Extracting hemodynamic activity with low-rank spatial signatures in functional ultrasound using tensor decompositions

Sofia-Eirini Kotti* and Borbála Hunyadi†
Signal Processing Systems, EEMCS, Delft University of Technology, Delft, The Netherlands
Email: *S.E.Kotti@tudelft.nl, †B.Hunyadi@tudelft.nl

***Abstract*—Functional ultrasound (fUS) is a neuroimaging modality that indirectly measures local neuronal activity by imaging cerebral blood volume fluctuations. However, accurately estimating neuronal activity from fUS measurements remains an open challenge. Hemodynamic changes are often modeled as the output of a system characterized by the hemodynamic response function (HRF), with neuronal activations as input. In this work, we propose a model for fUS measurements that assumes that hemodynamic activity has a low-rank spatial characterization. Starting from the tensor block term decomposition, we propose a method to estimate the spatial signatures, the HRF and the neuronal activation signals. This method is entirely data-driven and can be applied to entire fUS datasets. After an investigation using simulations, application to task experiment data of a mouse verified that activity that is spatially low rank and temporally correlated with the stimulus can be extracted in expected regions, which opens up the way to application on resting state data.**

***Index Terms*—Functional ultrasound, hemodynamic response function, tensor block term decomposition.**

I. Introduction

The continuing effort of humans to visualize and understand the brain's complex organization has led to significant discoveries in neuroscience and related fields [1], [2]. The ultimate goal is to provide a detailed view of the information processing architecture of cognitively engaged networks, in health or under various pathological conditions. A common objective in functional imaging is to characterize the activity in a particular brain region in terms of its interactions with other regions or with a behavioral state [3]. Hemodynamics-based imaging modalities, such as functional magnetic resonance imaging (fMRI) and functional ultrasound (fUS), measure brain activity in an indirect way through the neurovascular coupling [4]. Therefore, there is a need to invert this hemodynamic system and obtain an approximation of the neuronal signals. This need is motivated both neurobiologically, since brain interactions occur at the neuronal level, and mathematically, since modeling interactions at the hemodynamic level is not equivalent to modeling them neuronally [3].

The hemodynamic system is typically modeled as a linear time-invariant (LTI) system, described by the hemodynamic response function (HRF) [5]. There are two common assumptions in literature: a) a binary representation of the stimulus is used as a proxy for the neuronal signals, and b) the HRF is considered known a priori. These approaches can introduce bias in the results or cannot be applied to resting-state data.

In fMRI literature, multiple works investigate neuronal activations and the HRF. Total activation [6] estimates activity-inducing signals from noisy fMRI measurements, assuming a common, known HRF. The method uses sparse spatio-temporal priors that favor piecewise constant activity-inducing signals and coherent activation patterns in a priori atlas-defined brain regions. In [7], a low-rank matrix modeling of the fMRI signal is proposed: each extracted component consists of a temporal activation atom alongside an associated spatial map. These learned components correspond to distinctive functional networks in the brain. The temporal activation signals are assumed piecewise constant, whereas the HRF is fixed in the selected voxels to the canonical HRF model [8].

In [9], the authors propose an algorithm to jointly estimate the HRFs and the low-rank neuronal activity across the whole brain. The central assumption is that the observed fMRI signals result from convolving the weighted neuronal activation atoms with region-dependent HRFs. For this, a brain parcellation has to be provided. The region-specific HRF depends on a single dilation parameter δ_m as $h_{\delta_m}(t) = h_{\text{ref}}(\delta_m t)$, where $h_{\text{ref}}(t)$ is the canonical HRF. Thus, two distinct types of networks are defined: known disjoint anatomical regions based on an atlas and estimated overlapping functional networks.

In [10], the authors perform blind source separation utilizing tensor decompositions on multi-subject fMRI data tensors that have been unfolded into lower-order tensors. All considered tensor decompositions exploit a low rank constraint in the spatial domain, as motivated in [11].

Although the majority of research on HRF identification and the recovery of functional networks is focused on fMRI data, there have been recent works using fUS data, e.g., [4], [12], [13]. These methods, however, solve the problem for specific brain pixels, which are considered representative of regions of interest. In contrast, tensors and tensor decompositions naturally lend themselves to the representation and processing of original three-dimensional fUS data (time-varying 2D slices).

The contribution of this paper lies in applying tensor decompositions on experimental fUS measurements to uncover functional networks and subsequently performing deconvolution to estimate the neuronal activation signals and an HRF with a more flexible parametrization than in [9]. The proposed

This work is supported by the TU Delft AI Labs programme.

method follows a low-rank spatial model, similarly to [10], is entirely data-driven and, as such, does not require any a priori knowledge of the system.

II. Functional ultrasound system modeling

A. Signal model and assumptions

In this work, we assume that the hemodynamic system is an LTI system characterized by a single HRF, common everywhere in the imaged slice of the brain. The drivers of this system are the neuronal activation signals. These activation signals are analogous to the activation atoms in [7], [9] and the activity-inducing signals in [6], and, as such, they may correspond to activity under task performance or spontaneous activity. Following the approach outlined in these works, we assume these signals to be piecewise constant. Each activation signal has an associated spatial map containing the corresponding weights for all pixels. These spatial maps are two-dimensional, contrary to [7], [9], and are assumed to be of low rank, consistent with the assumptions in [10] and the multisubject dictionary learning probabilistic atlas of [11]. The fUS measurements (outputs of the hemodynamic system) capture the total activity per pixel owing to all components.

We consider an fUS data tensor $\mathcal{Y} \in \mathbb{R}^{N_z \times N_x \times N_t}$, where N_z is the number of pixels in the depth dimension, N_x is the number of pixels in the width dimension, and N_t is the number of time stamps. The 2D spatial maps are denoted $\mathbf{U}_r$, $r = 1, \ldots, R$, each with rank L_r. Thus, each can be rewritten as $\mathbf{U}_r = \mathbf{A}_r \mathbf{B}_r^{\mathrm{T}}$, where the matrices $\mathbf{A}_r \in \mathbb{R}^{N_x \times L_r}$ and $\mathbf{B}_r \in \mathbb{R}^{N_y \times L_r}$ are of full column rank L_r. With these assumptions, the signal model in Fig. 1 becomes

$$\mathcal{Y} = \sum_{r=1}^{R} \mathbf{U}_r \circ (\mathbf{h} * \mathbf{z}_r) = \sum_{r=1}^{R} \left(\mathbf{A}_r \mathbf{B}_r^{\mathrm{T}}\right) \circ (\mathbf{h} * \mathbf{z}_r), \quad (1)$$

where $\circ$ indicates the outer product operation, $*$ indicates the convolution operation, $\mathbf{h}$ is the HRF and $\mathbf{z}_r$ are the activation signals.

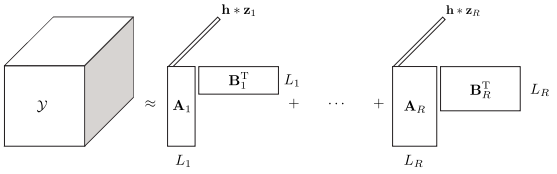

Fig. 1: Proposed fUS system model.

Furthermore, we assume that the HRF is parametrizable according to the following model [4], [12], [13]

$$h(t; \boldsymbol{\theta}) = \theta_1 (\Gamma(\theta_2))^{-1} \theta_3^{\theta_2} t^{\theta_2 - 1} \mathrm{e}^{-\theta_3 t}, \quad (2)$$

where $\boldsymbol{\theta} = [\theta_1,\ \theta_2,\ \theta_3]^{\mathrm{T}}$ and $\theta_1,\ \theta_2,\ \theta_3 > 0$. Vector $\mathbf{h}$ in (1) contains the samples of (2) at the given time stamps. This model is more flexible than the one used in [9] for fMRI and can capture large HRF variability.

B. Tensor block term decomposition

The model in (1) follows the form of a tensor $(L_r, L_r, 1)$-block term decomposition (BTD) [14], which writes a tensor $\mathcal{Y} \in \mathbb{R}^{N_x \times N_y \times N_t}$ as a sum of multilinear rank-$(L_r, L_r, 1)$ terms, $1 \leq r \leq R$, in the form

$$\mathcal{Y} = \sum_{r=1}^{R} \left(\mathbf{A}_r \mathbf{B}_r^{\mathrm{T}}\right) \circ \mathbf{c}_r, \quad (3)$$

where $\mathbf{A}_r$ and $\mathbf{B}_r$ are of full column rank L_r. In the sequel, the terms BTD and $(L_r, L_r, 1)$-BTD will be used interchangeably.

C. BTD-based fUS deconvolution

It is now easy to see that the model in (1) is that of an $(L_r, L_r, 1)$-BTD, whose time signatures are constrained to be the result of a convolution of $\mathbf{h}$ and $\mathbf{z}_r$. The question now is: how can we solve for this constrained BTD?

Directly estimating all quantities in (1) by minimizing a data-fitting error is a complex problem: the objective function would be jointly nonconvex in the large number of variables, but also nonconvex in the variable $\boldsymbol{\theta}$, when all other quantities are fixed. However, there is an alternative approach: if the BTD is (essentially) unique [14], it is possible to first calculate the $(L_r, L_r, 1)$-BTD of (3), potentially with a few simple constraints, and subsequently deconvolve the time signatures as $\mathbf{c}_r = \mathbf{h} * \mathbf{z}_r$, given the constraints on $\mathbf{h}$ and $\mathbf{z}_r$. If the BTD is unique, the interpretability of the components is also ensured [15], [16]. With this in mind, we propose a two-step strategy:

Step 1: Solve for the BTD spatial and temporal signatures, given the number of components R and the spatial ranks L_r

$$\begin{aligned} \min_{\mathbf{A}_r, \mathbf{B}_r, \mathbf{c}_r} \quad & \frac{1}{2} \left\| \mathcal{Y} - \sum_{r=1}^{R} \left(\mathbf{A}_r \mathbf{B}_r^{\mathrm{T}}\right) \circ \mathbf{c}_r \right\|_F^2 \\ \text{subject to} \quad & \|\mathbf{c}_r\|_2 = 1 \quad \text{and} \quad \mathbf{A}_r \geq 0, \mathbf{B}_r \geq 0. \end{aligned} \quad (4)$$

Step 2: Jointly estimate the HRF and the neuronal activation signals by solving

$$\begin{aligned} \min_{\boldsymbol{\theta}, \mathbf{z}_r} \quad & \frac{1}{2} \left\| \mathbf{C} - \mathbf{H}\mathbf{Z}^{\mathrm{T}} \right\|_{\mathrm{F}}^2 + \lambda \sum_{r=1}^{R} \|\mathbf{D}\mathbf{z}_r\|_1 \\ \text{subject to} \quad & \boldsymbol{\theta} = [\theta_1,\ \theta_2,\ \theta_3]^{\mathrm{T}} \\ & h(t; \boldsymbol{\theta}) = \theta_1 (\Gamma(\theta_2))^{-1} \theta_3^{\theta_2} t^{\theta_2 - 1} \mathrm{e}^{-\theta_3 t} \\ & 0.5 \leq \theta_1 \leq 5,\ \ 1 \leq \theta_2 \leq 6,\ \ 1 \leq \theta_3 \leq 10 \end{aligned} \quad (5)$$

where $\mathbf{C} = [\mathbf{c}_1\ \ldots\ \mathbf{c}_R]$, $\mathbf{H} = \text{Toeplitz}(\mathbf{h})$, $\mathbf{Z} = [\mathbf{z}_1\ \ldots\ \mathbf{z}_R]$ and $\mathbf{D}$ is the first-order difference operator.

A white Gaussian noise hypothesis leads to the least squares optimality criterion. The constraint $\|\mathbf{c}_r\|_2 = 1$ is added to deal with scale ambiguities. Constraining matrices $\mathbf{A}_r$ and $\mathbf{B}_r$ to be elementwise nonnegative ensures that the spatial maps $\mathbf{U}_r = \mathbf{A}_r \mathbf{B}_r^{\mathrm{T}}$ will also be nonnegative, as in [7], [9]. Piecewise constant activation signals are promoted through the total variation regularization term in (5). Finally, the parameters in $\boldsymbol{\theta}$ are constrained to lie in specific ranges to reduce potential indeterminacies and facilitate the optimization. These ranges are wide enough to cover the expected HRF shapes. The reader

is referred to [13] for the effect of $\boldsymbol{\theta}$ on the HRF shape. It should be noted that time-shift ambiguities in the estimation of $\mathbf{h}$ and $\mathbf{z}_r$ are unavoidable due to the convolution, and so are scaling ambiguities. We solve problems (4) and (5) using the structured data fusion (SDF) framework in Tensorlab [17].

III. Numerical Experiments

A. Simulation setup

For the numerical experiments, we defined a scenario with $R = 2$ components: one for a task-related source and one for spontaneous neuronal fluctuations. The former consisted of 2 nonzero blocks of random durations and amplitudes. A series of 8 Dirac impulses with random signs and amplitudes is generated for the latter. Their timings are drawn uniformly, and their intensity follows a Gaussian distribution $\mathcal{N}(0, 0.25)$. Integrating this signal over time produced the second activation signal. The temporal signals are generated with a sampling frequency of 8 Hz and a total length of 240 samples.

Subsequently, we defined two corresponding 2D spatial maps of size 20×20 pixels. Similarly to [9], each spatial map has a single activated squared region of size 5×5 pixels. Each active pixel has a randomly drawn weight between 1 and 100, and each map is normalized by its ℓ_1-norm. The signals resulting from the outer product of the activations with their spatial maps were then added together and convolved with a randomly defined HRF, with $\boldsymbol{\theta}$ in the ranges given in (5). We added white Gaussian noise to produce the observed signals at multiple signal-to-noise ratio (SNR) levels, as a final step.

The above procedure was repeated 35 times with randomly generated spatial maps, activation signals, and HRF. For solving (4), we assumed that the values of R and L are known. The regularization parameter was set using the `relweight` parameter of Tensorlab. A grid search approach led to choosing the `relweight` values of 30 and 1 for the two terms in (5), respectively. Each of the BTD and the deconvolution steps were run with 24 initializations. The initialization with the smallest reconstruction error was chosen for the final result, which aligns with the Gaussian noise assumption. The code for the simulations is available on GitHub[1].

B. Simulation results

The estimation results for the spatial maps, the HRF, and the activation signals are shown in the boxplots of Fig. 2(a)-(c), respectively, as obtained by solving (4) and (5). The black line indicates the median value. Regarding the spatial maps, the median relative error is generally low; even at 0 dB SNR, it is below 0.2. To compare the true and estimated HRF and activation signals point-by-point, the estimated HRF was shifted in time so as to have the maximum correlation with the true HRF. The estimated activation signals were shifted accordingly. Additionally, both were normalized to the same maximum amplitude as the true quantities, since the algorithm results suffer from scaling ambiguities. The median HRF error is low, around 0.2 even at -5 dB SNR. The error in the activation signals is larger, which was also the case in [7].

[1] https://github.com/sofigami/btd-based-fus-deconvolution

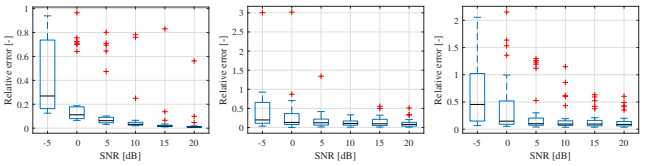

(a) Spatial maps. (b) HRF. (c) Activation signals.

Fig. 2: Boxplot of ℓ_2-norm relative error for different estimated quantities, defined as $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$ for quantity $\mathbf{x}$ and its estimate $\hat{\mathbf{x}}$, at different SNR levels, averaged over the $R = 2$ components and over the 35 runs.

It should be noted that all simulation results were included without removing cases where the BTD step did not converge to the global minimum; e.g., sometimes the resulting spatial maps were combinations of the true spatial maps. This explains the outliers and the large variance, mainly in the activation signals and the spatial maps at low SNR levels.

Example results for a single scenario in Fig. 3 show that, even at low SNR, the spatial maps are estimated very well. The estimation of the HRF and the activation signals in Fig. 4 is almost perfect at 20 dB SNR, whereas it is not sufficient at -5 dB (results are shown for Comp#1, but were similar for Comp#2).

The most challenging scenarios for the BTD are high-noise scenarios, especially when the activating regions overlap in space or are partly aligned. This, in turn, affects the estimation of the HRF and the activation signals. The deconvolution step, which is not convex in the $\boldsymbol{\theta}$ variable, adds further complexity. Overall, the BTD step of the proposed algorithm seems more robust than the deconvolution step.

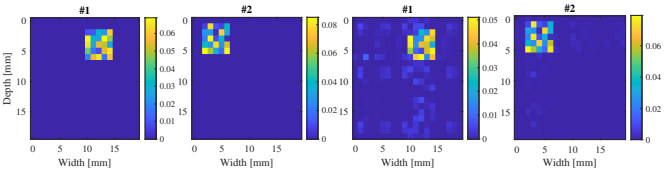

(a) True spatial signatures. (b) Estimated spatial signatures.

Fig. 3: Example BTD results at -5 dB SNR.

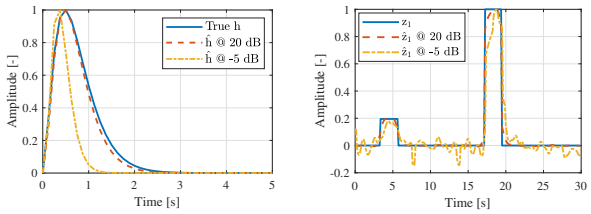

(a) True and estimated HRF. (b) True and estimated signal $\mathbf{z}_1$.

Fig. 4: Example deconvolution results at 20 dB and -5 dB SNR.

IV. Real Experiments

A. Data acquisition and preprocessing

Two visual experiments (on two brain slices of a single mouse) were conducted at the Center for Ultrasound and Brain Imaging at Erasmus University Medical Center (CUBE). The subject mouse was presented with randomly generated high-contrast images on two screens simultaneously, in 10 blocks of 4 seconds each. Each repetition was followed by a rest period

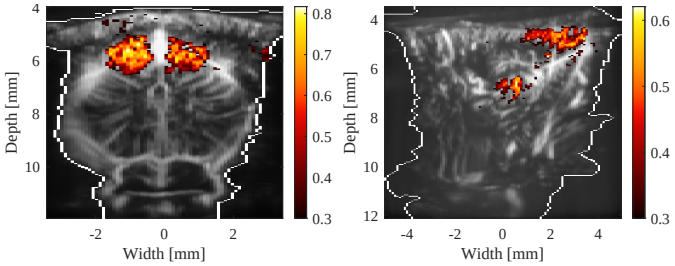

(a) Exp. 1. (b) Exp. 2.

Fig. 5: Correlation images for the two experiments.

lasting between 10 and 15 seconds. For the data acquisition procedure details, the reader is referred to [13].

The time series of a pixel corresponds to its power variation over the power Doppler image (PDI) stream [13]. The PDI rate was 4 Hertz. Outlier PDIs were removed and interpolated over, as in [18]. In Exp. 1, part of the SC was imaged, whereas in Exp. 2, parts of the LGN and V1, regions known to be involved in the processing of visual stimuli. The obtained correlation images can be found in Fig. 5, where pixels with a Pearson correlation coefficient (PCC) larger than 0.3 with the stimulus are indicated. The final step of data preprocessing included standardizing each pixel time series and band-pass filtering using a fifth-order Butterworth filter with a 0.01-0.3 Hertz passband in order to remove high-frequency noise components [19] and the slow baseline drift in the signals.

B. Implementation details

As discussed in Section II-C, the BTD has to be unique for the two-step method to be applicable. Solving (4) for the fUS data delivered factor matrices $\mathbf{A} = [\mathbf{A}_1 \ \ldots \ \mathbf{A}_R]$ and $\mathbf{B} = [\mathbf{B}_1 \ \ldots \ \mathbf{B}_R]$ that were full rank, but, unfortunately, matrix $\mathbf{C} = [\mathbf{c}_1 \ \ldots \ \mathbf{c}_R]$ had proportional columns; this means that the decomposition was not unique, as per the conditions in [14]. For this reason, in order to promote the uniqueness of the BTD, we imposed an orthogonality constraint on C, which can be easily added to Tensorlab. The result satisfied the uniqueness conditions of [14]. The orthogonality requirement means that the components must differ in their reaction to stimuli to be distinguishable from each other, as required in [13]; otherwise, they are extracted under a single component. Additionally, since the PDIs include pixels that do not belong to the brain (the rough brain limits in Fig. 5 were estimated from the data), the algorithm should be free to fit any value to those pixels of no interest. This was possible using `NaN` values and the `fmt` option of Tensorlab. Furthermore, we ran the algorithm with different values of R. Since we have no information on the spatial ranks, we assume that $L_r = L$ for all r, similarly to [10]. We tried 24 random initializations for each step. The best set of values for `relweight` was chosen using the stimulus timings.

C. Results

The spatial and temporal signatures for Exp. 1 when $R = 2$ and $L = 2, 5$ are shown in Fig. 6. We see that the spatial signature of component 2 in both cases resembles the correlation image of Fig. 5 (a). The obtained time signatures were very

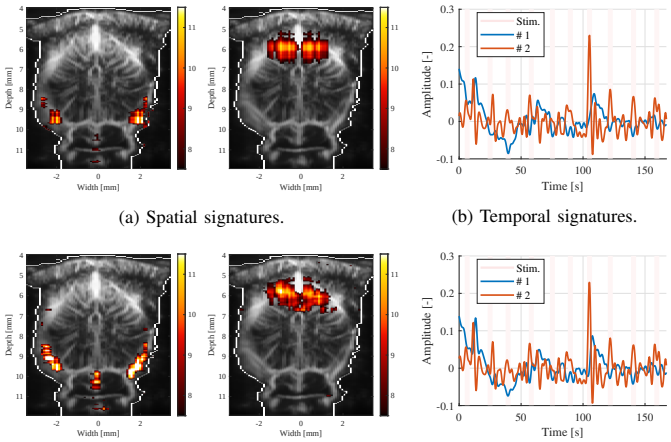

(a) Spatial signatures. (b) Temporal signatures.

(c) Spatial signatures. (d) Temporal signatures.

Fig. 6: Results for Exp. 1 when $R = 2$ and $L = 2$ (a), (b), $L = 5$ (c), (d).

similar for both values of L, and they had a PCC of 0.46 and 0.40 with the stimulus, respectively. It is clear that a higher spatial rank allows for more detailed shapes in the spatial maps. However, the overall shape of the SC is sufficiently captured for $L = 2$; therefore, we will only present results for $L = 2$ in the following.

When the number of components increases to 5 (Fig. 7), the SC is extracted at a much higher contrast (Comp#3). However, interpreting the rest of the components is challenging. Comp#5 mainly includes pixels on the big vessels near the brain limit. Comp#2 probably relates to an edge effect of the algorithm: its highest weights are for pixels outside the brain (not depicted), and most of its temporal signature samples are roughly 0. Comp#1 possibly captures meaningful information. For Exp. 2, two of the five extracted components are presented in Fig. 8. Comp#2 includes the LGN and V1 of Fig. 5 (b). Importantly, the BTD extracts the M1 region (reported in [20]) as well (Comp#3), which was not visible in Fig. 5 (b).

The deconvolution results for Exp. 1 can be found in Fig. 9. Comp#2 was not included in this step since it contained virtually no temporal variation. The activation signal for Comp#2 now clearly points towards the experimental paradigm timings. The reason the signal appears slightly ahead of the stimulus in time is the time-shift ambiguity shared with the HRF.

Some indicative results for various R are included in Table I. For Exp. 1, the maximum PCC with the stimulus significantly increases when the number of components goes from 2 to 5. No improvement is yielded as R increases further. In Exp. 2, additional components improve the extraction of the task-related source. In both experiments, increasing R from 2 to 20 slightly improves the BTD relative error. Note that the observed relative errors ($\sim$0.8) are comparable to those observed in the simulations for -5 dB SNR.

These results show that the component extraction heavily depends on the data: for Exp. 1, the task-related component can already be extracted with a high correlation for $R = 5$. This is not the case for $R = 2$, where we essentially divide components into two types, thus allowing part of the artifacts to mix into the task-related component. This is also clear from

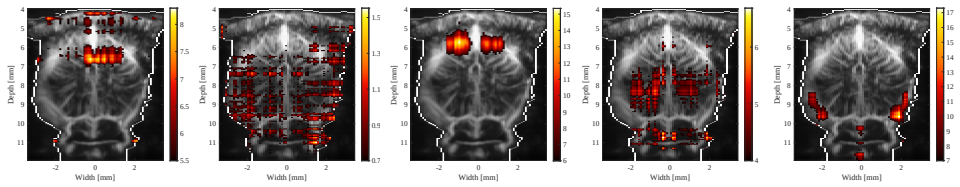

Fig. 7: Spatial maps for Exp. 1 when $R = 5$, $L = 2$.

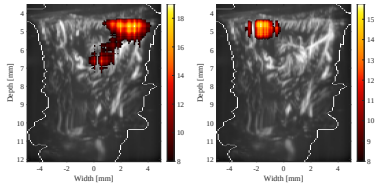

Fig. 8: Two spatial maps for Exp. 2 when $R = 5$, $L = 2$.

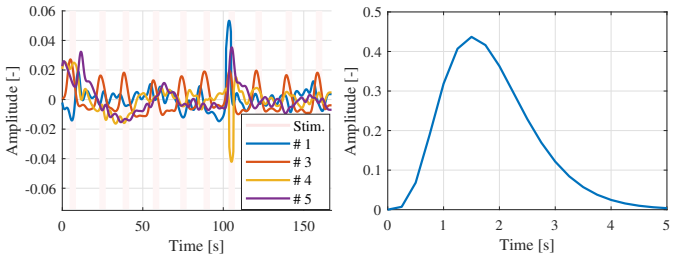

(a) Stimulus and estimated activ. signals. (b) Estimated HRF.

Fig. 9: Deconvolution results for Exp. 1.

the correlation results: the lag for the maximum correlation was not the same as the one used in Fig. 5 (a). The extraction of the task-related component was straightforward in Exp. 2 already for $R = 2$, albeit at a lower correlation, which was expected from the values in Fig. 5.

TABLE I: Results for both experiments and for $L = 2$.

Number of BTD components R	**Max. PCC**		**Relative error**	
	Exp. 1	Exp. 2	Exp. 1	Exp. 2
2	0.4594	0.5322	0.8886	0.8642
5	0.8667	0.5636	0.8600	0.7904
10	0.8678	0.6557	0.8029	0.7718
20	0.8711	0.7006	0.7866	0.7510

V. Conclusions

In this work, we extracted brain activity with low-rank spatial signatures in a two-step approach: using the BTD to estimate the spatial maps and semi-blind deconvolution to estimate the activation signals and the HRF in the imaged slice. The main advantages of this approach are that it is data-driven, acts on whole-slice data, and can be applied to resting-state data. As noted in [10], the BTD only relies on two assumptions: low spatial ranks and additive white Gaussian noise. In this real data implementation, we added the constraint of orthogonal time signatures to promote the uniqueness of the BTD, but other constraints can be used instead. The method does not come without consideration: the number of components and the ranks are generally unknown, so multiple values must be tested. Furthermore, the regularization parameter has to be finetuned. Finally, the extent to which the HRF can be considered constant in the slice remains an open question that has received diverse answers in literature [19], [21].

Acknowledgment

The authors would like to thank CUBE for providing the fUS data (www.ultrasoundbrainimaging.com).

References

[1] M. E. Raichle, "Functional brain imaging and human brain function," *Journal of Neuroscience*, vol. 23, no. 10, 2003.

[2] K. J. Friston, "Modalities, modes, and models in functional neuroimaging," *Science*, vol. 326, no. 5951, pp. 399–403, 2009.

[3] D. R. Gitelman, W. D. Penny, J. Ashburner, and K. J. Friston, "Modeling regional and psychophysiologic interactions in fMRI: the importance of hemodynamic deconvolution," *Neuroimage*, vol. 19, no. 1, 2003.

[4] A.-K. Aydin *et al.*, "Transfer functions linking neural calcium to single voxel functional ultrasound signal," *Nat. Commun.*, vol. 11, 06 2020.

[5] M. A. Lindquist, J. Meng Loh, L. Y. Atlas, and T. D. Wager, "Modeling the hemodynamic response function in fMRI: Efficiency, bias and mis-modeling," *NeuroImage*, vol. 45, no. 1, Supplement 1, 2009.

[6] F. I. Karahanoğlu, C. Caballero-Gaudes, F. Lazeyras, and D. Van De Ville, "Total activation: fMRI deconvolution through spatio-temporal regularization," *Neuroimage*, vol. 73, 2013.

[7] H. Cherkaoui, T. Moreau, A. Halimi, and P. Ciuciu, "fMRI BOLD signal decomposition using a multivariate low-rank model," in *2019 27th European Signal Processing Conference (EUSIPCO)*, IEEE, 2019.

[8] K. Friston *et al.*, "Event-related fMRI: Characterizing differential responses," *NeuroImage*, vol. 7, no. 1, pp. 30–40, 1998.

[9] H. Cherkaoui, T. Moreau, A. Halimi, C. Leroy, and P. Ciuciu, "Multivariate semi-blind deconvolution of fMRI time series," *NeuroImage*, vol. 241, p. 118418, 2021.

[10] C. Chatzichristos, E. Kofidis, M. Morante, and S. Theodoridis, "Blind fMRI source unmixing via higher-order tensor decompositions," *Journal of Neuroscience Methods*, vol. 315, 2019.

[11] G. Varoquaux, A. Gramfort, F. Pedregosa, V. Michel, and B. Thirion, "Multi-subject dictionary learning to segment an atlas of brain spontaneous activity," in *Information Processing in Medical Imaging: 22nd International Conference. Proceedings 22*, Springer, 2011.

[12] S.-E. Kotti, A. Erol, and B. Hunyadi, "Modeling nonlinear evoked hemodynamic responses in functional ultrasound," in *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, IEEE, 2023.

[13] A. Erol *et al.*, "Deconvolution of the functional ultrasound response in the mouse visual pathway using block-term decomposition," *Neuroinformatics*, vol. 21, no. 2, 2023.

[14] L. De Lathauwer, "Decompositions of a higher-order tensor in block terms—part II: Definitions and uniqueness," *SIAM Journal on Matrix Analysis and Applications*, vol. 30, no. 3, 2008.

[15] B. Hunyadi, P. Dupont, W. Van Paesschen, and S. Van Huffel, "Tensor decompositions and data fusion in epileptic electroencephalography and functional magnetic resonance imaging data," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 7, no. 1, 2017.

[16] T. Adali and V. D. Calhoun, "Reproducibility and replicability in neuroimaging data analysis," *Curr. Opin. Neurol.*, vol. 35, no. 4, 2022.

[17] N. Vervliet, O. Debals, L. Sorber, M. Van Barel, and L. De Lathauwer, "Tensorlab 3.0 [Online]. Available: https://www.tensorlab.net/," 2016.

[18] C. Brunner *et al.*, "Whole-brain functional ultrasound imaging in awake head-fixed mice," *Nature Protocols*, vol. 16, no. 7, 2021.

[19] A. O. Nunez-Elizalde *et al.*, "Neural correlates of blood flow measured by ultrasound," *Neuron*, vol. 110, no. 10, 2022.

[20] A. Erol, B. Generowicz, P. Kruizinga, and B. Hunyadi, "Evoked component analysis (ECA): Decomposing the functional ultrasound signal with GLM-regularization," *IEEE Transactions on Biomedical Engineering*, early access, 2024.

[21] D. A. Handwerker, J. M. Ollinger, and M. D'Esposito, "Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses," *NeuroImage*, vol. 21, no. 4, 2004.