al. “Functional specialization of seven mouse visual cortical areas”. In: *Neuron* 72.6 (2011), pp. 1040–1054.

[154] Allen Institute for Brain Science. *Allen Brain Atlas API*. Available from: brain-map.org/api/index.html. 2015.

[155] Pengfei Song et al. "Ultrasound Small Vessel Imaging With Block-Wise Adaptive Local Clutter Filtering". In: *IEEE Transactions on Medical Imaging* 36.1 (2017), pp. 251–262.

[156] R. Sala-Llonch et al. "Spatial Parcellations, Spectral Filtering, and Connectivity Measures in fMRI: Optimizing for Discrimination". In: *Hum. Brain Mapp.* 40.2 (2019), pp. 407–419.

[157] N. Correa et al. "Comparison of blind source separation algorithms for fMRI using a new Matlab toolbox: GIFT". In: *Proceedings. (ICASSP '05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.* Vol. 5. 2005, pp. v/401–v/404.

[158] A. T. Winder et al. "Weak correlations between hemodynamic signals and ongoing neural activity during the resting state". In: *Nat. Neurosci.* 20.12 (2017), pp. 1761–1760.

[159] A. K. Aydin et al. "Iliski, a software for robust calculation of transfer functions". In: *PLoS Comput. Biol.* 17.6 (2021), e1008614.

[160] K. B. J. Franklin and G. Paxinos. *The mouse brain in stereotaxic coordinates*. Second. Academic Press, 2001.

[161] D. A. Rahnev et al. "Prestimulus hemodynamic activity in dorsal attention network is negatively associated with decision confidence in visual perception". In: *J. Neurophysiol.* 108.5 (2012), pp. 1529–1536.

[162] K. J. Friston, L. Harrison, and W. Pennya. "Dynamic Causal Modelling". In: *Neuroimage* 19 (2003), pp. 1273–1302.

[163] Laura D Lewis et al. "Stimulus-dependent hemodynamic response timing across the human subcortical-cortical visual pathway identified through high spatiotemporal resolution 7T fMRI". In: *Neuroimage* 181 (2018), pp. 279–291.

[164] P. Kara, P Reinagel, and R. C. Reid. "Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons". In: *Neuron* 27.3 (2000), pp. 635–646.

[165] A. F. Meyer, J. O'Keefe, and J. Poort. "Two distinct types of eye-head coupling in freely moving mice". In: *Curr. Biol.*, 30 (2020), pp. 2116–2130.

[166] M. Gur and D. M. Snodderly. "Visual receptive fields of neurons in primary visual cortex (V1) move in space with the eye movements of fixation". In: *Vision Res.*, 37 (1997), pp. 257–265.

[167] Cristopher M Niell and Michael P Stryker. "Modulation of visual responses by behavioral state in mouse visual cortex". In: *Neuron* 65.4 (2010), pp. 472–479.

[168] Florentin Wörgötter et al. "State-dependent receptive-field restructuring in the visual cortex". In: *Nature* 396.6707 (1998), pp. 165–168.

[169] M. Ahmadlou, L. S. Zweifel, and J. A. Heimel. “Functional modulation of primary visual cortex by the superior colliculus in the mouse”. In: *Nat. Commun.* 9 (2018), p. 3895.

[170] Onkar S Dhande et al. “Contributions of retinal ganglion cells to subcortical visual processing and behaviors”. In: *Annual review of vision science* 1.1 (2015), pp. 291–328.

[171] Patrick J Drew. “Neurovascular coupling: motive unknown”. In: *Trends in neurosciences* 45.11 (2022), pp. 809–819.

[172] Ying Ma et al. “Resting-state hemodynamics are spatiotemporally coupled to synchronized and symmetric neural activity in excitatory neurons”. In: *Proceedings of the National Academy of Sciences* 113.52 (2016), E8463–E8471.

[173] John-Dylan Haynes and Geraint Rees. “Decoding mental states from brain activity in humans”. In: *Nature reviews neuroscience* 7.7 (2006), pp. 523–534.

[174] Amanda J Taylor, Jung Hwan Kim, and David Ress. “Characterization of the hemodynamic response function across the majority of human cerebral cortex”. In: *NeuroImage* 173 (2018), pp. 322–331.

[175] C. Lau et al. “BOLD temporal dynamics of rat superior colliculus and lateral geniculate nucleus following short duration visual stimulation”. In: *PLoS One* 6.4 (2011), e18914.

[176] J. Wang et al. “Visual receptive field properties of neurons in the superficial superior colliculus of the mouse”. In: *J. Neurosci.,* 30 (2010), pp. 16573–16584.

[177] S. Inayat et al. “Neurons in the most superficial lamina of the mouse superior colliculus are highly selective for stimulus direction”. In: *J. Neurosci.,* 35 (2015), pp. 7992–8003.

[178] X. Bai and B. He. “Estimation of Number of Independent Brain Electric Sources From the Scalp EEGs”. In: *IEEE Trans. Biomed. Eng.* 53.10 (2006), pp. 1883–1892.

[179] A. Erol et al. “Joint estimation of hemodynamic response and stimulus function in functional ultrasound using convolutive mixtures”. In: *Proc. Asilomar Conf. Signals, Syst. Comput.* (2020), pp. 246–250.

[180] Martin M Monti. “Statistical analysis of fMRI time-series: a critical review of the GLM approach”. In: *Frontiers in human neuroscience* 5 (2011), p. 28.

[181] Marta I Garrido et al. “Repetition suppression and plasticity in the human brain”. In: *Neuroimage* 48.1 (2009), pp. 269–279.

[182] Gabor Stefanics et al. “Timing of repetition suppression of event-related potentials to unattended objects”. In: *European Journal of Neuroscience* 52.11 (2020), pp. 4432–4441.

[183] T.-P. Jung et al. “Imaging brain dynamics using independent component analysis”. In: *Proc. IEEE* 89.7 (2001), pp. 1107–1122.

[184] Borbála Hunyadi et al. “ICA extracts epileptic sources from fMRI in EEG-negative patients: a retrospective validation study”. In: *PloS one* 8.11 (2013), e78796.

[185] Ludovica Griffanti et al. “Hand classification of fMRI ICA noise components”. In: *Neuroimage* 154 (2017), pp. 188–205.

[186] Gholamreza Salimi-Khorshidi et al. “Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers”. In: *Neuroimage* 90 (2014), pp. 449–468.

[187] Borbála Hunyadi et al. “A prospective fMRI-based technique for localising the epileptogenic zone in presurgical evaluation of epilepsy”. In: *Neuroimage* 113 (2015), pp. 329–339.

[188] Bastian S Generowicz et al. “Swept-3D Ultrasound Imaging of the Mouse Brain Using a Continuously Moving 1D-Array Part I: Doppler Imaging”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* (2023).

[189] Arne F Meyer et al. “Models of neuronal stimulus-response functions: elaboration, estimation, and evaluation”. In: *Frontiers in systems neuroscience* 10 (2017), p. 109.

[190] Michael X Cohen. “A tutorial on generalized eigendecomposition for denoising, contrast enhancement, and dimension reduction in multichannel electrophysiology”. In: *Neuroimage* 247 (2022), p. 118809.

[191] E Kelly Buchanan et al. “Penalized matrix decomposition for denoising, compression, and improved demixing of functional imaging data”. In: *BioRxiv* (2018), p. 334706.

[192] Amir A Khaliq, Ijaz M Qureshi, and Jawad A Shah. “Unmixing functional magnetic resonance imaging data using matrix factorization”. In: *International Journal of Imaging Systems and Technology* 22.4 (2012), pp. 195–199.

[193] Martin A Lindquist. “The statistical analysis of fMRI data”. In: (2008).

[194] Notger G Müller et al. “Repetition suppression versus enhancement—it’s quantity that matters”. In: *Cerebral cortex* 23.2 (2013), pp. 315–322.

[195] Ronald Sladky et al. “Slice-timing effects and their correction in functional MRI”. In: *Neuroimage* 58.2 (2011), pp. 588–594.

[196] David Parker, Xueqing Liu, and Qolamreza R Razlighi. “Optimal slice timing correction and its interaction with fMRI parameters and artifacts”. In: *Medical image analysis* 35 (2017), pp. 434–445.

[197] A Erol et al. “GLM-Regularized Low-Rank Factorization for Extracting Functional Response from Swept-3D Functional Ultrasound”. In: *2023 IEEE Data Science and Learning Workshop (DSLW)*. 2023.

[198] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*. Vol. 4. 4. Springer, 2006. Chap. 3.

[199] Valeria Della-Maggiore et al. “An empirical comparison of SPM preprocessing parameters to the analysis of fMRI data”. In: *Neuroimage* 17.1 (2002), pp. 19–28.

[200] David R Anderson, Dennis J Sweeney, and Thomas A Wlliams. “Statistics for business and economics 13th Edition”. In: (2011), p. 676.

[201] Jonas Larsson and Andrew T Smith. “fMRI repetition suppression: neuronal adaptation or stimulus expectation?” In: *Cerebral cortex* 22.3 (2012), pp. 567–576.

[202] Arvind Balachandrasekaran et al. “Reducing the effects of motion artifacts in fMRI: A structured matrix completion approach”. In: *IEEE Transactions on Medical Imaging* 41.1 (2021), pp. 172–185.

[203] V Calhoun, Xavier Golay, and Godfrey Pearlson. “Improved fMRI slice timing correction: interpolation errors and wrap around effects”. In: *Proceedings, ISMRM, 9th annual meeting, Denver*. Vol. 810. 2000.

[204] Fabien Cignetti et al. “Pros and cons of using the informed basis set to account for hemodynamic response variability with developmental data”. In: *Frontiers in neuroscience* 10 (2016), p. 322.

[205] Christos Chatzichristos et al. “Blind fMRI source unmixing via higher-order tensor decompositions”. In: *Journal of neuroscience methods* 315 (2019), pp. 17–47.

[206] Mansour Nejati et al. “Denoising by low-rank and sparse representations”. In: *Journal of Visual Communication and Image Representation* 36 (2016), pp. 28–39.

[207] Andrew P Bagshaw and Tracy Warbrick. “Single trial variability of EEG and fMRI responses to visual stimuli”. In: *Neuroimage* 38.2 (2007), pp. 280–292.

[208] Aniruddha Das, Kevin Murphy, and Patrick J Drew. “Rude mechanicals in brain haemodynamics: non-neural actors that influence blood flow”. In: *Philosophical Transactions of the Royal Society B* 376.1815 (2021), p. 20190635.

[209] Tracy Warbrick. “Simultaneous EEG-fMRI: what have we learned and what does the future hold?” In: *Sensors* 22.6 (2022), p. 2262.

[210] Gary H Glover. “Overview of functional magnetic resonance imaging”. In: *Neurosurgery Clinics* 22.2 (2011), pp. 133–139.

[211] Anders M Dale et al. “Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity”. In: *neuron* 26.1 (2000), pp. 55–67.

[212] Lim-Anna Sieu et al. “EEG and functional ultrasound imaging in mobile rats”. In: *Nature methods* 12.9 (2015), pp. 831–834.

[213] Vince D Calhoun et al. “Neuronal chronometry of target detection: fusion of hemodynamic and event-related potential data”. In: *Neuroimage* 30.2 (2006), pp. 544–553.

[214] Borbála Hunyadi et al. “Fusion of electroencephalography and functional magnetic resonance imaging to explore epileptic network activity”. In: *2016 24th European Signal Processing Conference (EUSIPCO)*. IEEE. 2016, pp. 240–244.

[215] John Martindale et al. “Long duration stimuli and nonlinearities in the neural–haemodynamic coupling”. In: *Journal of Cerebral Blood Flow & Metabolism* 25.5 (2005), pp. 651–661.

[216] Alexander Maier et al. “Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey”. In: *Nature neuroscience* 11.10 (2008), pp. 1193–1200.

[217] Mads Dyrholm, Scott Makeig, and Lars Kai Hansen. “Model selection for convolutive ICA with an application to spatiotemporal analysis of EEG”. In: *Neural computation* 19.4 (2007), pp. 934–955.

[218] Hoang-Lan Nguyen Thi and Christian Jutten. “Blind source separation for convolutive mixtures”. In: *Signal processing* 45.2 (1995), pp. 209–229.

[219] Alexey Ozerov and Cédric Févotte. “Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation”. In: *IEEE transactions on audio, speech, and language processing* 18.3 (2009), pp. 550–563.

[220] Yitong Tao. *Frequency Domain Joint Estimation of HRF and Stimulus from fUS Data.* [Master’s Thesis, Delft University of Technology]. 2022.

[221] Farnaz Sedighin et al. “Multimodal soft nonnegative matrix co-factorization for convolutive source separation”. In: *IEEE Transactions on Signal Processing* 65.12 (2017), pp. 3179–3190.

[222] Laura M Parkes, Marcel CM Bastiaansen, and David G Norris. “Combining EEG and fMRI to investigate the post-movement beta rebound”. In: *Neuroimage* 29.3 (2006), pp. 685–696.

[223] Lin Yang, Zhongming Liu, and Bin He. “EEG-fMRI reciprocal functional neuroimaging”. In: *Clinical Neurophysiology* 121.8 (2010), pp. 1240–1250.

[224] Ying Zheng et al. “A dynamic model of neurovascular coupling: Implications for blood vessel dilation and constriction”. In: *NeuroImage* 52.3 (2010), pp. 1135–1147.

[225] John HR Maunsell et al. “Visual response latencies of magnocellular and parvocellular LGN neurons in macaque monkeys”. In: *Visual neuroscience* 16.1 (1999), pp. 1–14.

[226] BS Generowicz et al. “Efficient and flexible spatiotemporal clutter filtering of high frame rate images using subspace tracking”. In: *2018 IEEE International Ultrasonics Symposium (IUS)*. IEEE. 2018, pp. 206–212.

[227] Sadaf Soloukey et al. “Subband processing of the Doppler signal in functional Ultrasound (fUS) brain imaging”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* (In submission).

[228] Natalia Petridou et al. “Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity”. In: *Human brain mapping* 34.6 (2013), pp. 1319–1329.

[229] Roger BH Tootell et al. “Functional analysis of primary visual cortex (V1) in humans”. In: *Proceedings of the National Academy of Sciences* 95.3 (1998), pp. 811–817.

[230] George H Denfield et al. “Attentional fluctuations induce shared variability in macaque primary visual cortex”. In: *Nature communications* 9.1 (2018), p. 2654.

[231] Charlie Demene et al. “Functional ultrasound imaging of brain activity in human newborns”. In: *Science translational medicine* 9.411 (2017), eaah6756.

GLOSSARY

NOTATION

SETS

$\mathbb{R}$	Real numbers.
$\mathbb{R}^I$	Real length-I vectors.
$\mathbb{R}^{I \times J}$	Real $I \times J$ matrices.
$\mathbb{R}^{I_1 \times I_2 \times \ldots \times I_D}$	Real $I_1 \times I_2 \times \ldots \times I_D$ tensors.

VECTORS, MATRICES AND TENSORS

x, X	Plain lowercase and uppercase letters denote scalar.
$\mathbf{x}$	Lowercase boldface letters denote vectors.
$\mathbf{X}$	Uppercase boldface letters denote matrices.
$\boldsymbol{\mathcal{X}}$	Calligraphic letters denote tensors.
$\mathbf{X}^{\mathrm{T}}$	Transpose of matrix $\mathbf{X}$.
$\mathbf{X}^{\dagger} := (\mathbf{X}^H \mathbf{X})^{-1} \mathbf{X}^H$	Pseudo inverse (or the left-inverse) of a full-column rank tall matrix $\mathbf{X}$.
$\mathbf{x} \circ \mathbf{y}$	Outer product between vectors $\mathbf{x}$ and $\mathbf{y}$.
$\boldsymbol{\mathcal{X}} \times_d \mathbf{Y}$	d-mode (tensor-matrix) product between tensor $\mathcal{X}$ and matrix $\mathbf{Y}$.

NORMS

$\|\mathbf{x}\|_1$	ℓ_1-norm of vector $\mathbf{x}$.
$\|\mathbf{x}\|_2$	Euclidean (or ℓ_2-)norm of vector $\mathbf{x}$.

STOCHASTIC PROCESSES

$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	Gaussian distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.

SUMMARY

The brain stands as the most powerful processor in the known universe. It generates a continuous stream of electrical and chemical signals that underpin every thought, sensation, and action. Our past efforts in decoding these signals have made it possible to diagnose and treat many neurological disorders, helped us gain a deeper understanding of cognitive processes and consciousness, and paved the way for brain-computer interfaces. To take another step forward in our long but rewarding journey of discovering the brain's complex organization, we rely on advances in imaging technologies and signal processing.

Functional ultrasound is a neuroimaging technique that has emerged in the last decade, and has gained remarkable attention since then. The popularity of this technique stems from its portability, high resolution and affordability. Functional ultrasound can detect subtle fluctuations in local blood dynamics, which serve as delayed indicators of the underlying changes in neuronal activity. The goal of this thesis is to develop novel signal models and processing algorithms that can reveal the spatial and temporal characteristics of hemodynamic activity induced by external stimuli using functional ultrasound.

Existing techniques that explore how the brain reacts in response to stimuli model the design variables using a linear time-invariant system with binarized input representations, marking when a stimulus is on or off. However, experimental evidence suggests that the brain reacts in a more intricate manner. While some regions exhibit consistent responses to repeated stimuli, others can show substantial variation even when exposed to the same stimulus seconds apart. In our in-vivo experiments, we particularly focus on key regions within the mouse visual processing pathway, which are analogous to those in the human brain. We track how visual information flows across these areas, and propose methods that can incorporate the spatiotemporal variability of brain responses when identifying evoked activity. Using these methods, we show that functional ultrasound can capture the dynamic nature of brain responses with high spatial and temporal resolution, and provide us with further insights into the functional organization of the brain. Future directions of this dissertation include multimodal processing of the functional ultrasound signal together with neuronal activity, aiming to enhance our understanding of neurovascular coupling.

SAMENVATTING

Het brein staat bekend als de krachtigste processor in het bekende universum. Het genereert een continue stroom van elektrische en chemische signalen die zorgen voor elke gedachte, sensatie, en actie. Onze eerdere inspanningen om deze signalen te ontcijferen hebben het mogelijk gemaakt om vele neurologische aandoeningen te diagnosticeren en behandelen, hebben ons geholpen een dieper inzicht te krijgen in cognitieve processen en bewustzijn, en hebben de weg vrijgemaakt voor brein-computer interfaces. Om een volgende stap te zetten in onze lange maar lonende reis van het ontdekken van de complexe organisatie van het brein, vertrouwen we op vooruitgang in beeldvormende technologieën en signaalverwerking.

Functionele echografie is een neuroimaging techniek die in het afgelopen decennium is opgekomen en sindsdien opmerkelijke aandacht heeft gekregen. De populariteit van deze techniek komt voort uit de draagbaarheid, hoge resolutie, en betaalbaarheid. Functionele echografie kan subtiele fluctuaties in lokale bloeddynamiek detecteren, die dienen als vertraagde indicatoren van de onderliggende veranderingen in neuronale activiteit. Het doel van dit proefschrift is om nieuwe signaalmodellen en verwerkingsalgoritmen te ontwikkelen die de ruimtelijke en temporele kenmerken van hemodynamische activiteit, geïnduceerd door externe stimuli, te onthullen met behulp van functionele echografie.

Bestaande technieken die kijken hoe het brein reageert op prikkels, gebruiken een eenvoudig model waarin wordt aangegeven of een prikkel aan of uit staat. Experimenteel bewijs suggereert echter dat het brein op een meer ingewikkelde manier reageert. Terwijl sommige hersengebieden consistente reacties vertonen op herhaalde stimuli, kunnen andere hersengebieden aanzienlijke variatie vertonen, zelfs wanneer ze worden blootgesteld aan dezelfde stimulus met slechts enkele seconden tussenpozen. In onze in-vivo experimenten richten we ons met name op belangrijke gebieden binnen de visuele verwerkingsroute van de muizenhersenen, die analoog zijn aan die in de menselijke hersenen. We volgen hoe visuele informatie door deze gebieden stroomt en stellen methoden voor die de spatiotemporele variabiliteit van hersenreacties kunnen incorporeren bij het identificeren van opgewekte activiteit. Met behulp van deze methoden tonen we aan dat functionele echografie de dynamische aard van hersenreacties kan vastleggen met hoge ruimtelijke en temporele resolutie, en ons verder inzicht kan geven in de functionele organisatie van de hersenen. Toekomstige richtingen van dit proefschrift omvatten multimodale verwerking van het functionele echografiesignaal samen met neuronale activiteit, met als doel ons begrip van neurovasculaire koppeling te verbeteren.

CURRICULUM VITÆ

AYBÜKE EROL was born in 1995 in Ankara, Turkey. She understands that there is a dispute over 1995, but she considers herself a proud Millennial. She earned her B.Sc. in Electrical and Electronics Engineering from Middle East Technical University (METU) in Ankara. During this time, she realized her passion towards signal processing, and obtained her M.Sc. in signal processing in METU in radar emitter identification while working as a teaching assistant in the meantime. One day, she decided to travel to The Netherlands as a tourist with her friends, which made her realize that she would like to live there for some time. The picture below is from that trip, which changed her life. Shortly after, she joined the Signal Processing Systems (SPS) group of Delft University of Technology as a Ph.D. student.

Although her general scientific interest lies in signal processing and developing mathematical algorithms, she feels particularly motivated when she sees an intrinsic value to her work. While she is aware that she helped no actual living being with her Ph.D., she was still motivated by the possibility of that happening. When writing this CV, she noticed that she does not enjoy writing about herself, but she does enjoy travelling, drawing and watching a TV series. Yes, she only watches one TV series over and over, and that TV series is Friends. She believes that she now looks more than 5 years older compared to the picture on the right, and she partially blames Ph.D. for this. If you disagree, you should tell her. That will make her happy.