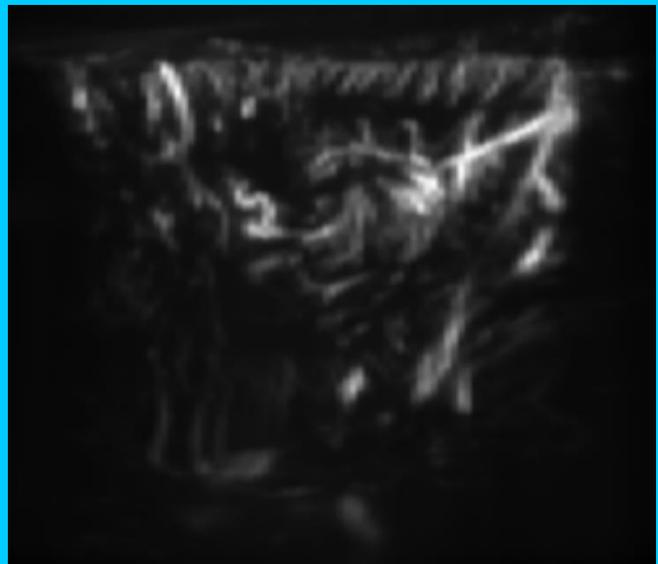


Advanced fUS Signal

Processing:

Exploiting Full-Detail Data and
Frequency Information Through
Tensor-Based Blind Source Separation

J.E. Bakker



Advanced fUS Signal Processing: Exploiting Full-Detail Data and Frequency Information Through Tensor-Based Blind Source Separation

by

Elise, J.E., Bakker

to obtain the degree of Master of Science in Biomedical Engineering, Track Medical Devices,
at the Delft University of Technology,
to be defended publicly on February 7, 2025 at 14:00 PM.

Student number: 4653599
Project duration: March 1, 2024 – February 7, 2025
Thesis committee: Dr. B. Hunyadi, TU Delft, supervisor
Ir. R. Wijnands, TU Delft, (daily) supervisor
Dr. Ir. P. Kruizinga, Erasmus Medical Center (CUBE)
Dr. Ir. R.C. Hendriks, TU Delft

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

In the last decade, the new functional neuroimaging technique functional ultrasound (fUS) has emerged as a potential new tool for clinical and neuroscientific applications. Unlike several conventional methods for functional brain imaging, fUS offers an unparalleled combination of submillimeter-subsecond spatiotemporal resolution and the ability to penetrate deep into brain tissue and capture large areas of interest. This makes fUS an exceptionally valuable tool for investigating brain function. The principal physiological mechanism utilized by fUS imaging is neurovascular coupling, which connects neuronal activity with local hemodynamic changes. Within the brain, increased neuronal activity results in the dilation of nearby blood vessels, subsequently leading to an increased local blood flow and volume. In fUS, the increase in blood volume is measured by the intensity of the Doppler signal caused by the Doppler effect.

To obtain relevant information from functional neuroimaging data, blind source separation (BSS) is often used. BSS is the separation of a set of source signals from a set of mixed signals, without the aid of information about the source signals or the mixing process. A commonly used BSS method for identifying brain networks and artifacts in functional neuroimaging techniques like fMRI is independent component analysis, which is a low-rank matrix decomposition. This work explores the use of canonical polyadic decomposition (CPD), a low-rank tensor decomposition, for BSS of fUS data. The CPD approximates a 3rd-order tensor, consisting of a space, frequency, and time dimension, by a sum of rank-1 tensors. The proposed method offers three key advantages that make it highly relevant to the field of fUS signal processing: 1) Tensor-based BSS by CPD allows for frequency as a third dimension to complement the spatial and temporal dimensions; 2) Tensor-based BSS by CPD allows for the use of all raw fUS data (compound images) rather than only power doppler images, exploiting more available information; and 3) Tensor-based BSS by CPD allows for a BSS method without strong constraints such as the case with matrix-based methods like independent component analysis.

The full processing pipeline developed in this work consists of four distinct stages. In stage one, pre-processing is carried out through SVD filtering, temporal demeaning, and pixel-based normalization. During the SVD filtering, normalization of the singular values is performed, which is novel. In stage two, compression is performed using truncated multilinear singular value decomposition. In stage three, BSS is carried out using CPD on the compressed data. In stage four, stable components are extracted by running the CPD 100 times and clustering the resulting components. These clusters were then averaged to create mean components. The CPD rank was estimated based on two quantifiable aspects: 1) Correlation between mean components and 2) Frequency of occurrence of mean components.

The method is entirely data-driven, can be applied to entire raw fUS datasets, and can run on a laptop with standard RAM. Application to task experiment data of a mouse verified that activity that is temporally correlated with the stimulus can be extracted in expected regions. The components also indicate expected functional connectivity, which is often not the case for matrix-based methods. Moreover, frequency spectra showed different characteristics for different types of components, which displays the relevance of this dimension for BSS, especially for nuisance components.

In conclusion, the proposed tensor-based blind source separation pipeline for raw fUS data was able to identify artifacts and meaningful neurological components based on distinctive characteristics in the temporal, spatial, and spectral domains. This work serves as a starting point for more advanced denoising, the identification of novel brain networks, and the comparison of brain networks across healthy and pathological conditions. Nevertheless, further research is necessary to ensure the utility of the method.

Acknowledgements

This thesis marks the end of my student days, which I hold onto as great memories. Starting out with a bachelor's in Clinical Technology, I have grown more into the engineering side, and ended with my thesis in the section of Signal Processing Systems. Nevertheless, I also very much enjoyed the 'side steps' towards (health) economics at Erasmus University and Leiden University during my minor and second master's. Along the way, I have met many inspiring people, lifelong friends, and my partner, Julian.

Bori and Ruben, I would like to thank you for your great guidance, help, and input during the project. I have learned a lot from you and always felt like I could come to you with my questions and ideas. I especially want to thank you for the inspiring discussions, your questions, and your enthusiasm for the project, which have kept me motivated.

Friends, family, and Julian, thank you for supporting me during my studies and thesis, giving advice, listening to all the things I was worried about, and providing necessary distractions.

I feel like this thesis really encompasses the field of Biomedical Engineering, combining advanced signal processing techniques with a large emphasis on biomedical interpretation. To me, the brain remains one of the most intriguing parts of the human body, and I hope to have contributed a tiny bit to understanding more of its secrets.

Elise
Delft, February 2025

Contents

1	Introduction	1
1.1	Why functional ultrasound?	1
1.2	Neurovascular coupling	1
1.3	Doppler effect.	2
1.4	Acquisition	2
1.4.1	Clutter filtering	2
1.5	Blind Source separation	3
1.5.1	Independent component analysis	3
1.6	Contribution.	3
1.6.1	Proposed model	3
1.6.2	Advantages	4
1.7	Problem statement	6
2	Methodology	7
2.1	Preliminaries	7
2.2	Data	7
2.3	Pipeline	8
2.3.1	Filtering	8
2.3.2	Temporal demeaning	8
2.3.3	Pixel-based normalization	9
2.3.4	Truncated multilinear singular value decomposition	9
2.3.5	Canonical polyadic decomposition.	10
2.3.6	Stable component extraction.	11
2.4	Validation and rank estimation	12
2.5	Assessing neuroscientific relevance.	13
3	Results	15
3.1	Identification of ROIs	15
3.2	Compression	15
3.3	CPD and rank estimation.	16
3.3.1	CPD iterations	16
3.3.2	Clustering	16
3.3.3	Normalization of singular values	16
3.3.4	Rank estimation	17
3.4	Neuroscientific analysis	19
3.4.1	Slice A.	20
3.4.2	Slice B.	24
3.4.3	Slice C	26
4	Discussion	29
5	Conclusion	33
A	Parameters clustering algorithm	39
B	SVD filtering threshold	41
C	Global Signal Regression	43
D	Spatial maps members cluster 1, slice A	45
E	Masking of slice A	47

List of Figures

1.1	Main functional imaging techniques	1
1.2	Acquisition pipeline of fUS	2
1.3	Brain voxel and corresponding frequency spectrum	3
1.4	Comparison between matrix-based and tensor-based approach	5
2.1	Set-up experiment	8
2.2	Processing pipeline	8
2.3	Rank estimation flow	13
3.1	Thresholded correlation images overlaid against the mean PDI for the different slices	15
3.2	Thresholded correlation images of the data after pre-processing and compression overlaid against the mean PDI for the different slices	16
3.3	Relative changes in the estimated factor matrices over the CPD iterations	16
3.4	Histograms of timecourse correlations between mean components for two different pipelines	17
3.5	Histograms of timecourse correlations between mean components for three different ranks	18
3.6	Histograms of timecourse correlations between mean components for three different ranks.	19
3.7	Histograms of timecourse correlations between mean components for three different ranks.	19
3.8	Mean component 1, slice A	20
3.9	Correlation between the delayed stimulus signal and the time course of mean component 1.	21
3.10	Mean component 2, slice A	21
3.11	Mean component 3, slice A	21
3.12	Correlation between the delayed stimulus signal and the time course of mean components 2 and 3	22
3.13	The spatial modes of the 11 cluster members belonging to mean component 3	22
3.14	Further analysis of component 3	23
3.15	Mean component 13, slice A	23
3.16	Mean component 56, slice A	24
3.17	Mean component 9, slice B	24
3.18	Mean component 50, slice B	25
3.19	Correlation between the delayed stimulus signal and the time course of mean components 9 and 50	25
3.20	Further analysis of component 9	25
3.21	Mean component 34, slice B	26
3.22	Mean component 20, slice B	26
3.23	Mean component 1, slice C	27
3.24	Mean component 6, slice C	27
3.25	Mean component 5, slice C	27
A.1	Histograms of timecourse correlations between mean components for different clustering settings for slice A	39
B.1	Thresholded correlation images overlaid against the mean PDI for different SVD filters	41
C.1	Mean component 15, regressed slice A	43
C.2	Mean component 16, regressed slice A	44
C.3	Correlation between the delayed stimulus signal and the time course of mean components 15 and 16	44
D.1	The spatial modes of the 30 cluster members belonging to mean component 1	46

E.2 Mean component 4, masked slice A 47

List of Tables

2.1	Matrix/tensor notations and symbols used in this work	7
3.1	Comparison storage complexities	15
3.2	Maximum frequency of occurrence and number of clusters for different ranks for slice A.	18
3.3	Maximum frequency of occurrence and number of clusters for different ranks for slice B.	18
3.4	Maximum frequency of occurrence and number of clusters for different ranks for slice C.	19

Introduction

First, an introduction on functional ultrasound (fUS) imaging and related concepts will be given, which is based on the comprehensive paper by Montaldo et al.[1]. Afterwards, the contribution of this work will be explained.

1.1. Why functional ultrasound?

In the last decade, the new functional neuroimaging technique fUS has emerged as a potential new tool for clinical and neuroscientific applications. Unlike several conventional methods for functional brain imaging, fUS offers an unparalleled combination of submillimeter-subsecond spatiotemporal resolution and the ability to penetrate deep into brain tissue and capture large areas of interest. This makes fUS an exceptionally valuable tool for investigating brain function. While local brain imaging techniques such as optical imaging and implanted EEG have high spatial resolution, they do not allow for investigating brain functions that span whole brain regions or even brain networks. In contrast, whole-brain imaging methods such as fMRI offer broader coverage but are constrained by lower temporal resolution, capturing only static "snapshots" of activity. fUS, however, bridges this gap by enabling dynamic, high-resolution visualizations of brain activity across large regions. See Figure 1.1 for a comparison of the characteristics with other modalities.

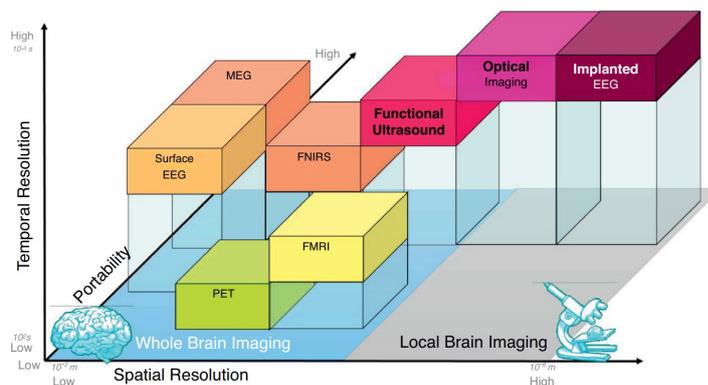


Figure 1.1: Main functional imaging techniques depicted based on temporal resolution, spatial resolution, and portability. Also, a division is made between local and whole-brain imaging, retrieved from [2].

1.2. Neurovascular coupling

The principal physiological mechanism utilized by fUS imaging is neurovascular coupling, which connects neuronal activity with local hemodynamic changes. Within the brain, increased neuronal activity results in the dilation of nearby blood vessels, subsequently leading to an increased local blood flow and volume. This response aims to meet the additional oxygen and glucose requirements of the activated

brain region. The hemodynamic response function (HRF) characterizes these changes, capturing the delayed increase in blood flow that typically peaks several seconds after the onset of neuronal activity.

1.3. Doppler effect

In fUS, the increase in blood volume is measured by the intensity of the Doppler signal caused by the Doppler effect. The Doppler effect is often explained as the shift in frequency that occurs when a sound source (i.e., an ambulance siren) is moving. Also, when a sound is emitted by a static source and reflected by a moving target, its frequency is shifted positively when the target is moving towards the source and shifted negatively when it is moving away from the source. In the context of brain imaging, the Doppler signal is the echo produced by moving tissue and blood cells when a short pulse is sent out to the brain by an ultrasound probe.

1.4. Acquisition

The fUS acquisition pipeline is depicted in Figure 1.2. fUS imaging uses angled plane waves sent to the brain through a cranial window. A linear array transducer is used, such that the backscattered signals (left side Figure 1.2), that are later beamformed and coherently compounded, constitute a 2D image of a given brain slice (middle Figure 1.2). Next, ensembles of adjacent compound frames are formed and filtered to reject undesired clutter from tissue artifacts. Finally, the 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 (right side Figure 1.2).

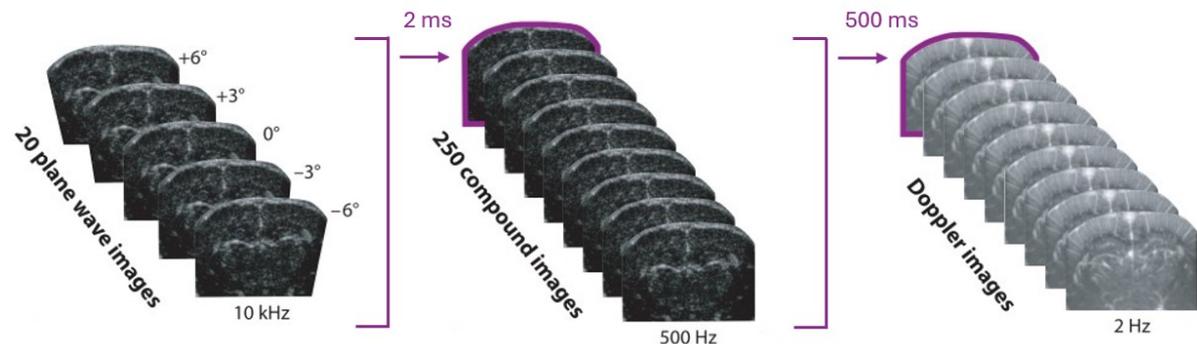
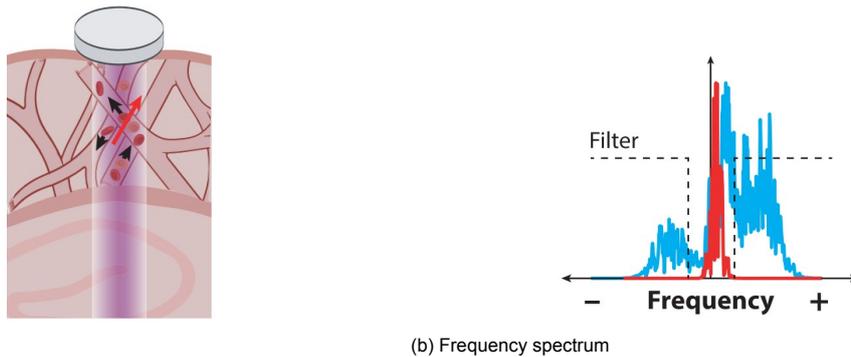


Figure 1.2: Acquisition pipeline of fUS, image based on Montaldo et al.[1].

1.4.1. Clutter filtering

The acquisition of fUS data is very heavily affected by noise. Effective separation of blood and tissue signals is essential due to the significantly stronger echoes produced by tissue, which are approximately two orders of magnitude higher, compared to those generated by blood. Specifically, echoes from red blood cells contribute only 1% of the overall signal amplitude. Initially, separation was done using high-pass filters, assuming brain tissue moves slowly (<4 mm/s) and produces lower-frequency signals than blood. However, this also removes signals from slower-flowing capillaries and arterioles (0.1–10 mm/s), which make up over 80% of cerebral blood volume and are the first to dilate during neuronal activation [3]. See Figure 1.3 for a visual explanation of this concept. Currently, the most commonly used method to perform clutter filtering is a singular value decomposition (SVD). This removes components with high spatiotemporal coherence, such as tissue noise.



(a) Brain

(b) Frequency spectrum

Figure 1.3: a) Realistic case of a brain voxel where multiple blood vessels and the tissue are mixed. The grey oval represents the ultrasound probe that sends a short pulse which produces an echo after hitting the structures. The red arrow represents the direction and amplitude of tissue motion and the black arrows represent the direction and amplitude of blood flow in the vessels. b) Frequency spectrum of the fUS signal. When multiple small vessels of different orientations are mixed in the same voxel, the spectra are added (blue). Opposite flow directions generate positive and negative frequencies. However, the intensity of the spectrum is proportional to the total number of cells in the voxel or to the blood volume. Brain tissue also moves but more slowly than the blood in most of the vessels, thus generating Doppler frequencies close to 0 (red). The tissue signal must be eliminated using filtering methods. Figures are obtained from [1]

1.5. Blind Source separation

To obtain relevant information from functional neuroimaging data, blind source separation (BSS) is often used [4]–[7]. BSS is the separation of a set of source signals from a set of mixed signals, without the aid of information about the source signals or the mixing process. A common example of a source separation problem is the cocktail party problem, where several people are speaking at the same time in a room (like at a cocktail party), and someone is trying to follow one of the conversations. In the context of this work, one would like to separate source signals related to artifacts and brain networks that are mixed in the fUS signal.

1.5.1. Independent component analysis

A commonly used BSS method for identifying brain networks and artifacts in functional neuroimaging techniques like fMRI is Independent Component Analysis (ICA) [8]. For ICA, a low-rank matrix decomposition, the following model is commonly assumed:

$$\mathbf{X} \approx \mathbf{A}\mathbf{S}, \quad (1.1)$$

where $\mathbf{X} \in \mathbb{R}^{I_N \times I_T}$ is the observed brain signal. The columns of $\mathbf{A} \in \mathbb{R}^{I_N \times R}$ represent the R spatial component maps, where I_N is the number of pixels of one spatial map. The rows of $\mathbf{S} \in \mathbb{R}^{R \times I_T}$ represent the time courses of the respective component maps, where I_T is the number of timepoints. Although meaningful components have been extracted using ICA-based approaches, they require either spatial or temporal independence, which is not realistic as brain networks can overlap and be correlated in both space and time. Also, a correlation between physiological noise and brain networks exists [7], [9].

1.6. Contribution

1.6.1. Proposed model

This work explores the use of canonical polyadic decomposition (CPD), a low-rank tensor decomposition, for BSS of fUS data. More specifically, the following model is proposed:

$$\underline{\mathbf{T}} \approx \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r, \quad (1.2)$$

where $\underline{\mathbf{T}} \in \mathbb{R}^{I_N \times I_F \times I_T}$ is the observed brain signal. I_F is the number of bins of the frequency spectrum. Furthermore, \mathbf{a}_r describes the r -th source in space, \mathbf{b}_r describes the r -th source in frequency, and \mathbf{c}_r describes the r -th source in time. Thus, CPD approximates the 3rd-order tensor by a sum of R (estimated number of sources) rank-1 tensors. CPD is generally unique under mild conditions. These

mild conditions are that the components should be sufficiently different and that their number is not unreasonably large. A deterministic condition that can tell us something about the uniqueness of CPD when the factor matrices have been computed is Kruskal's condition [10]. For this, Equation 1.2 can be rewritten in the matrix format

$$\underline{\mathbf{T}} \approx [[\mathbf{A}, \mathbf{B}, \mathbf{C}]] \quad (1.3)$$

where the columns of \mathbf{A} represent all \mathbf{a}_r spatial maps, the columns of \mathbf{B} represent all \mathbf{b}_r spectral maps, and the columns of \mathbf{C} represent all \mathbf{c}_r temporal maps. Then, according to the Kruskal's condition, if

$$k_A + k_B + k_C \geq 2R + 2$$

the CPD is essentially unique. k_A is the Kruskal rank of \mathbf{A} : the maximum value of k such that any k columns of \mathbf{A} are linearly independent. An essentially unique decomposition is subject to the trivial indeterminacies of permutation and scaling.

1.6.2. Advantages

The proposed method offers three key advantages that make it highly relevant to the field of fUS signal processing:

- Tensor-based BSS by CPD allows for frequency as a third dimension to complement the spatial and temporal dimensions.

The frequency spectra are expected to exhibit distinct characteristics among artifacts and brain regions due to the Doppler effect. For instance, static tissue noise, resulting from minimal motion, is likely to produce a symmetric spectrum centered around zero frequency, as there is no significant relative motion to cause a frequency shift.

In contrast, movements of the mouse, such as shifts in a specific direction, may produce an asymmetric spectrum due to the Doppler shift. When tissue or external objects move toward or away from the probe, the resulting frequency shift skews the spectrum to either the positive or negative side, depending on the direction of motion.

Similarly, for vessels feeding a brain region during activation, blood flows consistently in one direction, resulting in a spectrum with energy concentrated on one side—either positive or negative—depending on whether the flow is toward or away from the probe. The frequency range of the signal corresponds to the relative velocity between the probe and the blood flow.

The addition of the frequency dimension in the model allows for exploiting the diversity in the Doppler spectra to decompose the data. Tensor-based methods, such as CPD, inherently allow for more than two dimensions of data, in contrast to matrix-based methods like ICA.

- Tensor-based BSS by CPD allows for the use of all raw fUS data (compound images) rather than only PDIs, exploiting more available information.

When performing matrix-based BSS, such as ICA, on fUS data, the analysis is typically conducted on PDIs. These PDIs are ensemble averages of 200–250 compound images, resulting in a significant loss of detail. In theory, the compound images could be stacked to form a very wide, or "fat," matrix, where each column represents a compound image. However, this approach introduces problems related to both interpretability and computational complexity.

Interpretability issues arise because the time-related information resulting from ICA would correspond to each individual compound image rather than a larger, more meaningful group of images (like the average of 200–250 images in a PDI). This is problematic because individual compound images are noisy and less representative of overall patterns in the data. Since ICA would be blind to the fact that the compound images belong together in one ensemble, it can be compared to attempting to understand someone's speech by analyzing every single sound they make instead of looking at whole words or sentences. Tensor-based methods, like CPD, avoid this issue because they work with the entire 3D structure (space, time, frequency), preserving the relationships (ensembles) within the data and making the results easier to interpret.

From a computational perspective, performing ICA on a very fat matrix of compound images is more computationally expensive than performing ICA on a skinnier matrix of PDIs. ICA typically involves solving global optimization problems on the entire matrix, such as maximizing non-Gaussianity or performing eigenvalue decomposition on large covariance matrices. These operations scale poorly as the dataset size increases, especially with very wide matrices. CPD, on the other hand, explicitly assumes a low-rank model, meaning the tensor is approximated by R rank-1 components (outer products). This structured assumption simplifies the optimization, reducing the number of parameters to estimate and the size of intermediate computations.

Thus, the proposed method can use compound images while maintaining interpretability and limiting computational complexity, preserving the details in this "raw" data.

- Tensor-based BSS by CPD allows for a BSS method without strong constraints such as the case with matrix-based methods like ICA.

While ICA requires either spatial or temporal independence, which is unrealistic and could therefore produce invalid neuroscientific results, CPD only requires mild conditions. This means it is fully data-driven, allowing exploration of the data as it is.

To the best of the authors' knowledge, this is the first time that (tensor-based) BSS is performed on raw fUS data. For a graphic explanation of the difference between a conventional matrix-based approach and the proposed tensor-based approach, see Figure 1.4.

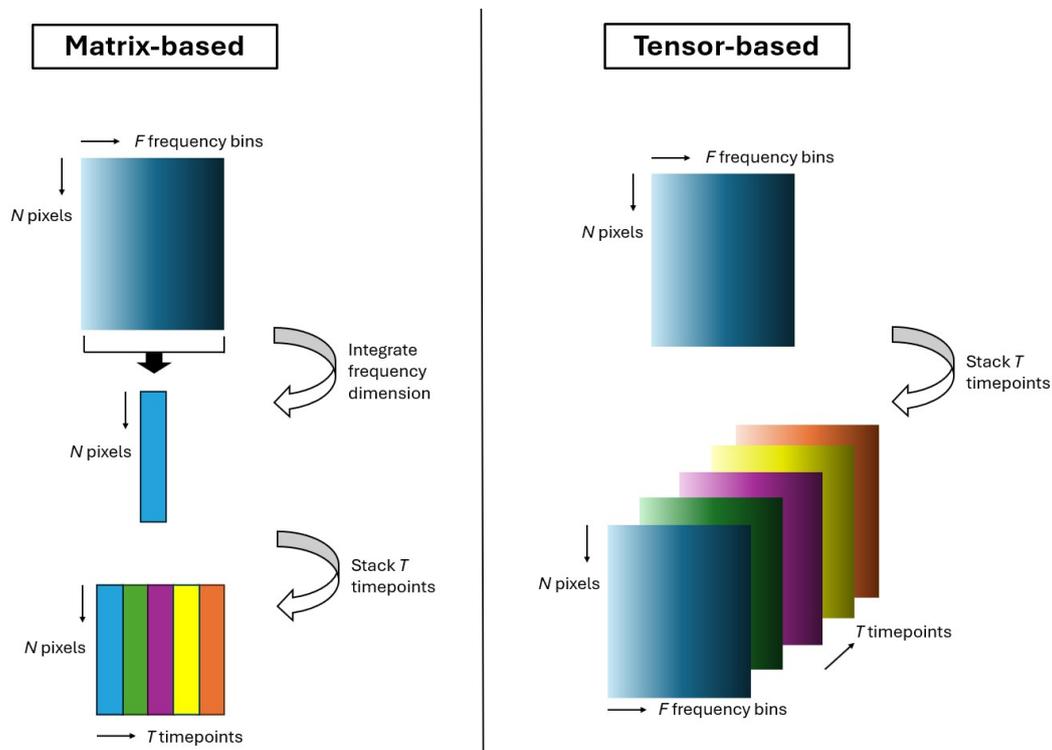


Figure 1.4: Comparison between matrix-based and tensor-based approach. The color gradient represents the detailed information present in the raw data. For visual purposes only five timepoints are shown.

1.7. Problem statement

The overall goal of this thesis is as follows:

To create a tensor-based blind source separation pipeline for raw fUS data to identify artifacts and meaningful neurological components based on distinctive characteristics in the temporal, spatial, and spectral domains.

This main goal can be captured in four research questions. The first question is as follows:

RQ 1. *What pre-processing techniques are needed on the raw functional ultrasound data to remove nuisance and ensure adequate and relevant BSS?*

When creating a pipeline for BSS on raw fUS data, an important aspect is handling the large size of such data. To ensure relevance for neuroscientific research, utilizing the pipeline should be possible with the random-access memory (RAM) of a general laptop. In this case, the RAM of a general laptop is defined as being around 8GB. To illustrate the importance of a compression step, the raw fUS data files used in this thesis are approximately 37GB. Therefore, the second question is as follows:

RQ 2. *How can raw functional ultrasound data be compressed such that CPD is possible with standard RAM?*

To ensure the validity of the pipeline it is important to assess if the resulting CPD components are unique and stable. Therefore, the third question is as follows:

RQ 3. *Does the pipeline result in unique, stable CPD components?*

And lastly, to ensure the relevance of the resulting CPD components, it is important to assess if they are physically meaningful. Therefore the fourth question is as follows:

RQ 4. *Are the CPD components meaningful from a neuroscientific perspective?*

2

Methodology

2.1. Preliminaries

Since notations and terminology used for tensors and tensor networks differ among scientific literature, an overview of the notations used in this work is presented in Table 2.1. This notation is adopted from [10].

<i>Notation</i>	<i>Meaning</i>
x	Scalar
\mathbf{x}	Vector
\mathbf{X}	Matrix
$\underline{\mathbf{X}}$	Tensor
I_n	Size of the n th dimension
$\mathbf{X}_{(n)}$	Mode- n matricization/unfolding
\times_n	Mode- n product
$\ \mathbf{X}\ _F$	Frobenius norm
\circ	Outer product
\odot	Khatri-Rao product
\otimes	Hadamard product
\dagger	Moore-Penrose pseudo-inverse
$R, (R_1, \dots, R_N)$	Tensor rank R and multilinear rank

Table 2.1: Matrix/tensor notations and symbols used in this work

2.2. Data

The data used in this thesis consists of a fUS experiment which was conducted at the Center for Ultrasound and Brain imaging at Erasmus MC (CUBE). In the experiment, a visual stimulus was displayed to a mouse in 20 blocks of 4 seconds in duration, see Figure 2.1 for a graphic representation. The visual stimulus consisted of a rectangular patch of randomly generated, high-contrast images, i.e. white “speckles” against a black background. Each stimulus epoch was followed by a random rest period of [10,15] seconds. Three slices of the mouse brain were imaged sagittally at Bregma -2.55 , $+2.15$, and -2.15 mm [11]. Complex-valued compound images of the mouse brain were made at a sampling rate of 800 Hz, and each ensemble consists of ~ 200 compound images. Each recording lasts ~ 6 minutes, resulting in 288000 successive compound images. These are used as input data for the processing pipeline. Conventional pipelines would integrate over the ensemble to create PDIs with a sampling rate of 4 Hz. This would result in 1440 successive PDIs.

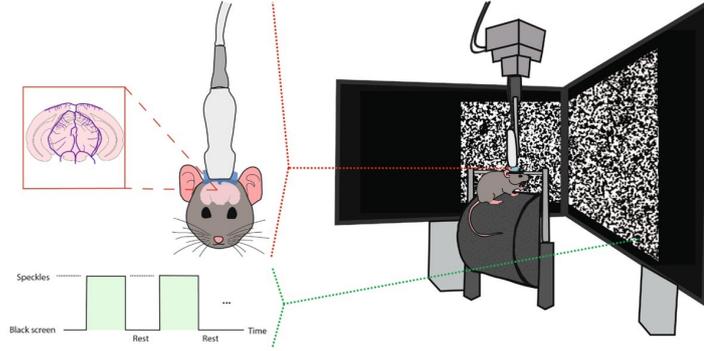


Figure 2.1: Set-up experiment, image retrieved from [12]

2.3. Pipeline

The full processing pipeline used on the data consists of four distinct stages. An overview of all steps is provided in Figure 2.2. In stage one, pre-processing is carried out through filtering, temporal demeaning, and pixel-based normalization. In stage two, compression is performed using truncated multilinear singular value decomposition (MLSVD), which will be explained in detail later. In stage three, BSS is carried out using CPD. In stage four, stable components are extracted to ensure reproducibility.

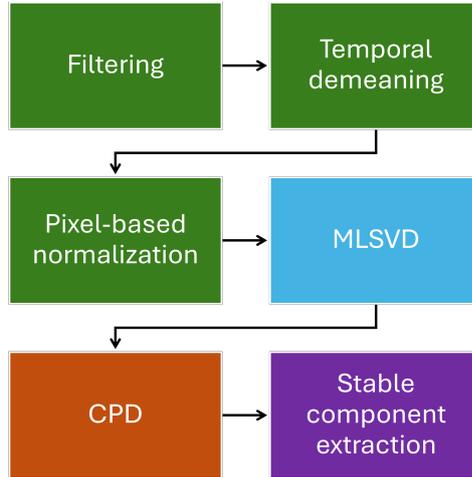


Figure 2.2: Processing pipeline

2.3.1. Filtering

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 whereas the latter removing noise. Also, the singular values were normalized to remove differences in energy between ensembles, potentially coming from unremoved tissue clutter. It is relevant to note that this normalization step is not common in fUS processing. Then, the Fourier transform of the ensemble was calculated to obtain the frequency spectrum, of which the magnitude was taken. Repeating this for every ensemble results in a tensor $\underline{\mathbf{T}} \in \mathbb{R}^{I_N \times I_F \times I_T}$, where I_N is the spatial dimension, I_F is the spectral dimension and I_T is the temporal dimension.

2.3.2. Temporal demeaning

Column-wise demeaning in the temporal mode unfolding of the tensor was performed such that $\mathbf{T}_{(3)} = \mathbf{T}_{(3)}(\mathbf{I} - \frac{1}{T}\mathbf{1}\mathbf{1}^T)$, where \mathbf{I} represents the identity matrix and $\mathbf{1}$ a column vector of ones. The purpose is to ensure that compression of the signal by MLSVD will not retain singular values containing the signal mean (offset), but rather singular values containing deviations from the mean which are of greater (neuroscientific) interest.

2.3.3. Pixel-based normalization

Matrix normalization of the N pixels was performed by taking the Frobenius norm $\|\underline{\mathbf{T}}_{[n,:;:]}\|_F$ where $\underline{\mathbf{T}}_{[n,:;:]}$, $\forall n \in [N]$ is the n -th horizontal slice of tensor $\underline{\mathbf{T}}$. Large vessels such as arteries and veins will contain fast flowing blood that produces a higher magnitude of the fUS signal compared to small arterioles and venules. Nevertheless, especially those smaller vessels are central to the neurovascular coupling effect, as they constitute over 80% of the total cerebral blood volume and are the initial vessels to undergo dilation in response to neuronal activation [1]. CPD will fit the signal with the highest variance, which creates a bias towards large vessels. Therefore, normalization is necessary.

2.3.4. Truncated multilinear singular value decomposition

Compression of the data was performed by truncated MLSVD. MLSVD is a factorization of a tensor into a relatively small size core tensor and factor matrices, expressed as:

$$\underline{\mathbf{T}} \approx \underline{\mathbf{G}} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times_3 \mathbf{U}^{(3)} \quad (2.1)$$

where $\underline{\mathbf{T}} \in \mathbb{R}^{I_N \times I_F \times I_T}$ is the tensor, $\mathbf{U}^{(1)} \in \mathbb{R}^{I_N \times R_1}$, $\mathbf{U}^{(2)} \in \mathbb{R}^{I_F \times R_2}$, $\mathbf{U}^{(3)} \in \mathbb{R}^{I_T \times R_3}$ are the orthogonal factor matrices, and $\underline{\mathbf{G}} \in \mathbb{R}^{R_1 \times R_2 \times R_3}$ is the all-orthogonal core tensor. The core tensor is ordered in the sense that the Frobenius norms of the slices in each mode are decreasing with the increase in the running index. Therefore, a tensor can be approximated by discarding the columns of the factor matrices and the corresponding slices of the core tensor with small singular values. Truncated MLSVD can be obtained by Algorithm 1, which makes use of SVD [10].

Algorithm 1 Truncated multilinear singular value decomposition of a third-order tensor

Input: Tensor $\underline{\mathbf{T}} \in \mathbb{R}^{I_N \times I_F \times I_T}$, truncation (R_1, R_2, R_3)

Output: Low multilinear rank approximation $\underline{\mathbf{T}} \approx \underline{\mathbf{G}} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times_3 \mathbf{U}^{(3)}$

- 1: **for all** $n \in \{1, 2, 3\}$ **do**
 - 2: $\mathbf{U}_n, \mathbf{S}_n, \mathbf{V}_n^T \leftarrow \text{svd}(\mathbf{T}_{(n)})$
 - 3: **end for**
 - 4: $\mathbf{U}^{(1)} \leftarrow$ first R_1 columns of \mathbf{U}_1
 - 5: $\mathbf{U}^{(2)} \leftarrow$ first R_2 columns of \mathbf{U}_2
 - 6: $\mathbf{U}^{(3)} \leftarrow$ first R_3 columns of \mathbf{U}_3
 - 7: $\underline{\mathbf{G}} \leftarrow \underline{\mathbf{T}} \times_1 \mathbf{U}^{(1)T} \times_2 \mathbf{U}^{(2)T} \times_3 \mathbf{U}^{(3)T}$
-

However, performing SVD on large matrices that can not be stored in (standard-sized) RAM is not feasible. To solve this, randomized SVD with power iterations was used. The implementation such as depicted in Algorithm 2 was used, with the first mode unfolding of tensor $\underline{\mathbf{T}}$ serving as an example. Randomized SVD uses random sampling to identify a subspace that captures most of the content of a matrix. The input matrix is then compressed to this subspace, and the reduced matrix is used to obtain the SVD. In cases where the rank is known, randomized SVD will estimate an adequate subspace as long as R is equal or larger than this true rank. In the context of this thesis, the rank is unknown, and therefore so-called power iterations are needed. These power iterations involve multiplying the current subspace approximation by the original matrix and orthogonalizing the result. These iterations serve to improve the quality of the subspace estimate by focusing more on the directions that capture the most important variance of the matrix. By iteratively refining the approximation, the power iterations compensate for the uncertainty in the rank. It can adapt to capture the most significant components of the matrix, even if the exact rank is not known.

Algorithm 2 is a slight adaption from Halko et al.[13]. Specifically, in the original algorithm an extra orthonormalization (qr) was performed in step 4, while here this was only done in step 5. These orthonormalizations are practically used to alleviate round-off errors in the floating-point computation. The adaptation was made due to memory constraints, as the available RAM would not suffice for the extra orthonormalization step. However, as noted by Feng et al.[14], this adjustment has minimal impact on accuracy. Furthermore, although not stated in Algorithm 2, all matrix multiplications (step 2, 4, 5, 7) were performed block-wise to limit RAM usage.

Algorithm 2 Randomized SVD with power iterations

Input: $\mathbf{T}_{(1)} \in \mathbb{R}^{I_N \times I_F I_T}$, rank parameter R , maximum number of power iterations P

Output: $\mathbf{U} \in \mathbb{R}^{I_N \times R}$, $\mathbf{S} \in \mathbb{R}^{R \times R}$, $\mathbf{V}^T \in \mathbb{R}^{I_F I_T \times R}$

1: $\Omega \leftarrow \text{randn}(I_F I_T, R)$

2: $\mathbf{Q} \leftarrow \text{qr}(\mathbf{T}_{(1)} \Omega)$

3: **while** not converged or power iteration limit is not reached **do**

4: $\mathbf{K} \leftarrow \mathbf{T}_{(1)}^T \mathbf{Q}$

5: $\mathbf{Q} \leftarrow \text{qr}(\mathbf{T}_{(1)} \mathbf{K})$

6: **end while**

7: $\mathbf{W} \leftarrow \mathbf{Q}^T \mathbf{T}_{(1)}$

8: $[\hat{\mathbf{U}}, \mathbf{S}, \mathbf{V}^T] \leftarrow \text{svd}(\mathbf{W}, 'econ')$

▷ SVD is performed on a matrix with dimensions $R \times I_F I_T$

9: $\mathbf{U} \leftarrow \hat{\mathbf{U}}$

2.3.5. Canonical polyadic decomposition

To perform BSS on the compressed tensor, CPD is used. The proposed algorithm follows from the framework described by Vervliet et al.[15]. The goal is to decompose $\underline{\mathbf{T}}$ into R rank-1 tensors such that

$$\underline{\mathbf{T}} \approx \underline{\mathbf{L}} \times_1 \mathbf{A} \times_2 \mathbf{B} \times_3 \mathbf{C} \quad (2.2)$$

where $\underline{\mathbf{L}} \in \mathbb{R}^{R \times R \times R}$ is a superdiagonal tensor of rank R , and $\mathbf{A} \in \mathbb{R}^{I_N \times R}$, $\mathbf{B} \in \mathbb{R}^{I_F \times R}$, $\mathbf{C} \in \mathbb{R}^{I_T \times R}$ the factor matrices. In further notations, the diagonal weights in $\underline{\mathbf{L}}$ are absorbed into the factor matrices. An alternating least squares (ALS) approach was followed to obtain \mathbf{A} , \mathbf{B} , and \mathbf{C} . In order to solve for the first CPD factor matrix (spatial), the following cost function has to be minimized

$$\begin{aligned} f(\mathbf{A}, \mathbf{B}, \mathbf{C}) &= \|\mathbf{T}_{(1)} - \mathbf{A}(\mathbf{C} \odot \mathbf{B})^T\|_F^2 \\ &= \text{Tr}((\mathbf{T}_{(1)})^T \mathbf{T}_{(1)}) - 2\text{Tr}((\mathbf{T}_{(1)})^T \mathbf{A}(\mathbf{C} \odot \mathbf{B})^T) \\ &\quad + \text{Tr}((\mathbf{C} \odot \mathbf{B}) \mathbf{A}^T \mathbf{A}(\mathbf{C} \odot \mathbf{B})^T) \end{aligned} \quad (2.3)$$

Taking the derivative of $f(\mathbf{A}, \mathbf{B}, \mathbf{C})$ with respect to matrix \mathbf{A} gives:

$$\frac{\partial f(\mathbf{A}, \mathbf{B}, \mathbf{C})}{\partial \mathbf{A}} = -2\mathbf{T}_{(1)}(\mathbf{C} \odot \mathbf{B}) + 2\mathbf{A}(\mathbf{C}^T \mathbf{C} \odot \mathbf{B}^T \mathbf{B}). \quad (2.4)$$

Now, $\partial f(\mathbf{A}, \mathbf{B}, \mathbf{C}) / \partial \mathbf{A} = 0$ is set to get a closed-form solution for \mathbf{A} . Also, the MLSVD representation of $\underline{\mathbf{T}}$ (Equation 2.1) in its first mode unfolding is plugged in, which is $\mathbf{T}_{(1)} = \mathbf{U}^{(1)} \mathbf{G}_{(1)} (\mathbf{U}^{(3)} \otimes \mathbf{U}^{(2)})^T$. This gives:

$$\begin{aligned} \mathbf{A} &= \mathbf{U}^{(1)} \mathbf{G}_{(1)} (\mathbf{U}^{(3)} \otimes \mathbf{U}^{(2)})^T (\mathbf{C} \odot \mathbf{B}) (\mathbf{C}^T \mathbf{C} \odot \mathbf{B}^T \mathbf{B})^\dagger \\ &= \mathbf{U}^{(1)} \mathbf{G}_{(1)} (\mathbf{U}^{(3)T} \mathbf{C} \odot \mathbf{U}^{(2)T} \mathbf{B}) (\mathbf{C}^T \mathbf{C} \odot \mathbf{B}^T \mathbf{B})^\dagger, \end{aligned} \quad (2.5)$$

where the 'trick'

$$(\mathbf{U}^{(3)} \otimes \mathbf{U}^{(2)})^T (\mathbf{C} \odot \mathbf{B}) = \mathbf{U}^{(3)T} \mathbf{C} \odot \mathbf{U}^{(2)T} \mathbf{B} \quad (2.6)$$

was used to minimize the number of computations. Following the same steps for the third factor matrix (temporal), the following update was obtained:

$$\mathbf{C} = \mathbf{U}^{(3)} \mathbf{G}_{(3)} (\mathbf{U}^{(2)T} \mathbf{B} \odot \mathbf{U}^{(1)T} \mathbf{A}) (\mathbf{B}^T \mathbf{B} \odot \mathbf{A}^T \mathbf{A})^\dagger \quad (2.7)$$

For the second factor matrix, which entails the frequency spectra, a non-negativity constraint was added to the cost function. This constraint matches the prior knowledge that the frequency spectra are absolute, and therefore enhances interpretability. One can see that the problem

$$\begin{aligned} \min_{\mathbf{B}} \quad & \|\mathbf{T}_{(2)} - \mathbf{B}(\mathbf{C} \odot \mathbf{A})^T\|_F^2 \\ \text{s.t.} \quad & \mathbf{B} \geq 0, \end{aligned} \quad (2.8)$$

can be split into I_F sub-problems, which could be parallelized:

$$\min_{\bar{\mathbf{b}}_j} \|\bar{\mathbf{t}}_{(2)j} - \mathbf{P}\bar{\mathbf{b}}_j\|_2^2 \quad \text{s.t.} \quad \bar{\mathbf{b}}_j \geq 0, \quad (2.9)$$

where $\mathbf{P} = (\mathbf{C} \odot \mathbf{A})$, and column-vectors $\bar{\mathbf{b}}_j$ and $\bar{\mathbf{t}}_{(2)j}$ are the j 'th row of \mathbf{B} and $\mathbf{T}_{(2)}$, respectively. The fast non-negative least squares (FNNLS) algorithm proposed in [16] is used. The algorithm takes two inputs $\mathbf{P}^T \mathbf{P}$ and $\mathbf{P}^T \bar{\mathbf{t}}_{(2)j}$. Since our data is implicit, $\mathbf{P}^T \bar{\mathbf{t}}_{(2)j}$ is obtained by substituting the second mode unfolding of the MLSVD representation, $\mathbf{T}_{(2)} = \mathbf{U}^{(2)} \mathbf{G}_{(2)} (\mathbf{U}^{(3)} \otimes \mathbf{U}^{(1)})^T$, and use (2.6) again to compute

$$\begin{aligned} \mathbf{P}^T (\mathbf{T}_{(2)})^T &= (\mathbf{T}_{(2)} \mathbf{P})^T \\ &= (\mathbf{C} \odot \mathbf{A})^T (\mathbf{U}^{(3)} \otimes \mathbf{U}^{(1)}) (\mathbf{G}_{(2)})^T \mathbf{U}^{(2)T} \\ &= (\mathbf{U}^{(3)T} \mathbf{C} \odot \mathbf{U}^{(2)T} \mathbf{A})^T (\mathbf{G}_{(2)})^T \mathbf{U}^{(2)T} \end{aligned} \quad (2.10)$$

at once. Subsequently, one can take a column j , which is equal to $\mathbf{P}^T \bar{\mathbf{t}}_{(2)j}$. Furthermore, $\mathbf{P}^T \mathbf{P} = \mathbf{C}^T \mathbf{C} \circledast \mathbf{A}^T \mathbf{A}$. Thus, now FNNLS can be run within the context of our problem. A summary of the implementation can be found in Algorithm 3. Note that the factor matrices of the MLSVD and CPD must have the same rank R .

Algorithm 3 ALS for the CPD of a third-order tensor in MLSVD representation with a non-negativity constraint on the second mode

Input: the factor matrices $\mathbf{U}^{(1)} \in \mathbb{R}^{I_N \times R}$, $\mathbf{U}^{(2)} \in \mathbb{R}^{I_F \times R}$, $\mathbf{U}^{(3)} \in \mathbb{R}^{I_T \times R}$, and core tensor $\mathbf{G} \in \mathbb{R}^{R \times R \times R}$ of the low multilinear rank approximation of tensor \mathbf{T} , maximum number of iterations M

Output: $\mathbf{A} \in \mathbb{R}^{I_N \times R}$, $\mathbf{B} \in \mathbb{R}^{I_F \times R}$, $\mathbf{C} \in \mathbb{R}^{I_T \times R}$

- 1: Random initialization of $\mathbf{A}, \mathbf{B}, \mathbf{C}$
 - 2: **while** iteration limit is not reached **do**
 - 3: $\mathbf{A} \leftarrow \mathbf{U}^{(1)} \mathbf{G}_{(1)} (\mathbf{U}^{(3)T} \mathbf{C} \odot \mathbf{U}^{(2)T} \mathbf{B}) (\mathbf{C}^T \mathbf{C} \circledast \mathbf{B}^T \mathbf{B})^\dagger$
 - 4: Normalize column vectors of \mathbf{A}
 - 5: $\mathbf{P}^T \mathbf{P} \leftarrow \mathbf{C}^T \mathbf{C} \circledast \mathbf{A}^T \mathbf{A}$
 - 6: $\mathbf{P}^T \bar{\mathbf{t}}_{(2)} \leftarrow (\mathbf{U}^{(3)T} \mathbf{C} \odot \mathbf{U}^{(2)T} \mathbf{A})^T (\mathbf{G}_{(2)})^T \mathbf{U}^{(2)T}$
 - 7: **for all** columns j **do**
 - 8: $\mathbf{B}[:, j] \leftarrow \text{FNNLS}(\mathbf{P}^T \mathbf{P}, \mathbf{P}^T \bar{\mathbf{t}}_{(2)[:, j]})$
 - 9: **end for**
 - 10: Normalize column vectors of \mathbf{B}
 - 11: $\mathbf{C} \leftarrow \mathbf{U}^{(3)} \mathbf{G}_{(3)} (\mathbf{U}^{(2)T} \mathbf{B} \odot \mathbf{U}^{(1)T} \mathbf{A}) (\mathbf{B}^T \mathbf{B} \circledast \mathbf{A}^T \mathbf{A})^\dagger$
 - 12: Normalize column vectors of \mathbf{C} , store the norms in vector λ
 - 13: **end while**
 - 14: **return** $\mathbf{A}, \mathbf{B}, \mathbf{C}$ and λ
-

2.3.6. Stable component extraction

For many matrix and tensor factorizations, including the CPD, the associated objective functions are nonconvex. Consequently, the optimization algorithm used may converge to local optima rather than the global solution. To improve the likelihood of finding a stable outcome, it is common practice to run the optimization multiple times, each time starting from a different initialization. The final solution is then selected based on some criterion, often the one with the lowest cost value. However, in noise-sensitive applications, this approach does not guarantee the most meaningful result. In those applications, when fitting the CPD model by minimizing the mean squared error, one or more of the CPD components may model variance that is due to the noise or strong artifacts, neglecting ‘true’ sources of interest that remain undiscovered.

Since fUS data is highly affected by noise, assessment of the reproducibility of the CPD is needed. [17] has proposed an algorithm for this purpose that investigates the similarity of components across all repetitions of the decomposition with different initializations and clusters them accordingly. Although related to the well-known ICASSO, this method does not depend on a user-defined number of clusters, which enhances reproducibility. The steps of the clustering approach are as follows:

1. Run CPD 100 times with random initialization, and store the resulting factor matrices with each R components.
2. Using the factor matrices, run the algorithm of Van Eyndhoven et al., which will construct a similarity matrix of the components and compute the Eigenvalues of this matrix. It will suggest the number of significant eigenvalues (clusters) based on the largest gap between consecutive eigenvalues.
3. Select the number of clusters as suggested.
4. Run the second part of the algorithm of Van Eyndhoven et al., which will cluster the components by a user-defined clustering method. Clusters are ordered based on the number of components they contain. Clusters that have components from less than 5% of the runs are discarded.
5. Compute the spatial, temporal, and spectral mean of each cluster. Due to the sign ambiguity of CPD, the modes are sign-flipped such that they all positively correlate with each other to prevent cancelling each other out.

By using this clustering approach, the resulting mean components will represent components that occur frequently over different runs.

2.4. Validation and rank estimation

The success of CPD for a certain application is based on the match of the model to the way the observations are generated [18]. Selecting the right model order (rank) plays a key role in model match. Similar to the initialization, this is often done by choosing the rank which results in the lowest cost function. However, in this noise-sensitive context, it does not guarantee the most meaningful result. While it remains ambiguous, in this work the rank was estimated based on two quantifiable aspects:

1. Correlation between mean components.

To allow for useful interpretability of the mean components, there should be enough diversity among them. If they are all highly correlated with each other, attributing physical significance to them would make little sense. Therefore, for different ranks, the correlation between mean components is calculated and the distribution is visualized in histograms. Ranks that exhibited limited diversity, particularly those with high correlation values dominating the distribution, were deemed unsuitable.

2. Frequency of occurrence of mean components.

A good model match enhances the stability of the solutions, reflecting better agreement with the way observations are generated. Since higher stability leads to components that occur more frequently over different runs, the occurrence of components can be used as a proxy for choosing the correct rank.

For ranks up to 40, stable components were extracted, and both the correlation between mean components and their frequency of occurrence were analyzed. The rank estimation was guided by these metrics. Although the maximum rank was chosen somewhat arbitrarily, having more than 40 components does not appear to be neuroscientifically meaningful. See Figure 2.3 for the process flow of the rank estimation. While these metrics provide valuable insights, they lack the sensitivity to detect subtle variations in model fit, highlighting the estimation's approximate nature.

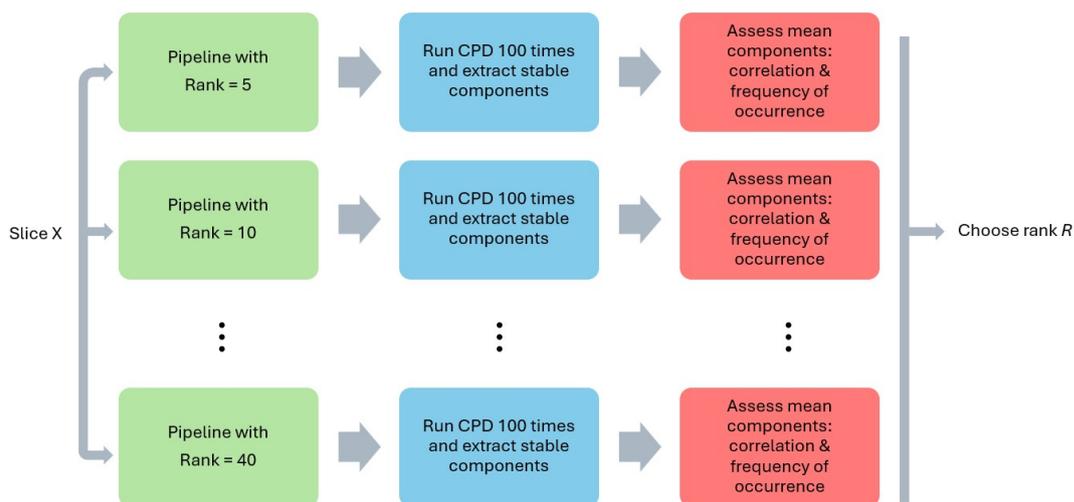


Figure 2.3: Rank estimation flow

2.5. Assessing neuroscientific relevance

Since there is no ground truth or reference available to assess if the CPD components represent actual neurological information and/or artifacts, it is difficult to assess the neuroscientific relevance. Nevertheless, there is some prior knowledge which can be used to assess the results. The most important prior knowledge is the timecourse of the visual stimulus that was displayed to the mouse. The timecourse can be correlated with the PDIs of the data to understand which areas are highly correlated with the task. This is a so-called correlation analysis and helps to identify regions of interest (ROIs). One way to assess the validity of the CPD components is to see if these ROIs are also present in (some of) the spatial maps.

Furthermore, the timecourses of the components can be correlated with the task timecourse. Components that are expected to be associated with the task should have a high correlation with the task timecourse.

3

Results

3.1. Identification of ROIs

The correlation analysis reveals activations in two ROIs: visual cortex (VIS) and lateral geniculate nucleus (LGN), see Figure 3.1. Both are known to be involved in the visual pathway. Although all three slices show activations in these regions, the VIS and LGN in slice A have a higher correlation with the task (more yellow) than slice B and C.

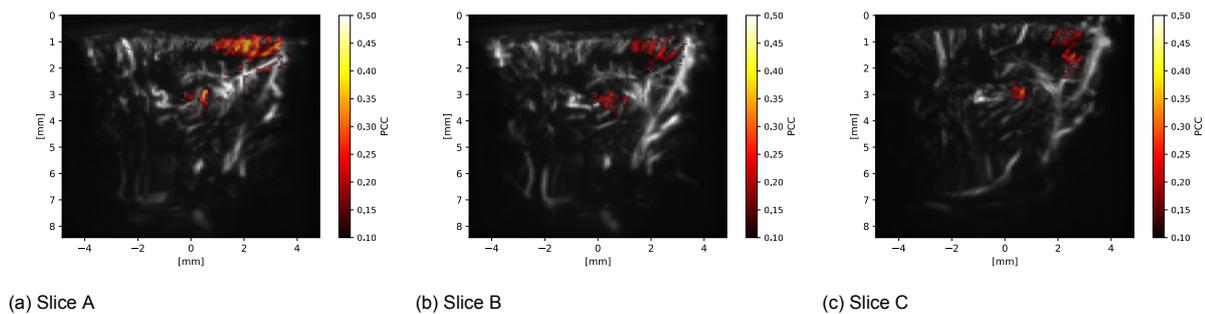


Figure 3.1: Thresholded correlation images ($PCC > 0.10$) overlaid against the mean PDI for the different slices, displaying LGN (bottom region) and VIS (top region). The correlations are calculated with a 2.5 s delay of the stimulus. The slices were imaged sagittally at Bregma -2,15, +2,15, and -2,55 mm for slice A, B, and C respectively. PCC = Pearson Correlation Coefficient

3.2. Compression

See Table 3.1 for a comparison between the storage complexity of the tensor in original format and the truncated MLSVD representation. Based on the number of entries, such as depicted in the last column of Table 3.1, the achieved data compression ratio is 1:5 898.

Representation	Complexity (\mathcal{O})	Complexity (entries)
Original tensor	$\mathcal{O}(I^3)$	$4.65 * 10^9$
Truncated MLSVD representation	$\mathcal{O}(3IR + R^3)$	$7.89 * 10^5$

Table 3.1: Comparison between the storage complexity of the original 3th-order tensor $T \in \mathbb{R}^{I_N \times I_F \times I_T}$ and the truncated MLSVD representation, where $I = \max(I_N, I_F, I_T)$. The values in the last column are based on the parameters $I_N = 16490$, $I_F = 197$, $I_T = 1432$, and $R = 40$. These parameters resemble the dimensions of the slices and the maximum rank used for compression.

To gain insight into how the functional information is preserved, a correlation analysis was done on the data after compression, see Figure 3.2. To perform this, the tensor was reconstructed and the frequency dimension was collapsed to form PDIs. In all three slices, the correlations in the VIS are higher (more yellow) than in the initial PDIs (Figure 3.1). However, the correlations in the LGN seem vanished. One should realize that these correlations are observed after compression *and* pre-processing of the data.

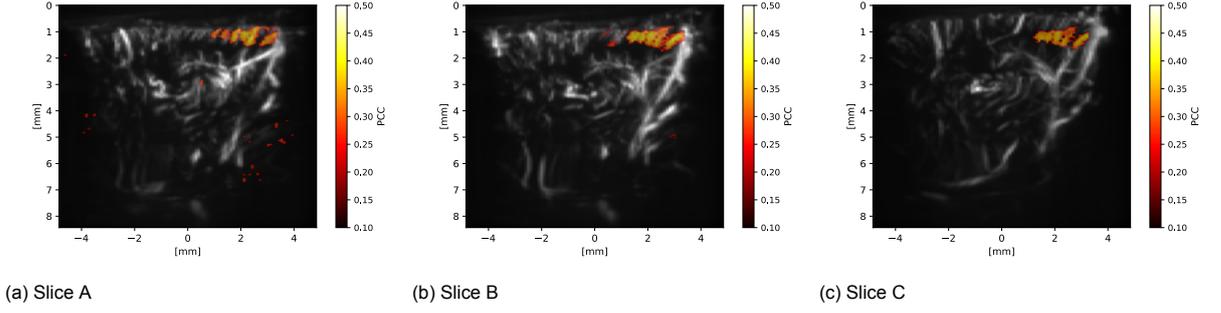


Figure 3.2: Thresholded correlation images of the data after pre-processing and compression overlaid against the mean PDI for the different slices. The correlations are calculated with a 2.5 s delay of the stimulus. PCC = Pearson Correlation Coefficient

3.3. CPD and rank estimation

3.3.1. CPD iterations

For the CPD, a maximum number of iterations of 10 was chosen. This is based on the relative changes in the estimated factor matrices over the CPD iterations, such as presented in Figure 3.3. For all three slices, the similarity converges after about four iterations, implying 10 will suffice.

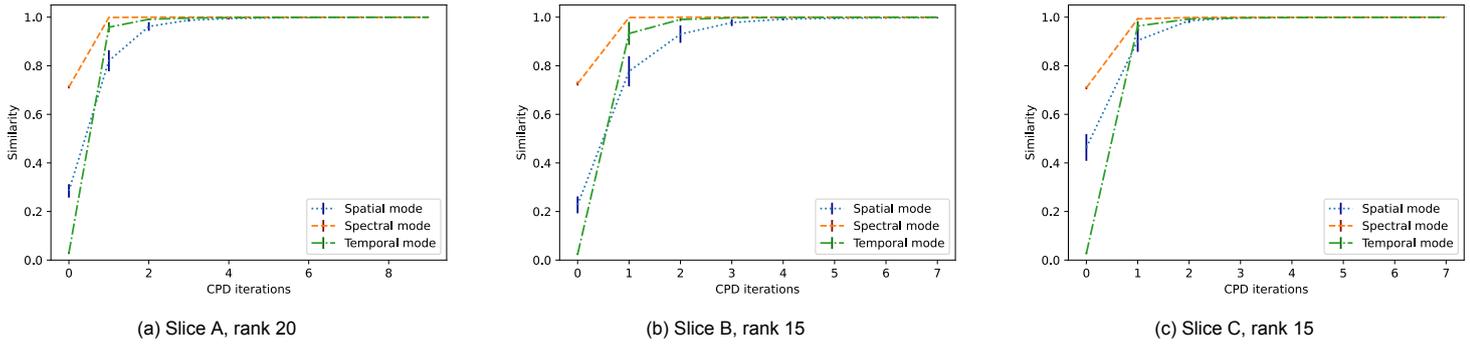


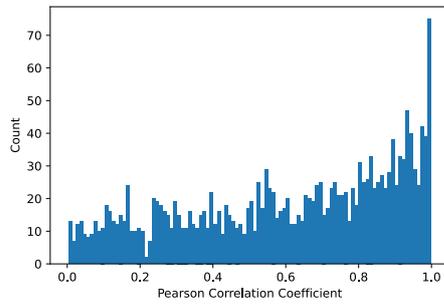
Figure 3.3: Relative changes in the estimated factor matrices over the CPD iterations. Similarity was calculated by $\sum \sigma / R$, where σ is the vector of singular values resulting from the matrix multiplication of two consecutive factor matrices ($\mathbf{A}_{k-1}^T \mathbf{A}_k$). The data points are the averaged similarities of the 100 runs used for the extraction of stable components. The variability is depicted by the error bars. The displayed ranks are the same ones as used for the neuroscientific analysis.

3.3.2. Clustering

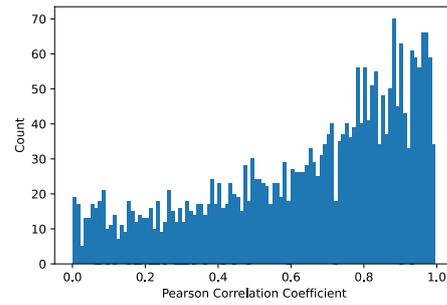
For the clustering algorithm of Van Eyndhoven et al.[17], different options can be picked for the construction of the similarity matrix as well as the clustering itself. To choose the most useful settings the correlations between the time courses of the resulting mean components were calculated to assess diversity. Based on these results, see Appendix A, an appropriate clustering method was picked.

3.3.3. Normalization of singular values

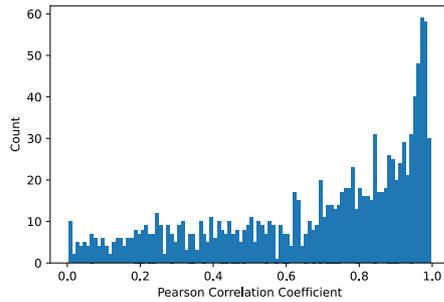
Since the normalization of the singular values in the filtering step is rather unconventional, this processing step was evaluated separately. In Figure 3.4 one can see the differences between the pipeline with and without this step. For all three slices, the maximum frequency of occurrence is higher for the pipeline with the normalization step. Also, for slice C, the distribution of the correlations is very concentrated at the high correlations for the pipeline without the normalization, but there is much more variety when the normalization is included (Figure 3.4e v.s. 3.4f).



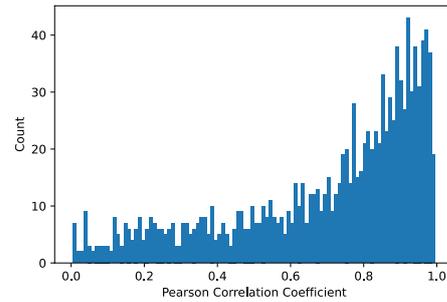
(a) Slice A; Proposed pipeline with rank 20.
Maximum frequency of occurrence = 30. Number of clusters = 62.



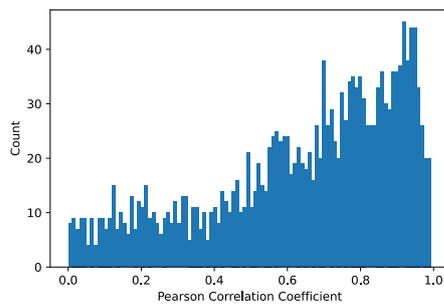
(b) Slice A; Pipeline without normalization of singular values with rank 20.
Maximum frequency of occurrence = 24. Number of clusters = 74.



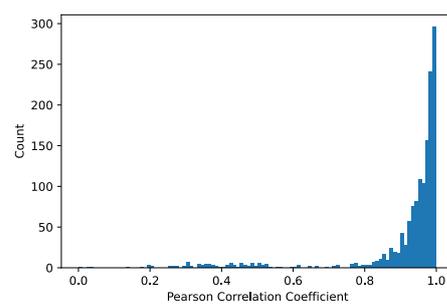
(c) Slice B; Proposed pipeline with rank 15.
Maximum frequency of occurrence = 24. Number of clusters = 50.



(d) Slice B; Pipeline without normalization of singular values with rank 15.
Maximum frequency of occurrence = 16. Number of clusters = 49.



(e) Slice C; Proposed pipeline with rank 15.
Maximum frequency of occurrence = 26. Number of clusters = 61.



(f) Slice C; Pipeline without normalization of singular values with rank 15.
Maximum frequency of occurrence = 16. Number of clusters = 54.

Figure 3.4: Histograms of timecourse correlations between mean components for two different pipelines

3.3.4. Rank estimation

To achieve the best model fit the correlation between the mean components as well as the frequency of occurrence of the mean components was assessed. These choices for all three slices are elaborated on in the following subsections.

Slice A

For slice A, $R = 20$ was chosen based on the analysis of the correlation between the mean components and the frequency of occurrence of the mean components.

In Table 3.2, an overview of the number of clusters and the maximum frequency of occurrence of these clusters is given for different ranks. One can observe a peak in the max frequency of occurrence at rank 20. The correlations between the mean components are also adequately spread out for rank 20. To illustrate differences between ranks, the histograms of the correlations for ranks 15, 20, and 25 are depicted in Figure 3.5. One can see that while for rank 15 the distribution is concentrated at the high correlations, it is more spread out for rank 20. While it is even more spread out for rank 25, this is not a measure of goodness of fit. Rather, having a somewhat diverse histogram is a necessary requirement for useful interpretability. Thus, based on the peak in occurrence (Table 3.2) and satisfying the requirement of diversity in correlations, rank 20 was considered a good fit for the model.

<i>Rank</i>	<i>Maximum frequency of occurrence</i>	<i>Number of clusters</i>
5	12	18
10	25	40
15	22	55
20	30	62
25	15	82
30	15	100
35	20	115
40	14	138

Table 3.2: Maximum frequency of occurrence and number of clusters for different ranks for slice A.

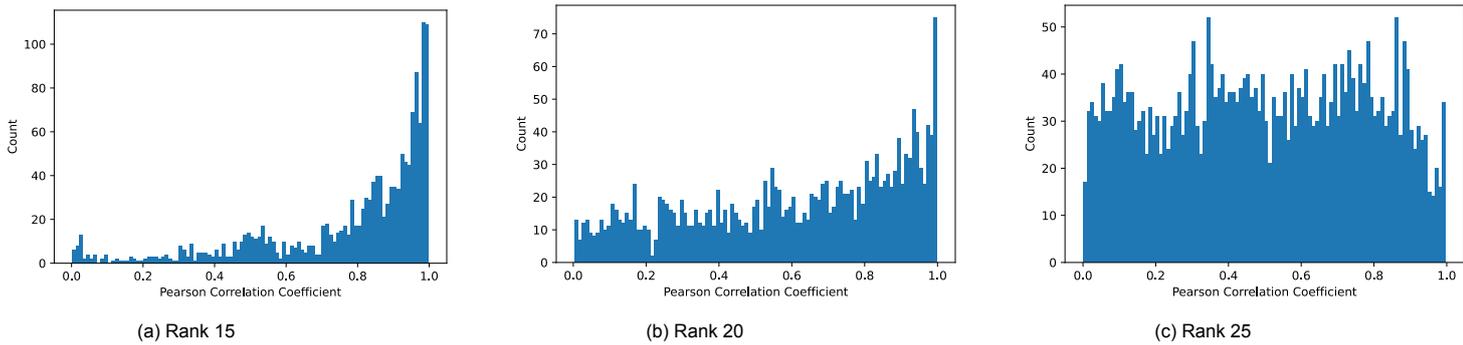


Figure 3.5: Histograms of timecourse correlations between mean components for three different ranks

Slice B

For slice B, $R = 15$ was chosen based on the analysis of the correlation between the mean components and the frequency of occurrence of the mean components.

In Table 3.3, an overview of the number of clusters and the maximum frequency of occurrence of these clusters is given for different ranks. The maximum occurrence is 24, which is reached at rank 10 and 15.

<i>Rank</i>	<i>Maximum frequency of occurrence</i>	<i>Number of clusters</i>
5	20	18
10	24	34
15	24	50
20	14	65
25	17	63
30	18	79
35	13	116
40	14	129

Table 3.3: Maximum frequency of occurrence and number of clusters for different ranks for slice B.

When comparing the histograms of the correlations (Figure 3.6), one can see a clear difference between rank 10 and 15. For rank 10, the distribution has a sharp peak at the high correlations (PCC~0.95), and (almost) no correlations are present below 0.4. For rank 15, the distribution is more spread out. Thus, based on the peak in occurrence (Table 3.3) and the diversity in correlations, rank 15 was considered a good fit for the model.

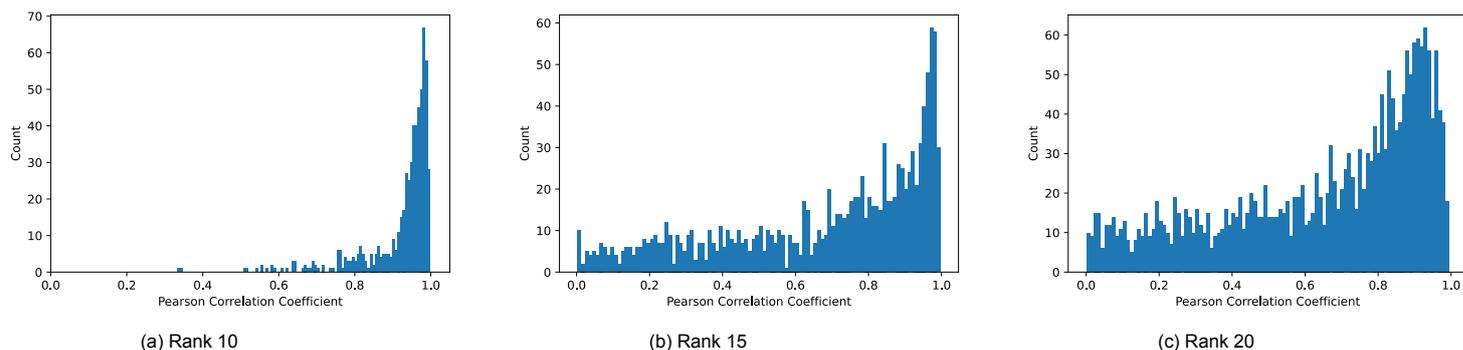


Figure 3.6: Histograms of timecourse correlations between mean components for three different ranks.

Slice C

For slice A, $R = 15$ was chosen based on the analysis of the correlation between the mean components and the frequency of occurrence of the mean components.

In Table 3.4, an overview of the number of clusters and the maximum frequency of occurrence of these clusters is given for different ranks. One can observe a peak in the max frequency of occurrence at rank 15.

Rank	Maximum frequency of occurrence	Number of clusters
5	7	17
10	16	49
15	26	61
20	21	76
25	19	101
30	19	114
35	19	135
40	18	153

Table 3.4: Maximum frequency of occurrence and number of clusters for different ranks for slice C.

The correlations between the mean components are also adequately spread out for rank 15. To illustrate differences between ranks, the histograms of the correlations for ranks 15, 20, and 25 are depicted in Figure 3.7. Thus, based on the peak in occurrence (Table 3.4) and satisfying the requirement of diversity in correlations, rank 15 was considered a good fit for the model.

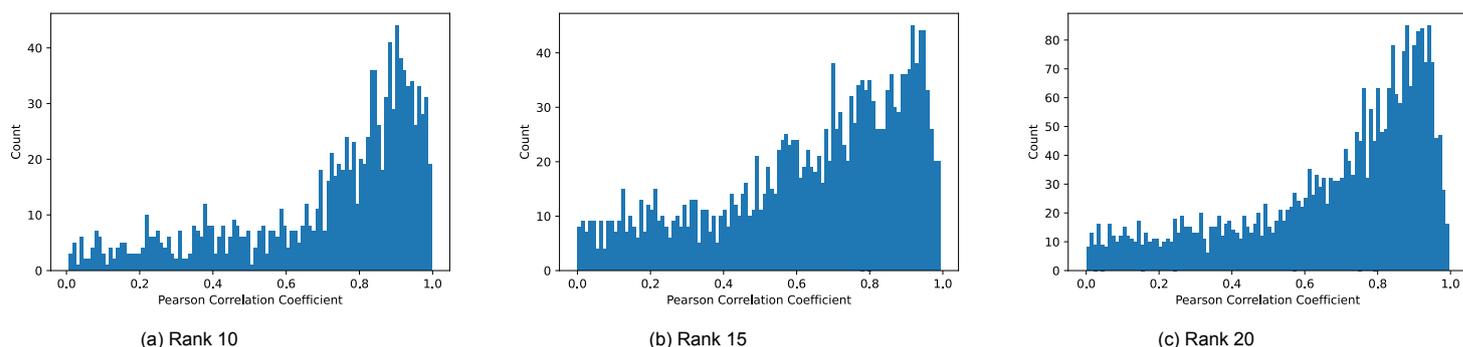


Figure 3.7: Histograms of timecourse correlations between mean components for three different ranks.

3.4. Neuroscientific analysis

In this section several relevant components of the three slices will be presented and further analyzed. Nevertheless, given the large number of components, these will be limited to components that contribute to answering the research question.

Given the ambiguity in scaling of the estimated factor matrices, visually interpreting the frequency spectra can be challenging. To address this, three metrics were calculated for the spectra of all components. The first metric is the ratio of the maximum energy of the highest peak on one side of the spectrum (+ or -) to the highest peak on the other side of the spectrum, providing a measure of asymmetry. The second metric, calculated for the side of the spectrum (+ or -) containing the highest peak, is the ratio of the peak energy to the energy at the maximum absolute frequency (400 Hz or -400 Hz). Since it was observed that these extremes contain the minimum energies, this metric provides a relative indication of the peak's prominence compared to the rest of the spectrum. The third metric involves calculating the half bandwidth of the peak on the side of the spectrum (+ or -) with the highest peak. The half bandwidth represents the width of the frequency range at half the peak's maximum power, indicating the sharpness of the peak and whether the energy is concentrated around a specific frequency or more broadly distributed.

3.4.1. Slice A

For slice A the mean asymmetry ratio of the components was 1.278 (SD: 0.902), the mean relative height was 3.83 (SD: 0.95), and the mean half bandwidth was 26.4 Hz (SD: 4.61 Hz). In general, the frequency spectra of the components contain two peaks, one on the negative side between -100 and -200 Hz, and one on the positive side between 100 and 200 Hz.

The mean component that has components from the highest number of runs is depicted in Figure 3.8. The spatial mode displays a rather uniform activation of most areas in the brain, except for a large vessel close to the VIS. Note that, considering the ambiguity in sign of the estimated factor matrices, it is not possible to define if this activation is positive or negative in reality based on this figure.

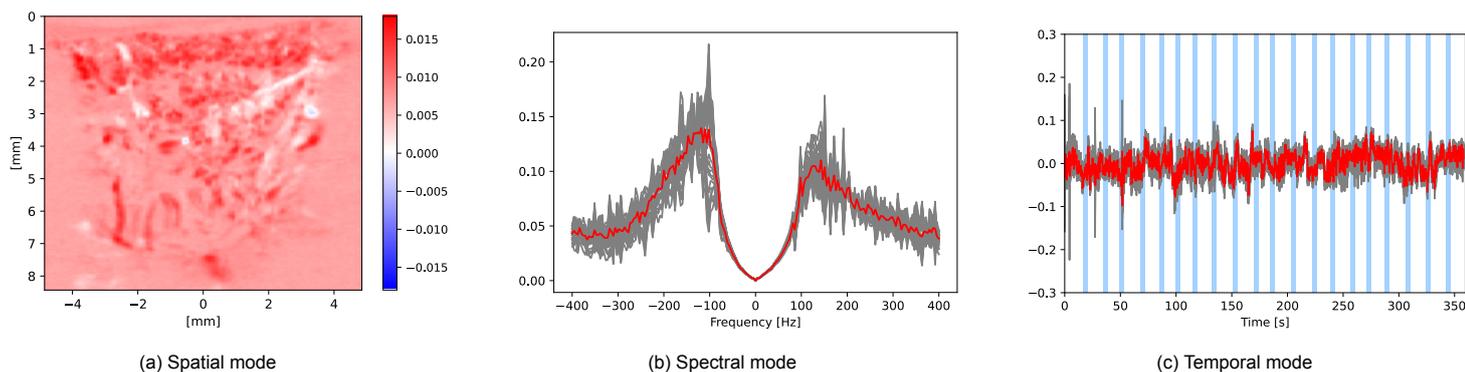
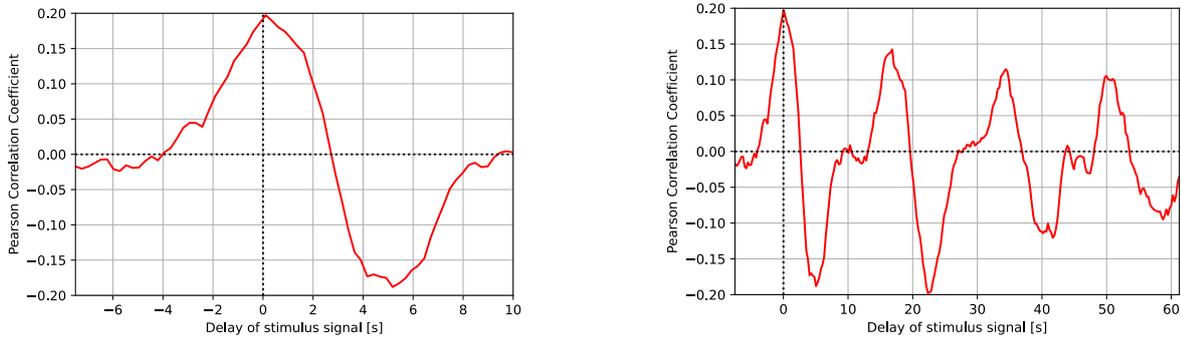


Figure 3.8: Mean component 1, which is the component that was present in the most runs (30 times). The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.267, the relative height was 3.24 and the half bandwidth was 29.6 Hz.

When the time course of the component (Figure 3.8c) is correlated with the time course of the task with a varying delay, the correlations such as depicted in Figure 3.9a can be observed. From a correlation of approximately 0 before the stimulus, the correlation becomes positive with a peak at ~ 0.1 s after the stimulus. Then, the correlation switches from positive to negative at ~ 2.5 s and peaks around ~ 5 s. The correlations with the delayed stimulus for a duration of multiple stimuli are displayed in Figure 3.9b. The described correlation pattern is repeated at similar intervals as that of the visual stimuli. Keep in mind that the interval between stimuli varies slightly each time, probably explaining the decaying height of the correlations.

The mean components that have the second and third highest frequency of occurrence are depicted in Figure 3.10 and Figure 3.13. Both components contain the VIS and LGN in their spatial map, but differences can be observed as well. While for mean component three activation is especially prominent in the VIS and LGN, for mean component two this is more spread out to other brain regions in the top area as well. The two timecourses are very similar (PCC of 0.99). In Figure 3.20, one can observe the correlations with the stimulus signal for a range of delays. Both graphs show a similar trend, with peaks located at ~ 2.5 s delay. However, the peak of component three is somewhat higher than for

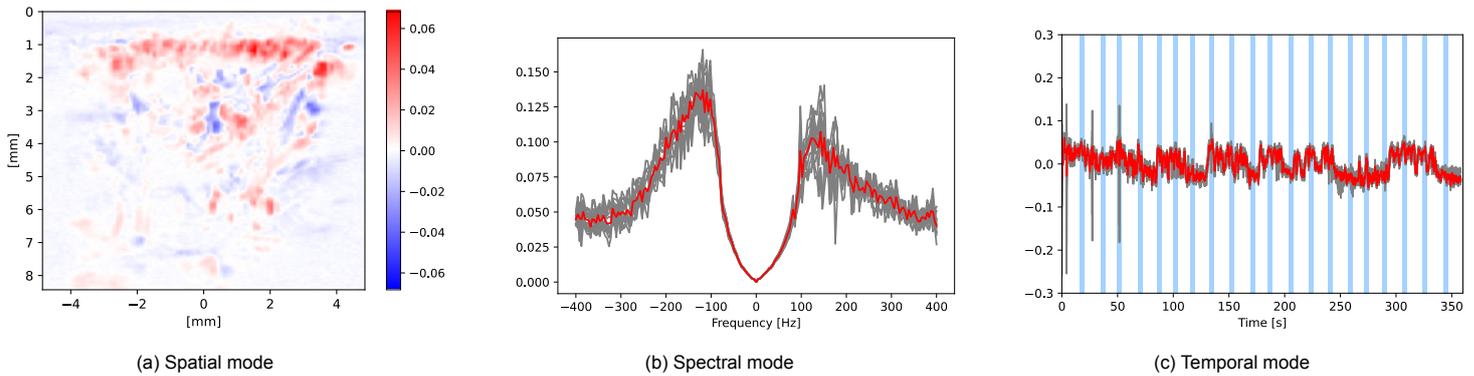
component two. It is relevant to note that almost all members belonging to cluster 2 and 3 come from different CPD runs. More specifically, only one component is member of both cluster 2 and 3. This means that both components likely represent the same underlying source, only the source is extracted slightly differently depending on the CPD run.



(a) Correlation between the delayed stimulus signal and the time course of mean component 1

(b) Correlation between the delayed stimulus signal and the time course of mean component 1 that includes multiple stimuli

Figure 3.9: Correlation between the delayed stimulus signal and the team course of mean component 1.

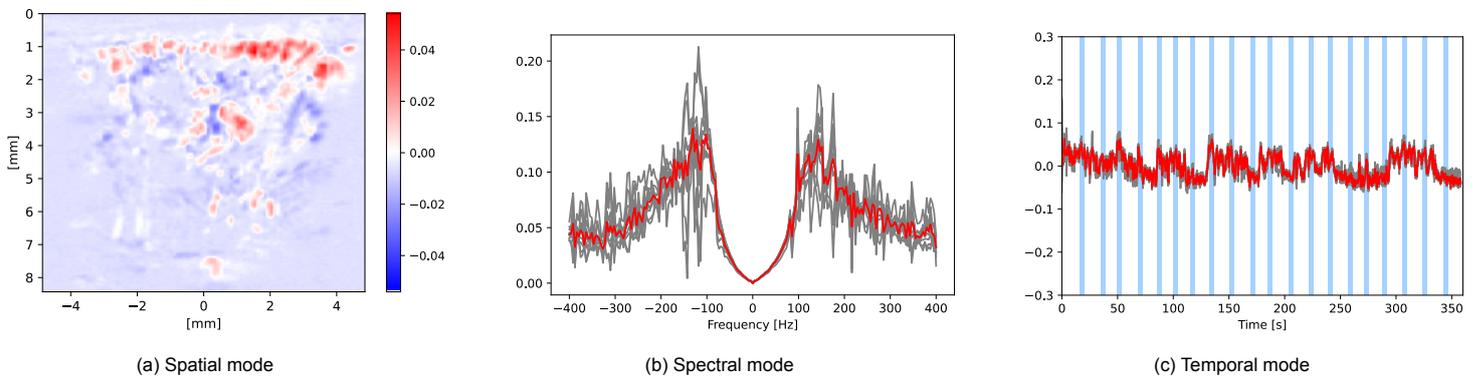


(a) Spatial mode

(b) Spectral mode

(c) Temporal mode

Figure 3.10: Mean component 2, which occurred 13 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.277, the relative height was 3.06 and the half bandwidth was 29.7 Hz.

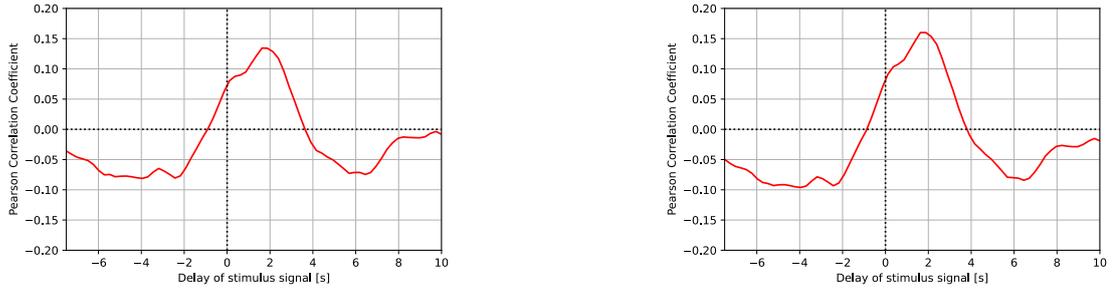


(a) Spatial mode

(b) Spectral mode

(c) Temporal mode

Figure 3.11: Mean component 3, which occurred 11 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.078, the relative height was 3.17 and the half bandwidth was 29.6 Hz.



(a) Correlation with component 2

(b) Correlation with component 3

Figure 3.12: Correlation between the delayed stimulus signal and the time course of mean components 2 and 3

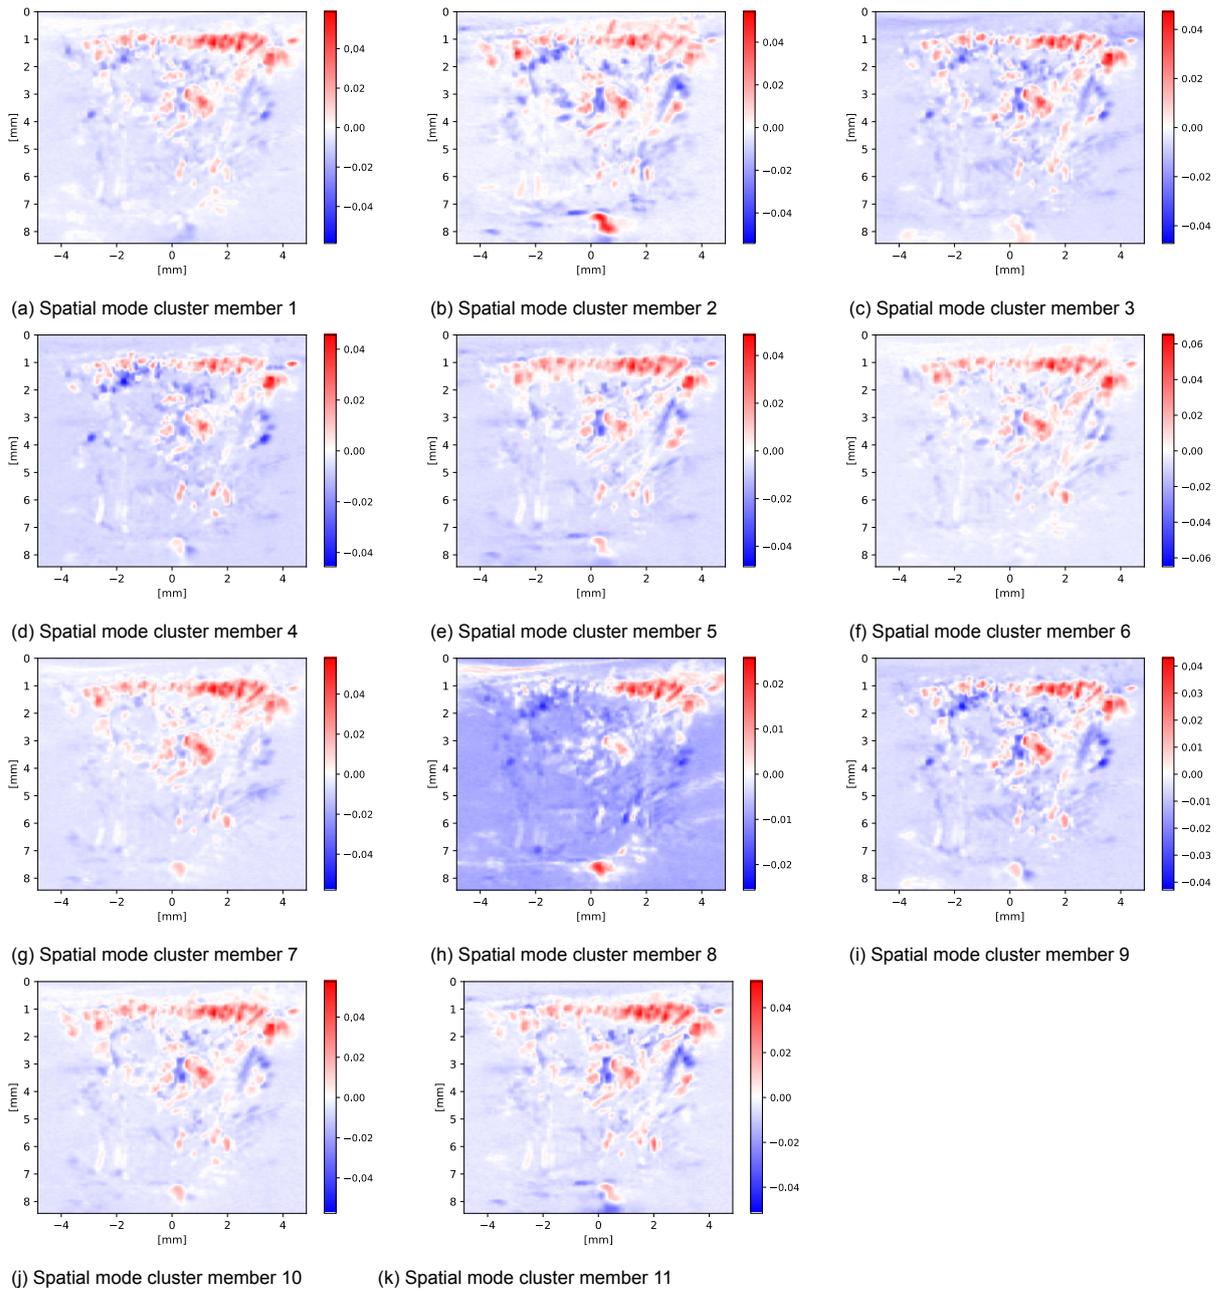


Figure 3.13: The spatial modes of the 11 cluster members belonging to mean component 3

To show the variance in the spatial map of component 3, the spatial maps of all cluster members of component 3 are depicted in Figure 3.13. Overall, one can observe very similar maps with some differences, especially for members 2 and 8 (D.1b, D.1h).

To aid the interpretation of component 3, the visual map was thresholded with 2.5 times the SD and overlaid against the mean PDI, see Figure 3.14a. One can now more easily see that the activated regions correspond to the identified ROIs (Figure 3.1a). The assumption that this component is related to the task signal is further strengthened by observing the correlations with the delayed stimulus for a duration of multiple stimuli, see Figure 3.14b. Peaks in the correlation can be observed at similar intervals as that of the visual stimuli.

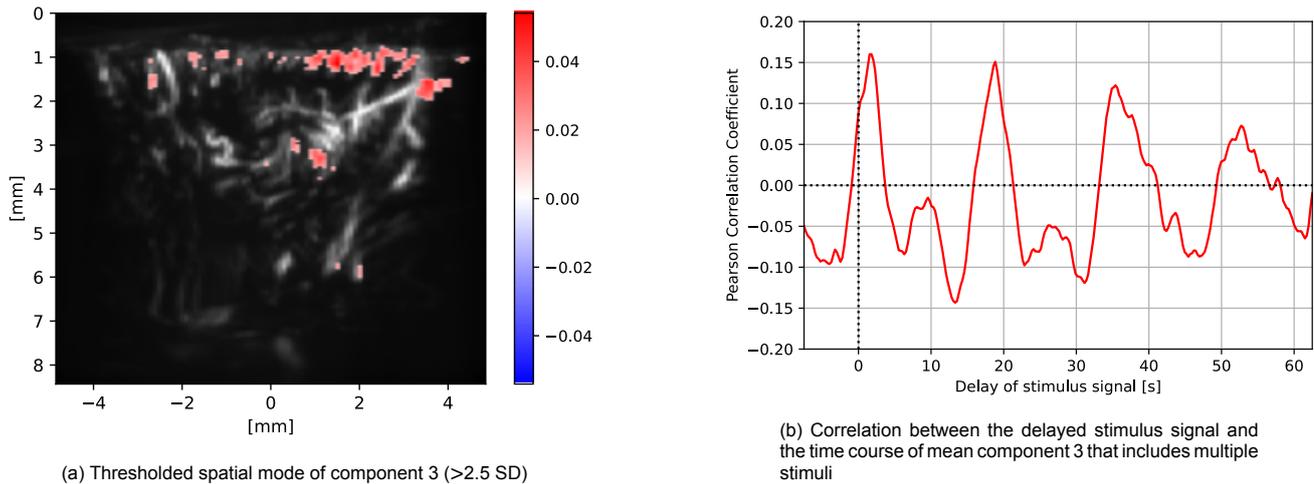


Figure 3.14: Further analysis of component 3

In Figure 3.15 and Figure 3.16 two artifactual appearing mean components can be observed. In the spatial map of component 13, a noisy hue can be observed around the brain. In the spatial map of component 56, a more demarcated artifact is visible in the bottom. When comparing the frequency spectra, component 13 has rather round peaks while component 56 has sharp, well-defined peaks. This is also reflected in the values of the relative height (2.83 vs 4.47) and half bandwidth (31.6 vs 23.1). Another interesting observation for component 13 is the sharp peaks in the time course which are visible in the first ~50 s.

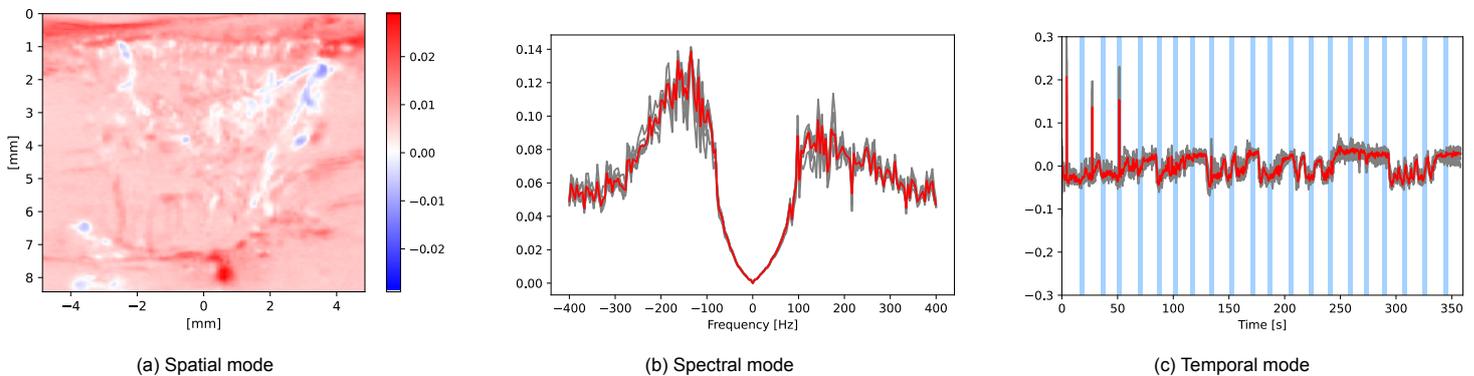


Figure 3.15: Mean component 13, which occurred 8 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.418, the relative height was 2.83 and the half bandwidth was 31.6 Hz.

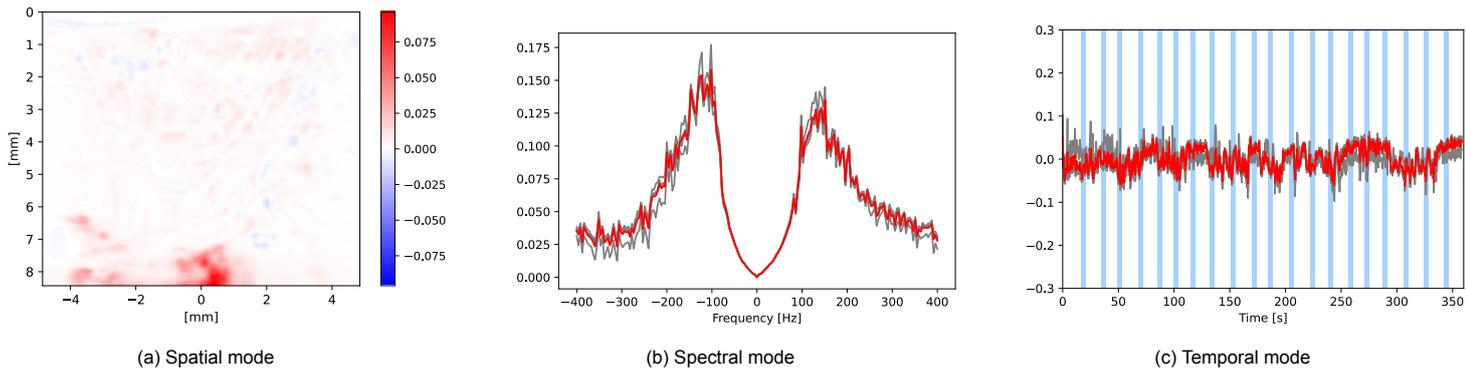


Figure 3.16: Mean component 56, which occurred 5 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.176, the relative height was 4.47 and the half bandwidth was 23.1 Hz.

3.4.2. Slice B

For slice B the mean asymmetry ratio of the components was 1.644 (SD: 0.224), the mean relative height was 4.79 (SD: 1.11), and the mean half bandwidth was 19.8 Hz (SD: 4.82 Hz). Similar to slice A, the frequency spectra of the components contain two peaks, one on the negative side between -100 and -200 Hz, and one on the positive side between 100 and 200 Hz. An interesting observation that was present in almost all temporal maps is the sudden change of the signal around ~ 200 s.

Two components that contain the VIS and LGN are 9 and 50, see Figure 3.17 and Figure 3.18. Nevertheless, activations in other areas are also present. It is relevant to note that all members belonging to cluster 9 and 50 come from different CPD runs. Thus, no components are members of both cluster 9 and 50. When the timecourses are correlated with the timecourse of the task with varying delays, the correlations such as depicted in Figure 3.19 can be observed. For both components, a peak in the correlation can be observed after ~ 2.5 s delay. However, the peak of component 50 is somewhat higher than for component 9.

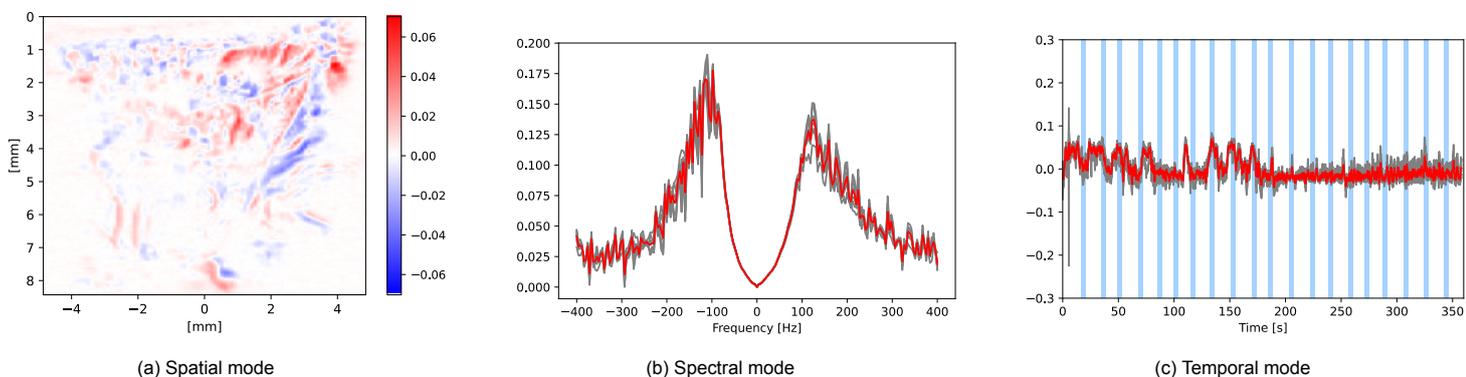


Figure 3.17: Mean component 9, which occurred 9 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.295, the relative height was 4.26 and the half bandwidth was 20.4 Hz.

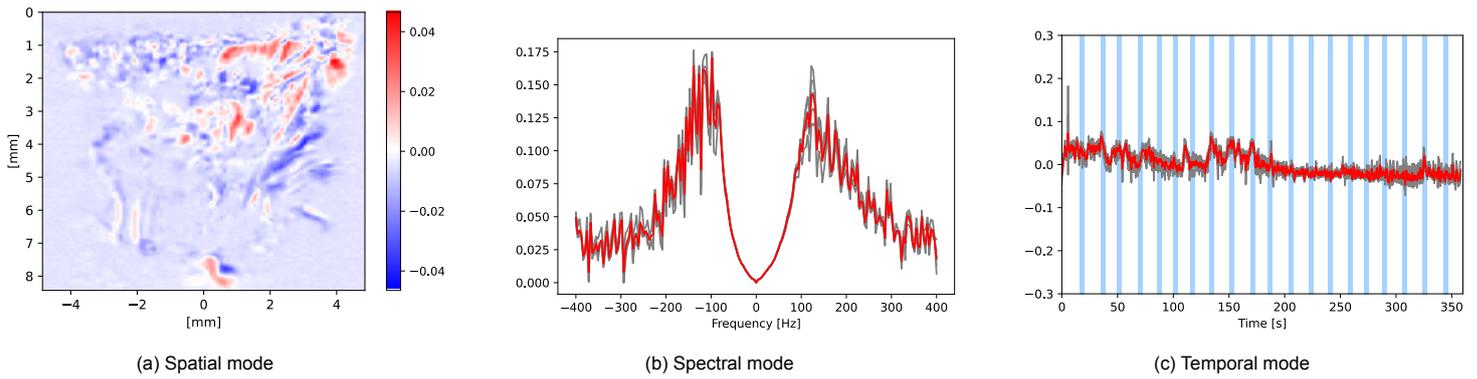


Figure 3.18: Mean component 50, which occurred 5 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.194, the relative height was 3.46 and the half bandwidth was 11.8 Hz.

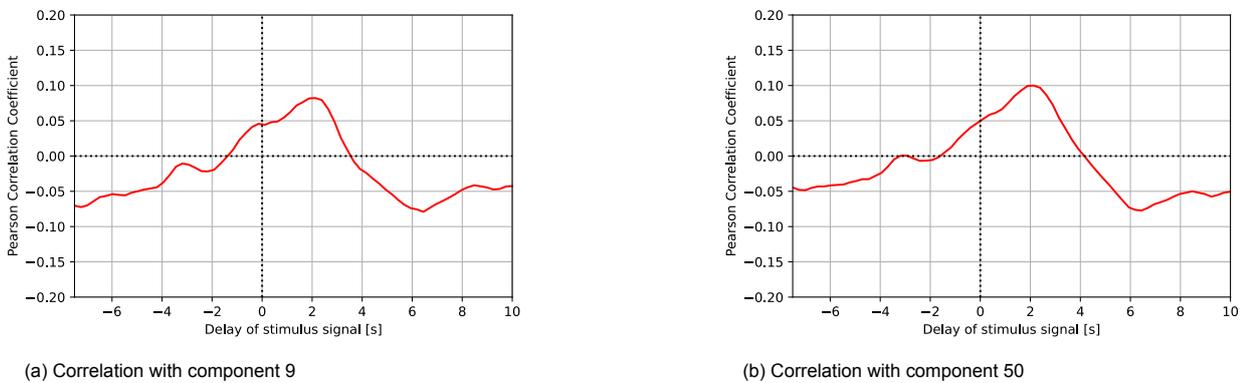


Figure 3.19: Correlation between the delayed stimulus signal and the time course of mean components 9 and 50

To facilitate the interpretation of component 9, the visual map was thresholded at 2.5 times the standard deviation (SD) and overlaid on the mean PDI, as shown in Figure 3.20a. This approach allows for easier identification of the activated regions, which correspond to the identified ROIs (Figure 3.1b). The assumption that this component is related to the task signal is further supported by examining its correlations with the delayed stimulus over multiple stimuli durations (Figure 3.20b). Peaks in the correlation can be observed at similar intervals as that of the visual stimuli.

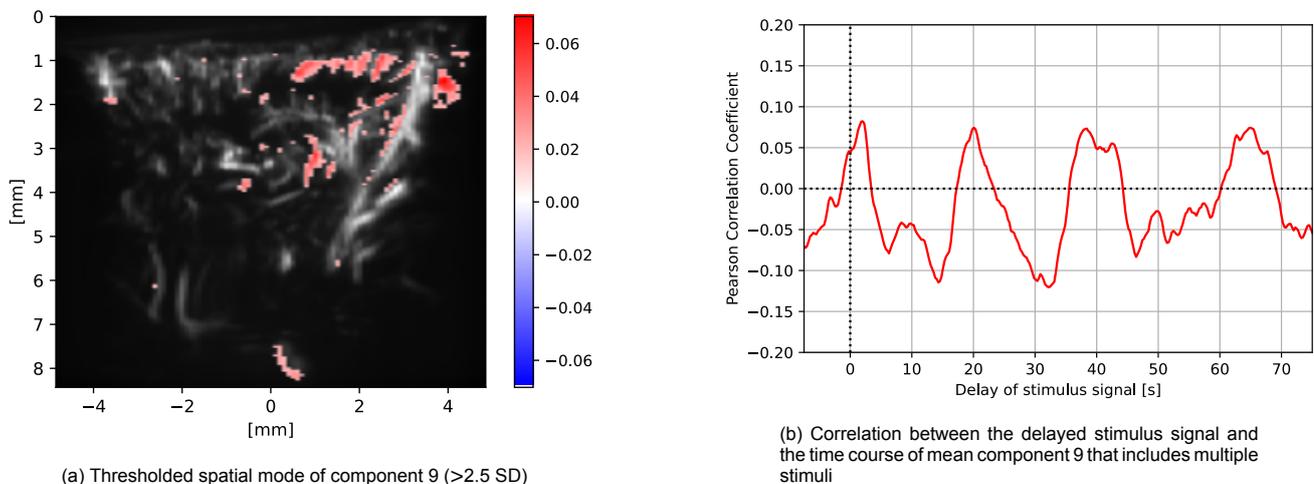


Figure 3.20: Further analysis of component 9

In Figure 3.21 and Figure 3.22 two artifactual appearing mean components can be observed. In the spatial map of component 20, a noisy hue can be observed around the brain, predominantly in the upper part of the spatial map, similar to component 13 of slice A (Figure 3.15). Also similar to component 13 of slice A, a sharp peak in the time course can be observed around ~ 5 s. In the spatial map of component 56, a more demarcated artifact is visible in the bottom, similar to component 56 of slice A (Figure 3.16). Although less pronounced than for slice A, differences in the frequency spectra of the artifactual components can be observed as well. While component 34 has quite round peaks, component 20 has sharp, well-defined peaks. This is also reflected in the values of the relative height (4.11 vs 5.52) and half bandwidth (23.9 vs 20.1).

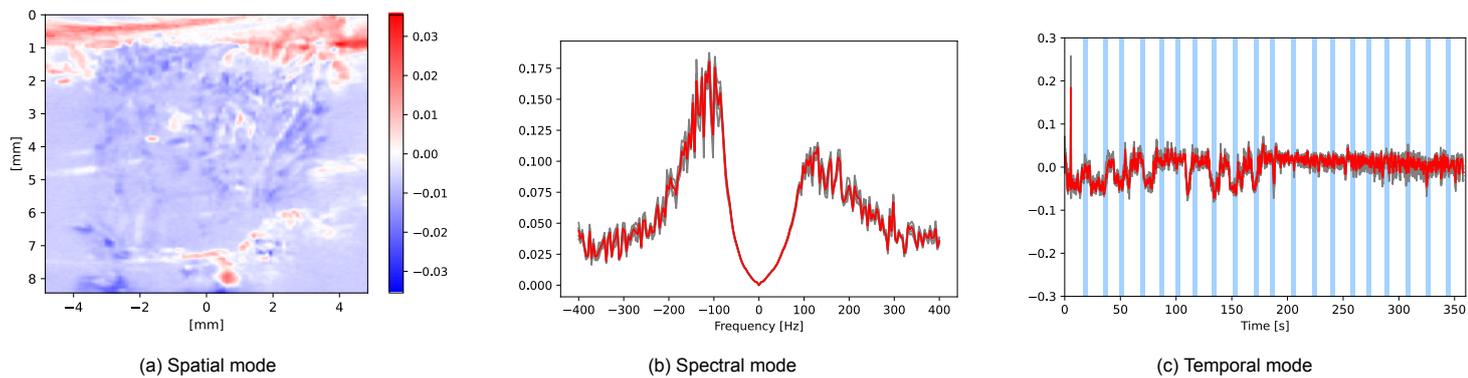


Figure 3.21: Mean component 34, which occurred 5 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.622, the relative height was 4.11 and the half bandwidth was 23.9 Hz.

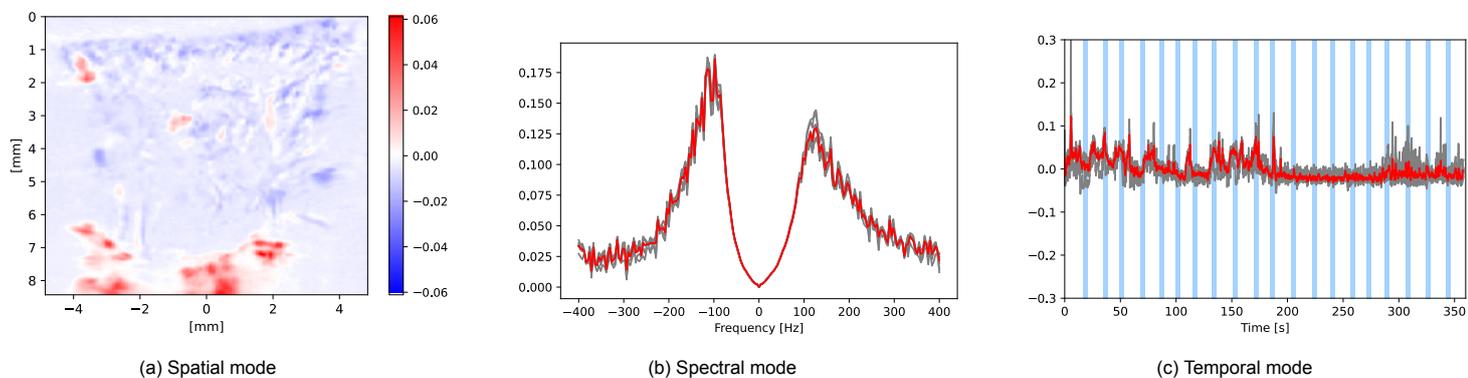


Figure 3.22: Mean component 20, which occurred 7 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines. The asymmetry ratio was 1.428, the relative height was 5.52 and the half bandwidth was 20.1 Hz.

3.4.3. Slice C

For slice C, the components primarily seem to contain noise, and based on visual inspection no spatial map was found that seemed to contain the VIS and LGN. Most components have a spatial map that contains an artifact on the bottom, similar to component 20 of slice B (Figure 3.22) and component 56 of slice A (Figure 3.16). Two examples are component 1, the component with the highest occurrence, and component 6, which are depicted in Figure 3.23 and Figure 3.25.

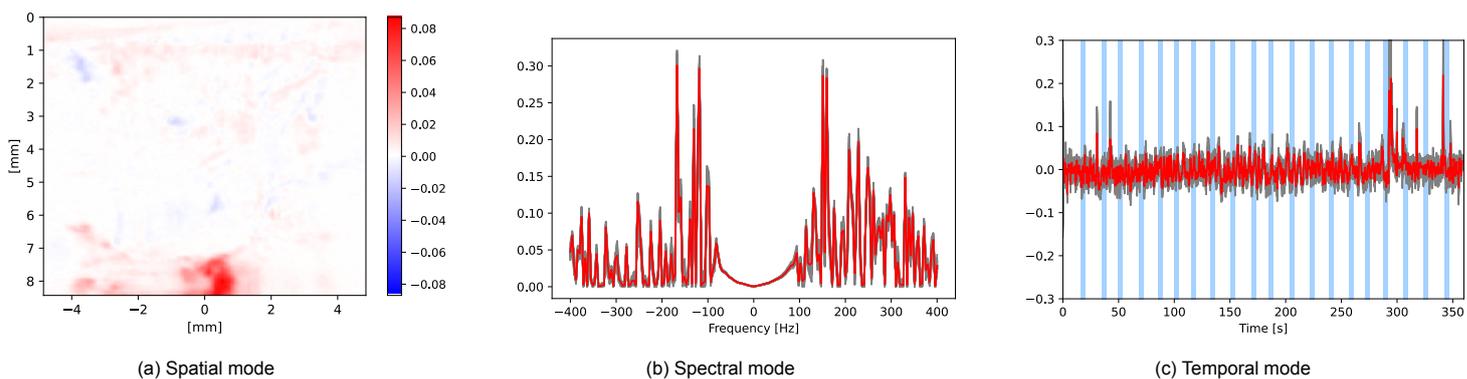


Figure 3.23: Mean component 1, which is the component that was present in the most runs (26 times). The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines.

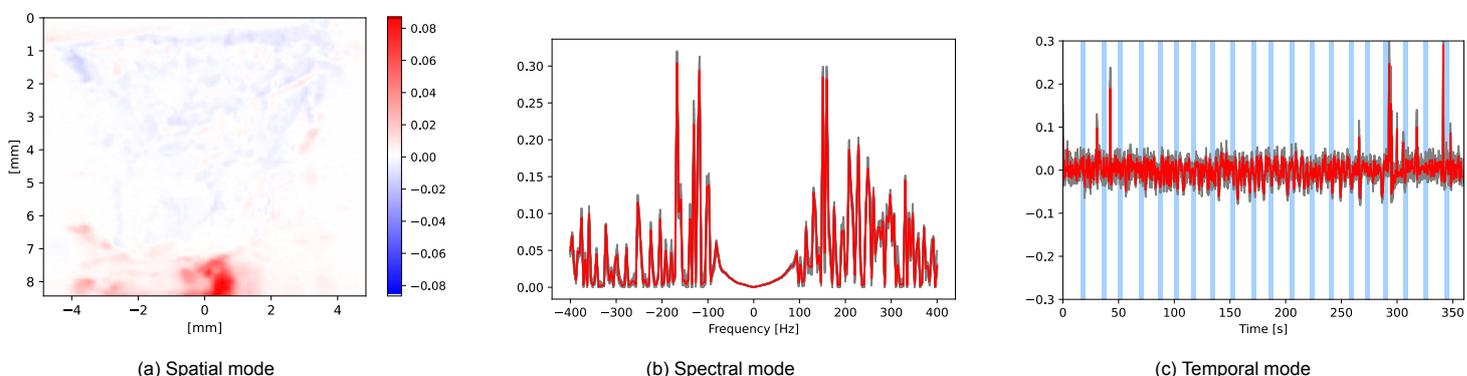


Figure 3.24: Mean component 6, which occurred 12 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines.

Some components also show activations in the upper part of the brain. An example is component 5, depicted in Figure 3.25. In this component, while the artifactual area shows activation (red), the brain area shows deactivation (blue), suggesting anticorrelation. However, this could also be the other way around, given the sign ambiguity of CPD.

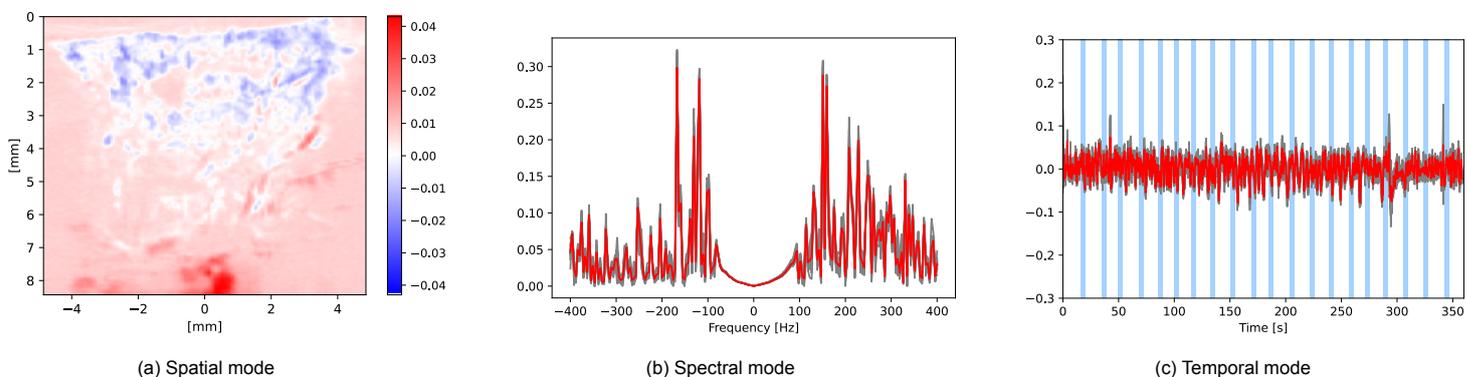
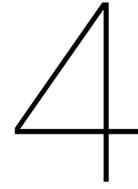


Figure 3.25: Mean component 5, which occurred 14 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines.

The frequency spectra of component 1, 6, and 5 show very sharp peaks and the overall shape seems very rough, as compared to the frequency spectra of component A and B. Based on this, the previously defined metrics do not seem interpretable, and will therefore not be discussed for this slice.



Discussion

This chapter discusses the results and evaluates how they answer the research questions posed in Section 1.7.

Pre-processing

Using an iterative design process, three pre-processing steps were identified to ensure adequate and relevant BSS: filtering, temporal demeaning, and pixel-based normalization.

In the filtering step, SVD filtering was performed using conventional thresholds that have been shown to produce useful results in the literature [19]. The goal of this step is to remove tissue clutter. Given the overlap in subspaces between the clutter and blood signals, the aim was to discard enough of the tissue subspace to ensure that the clutter does not dominate the signal—and, by extension, the CPD components—without removing so much that the blood signal is also lost. By retaining the portion of the clutter subspace that overlaps with the blood signal, these two subspaces could be separated using CPD. For slices A and B, the chosen thresholds appeared to suffice: some components contained clutter, while others contained task-related blood signals. This indicates that the discarded subspace was neither excessive nor insufficient. However, for slice C, the CPD components were dominated by clutter, suggesting inadequate source separation. Possible explanations could be insufficient subspace removal, higher noise variance relative to the blood signal variance in the remaining subspace, or a combination of both. In general, it is ideal to use high-quality data with minimal clutter caused by motion or other artifacts. The PDIs correlated with the task signal (Figure 3.1) suggest that the task signal is less prominent in slice C compared to, for example, slice A, indicating lower data quality. Nonetheless, it may be worth exploring whether slices could benefit from different SVD thresholds. However, manually selecting such thresholds is challenging and not in line with the data-driven nature of the method. Also, the size of the clutter subspace is probably not fixed, indicating a dynamic threshold is needed. In literature, algorithms have been described that allow for automatic, dynamic thresholding, but no 'golden standard' exists [20], [21]. Also, since these algorithms were not designed for the goal of BSS, they might not be useful. A paper by Baranger et al.[22] that compares 14 different estimators for the optimal threshold, concludes that the one that uses the spatial similarity matrix (SSM) outperforms all others based on the contrast-to-noise ratio. To check if this algorithm would also be useful from a functional connectivity perspective, the raw data of slice C was SVD filtered with thresholds based on the proposed algorithm[23], and the resulting correlation image was compared to the one following from the previously defined, fixed threshold, see Appendix B. However, the (preliminary) result is not convincing, since the correlation image seems to have lower correlations.

Something that is not common practice in fUS processing is the normalization of the singular values performed in the filtering step. Based on the results presented in Figure 3.4, this step seems to improve the results. For all three slices, the maximum frequency of occurrence increases. Especially for slice C a large difference can be observed between the pipeline with and without this step. Therefore, the normalization of the singular values is deemed a valuable addition to the pipeline.

In the other two pre-processing steps the data was further optimized for BSS. In the temporal demeaning step, the offset was removed, and in the pixel-based normalization step, the influence of large

vessels was reduced.

The primary objective of the pre-processing steps in the pipeline is to generate CPD components that align with the goal of this study, namely, *identifying artifacts and meaningful neurological components*. In data where such meaningful neurological signals are not very strong (e.g. slice C) other signals tend to be more dominant, making it challenging to extract the desired signals, e.g. task-related ones. While it can be argued that this reflects the nature of the actual data, it is less relevant from a neuroscientific perspective. A straightforward approach would be to only use data wherein the (strong) presence of certain desired signals are known. However, this may not always be predefined or easy to verify. Introducing additional preprocessing steps to enhance relevant signals and suppress irrelevant ones could therefore be beneficial. That said, achieving this in a data-driven manner may be difficult since it potentially requires making assumptions about the source signals — an approach that could undermine the intended purpose. A method that is very controversial but often used in functional neuroimaging, specifically for fMRI, is global signal regression (GSR) [24]. Here, a global signal is assumed to be present in the data, which should be removed to emphasize relevant signals. The global signal is removed from the time series of each voxel through linear regression. The global signal can be defined as the signal obtained from averaging all timeseries along the spatial dimension. Different researchers have different ideas as to whether GSR should be used in the processing of fMRI. While some reason that GSR can remove unwanted global confounds, others conclude that it can introduce artefactual anti-correlations that were not originally present in the modeled data [24]. Although GSR requires the assumption that a certain 'global signal' is present in the fUS data of this work, it could be a useful addition. Therefore, this step was added to the pipeline and the results were compared to the results of the pipeline without this step. However, the results were not convincing, and the pipeline seemed to produce worse results, see Appendix C.

Compression

Successful compression was achieved using MLSVD, enabling the pipeline to operate within standard RAM constraints. MLSVD retains the information that captures the highest variance in the data, which distinguishes it from approaches such as (random) sampling. However, compression inherently results in the loss of certain parts of the data. Whether this is disadvantageous in this specific context remains unclear. Ideally, the discarded subspace primarily contains noise, in which case MLSVD not only compresses but also denoises the data. Nevertheless, the discarded subspace might also include neuronal information with low variance, for example, task-specific resting state neural activity.

While (computational) efficiency was not within the scope of this work, two relevant remarks can be made regarding this. First, since both the filtering and compression steps involve SVD, it may be possible to integrate them into a single step, thereby improving the efficiency of the pipeline. However, because the filtering step occurs before temporal demeaning and pixel-based normalization, while the MLSVD is performed afterward, integrating these steps is not straightforward. Second, given the large raw data tensor and therefore large matrices that are used as input to the randomized SVD, the compression (and pre-processing) is very computationally costly, and takes multiple days (~2). While optimized to a certain extent, there are likely more optimizations possible to reduce this run time. Nevertheless, one should realize that pre-processing and compression only need to be performed once for a sufficiently high rank. After this, the MLSVD representation can simply be truncated to other (smaller) ranks and be used for CPD, which (due to the compression) has a runtime of only a few seconds.

In general, both the issues of discarding information by compression and the computational cost could be improved upon by using hardware that has a (much) higher RAM and allows for faster processing. Nevertheless, this would make it much less accessible for neuroscientific research.

Stability CPD components

To identify stable components, 100 CPD iterations were performed, and the results were clustered to account for the non-convexity of the cost function. A high occurrence of a cluster indicates that the corresponding component is highly reproducible or "stable," as it appears in the factor matrices across many iterations. By using the rank estimation process such as defined in Figure 2.3, it was possible to choose optimal ranks based on the pre-defined metrics (frequency of occurrence and correlation).

For every slice a peak in frequency of occurrence was found for a rank that also had variety in the distribution of the correlations. This strengthens the idea that the proposed method for rank estimation is useful.

In the ideal situation, the factorizations consistently converge to the same rank R solutions in every run, and thus the number of clusters would equal R and each cluster would occur 100 times. For slice A, 62 clusters were identified, and the most stable cluster occurred 30 times. While 30 out of 100 repetitions might not seem particularly frequent, it is worth noting that visual inspection revealed many similar clusters. For instance, multiple clusters resembled a similar artifact or brain network across different runs, likely resulting from different local optima reached during the decompositions. Examples are the components in Figure 3.23 and 3.24 and the components in Figure 3.17 and 3.18. In general, greater trust should be placed in components belonging to large clusters (with higher occurrences), as smaller clusters are more likely to represent components tied to individual, potentially suboptimal, local minima. One way to increase the number of occurrences per cluster and reduce the total number of clusters could be to modify the initialization of the CPD. In this work, random initialization was used to maintain a data-driven approach. However, using a non-random initialization might restrict the process to a smaller subset of local minima, potentially increasing the number of occurrences per cluster and decreasing the overall number of clusters. Nevertheless, selecting the right initialization would be challenging and would inherently require making assumptions, which could undermine the data-driven nature of the approach and introduce bias.

Neuroscientific relevance

Overall, the proposed method appears to generate components that are neuroscientifically meaningful. For both slices A and B, components were identified that include the VIS and LGN and exhibit a peak correlation with the stimulus approximately 2.5 seconds after its onset. These brain regions and timing are consistent with expectations for a visual stimulus and align with findings in the literature [12], [25]. Additionally, the timing of the peak correlations with the stimulus follows a pattern consistent with expectations, occurring at intervals similar to the spacing between the stimuli. Furthermore, especially for slice A, it is worth emphasizing that these components have a relatively high occurrence (stability), which underscores the relevance of the results and therefore the proposed method.

A less straightforward component is component 1 of slice A, which occurs the most and is thus the most stable. The spatial map seems to indicate whole-brain activation and a relatively lower activation of the background. Moreover, the timecourse has relatively high correlations with the stimulus. There are two potential explanations for this phenomenon. The first is based on neuroscientific literature that supports the idea that visual stimuli can elicit brain-wide activation by engaging both sensory processing and arousal-related mechanisms [26]–[28]. Arousal-related factors, such as running, whisking, and pupil diameter, influence neuronal activity and modulate the activity of sensory cortices, including the visual cortex. This behaviorally related neuronal activity extends across the brain [28]. If this component indeed reflects whole-brain activation caused by arousal, it could also account for its correlation with the stimulus. The hypothesis, in this case, would be that the mouse anticipates the visual stimulus, entering a state of arousal shortly before the stimulus (<2.5 s delay) and returning from this state shortly afterward (>2.5 s delay). However, this explanation does not account for the relatively strong activation of the background, which appears to be an artifactual baseline. This leads to a second potential explanation: the component might represent an artifact introduced by the pipeline. It might be possible that one of the preprocessing steps introduces an artificial baseline. Alternatively, many of the CPD factorizations could include a component representing nonspecific, residual variance, which may then cluster together during the clustering step. For completeness, the spatial maps of all members of component 1 are presented in Appendix D. To verify if the whole-brain activation remains present when the background is removed, a masking procedure was performed. See Appendix E for details. The whole-brain activation is observed in the components following from both the masked and unmasked data. Therefore, this suggests that the whole-brain activation is not (entirely) artifactually driven by pixels outside of the brain, indicating potential physiological relevance.

For all three slices, components containing nuisance signals were identified. The shape of these nuisances is generally well-defined, with limited activations within the brain. It is particularly convincing that the same types of artifacts consistently appear across all three slices (e.g., Figures 3.16, 3.22, 3.23, 3.24 and 3.15, 3.21). The slices were acquired at (somewhat) different time points, and thus the nuisance caused by mouse movements is not identical. However, the shape of such nuisance signals

can be expected to remain similar.

Based on the obtained frequency spectra, which are especially novel about the proposed method, some interesting observations could be made. In general, a 'gap' in energy around 0 Hz is present in the frequency spectra, which is to be expected due to the SVD filtering. Furthermore, for all spectra, two peaks in energy are visible between +/-100 and +/-200 Hz. This matches the observations found in the thesis of Soloukey Tbalvandany[29], where a correlation analysis showed that correlations of the VIS with the stimulus are particularly high for the frequency subband 100-200 Hz (Figure 12.3). Apart from the similarities across frequency spectra of the components, differences could be observed as well. Given these differences, one can assume the frequency dimension contributes to the BSS, i.e. helps separate the underlying sources. Based on the defined metrics, some patterns can be observed across the components. Especially specific characteristics were found for 'local' artifacts, such as Figure 3.16 and 3.22, and more global artifacts, such as Figure 3.15 and 3.21. When comparing the two, the former has rather round peaks and low relative height, while the latter has sharper, more well-defined peaks. This information regarding the frequency content of underlying sources is new and is not possible to obtain with matrix-based methods that only include space and time information. Initially, the author suspected that components displaying a task-related activation would have a frequency spectrum with energy concentrated on one side of the spectrum. However, for Figure 3.10, 3.13, 3.17 and 3.18 this does not seem to be the case, as a peak on the positive and negative side is visible with relatively similar energy. An explanation could be that although the blood does all flow towards a certain task-related brain area, i.e. feeding vessels, this blood does not necessarily all flow towards or from the fUS probe. For example, blood that flows through vessels that are perpendicular to the imaging plane or have certain bends that cross the imaging plane will result in both positive and negative frequencies, while they still flow towards the task-related brain area. This is similar to the visual representation of the brain in Figure 1.3a.

When comparing the results of the proposed method with results generated by matrix-based methods, e.g. ICA, some interesting differences can be observed. In spatial ICA, where spatial independence is assumed, brain regions are often quite well demarcated, with very local activations and little to no activations elsewhere in the brain. This eases the interpretation of such components, and people with limited prior experience can interpret them. For the results of the method proposed in this work, a deeper understanding is needed to interpret the results. For example for the component in Figure 3.13, multiple areas are activated (VIS and LGN), and a deeper understanding can be achieved by analyzing the correlation with the task signal. However, the author argues that ICA gives too simplified results, ignoring functional connectivity. More specifically, in ICA decompositions, the VIS and LGN would show up in two separate components, while for the method in this work the VIS and LGN show up simultaneously in both slice A and slice B. Therefore, functional connectivity is suggested in the method of this work, while this is not the case for ICA. Since the overall goal of functional neuroimaging, and BSS specifically, is to unravel the underlying brain dynamics, being able to represent such functional connectivities in the components is critical. Similar distinctions can be observed in the handling of artifactual components. In ICA, spatially distributed noise, such as that seen in Figure 3.21, is often split into multiple components. This fragmentation limits the observation of global dynamics and complicates attempts to infer the origins of such artifacts.

5

Conclusion

In conclusion, the proposed tensor-based blind source separation pipeline for raw fUS data was able to identify artifacts and meaningful neurological components based on distinctive characteristics in the temporal, spatial, and spectral domains.

Application to task experiment data of a mouse verified that activity that is temporally correlated with the stimulus can be extracted in expected regions. The components also indicate expected functional connectivity, which is often not the case for matrix-based methods. Moreover, frequency spectra showed different characteristics for different types of components, which displays the relevance of this dimension for BSS, especially for nuisance components. The method is entirely data-driven, can be applied to entire raw fUS datasets, and can run on a laptop with standard RAM.

Nevertheless, further research is necessary to ensure the utility of the method. Data of other experiments and cranially imaged data should be tested to see if relevant components can be extracted.

Future work

This work serves as a starting point for more advanced denoising, the identification of novel brain networks, and the comparison of brain networks across healthy and pathological conditions. To achieve these goals, it is essential to obtain the most optimal CPD factorizations for a given dataset. Here, "optimal" refers to reaching the local optimum that is the most neurologically plausible. Achieving this likely requires further optimization of the pipeline, particularly in terms of CPD initialization and pre-processing, both of which depend on the specific experimental context.

To define an appropriate initialization, the pipeline could initially be run with random initializations. The resulting components that are deemed relevant could then serve as a starting point for a more informed initialization that leads to an optimal CPD factorization.

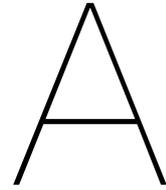
For denoising, once such an optimal CPD factorization is achieved, noise-containing components can be discarded, and the data can be reconstructed to produce a denoised version. It is important to note that, in this work, clusters are averaged to form mean components. However, these mean components cannot be used for reconstruction since they are not part of the actual CPD factorization. However, similarity metrics can be used to identify a component within a cluster that is most similar to other members of the cluster, which can then be used for reconstruction.

Bibliography

- [1] G. Montaldo, A. Urban, and E. Macé, “Annual Review of Neuroscience Functional Ultrasound Neuroimaging,” 2022. DOI: 10.1146/annurev-neuro-111020. [Online]. Available: <https://doi.org/10.1146/annurev-neuro-111020->.
- [2] T. Deffieux, C. Demene, M. Pernot, and M. Tanter, “Functional ultrasound neuroimaging: a review of the preclinical and clinical state of the art,” *Current Opinion in Neurobiology*, vol. 50, pp. 128–135, Jun. 2018, ISSN: 09594388. DOI: 10.1016/j.conb.2018.02.001.
- [3] A. Y. Shih, P. Blinder, P. S. Tsai, *et al.*, “The smallest stroke: occlusion of one penetrating vessel leads to infarction and a cognitive deficit,” *Nature Neuroscience*, vol. 16, no. 1, pp. 55–63, Jan. 2013, ISSN: 1097-6256. DOI: 10.1038/nn.3278.
- [4] M. F. Glasser, T. S. Coalson, J. D. Bijsterbosch, *et al.*, “Using temporal ICA to selectively remove global noise while preserving global signal in functional MRI data,” *NeuroImage*, vol. 181, pp. 692–717, Nov. 2018, ISSN: 1095-9572. DOI: 10.1016/j.neuroimage.2018.04.076.
- [5] Z. Chen and V. D. Calhoun, “Task-evoked brain functional magnetic susceptibility mapping by independent component analysis (ICA),” *Journal of Neuroscience Methods*, vol. 261, pp. 161–171, Mar. 2016, ISSN: 01650270. DOI: 10.1016/j.jneumeth.2016.01.007.
- [6] N. Wang, W. Zeng, Y. Shi, *et al.*, “WASICA: An effective wavelet-shrinkage based ICA model for brain fMRI data analysis,” *Journal of Neuroscience Methods*, vol. 246, pp. 75–96, May 2015, ISSN: 01650270. DOI: 10.1016/j.jneumeth.2015.03.011.
- [7] L. Griffanti, G. Salimi-Khorshidi, C. F. Beckmann, *et al.*, “ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging,” *NeuroImage*, vol. 95, pp. 232–247, Jul. 2014, ISSN: 10538119. DOI: 10.1016/j.neuroimage.2014.03.034.
- [8] M. J. McKeown, L. K. Hansen, and T. J. Sejnowski, “Independent component analysis of functional MRI: what is signal and what is noise?” *Current opinion in neurobiology*, vol. 13, no. 5, pp. 620–9, Oct. 2003, ISSN: 0959-4388. DOI: 10.1016/j.conb.2003.09.012.
- [9] A. M. Golestani, C. Chang, J. B. Kwinta, Y. B. Khatamian, and J. Jean Chen, “Mapping the end-tidal CO₂ response function in the resting-state BOLD fMRI signal: Spatial specificity, test–retest reliability and effect of fMRI sampling rate,” *NeuroImage*, vol. 104, pp. 266–277, Jan. 2015, ISSN: 10538119. DOI: 10.1016/j.neuroimage.2014.10.031.
- [10] A. Cichocki, N. Lee, I. Oseledets, A.-H. Phan, Q. Zhao, and D. P. Mandic, “Tensor Networks for Dimensionality Reduction and Large-scale Optimization: Part 1 Low-Rank Tensor Decompositions,” *Foundations and Trends® in Machine Learning*, vol. 9, no. 4-5, pp. 249–429, 2016, ISSN: 1935-8237. DOI: 10.1561/22000000059.
- [11] G. Paxinos and K. B. Franklin, *The Mouse Brain in Stereotaxic Coordinates*, 2nd Edition. 2001, ISBN: 9780125476362.
- [12] A. Erol, C. Soloukey, B. Generowicz, *et al.*, “Deconvolution of the Functional Ultrasound Response in the Mouse Visual Pathway Using Block-Term Decomposition,” *Neuroinformatics*, vol. 21, no. 2, pp. 247–265, Apr. 2023, ISSN: 15590089. DOI: 10.1007/s12021-022-09613-3.
- [13] N. Halko, P. G. Martinsson, and J. A. Tropp, “Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions,” Tech. Rep., 2010.
- [14] X. Feng, W. Yu, and Y. Xie, “Pass-Efficient Randomized SVD with Boosted Accuracy,” in 2023, pp. 3–20. DOI: 10.1007/978-3-031-26387-3_{_}1.
- [15] N. Vervliet, O. Debals, and L. De Lathauwer, “Exploiting Efficient Representations in Large-scale Tensor Decompositions,” Tech. Rep.

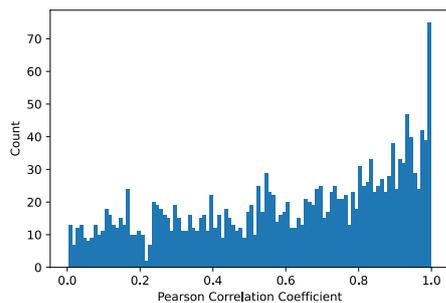
- [16] R. Bro and S. De Jong, "A fast non-negativity-constrained least squares algorithm," *Journal of Chemometrics*, vol. 11, no. 5, pp. 393–401, Sep. 1997, ISSN: 0886-9383. DOI: 10.1002/(SICI)1099-128X(199709/10)11:5<393::AID-CEM483>3.0.CO;2-L.
- [17] S. V. Eyndhoven, N. Vervliet, L. D. Lathauwer, and S. V. Huffel, "Identifying Stable Components of Matrix /Tensor Factorizations via Low-Rank Approximation of Inter-Factorization Similarity," in *2019 27th European Signal Processing Conference (EUSIPCO)*, IEEE, Sep. 2019, pp. 1–5, ISBN: 978-9-0827-9703-9. DOI: 10.23919/EUSIPCO.2019.8902954.
- [18] T. Adali, F. Kantar, M. A. B. S. Akhonda, S. Strother, V. D. Calhoun, and E. Acar, "Reproducibility in Matrix and Tensor Decompositions: Focus on model match, interpretability, and uniqueness," *IEEE Signal Processing Magazine*, vol. 39, no. 4, pp. 8–24, Jul. 2022, ISSN: 15580792. DOI: 10.1109/MSP.2022.3163870.
- [19] C. Demene, T. Deffieux, M. Pernot, *et al.*, "Spatiotemporal Clutter Filtering of Ultrafast Ultrasound Data Highly Increases Doppler and fUltrasound Sensitivity," *IEEE Transactions on Medical Imaging*, vol. 34, no. 11, pp. 2271–2285, Nov. 2015, ISSN: 0278-0062. DOI: 10.1109/TMI.2015.2428634.
- [20] G. Lacoïn, I. Zemmoura, J.-L. Gennisson, D. Kouamé, and J.-P. Remenieras, "Multi-layered adaptive neoangiogenesis Intra-Operative quantification (MANIOQ)," *Journal of Cerebral Blood Flow & Metabolism*, vol. 43, no. 9, pp. 1557–1570, Sep. 2023, ISSN: 0271-678X. DOI: 10.1177/0271678X231170504.
- [21] B. Fang, F. Meng, Y. Chen, J. Luo, and X. Luo, "A Competitive Swarm Optimized SVD-based Clutter Filter," in *2023 IEEE International Ultrasonics Symposium (IUS)*, IEEE, Sep. 2023, pp. 1–4, ISBN: 979-8-3503-4645-9. DOI: 10.1109/IUS51837.2023.10306441.
- [22] J. Baranger, B. Arnal, F. Perren, O. Baud, M. Tanter, and C. Demene, "Adaptive Spatiotemporal SVD Clutter Filtering for Ultrafast Doppler Imaging Using Similarity of Spatial Singular Vectors," *IEEE Transactions on Medical Imaging*, vol. 37, no. 7, pp. 1574–1586, Jul. 2018, ISSN: 0278-0062. DOI: 10.1109/TMI.2018.2789499.
- [23] J. Baranger, J. Aguet, and O. Villemain, "Fast Thresholding of SVD Clutter Filter Using the Spatial Similarity Matrix and a Sum-Table Algorithm," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 70, no. 8, pp. 821–830, Aug. 2023, ISSN: 0885-3010. DOI: 10.1109/TUFFC.2023.3289235.
- [24] K. Murphy and M. D. Fox, "Towards a consensus regarding global signal regression for resting state functional connectivity MRI," *NeuroImage*, vol. 154, pp. 169–173, Jul. 2017, ISSN: 10538119. DOI: 10.1016/j.neuroimage.2016.11.052.
- [25] H. Lambers, M. Segeroth, F. Albers, L. Wachsmuth, T. M. van Alst, and C. Faber, "A cortical rat hemodynamic response function for improved detection of BOLD activation under common experimental conditions," *NeuroImage*, vol. 208, p. 116446, Mar. 2020, ISSN: 10538119. DOI: 10.1016/j.neuroimage.2019.116446.
- [26] D. Shimaoka, K. D. Harris, and M. Carandini, "Effects of Arousal on Mouse Sensory Cortex Depend on Modality.," *Cell reports*, vol. 22, no. 12, pp. 3160–3167, Mar. 2018, ISSN: 2211-1247. DOI: 10.1016/j.celrep.2018.02.092.
- [27] M. Vinck, R. Batista-Brito, U. Knoblich, and J. A. Cardin, "Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding," *Neuron*, vol. 86, no. 3, pp. 740–754, May 2015, ISSN: 08966273. DOI: 10.1016/j.neuron.2015.03.028.
- [28] C. Stringer, M. Pachitariu, N. Steinmetz, C. B. Reddy, M. Carandini, and K. D. Harris, "Spontaneous behaviors drive multidimensional, brainwide activity," *Science*, vol. 364, no. 6437, Apr. 2019, ISSN: 0036-8075. DOI: 10.1126/science.aav7893.
- [29] S. Soloukey Tbalvandany, "Everything flows: Functional ultrasound imaging and neuromodulation of the brain and spinal cord," Ph.D. dissertation, Erasmus University Rotterdam, 2023, pp. 293–318.
- [30] S. Van Eyndhoven, P. Dupont, S. Tousseyn, *et al.*, "Augmenting interictal mapping with neurovascular coupling biomarkers by structured factorization of epileptic EEG and fMRI data," *NeuroImage*, vol. 228, p. 117652, Mar. 2021, ISSN: 10538119. DOI: 10.1016/j.neuroimage.2020.117652.

-
- [31] S.-E. Kotti and B. Hunyadi, "Extracting Hemodynamic Activity with Low-Rank Spatial Signatures in Functional Ultrasound Using Tensor Decompositions," in *2024 32nd European Signal Processing Conference (EUSIPCO)*, IEEE, Aug. 2024, pp. 1347–1351, ISBN: 978-9-4645-9361-7. DOI: 10.23919/EUSIPCO63174.2024.10714979.

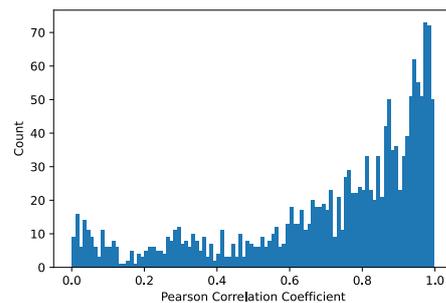


Parameters clustering algorithm

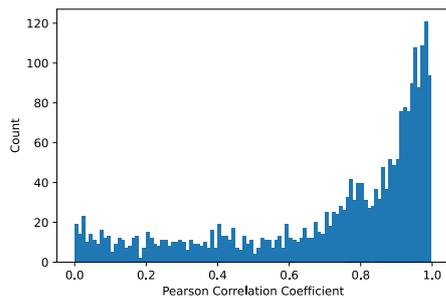
In the clustering algorithm of Van Eyndhoven et al.[17], different algorithms are available for the construction of the similarity matrix as well as the clustering itself. The differences among the combinations of settings can be observed in Figure A.1. Given the dominant high temporal correlations of options *b* and *c*, these options were deemed inferior. Although based on visual inspection the components of options *a* and *d* were similar, option *a* was chosen because of the higher variety in the correlations. For more information regarding the settings, see the associated code of [17] and [30] on GitHub (<https://github.com/svaneynd/structured-cmtf>).



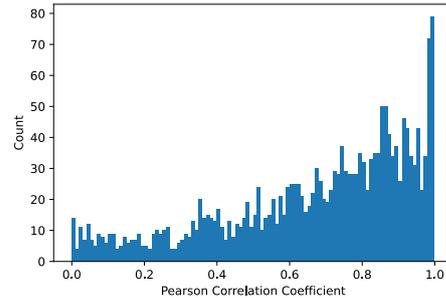
(a) Combination of Hungarian and Hungarian algorithm.
Maximum frequency of occurrence = 30. Number of clusters = 62.



(b) Combination of greedy and average linkage algorithm.
Maximum frequency of occurrence = 51. Number of clusters = 57.



(c) Combination of greedy and Hungarian algorithm.
Maximum frequency of occurrence = 14. Number of clusters = 69.



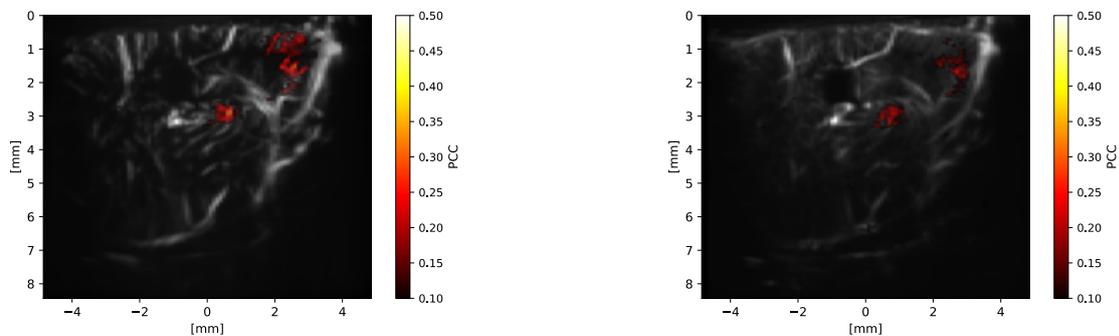
(d) Combination of Hungarian and average linkage algorithm.
Maximum frequency of occurrence = 49. Number of clusters = 63.

Figure A.1: Histograms of timecourse correlations between mean components for different clustering settings for slice A

B

SVD filtering threshold

Slice C was SVD filtered based on the algorithm by Barranger et al.[23], see Figure B.1. The correlations are more pronounced for the conventional threshold than for the dynamic threshold (Figure B.1a v.s. B.1b).



(a) Conventional threshold

(b) Algorithm by Barranger et al.

Figure B.1: Thresholded correlation images ($PCC > 0.10$) of slice C overlaid against the mean PDI for the fixed threshold used in this work and the dynamic, automatically chosen threshold of the algorithm of Barranger et al.[23]. PCC = Pearson Correlation Coefficient

C

Global Signal Regression

GSR was added as a last pre-processing step before compression. For slice A, rank 15 was chosen based on the frequency of occurrence and correlations between mean components. In Figure C.1 and Figure C.2, two components resembling components 2 and 3 (Figure 3.10 and 3.13) of the pipeline without GSR are depicted. In Figure C.3, their respective correlations with the stimulus are presented. Compared to the components of the unregressed pipeline, the components have a lower frequency of occurrence and a less defined VIS. Furthermore, artifactual spikes are visible in the timecourse of mean component 16. The correlation with the stimulus of component 15 is slightly higher than for the components from the unregressed pipeline. However, the highest correlation with the stimulus of component 16 is lower and has an altered timing (~ 1 s). Based on these results, adding the GSR step was not deemed beneficial.

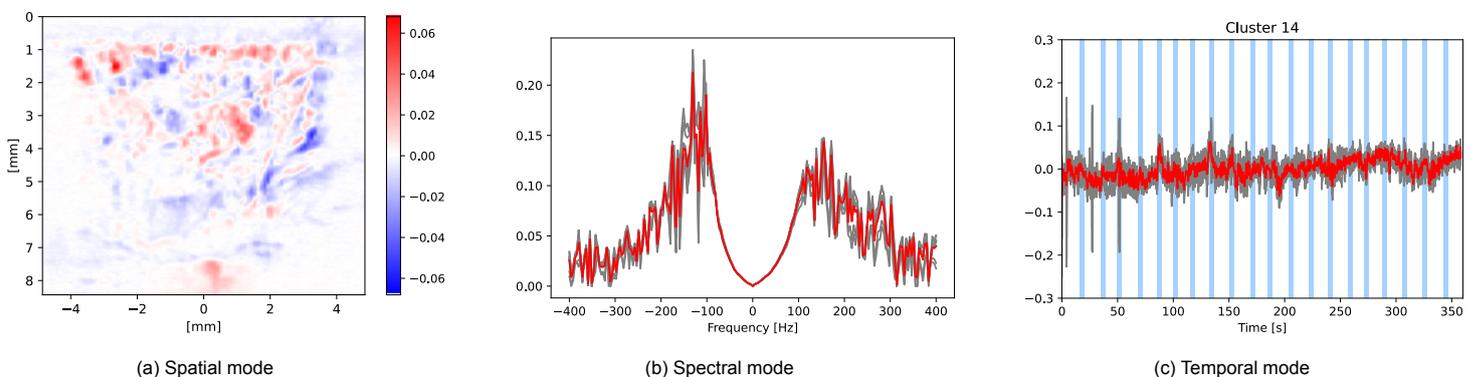


Figure C.1: Mean component 15, which occurred 8 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines.

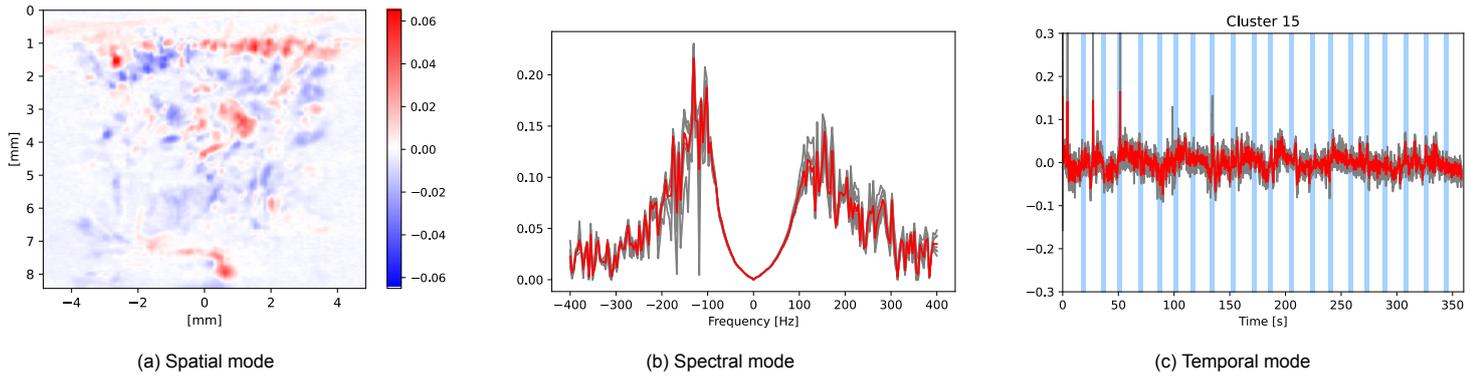


Figure C.2: Mean component 16, which occurred 7 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines.

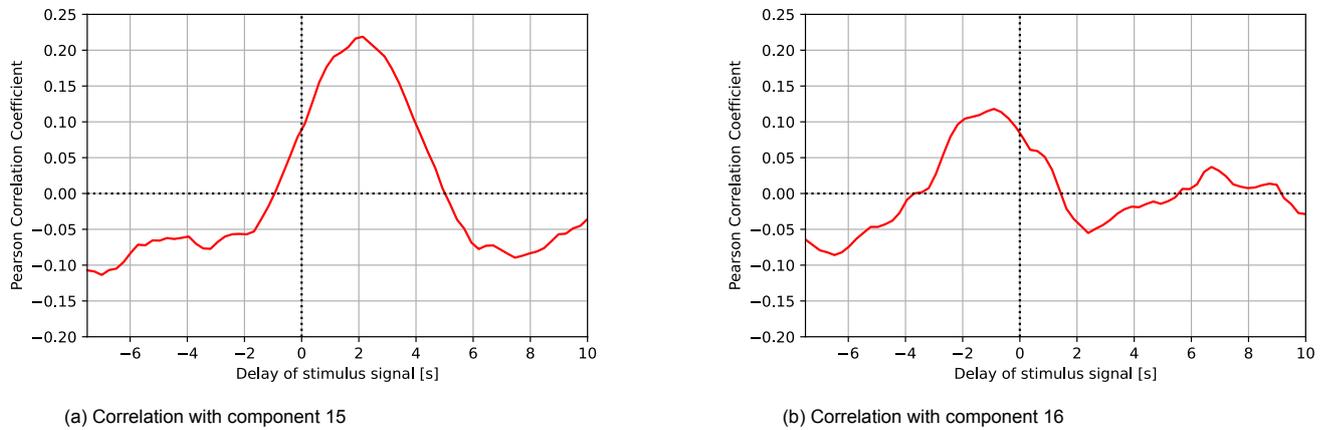


Figure C.3: Correlation between the delayed stimulus signal and the time course of mean components 15 and 16

D

Spatial maps members cluster 1, slice A

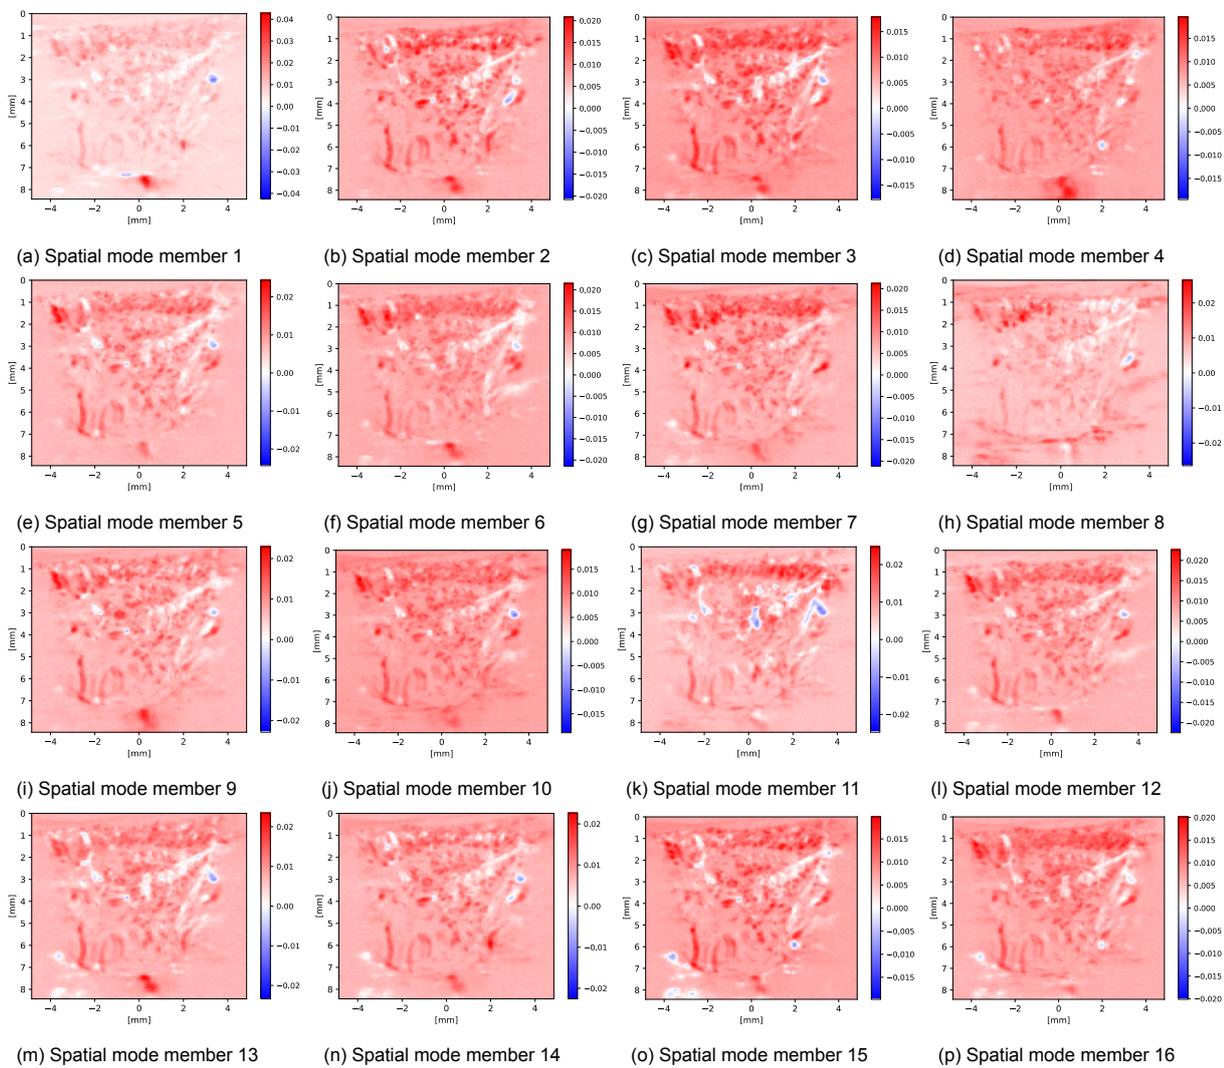


Figure continues on the next page

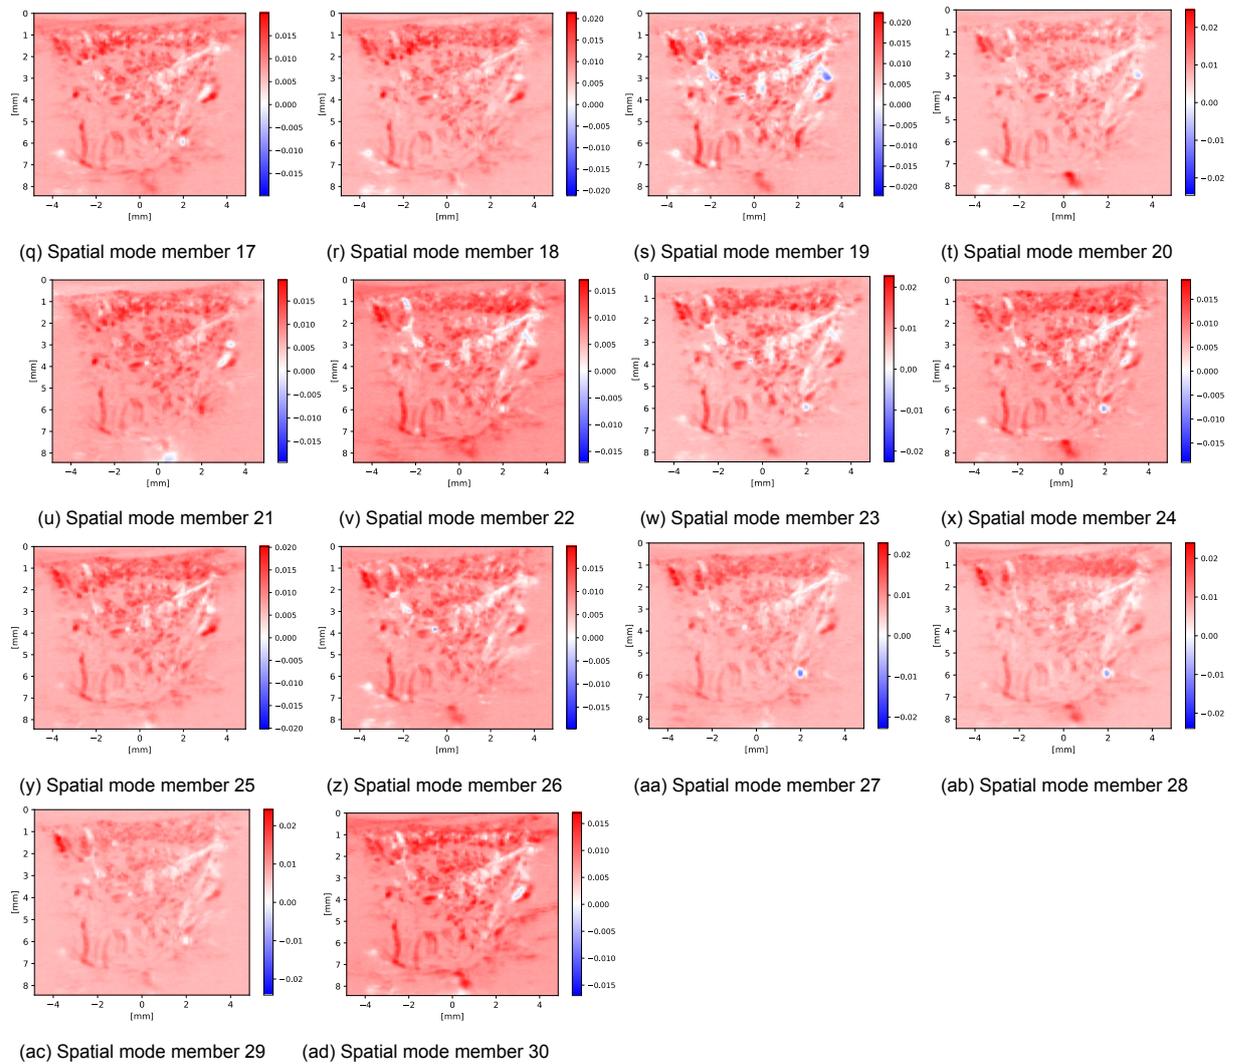
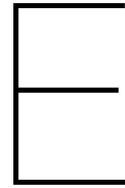


Figure D.1: The spatial modes of the 30 cluster members belonging to mean component 1



Masking of slice A

A mask containing (primarily) the pixels within the brain was created. This was done by a function that thresholded the activation level of the pixels. Subsequently, the edge of the activated pixels was taken, and extensive manual refinement was performed to create a continuous brain boundary. The function was developed by Sofia Kotti, also see [31]. In Figure E.1, both the masked and unmasked PDI are depicted.

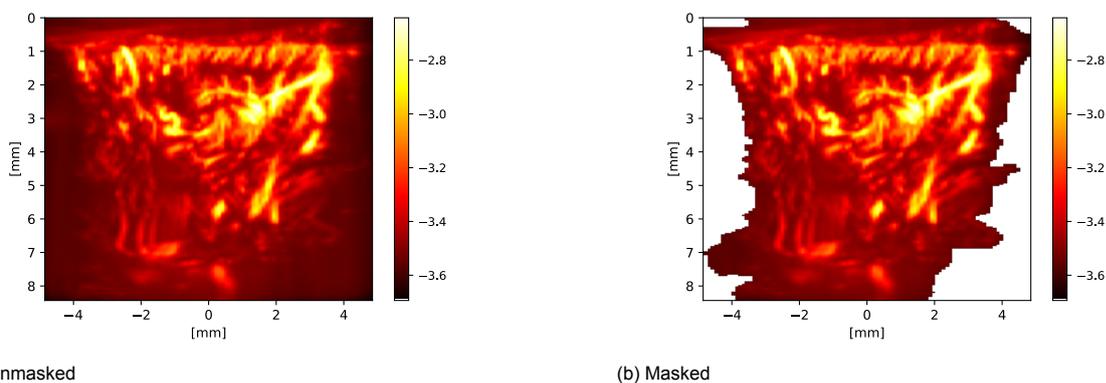


Figure E.1: The masked and unmasked version of the PDI of slice A. The log of the PDI was taken to better visualize activations in the background.

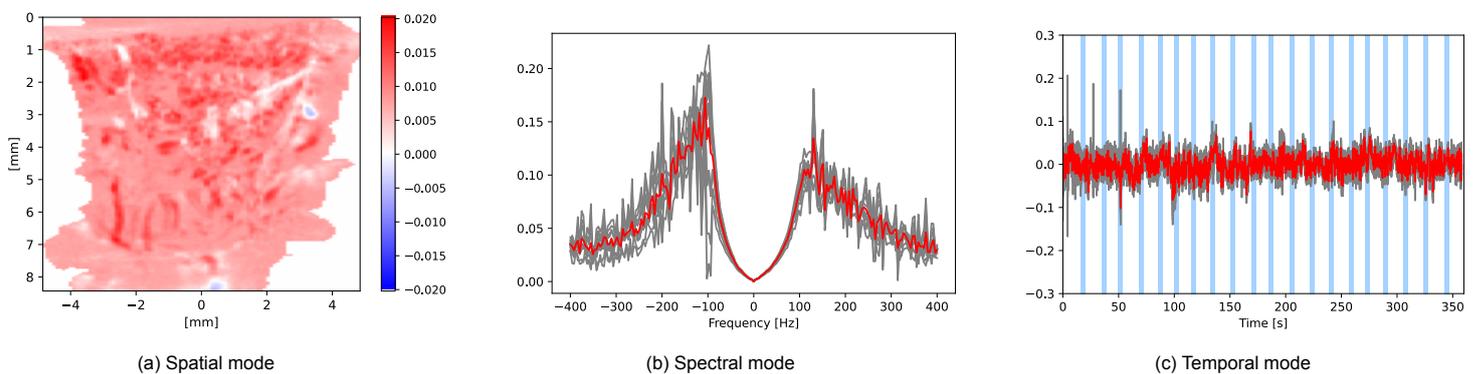


Figure E.2: Mean component 4 resulting from the masked version of slice A, which occurred 11 times. The spatial mode (a) is the mean spatial signature of the components belonging to this cluster. The grey is the plotted spectra (b) and time courses (c) of the components belonging to this cluster. The mean spectra and time courses are superimposed in red. The timepoints where the stimulus is 'on' are visualized as vertical blue lines.

In Figure E.2, a mean component following from the masked data similar to the one in Figure 3.8

is presented. In the spatial map, one can clearly see that the pixels outside of the mask have no activation (white). The whole-brain activation is present in both the masked and unmasked components. Therefore, this suggests that the whole-brain activation is not (entirely) artifactually driven by pixels outside of the brain, indicating potential physiological relevance.