Evoked Component Analysis (ECA)

Decomposing the Functional Ultrasound Signal with GLM-Regularization

Erol, Aybuke; Generowicz, Bastian; Kruizinga, Pieter; Hunyadi, Borbala

DOI
10.1109/TBME.2024.3395154

Publication date
2024

Document Version
Final published version

Published in
IEEE Transactions on Biomedical Engineering

Evoked Component Analysis (ECA): Decomposing the Functional Ultrasound Signal With GLM-Regularization

Aybüke Erol, *Student Member, IEEE*, Bastian Generowicz, Pieter Kruizinga, and Borbála Hunyadi, *Senior Member, IEEE*

***Abstract*—Analysis of functional neuroimaging data aims to unveil spatial and temporal patterns of interest. Existing analysis methods fall into two categories: fully data-driven approaches and those reliant on prior information, e.g. the stimulus time course. While using the stimulus signal directly can help identify the activated brain areas, it is known that the relationship between stimuli and the brain's response exhibits nonlinear and time-varying characteristics. As such, relying completely on the stimulus signal to describe the brain's temporal response leads to a restricted interpretation of the brain function. In this paper, we present a new technique called Evoked Component Analysis (ECA), which leverages prior information up to a defined extent. This is achieved by including the general linear model (GLM) design matrix as a regulatory term and estimating the factor matrices in both space and time through an alternating minimization approach. We apply ECA to 2D and swept-3D functional ultrasound (fUS) experiments conducted with mice. When decomposing 2D fUS data, we employ GLM regularization at various intensities to emphasize the role of prior information. Furthermore, we show that incorporating multiple hemodynamic response functions within the design matrix can provide valuable insights into region-specific characteristics of evoked activity. Finally, we use ECA to analyze swept-3D fUS data recorded from five mice engaged in two distinct visual tasks. Swept-3D fUS images the 3D brain sequentially using a moving probe, resulting in different slice acquisition time instants. We show that ECA can estimate factor matrices with a fine resolution at each slice acquisition time instant and yield higher t-statistics compared to GLM and correlation analysis for all subjects.**

***Index Terms*—Functional ultrasound, general linear model, regularized low-rank factorization.**

I. Introduction

FUNCTIONAL ultrasound (fUS) is a neuroimaging modality that uses plane-wave irradiation to generate 2D images of changing local blood dynamics [1]. Owing to its low-cost and portable design, as well as its impressive spatio-temporal resolution, fUS has been used for numerous neuroimaging applications in the last decade [2], [3]. Similar to functional magnetic resonance imaging (fMRI), fUS records hemodynamic changes induced by the local alterations in neuronal activity. In other words, the fUS signal reflects neuronal activity indirectly through the impulse response of the underlying neurovascular system, known as the hemoydnamic response function (HRF) [4].

For studying the relation between brain's hemodynamic response and external stimuli, two methods come forward in the literature. First one of these methods is correlation analysis, where temporal synchrony between the stimulus signal and voxel time series is explored by computing the Pearson Correlation Coefficient (PCC) per voxel. The second method is the general linear model (GLM). GLM can be considered as an extension of correlation analysis which allows to define multiple design variables for modelling the observed neuroimaging data in a multivariate regression problem [5]. Typically, a design variable is calculated as the convolution between a stimulus signal and an HRF, reflecting the expected response pattern when a brain voxel is activated by the stimulus. GLM can incorporate multiple stimuli and HRF shapes [6] as additional columns in the design matrix. While the former enables the extraction of distinct spatial activation maps per stimulus, the latter models HRF variations across voxels.

Both of these methods are quite intuitive and can provide valuable insight for identification of activated brain areas. However, they both rely on an overly-simplified representation of the stimulus signal. For example, it is known that the brain can habituate to repeated stimuli, resulting in weakened responses over time – a phenomenon known as Repetition Suppression (RS) [7]. These epoch-dependent variations would cause lower correlation values and lower GLM coefficients to be estimated in responsive brain areas, as the overall match between the

Manuscript received 26 October 2023; revised 4 March 2024; accepted 22 April 2024. Date of publication 30 April 2024; date of current version 26 August 2024. This work was supported by the Synergy Grant of Department of Microelectronics of Delft University of Technology and the Delft Technology Fellowship. *(Corresponding author: Borbála Hunyadi.)*

Aybüke Erol is with the Signal Processing Systems (SPS) Group, Department of Microelectronics, Delft University of Technology, The Netherlands.

Bastian Generowicz and Pieter Kruizinga are with the Center for Ultrasound and Brain Imaging at Erasmus MC (CUBE), Department of Neuroscience, Erasmus University Medical Center, The Netherlands.

Borbála Hunyadi is with the Signal Processing Systems (SPS) Group, Department of Microelectronics, Delft University of Technology Delft 2628 CD, The Netherlands (e-mail: b.hunyadi@tudelft.nl).

This article has supplementary downloadable material available at https://doi.org/10.1109/TBME.2024.3395154, provided by the authors.

Digital Object Identifier 10.1109/TBME.2024.3395154

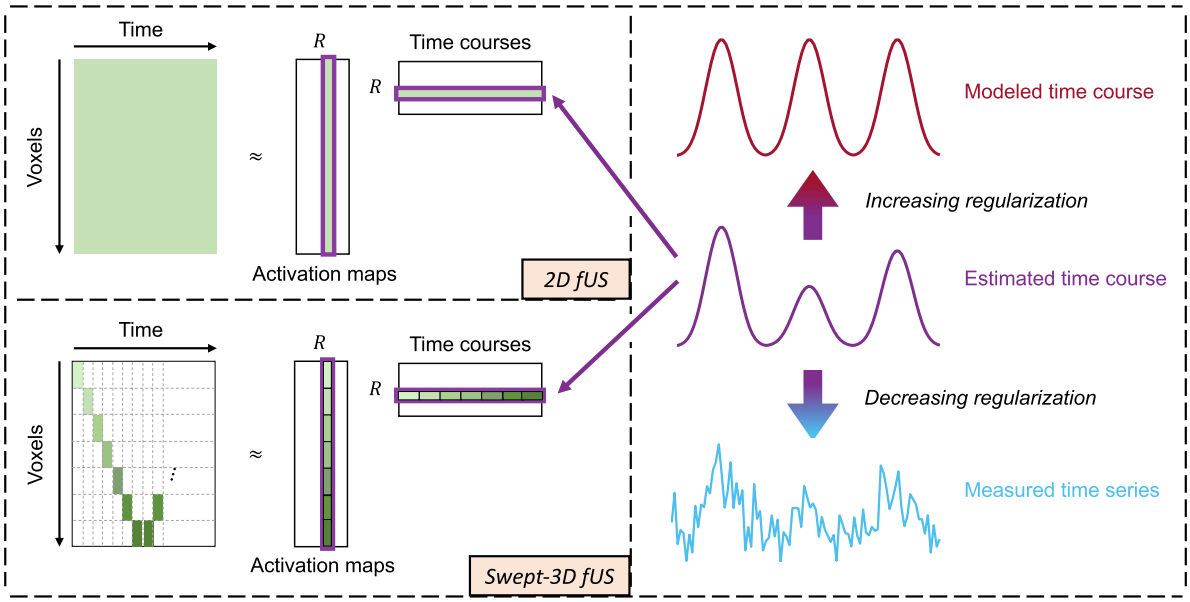

Fig. 1. Illustration of ECA for decomposing neuroimaging data into its temporal and spatial factor matrices of rank R. Modeled time courses are the design variables, obtained by convolving the known stimulus signal(s) with an HRF. By regularizing the temporal factor matrix to the design matrix, ECA achieves a non-restricting estimation of the brain's evoked response. In case of swept-3D fUS, where slice timings are different as a result of probe motion, ECA can recover the evoked response at the whole resolution without necessitating interpolation to a common time axis.

corresponding voxel time series and the stimulus signal will be reduced [8]. To understand why low activations are observed would require a subsequent step of choosing ROIs (either by anatomical labelling or by thresholding the spatial activation maps obtained by either method) and investigating their time courses thoroughly.

On the other hand, there are also fully data-driven approaches, such as independent component analysis (ICA), which do not make use of any prior information of the stimulus, and returns spatio-temporal signatures for each independent component [9]. Using a data-driven approach is crucial when no prior knowledge is available on the expected activation time courses, such as in case of resting-state experiments [10], or in presurgical electroencephalogram (EEG)-fMRI studies in epilepsy, when no epileptic spikes are observed in the EEG [11]. However, interpreting the large number of resulting independent components requires time-consuming manual classification [12] or subsequent machine learning in order to identify artifact sources and components of interest [13], [14].

Overall, model-based approaches offer an undeniable ease in detection and interpretation of evoked activities of interest, but fall short in capturing sources of variability, whereas the opposite holds for data-driven methods. In this work, we propose a novel method called as Evoked Component Analysis (ECA) to balance between the two extremities by utilizing prior information only as a guiding factor. We achieve this by incorporating the GLM design matrix as a regulatory term in a low-rank decomposition framework instead of enforcing it. This flexibility allows for extracting time courses that follow the stimulus onsets but still get nourished from the measured data. An illustration of ECA is provided in Fig. 1.

We apply ECA on mouse-based 2D and 3D fUS experiments. In 2D fUS data, we investigate the impact of algorithmic parameters, namely the rank and regularization coefficient. For 3D ultrasound acquisition, we employ a new technique known as swept-3D fUS [15], [16], where a full volumetric scan of the brain is obtained by moving the ultrasound probe across the exposed brain. In other words, swept-3D fUS does not capture the whole brain at once but instead images it sequentially – resulting in an incomplete data array. We demonstrate how ECA can also be used to decompose such 3D neuroimaging data at the full resolution without relying on slice timing correction techniques. In our previous study [17], we presented GLM-regularized low-rank factorization only in the setting of swept-3D fUS. Within this work, we generalize our problem definition to extracting evoked activity both in 2D and 3D neuroimaging data. We explore the effect of GLM regularization thoroughly at different intensities and assess the performance of ECA at higher ranks, various stimulus conditions and multiple subjects.

The rest of this paper is organized as follows. We first explain our problem formulation, accompanied by the proposed solution. We then describe the fUS data acquisition pipeline for both standard 2D-fUS and swept-3D fUS. Next, we present our results on in-vivo 2D and swept-3D fUS experiments. We compare our findings with correlation and GLM analysis. Finally, we conclude our paper with discussion and possible future extensions.

II. Problem Formulation

The objective of this work is to estimate evoked brain activity both in space and time (i.e. its anatomical location and the activity time course), considering that the stimuli inducing this activity are known. However, stimuli might not always translate to brain activity as expected. Indeed, it is known that the relationship between stimuli and the brain's response is quite complex and can exhibit nonlinear and time-varying traits [18],

[19]. Hence, we aim to devise a novel decomposition technique that exploits prior information while not completely relying on it.

Matrix decompositions are used for a variety of purposes in neuroimaging, including denoising [20], compression [21] and demixing [22]. In general, decomposition of neuroimaging data reveals underlying patterns of interest that may not be immediately apparent in its raw form [23]. The goal of matrix decompositions is to express the input data matrix $\mathbf{Y}$ as a product of factor matrices $\mathbf{U}$ and $\mathbf{V}$:

$$\mathbf{Y} \approx \mathbf{U}\mathbf{V}^{\mathrm{T}}. \tag{1}$$

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

Different decomposition techniques make different assumptions on the factor matrices. For instance, SVD achieves an exact decomposition of $\mathbf{Y}$ by extracting orthogonal factor matrices [24], while ICA attempts to find a set of maximally independent components [25]. While both SVD and ICA are fully data-driven, GLM specifies the factor matrix in time by defining design variables using the known time traces of triggering events (stimuli, motion, etc.) and HRFs. This way, a spatial activation map is estimated for each design variable. By definition, GLM assumes the same response to different repetitions of the same stimulus, although in reality, brain response exhibits dynamic characteristics [26].

A. 2D fUS

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

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

where $\mathbf{X} \in \mathbb{R}^{T\times K}$ is the GLM design matrix, K is the number of design variables and $j \in \Psi$ denotes the columns of $\mathbf{V}$ that are associated with the design matrix. The regularization coefficient λ adjusts the influence of design variables on these time courses. The remaining time courses (i.e. rest of $\mathbf{V}$'s columns) are kept unregularized to model the content that can not be pre-defined in a design matrix, such as background hemodynamic activity. In this work, we obtain the design variables by convolving each stimulus signal with an HRF, while in standard GLM, it is quite common to accommodate nuisance regressors in the design matrix as well [5]. These regressors, called as covariates, are used to suppress confounding effects, but are generally of no particular interest by themselves. Our decomposition framework allows for modelling of all such components in the non-regularized column(s) of $\mathbf{V}$, therefore we keep the design matrix simple by including only the stimuli of interest. Lastly, we incorporate a sparse derivative regularization on the same column set of $\mathbf{V}$ ($j \in \Psi$) to avoid sudden temporal changes in the components of interest, which can be induced by motion. For this purpose, we define the first order difference operator $\mathbf{D}$ [27] and a constant $\boldsymbol{\eta}$ to control the desired smoothness of the functional response.

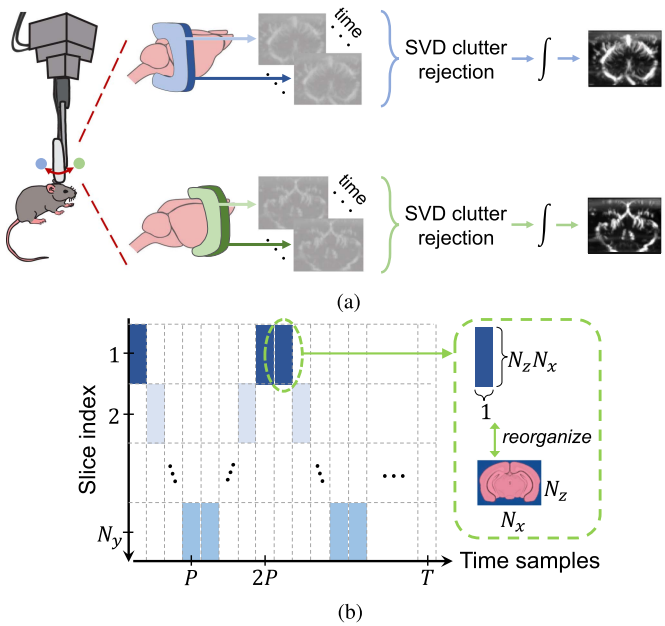

Fig. 2. Illustration for swept-3D fUS. (a) The ultrasound probe is continuously moved back and forth along the mouse brain. Beamformed echo frames are clutter-filtered with SVD in ensembles and later integrated over the ensemble dimension to create PDIs. (b) Due to the movement of the probe, swept-3D fUS creates an incomplete data matrix. Each colored block stands for an imaged brain slice in vectorized format, whereas the remaining (i.e. blank) blocks refer to points of no acquisition. P denotes the time for one sweep.

B. Swept-3D fUS

Swept-3D fUS does not image the whole brain at once, instead, captures it sequentially slice-by-slice at each sweep, leading to a slice timing offset (STO) problem. Correction of slice timings is a well-known challenge for fMRI as well, and is often tackled using interpolation techniques to re-align data from all slices to a common reference time axis [28]. However, the true interpolating function used for transferring neuroimaging data to a new time axis is actually unknown, and it is shown that its selection can have a major impact on further processing [29]. We instead propose to treat the data as it is, considering that one slice is measured at a time. This leads to a data array with missing values as shown in Fig. 2(b) [17]. Namely, an $(N_z \times N_x)$ image of a brain slice is acquired at each time instant, where N_z and N_x is the number of voxels in depth and width respectively. Within a sweep, N_y of such slices are acquired with slight shifts in time, creating a 3D image of the brain. By vectorizing the depth-width dimension, the swept-3D fUS data can be represented as an incomplete space-time matrix $\mathbf{Y}$.

Algorithm 1: Steps of the Proposed AM Algorithm.

1: Inputs: $\mathbf{Y}, \mathbf{X}$
2: Initialize:
$\mathbf{V}^{(0)}_{:,j\in\Psi} \leftarrow \mathbf{X}, \mathbf{V}^{(0)}_{:,j\notin\Psi} \leftarrow \mathcal{N}(0,1), k=0, \lambda, \eta, \epsilon$
3: **while** $\frac{\|\mathbf{U}^{(k+1)}-\mathbf{U}^{(k)}\|_F^2}{\|\mathbf{U}^{(k+1)}\|_F^2} + \frac{\|\mathbf{V}^{(k+1)}-\mathbf{V}^{(k)}\|_F^2}{\|\mathbf{V}^{(k+1)}\|_F^2} > \epsilon$ **do**

$$\mathbf{U}^{(k+1)} \leftarrow \underset{\mathbf{U}}{\operatorname{argmin}} \sum_{(i,j)\in\Omega} [\mathbf{Y}_{ij} - (\mathbf{U}\mathbf{V}^{(k)^{\mathrm{T}}})_{ij}]^2 \quad (4)$$

$$\mathbf{V}^{(k+1)} \leftarrow \underset{\mathbf{V}}{\operatorname{argmin}} \sum_{(i,j)\in\Omega} [\mathbf{Y}_{ij} - (\mathbf{U}^{(k+1)}\mathbf{V}^{\mathrm{T}})_{ij}]^2 + \lambda\|\mathbf{V}_{:,j\in\Psi} - \mathbf{X}\|_F^2 + \eta\|\mathbf{D}\mathbf{V}_{:,j\in\Psi}\|_1 \quad (5)$$

4: $k \leftarrow k+1$
5: **end while**
6: Outputs: $\mathbf{U}, \mathbf{V}$

Due to missing slice information, the optimization problem given in (2) should be solved only over the observed entries of $\mathbf{Y}$ for swept-3D fUS. This way, ECA also recovers the evoked activity of interest at each acquisition point. The minimization problem for swept-3D fUS data can be expressed as follows:

$$(\hat{\mathbf{U}}, \hat{\mathbf{V}}) = \underset{(\mathbf{U},\mathbf{V})}{\operatorname{argmin}} \sum_{(i,j)\in\Omega} [\mathbf{Y}_{ij} - (\mathbf{U}\mathbf{V}^{\mathrm{T}})_{ij}]^2 + \lambda\|\mathbf{V}_{:,j\in\Psi} - \mathbf{X}\|_F^2 + \eta\|\mathbf{D}\mathbf{V}_{:,j\in\Psi}\|_1. \quad (3)$$

III. Proposed Method

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

IV. fUS Data Acquisition

FUS imaging uses angled plane waves sent to the brain through a cranial window. In the regular (i.e., 2D) case, a linear array transducer is used, such that the backscattered signals, that are later beamformed and coherently compounded, constitute a 2D image of a given brain slice. Next, ensembles of adjacent compound frames are formed and SVD-filtered to reject undesired tissue artifacts. Finally, the SVD-filtered frames are integrated over the ensemble to create power-Doppler images (PDIs), whose pixel amplitude varies in proportion to the changes in local cerebral blood volume [30].

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

Although initially developed for acquisition of a single brain slice, various 3D-volumetric extensions of fUS imaging have been employed since then. One such extension is obtained by concatenating 2D fUS images of different slices by repeating the same experiment at each slice [31]. Alternatively, a 2D matrix array can be used for ultrasound transmission [32], which requires expensive hardware and acquired volumes suffer from lower sensitivity. Recently, another solution known as swept-3D fUS (Fig. 2(a)) has been proposed, which uses a continuously-moving ultrasound probe. The probe is moved continuously back-and-forth over the craniotomy during the experiment, resulting in a 3D volume after every half-cycle of the probe's movement (called as a *sweep*). Eventually, a full PDI volume of the mouse brain (roughly from Bregma −4 mm to +2 mm) is completed in $\sim$ 1 second. Further information on our acquisition setup can be found in Supplementary Materials Section I-B. For a more comprehensive understanding of swept-3D fUS as an imaging technique, including how the sweeping speed affects the resulting PDIs, we refer the reader to [15], [16].

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

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

Prior to applying ECA, we standardized each PDI voxel time series to zero-mean and unit variance and applied spatial smoothing with a Gaussian kernel with standard deviation of a half-voxel in size. Lastly, we registered the corresponding slices of Allen Brain Atlas on the mean PDIs to locate the regions-of-interest (ROIs) [33].

V. Results

We first present our results on 2D fUS data and thoroughly examine the impact of hyperparameter selection in ECA. Next, we apply ECA on swept-3D fUS data and use t-statistics for comparing our findings to correlation and GLM analyses. Note that, while ECA allows the observation matrix to encompass

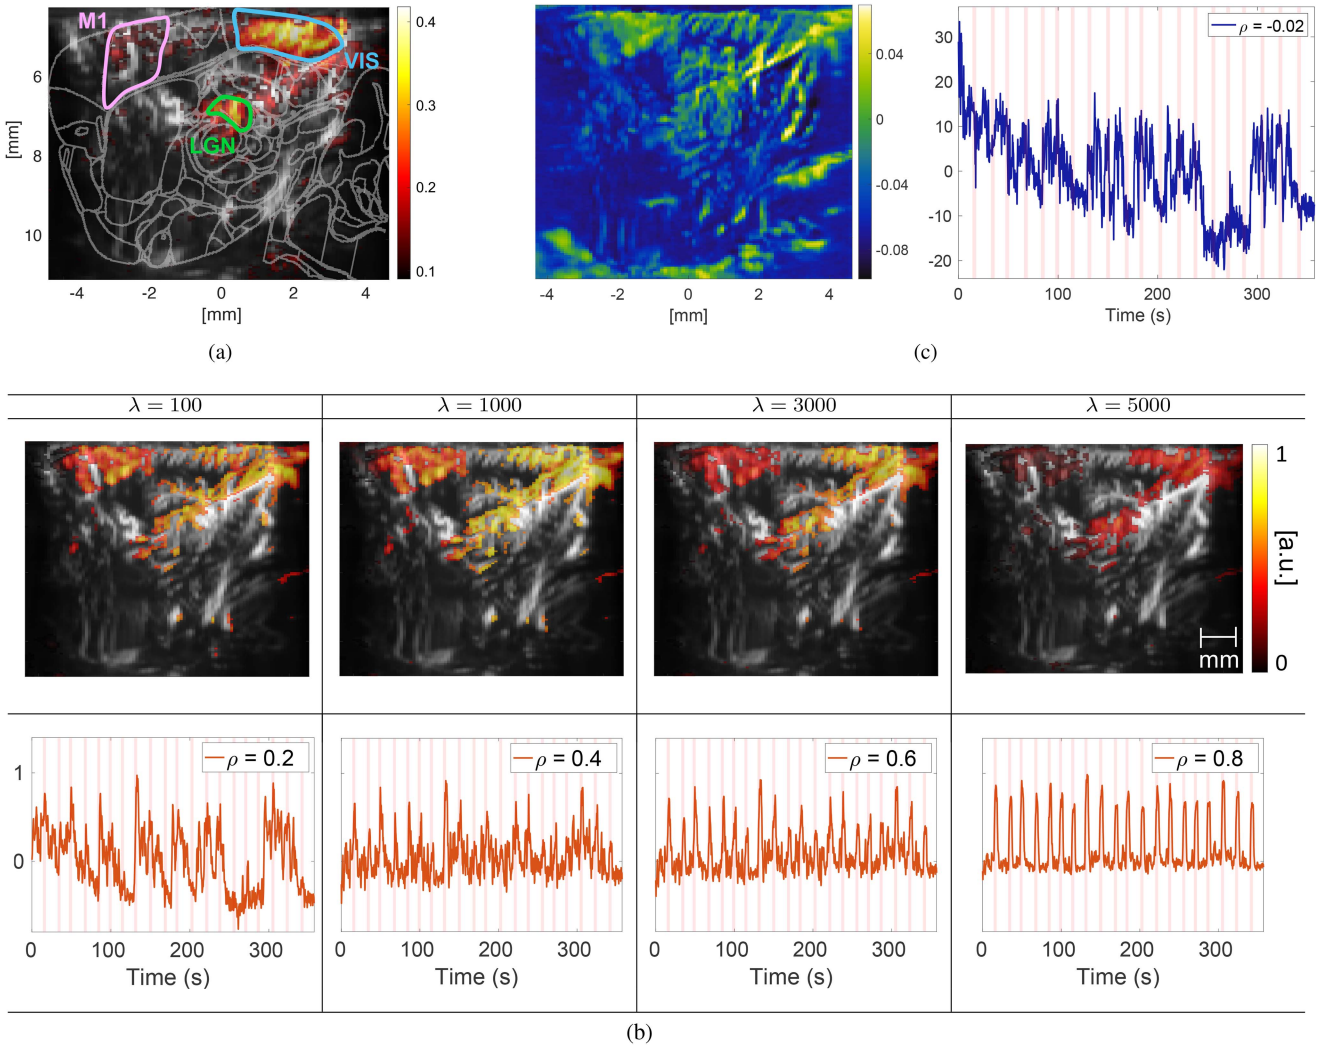

Fig. 3. Analysis of 2D fUS data. (a) PCC image overlaid against the mean PDI. The ROIs are highlighted over the registered atlas (VIS: visual cortex, M1: primary motor area, LGN: lateral geniculate nucleus). (b) ECA of 2D fUS data with different degrees of regularization. The PCC (ρ) between each estimated time course and the stimulus signal is indicated in parenthesis at the bottom plots. The estimated spatial maps (top plots) display voxels with z-score $\geq$1.5 overlaid against the mean PDI. Due to the ambiguity in sign and scaling of the estimated factor matrices, we normalized each time course to have a maximum amplitude of 1 and applied a reversed scaling on the corresponding spatial map for a fair comparison of voxel activation levels across various regularizations. Afterward, we used the same color range for all spatial maps. (c) Estimated spatial activation map and time course of the non-regularized component at $\lambda = 1000$. Notice that, z-score thresholding was not applied to the non-regularized component to showcase the vascular structure with full contrast without overlaying of the mean-PDI.

missing entries, accounting for slice timing differences illustrated in Fig. 2(b), slice measurements should be carried to the same reference time axis for the standard GLM formulation. To that end, we employ sinc interpolation as offered by statistical parametric mapping (SPM) [34] for slice timing correction while applying GLM.

A. Results on 2D fUS

We first performed a conventional correlation analysis (Fig. 3(a)) on the 2D fUS data, which reveals activations in three ROIs: visual cortex (VIS), lateral geniculate nucleus (LGN) and primary motor area (M1). The PCC image suggests a clear response to the stimulus in LGN and VIS, and a lower level of activation in M1.

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

We applied ECA at 4 different values of λ. The estimated spatial maps and corresponding time courses are shown in Fig. 3(b). When $\lambda = 5000$, the method attempts at discovering voxels whose response aligns almost perfectly with the stimulus signal, which is only observed in parts of LGN and VIS with not as high activation levels as in less strict cases of regularization. As λ is decreased, epoch-dependent variations appear in the estimated time course, and activation levels rise in LGN, VIS and M1. LGN and VIS appear brightest at $\lambda = 1000$, meaning that

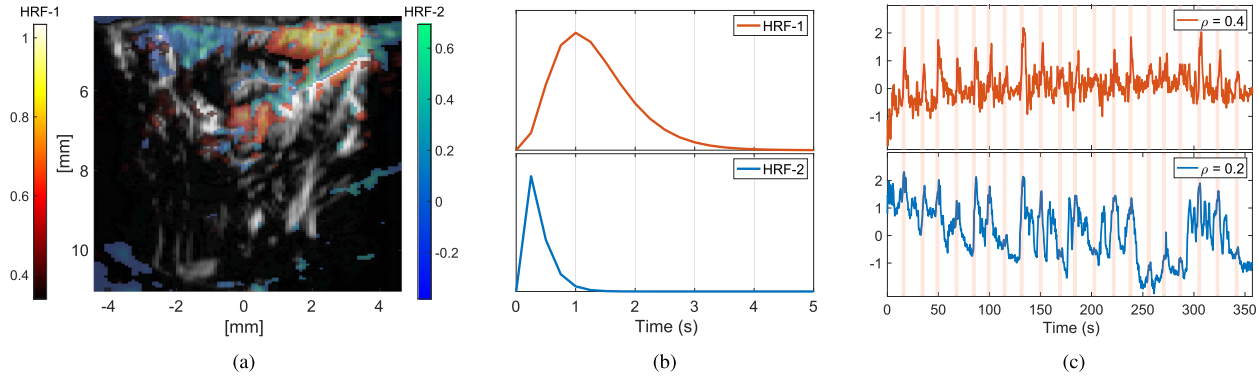

Fig. 4. Decomposition of 2D fUS data at $R = 3$ with two regularized components. Only the voxels that are found significantly active (z-score $\geq$1.5) for eact spatial component are displayed in (a). The newly introduced regularized component is shown in blue color. (a) Estimated spatial maps of the regularized components. (b) HRFs used in the design matrix. (c) Estimated time courses of the regularized components.

the time course obtained at $\lambda = 1000$ is the most representative for the response of these regions.

When λ is reduced further to 100, activation levels in LGN and VIS drop, while M1 reaches its highest activation. As a result of capturing of the common fluctuations in these regions instead of the stimulus-based design variable, the contrast between LGN/VIS and M1 is almost vanished at $\lambda = 100$. The changing temporal patterns can also be visualized epoch-by-epoch to highlight epoch-specific variations of the brain's functional response (Fig. S1). Our results indicate that the response of M1 to the stimulus is not consistent across all epochs, hence activations in this area become more prominent at low regularizations. Notice that even when $\lambda = 0$, the estimated time course can still reflect a link to the stimulus onsets (Fig. S2). This link stems from the algorithm's initialization approach, wherein $\mathbf{V}$'s column(s) of interest are initialized directly as the design variable(s). It should as well be noted that the stopping criterion is fixed at $\epsilon = 0.1$ for all results presented in this paper, with an example convergence plot provided in Fig. S3.

On the other hand, the estimated time course of the non-regularized component (Fig. 3(c)) manifests great baseline shifts, and is unrelated from the stimulus signal. Likewise, the associated spatial map exhibits much higher activity in large brain vessels and outside the brain instead of functional brain areas, outlining the brain's vasculature. The non-regularized component was observed to be less affected by varying λ (Fig. S4).

2) Rank Selection: In this section, we will discuss the effect of rank on the estimated factors. When a single regularized component is assumed ($R = 1$), we observed that this component is severely affected by noise compared to the results in Fig. 3. While adding at least one non-regularized component significantly helps with the noise suppression in task-relevant component(s) of interest, introducing more non-regularized terms cause no remarkable change (Fig. S5). Therefore, we kept the number of non-regularized components as 1 in the results to follow.

Subsequently, we applied ECA with two task-relevant components (both regularized at $\eta = 10$ and $\lambda = 200$, selected heuristically) and one non-regularized component. Notice that, we reduced the value of λ here compared to its ideal value found in rank-2 analysis. We observed that keeping λ in the order of thousands in the high-rank case outputs time courses that are very close to what was hypothesised in the design matrix (Fig. S7). Increasing the rank allows for a more detailed representation of the data through the $\mathbf{UV}$ factorization, thus part of the cost function corresponding to the data reconstruction error is reduced. Hence, for preserving the balance between data-fitting (first term in (2)) and model-fitting (second term in (2)), λ was decreased as well [35].

For the newly introduced task-relevant component, we utilized a different HRF shape in the design matrix. The results are provided in Fig. 4. To start with, we can observe the voxels that were mostly captured together in the lower-rank case being divided into two groups. Particularly, M1 and the vessels surrounding VIS and LGN are revealed to exhibit a more dynamic response to the stimulus, resulting in their separation from VIS and LGN. These results are in accordance with our observations in rank-2 analysis, where the highest activation of M1 was detected in the low regularization case, with again a PCC of 0.2 shared with the stimulus signal.

The non-regularized component (Fig. S6) similarly highlights the brain's vasculature as in the rank-2 case. However, the activations seem to have receded from the pronounced influence of M1 and blood vessels surrounding VIS, since now their stimulus-evoked content is captured by the second HRF. Moreover, almost the entirety of the response in VIS is attributed to the regularized components, resulting in activation values near 0 for VIS in the non-regularized component.

B. Results on Swept-3D fUS

Our goal in this experiment was to observe if ECA can capture activity evoked by different tasks, associated with the position of the LED stimulus (either left or right side). For this purpose, we constructed two design variables based on the stimulus onsets of each task. We assumed a total rank of 3, containing two task-regularized components and one non-regularized component to model spatially structured fluctuations due to background activity or noise. We set $\lambda = 1000$ and $\eta = 10$. Our results from one subject are displayed in Fig. 5 (time courses and corresponding

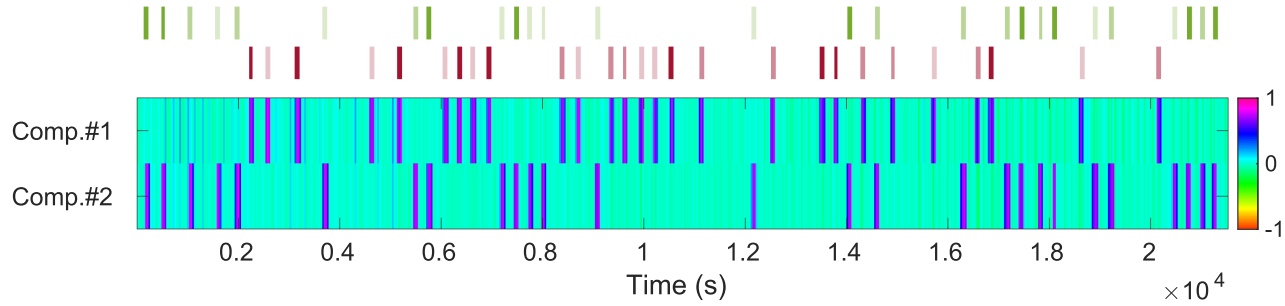

(a)

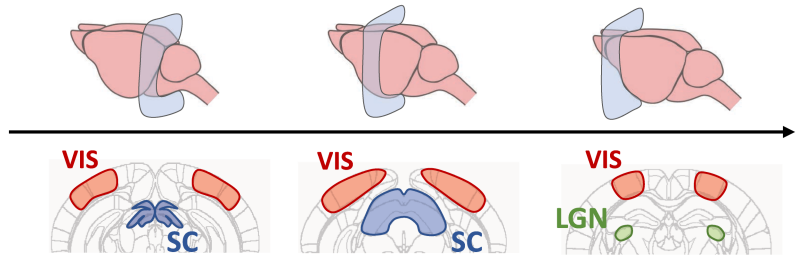

(b)

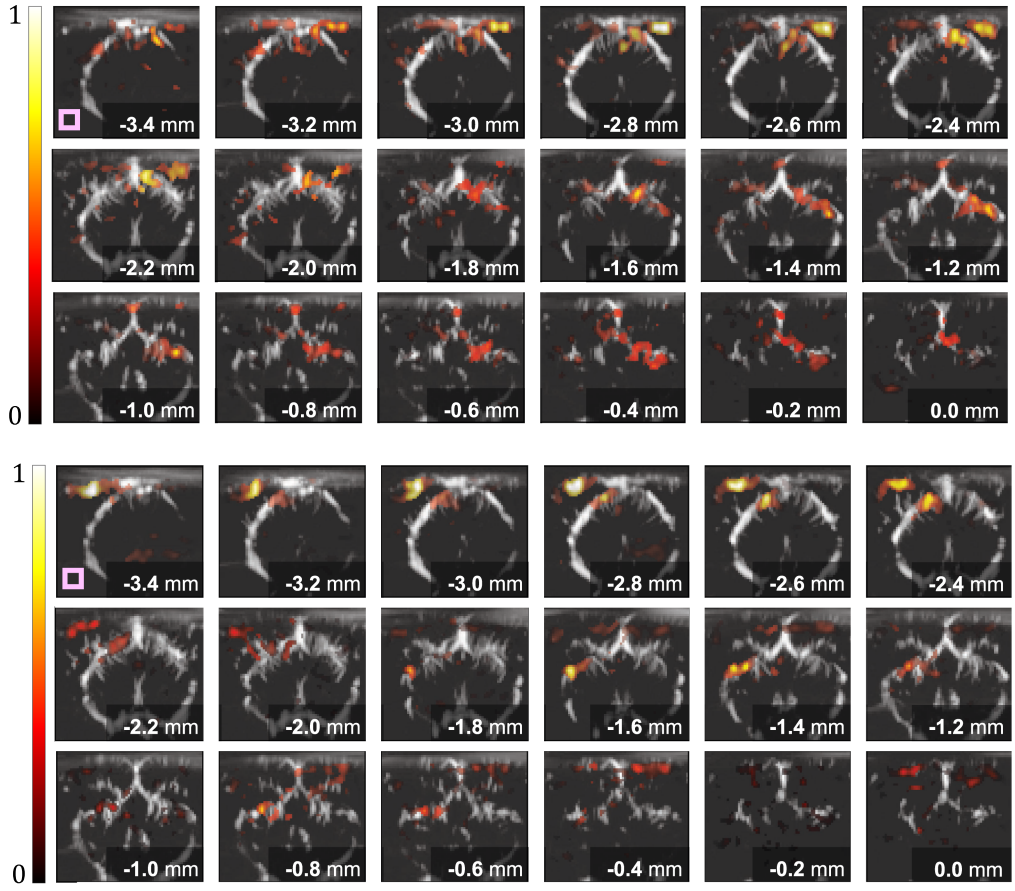

(c)

Fig. 5. Estimated time courses (a) and associated spatial maps (c) for one subject (N=1). The ROIs shown over the mouse brain atlas (b) are found to be activated in (c), where only voxels with z-score ≥ 1.5 are displayed against the mean PDI. The selected control region is shown only in the first slice, but was in fact defined over the 3D brain by repeating the same mask at every slice. (a) Estimated time courses of the two regularized components. The red and green colored bars on top indicate the left and right-side LED stimuli respectively, with different shades referring to the LED brightness level. The colorbar denotes the amplitude changes of the time courses. (b) Illustration of the probe's movement and evolution of ROIs within the mouse visual pathway. (c) Estimated spatial maps corresponding to left-side (top) and right-side (bottom) stimuli.

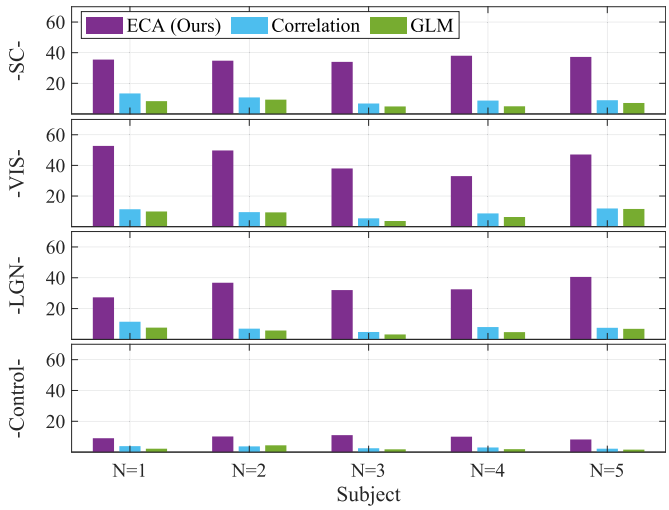

Fig. 6. Mean of top 20 t-statistics per region for all subjects.

spatial maps of the GLM-regularized components). Although we did not discover any significant brightness level-dependent changes in the magnitude of estimated time courses, we observed an overall decreasing trend in their magnitudes across epochs (Fig. 5(a)). Compared to the 2D fUS experiment with on-screen stimuli, we found the LED stimulus to elicit a more consistent brain response, with less epoch-specific highs or lows.

For assessing the statistical significance of voxel activation levels estimated with ECA as opposed to standard analyses, we checked the mean of top 20 t-statistics [29] achieved at the ROIs, namely superior colliculus (SC), LGN and VIS, and a control region. We selected the control region outside of the functional areas to ensure that the higher statistics obtained in the ROIs by either method is indeed a result of the task-evoked activity in the ROIs. The t-statistic evaluates the accuracy of a model fit by computing the standard error (SE) based on the variance of the residual between the modeled time series and measured data [36].

Specifically, for each voxel in the brain, the null hypothesis asserts that there is no effect of the predictor variables (i.e., columns of the temporal factor matrix) on the response variable (i.e., voxel time series) by assuming that the true coefficient value for the predictor variable is zero, generating a high residual. For GLM (with sinc interpolation), this residual corresponds to the difference between the modeled design variables and voxel time series. For Pearson correlation, the t-statistic is given by $t = \frac{\rho\sqrt{n-2}}{1-\rho^2}$, where n denotes the sample size [37]. For ECA, the predictor variables correspond to the estimated time courses, which are regularized by, but not solely based on stimulus information. As such, ECA time courses constitute a much better fit for the actual data, resulting in very low SE values and much higher t-statistics for all subjects, as shown in Fig. 6.

To further support our conclusions, we provide a visual comparison between ECA and the PCC maps obtained from the same subject (Fig. S8) by computing the difference in voxel activations (normalized to unit energy) as estimated by both methods in Fig. S9. These results confirm that the activations in regions of interest found by ECA are stronger than those found by PCC. In addition, we share the results of ECA on swept-3D fUS data at lower λ in Fig. S10, which aligns with our prior observations in the 2D fUS case as the estimated factor matrices in both space and time absorb more of the task-irrelevant content. Consequently, the contrast of the regions of interest in the spatial activation maps is somewhat diminished compared to Fig. S5. However, even at lower or higher values of λ, ECA retains higher t-statistics than conventional methods in regions of interest in contrast to the control region, indicated by the ratio between activations in ROIs and the control region as provided in Fig. S11, preserving its favorable stance for analyzing evoked components within the brain.

VI. Discussion

In this study, we introduced ECA as a novel decomposition technique for analyzing neuroimaging data. Existing decompositions in the literature are typically either data-driven or completely reliant on prior information. This prior information is commonly given by the stimulus time course that entails when a stimulus is shown, represented as a boxcar function. However, the brain response may not always vary in sync with stimuli due to various factors including neuronal adaptation and stimulus expectation [38]. In our approach, we use the stimulus time course only as a guiding term, striking a balance between expectation (design variables defined based on stimuli) and reality (measured data). This balance is controlled through regularization of the temporal factor matrix with respect to the design variables.

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

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

Subsequently, we applied our method on swept-3D fUS data that is subject to a similar slice timing offset problem as in fMRI. Instead of resolving the problem by interpolation (as commonly done in fMRI), we proposed to treat the data as it was acquired and implicitly recover the evoked activity at each acquisition point. We showed that ECA more precisely describes temporal and spatial brain responses than standard approaches. To elaborate, we experimented on five mice and compared the resulting t-statistics to those obtained via GLM (with sinc interpolation, [34]) and correlation analysis. While applying ECA, we assumed a rank of 3, where the first component was regularized according to stimuli presented to the left eye, the second component towards stimuli shown to the right eye, and the last component was not subject to any regularization. We observed that the estimated spatial maps showed clear differences in the activated areas associated with left and right side LED stimulus. Our approach yielded significantly higher t-statistics in the visual processing pathway of the mouse brain. Note that, low-rank completion has been proposed for recovery of data points corrupted with artifacts and slice-timing correction before for fMRI [39]. This approach is based on the assumption that the temporal signal of a voxel at any time point can be expressed as a linear combination of its previous samples, and that these linear weights are shared within all voxels of functionally connected regions. To the best of our knowledge, ours is the first study to achieve recovery of evoked activity at the full scale, i.e. for all slice acquisition points using only the known stimulus onsets. This perspective can be beneficial for future work on task-based swept-3D fUS and fMRI studies, since it does not require selecting an appropriate interpolating function, which is shown to have a significant effect in processing of the data [40].

Notice that ECA necessitates the use of a design matrix as input for regularization, which in turn requires the specification of an HRF shape (or shapes). Indeed, similar specifications are unavoidable for GLM and correlation-based analyses as well. To elaborate, standard GLM also requires a design matrix with user-defined HRF shapes, commonly picked as only the canonical form or canonical form with its derivative(s) [41]. For correlation analysis, the stimulus time course (or if known, neuronal activity) is either subjected to a certain amount of delay or convolved with an HRF [42], which is again a decision left to the user. In fact, ECA provides more flexibility than GLM or correlation analyses since it does not absolutely depend on the given prior information, including the pre-specified HRF shape(s). At the same time, it is worth emphasizing that the results of ECA are influenced by the regularization intensity, λ. While we suggest $\lambda = 1000$ as a point of reference, its choice is indeed data and purpose-dependent, such as whether the application requires a more model-based (high λ) analysis or a data-driven (low λ) one. For example, while Fig. S11 reveals a lower contrast in regions of interest compared to the control region at lower λ, the better capturing of shared fluctuations also result in a more accurate portrayal of the brain's activity along time for all regions. In other words, each stage of Fig. 3 provides us with valuable insights into evoked and inherent activity of the brain, yet the full picture comes together when inspecting the data at various λ values, which might be favorable for exploratory analyses. Although the optimal λ is indeed purpose-dependent and influences the relative contrast between interest and control regions, note that ECA, regardless of the choice of λ, provides better contrast than alternative methods (Fig. S11).

Another important aspect to address is the computational efficiency of the algorithm. Unlike GLM or correlation analysis, which assume known temporal regressors and only predict their spatial counterparts, ECA estimates time courses as well. This flexibility allows for capturing the time-varying characteristics of the brain response, but comes at the expense of computational time. The estimation of spatial activation maps in 2D fUS takes up a negligible amount of time, as (4) becomes equivalent to least-squares. On the contrary, due to missing entries within our swept-3D fUS matrix formulation, again a point-wise minimization is required in space. Ultimately, the complete decomposition of the swept-3D fUS data (of size 71280×21402) requires approximately half an hour when executed using MATLAB 2021a on a high-performance computing system running Linux, equipped with two AMD EPYC 32-Core Processors and 528 GB of memory. To enhance the algorithm's efficiency, additional structures can be imposed in space or time. For instance, assuming that the estimated spatial maps can be reconstructed from low-rank factors in each space dimension (depth, width and height) can help speed up the process [43], [44].

VII. Conclusion

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

References

[1] E. Macé et al., "Functional ultrasound imaging of the brain," *Nature Methods*, vol. 8, no. 8, pp. 662–664, 2011.

[2] S. Soloukey et al., "Functional ultrasound (fUS) during awake brain surgery: The clinical potential of intra-operative functional and vascular brain mapping," *Front. Neurosci.*, vol. 13, 2020, Art. no. 1384.

[3] M. Gesnik et al., "3D functional ultrasound imaging of the cerebral visual system in rodents," *Neuroimage*, vol. 149, pp. 267–274, 2017.

[4] A. K. Aydin et al., "Transfer functions linking neural calcium to single voxel functional ultrasound signal," *Nature Commun*, vol. 11, 2020, Art. no. 2954.

[5] M. M. Monti, "Statistical analysis of fMRI time-series: A critical review of the GLM approach," *Front. Hum. Neurosci.*, vol. 5, 2011, Art. no. 28.

[6] 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, pp. 1639–1651, 2004.

[7] M. I. Garrido et al., "Repetition suppression and plasticity in the human brain," *Neuroimage*, vol. 48, no. 1, pp. 269–279, 2009.

[8] G. Stefanics et al., "Timing of repetition suppression of event-related potentials to unattended objects," *Eur. J. Neurosci.*, vol. 52, no. 11, pp. 4432–4441, 2020.

[9] T.-P. Jung et al., "Imaging brain dynamics using independent component analysis," *Proc. IEEE*, vol. 89, no. 7, pp. 1107–1122, Jul. 2001.

[10] K. Smitha et al., "Resting state fMRI: A review on methods in resting state connectivity analysis and resting state networks," *Neuroradiol. J.*, vol. 30, no. 4, pp. 305–317, 2017.

[11] B. Hunyadi et al., "ICA extracts epileptic sources from fMRI in EEG-negative patients: A retrospective validation study," *PLoS One*, vol. 8, no. 11, 2013, Art. no. e78796.

[12] L. Griffanti et al., "Hand classification of fMRI ICA noise components," *Neuroimage*, vol. 154, pp. 188–205, 2017.

[13] G. Salimi-Khorshidi et al., "Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers," *Neuroimage*, vol. 90, pp. 449–468, 2014.

[14] B. Hunyadi et al., "A prospective fMRI-based technique for localising the epileptogenic zone in presurgical evaluation of epilepsy," *Neuroimage*, vol. 113, pp. 329–339, 2015.

[15] B. S. Generowicz et al., "Swept-3D ultrasound imaging of the mouse brain using a continuously moving 1D-array part I: Doppler imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 70, no. 12, pp. 1714–1725, Dec. 2023.

[16] B. S. Generowicz et al., "Swept-3D ultrasound imaging of the mouse brain using a continuously moving 1D-array part II: Functional imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 70, no. 12, pp. 1714–1725, Dec. 2023.

[17] A. Erol et al., "GLM-Regularized low-rank factorization for extracting functional response from Swept-3D functional ultrasound," in *Proc. IEEE Data Sci. Learn. Workshop*, 2023, pp. 1–5.

[18] A. F. Meyer et al., "Models of neuronal stimulus-response functions: Elaboration, estimation, and evaluation," *Front. Syst. Neurosci.*, vol. 10, 2017, Art. no. 109.

[19] A. L. Vazquez and D. C. Noll, "Nonlinear aspects of the bold response in functional MRI," *Neuroimage*, vol. 7, no. 2, pp. 108–118, 1998.

[20] M. X. Cohen, "A tutorial on generalized eigendecomposition for denoising, contrast enhancement, and dimension reduction in multichannel electrophysiology," *Neuroimage*, vol. 247, 2022, Art. no. 118809.

[21] E. A. Mukamel, A. Nimmerjahn, and M. J. Schnitzer, "Automated analysis of cellular signals from large-scale calcium imaging data," *Neuron*, vol. 63, no. 6, pp. 747–760, 2009.

[22] A. A. Khaliq, I. M. Qureshi, and J. A. Shah, "Unmixing functional magnetic resonance imaging data using matrix factorization," *Int. J. Imag. Syst. Technol.*, vol. 22, no. 4, pp. 195–199, 2012.

[23] M. A. Lindquist, "The statistical analysis of fMRI data," *Stat. Sci.*, vol. 23, no. 4, pp. 439–464, 2008.

[24] K. Li et al., "Review of methods for functional brain connectivity detection using fMRI," *Comput. Med. Imag. Graph.*, vol. 33, no. 2, pp. 131–139, 2009.

[25] V. D. Calhoun and T. Adali, "Unmixing fMRI with independent component analysis," *IEEE Eng. Med. Biol. Mag.*, vol. 25, no. 2, pp. 79–90, Mar./Apr., 2006.

[26] N. G. Müller et al., "Repetition suppression versus enhancement–it's quantity that matters," *Cereb. Cortex*, vol. 23, no. 2, pp. 315–322, 2013.

[27] F. I. Karahanoğlu et al., "Total activation: FMRI deconvolution through spatio-temporal regularization," *Neuroimage*, vol. 73, pp. 121–134, 2013.

[28] R. Sladky et al., "Slice-timing effects and their correction in functional MRI," *Neuroimage*, vol. 58, no. 2, pp. 588–594, 2011.

[29] D. Parker, X. Liu, and Q. R. Razlighi, "Optimal slice timing correction and its interaction with fMRI parameters and artifacts," *Med. Image Anal.*, vol. 35, pp. 434–445, 2017.

[30] T. Deffieux, C. Demené, and M. Tanter, "Functional ultrasound imaging: A new imaging modality for neuroscience," *Neuroscience*, vol. 474, pp. 110–121, 2021.

[31] É. Macé et al., "Whole-brain functional ultrasound imaging reveals brain modules for visuomotor integration," *Neuron*, vol. 100, no. 5, pp. 1241–1251, 2018.

[32] C. Rabut et al., "4D functional ultrasound imaging of whole-brain activity in rodents," *Nature Methods*, vol. 16, no. 10, pp. 994–997, 2019.

[33] Allen Institute for Brain Science., "Allen brain atlas API," 2015. [Online]. Available: brain-map.org/api/index.htm

[34] W. D. Penny et al., *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Amsterdam, The Netherlands: Elsevier, 2011.

[35] C. M. Bishop and N. M. Nasrabadi, *Pattern Recognition and Machine Learning*, vol. 4, Berlin, Germany: Springer, 2006.

[36] V. Della-Maggiore et al., "An empirical comparison of SPM preprocessing parameters to the analysis of fMRI data," *Neuroimage*, vol. 17, no. 1, pp. 19–28, 2002.

[37] D. R. Anderson, D. J. Sweeney, and T. A. Wlliams, *Statistics for Business and Economics*, 13th edition, Boston, MA, USA: Cengage Learning, Ch. 14, 2011, p. 676.

[38] J. Larsson and A. T. Smith, "fMRI repetition suppression: Neuronal adaptation or stimulus expectation," *Cereb. Cortex*, vol. 22, no. 3, pp. 567–576, 2012.

[39] A. Balachandrasekaran et al., "Reducing the effects of motion artifacts in fMRI: A structured matrix completion approach," *IEEE Trans. Med. Imag.*, vol. 41, no. 1, pp. 172–185, Jan. 2022.

[40] V. Calhoun, X. Golay, and G. Pearlson, "Improved fMRI slice timing correction: Interpolation errors and wrap around effects," in *Proc. ISMRM, 9th Annu. Meeting*, 2000, p. 810.

[41] F. Cignetti et al., "Pros and cons of using the informed basis set to account for hemodynamic response variability with developmental data," *Front. Neurosci.*, vol. 10, 2016, Art. no. 322.

[42] H. Laufs et al., "EEG-correlated fMRI of human alpha activity," *Neuroimage*, vol. 19, no. 4, pp. 1463–1476, 2003.

[43] A. Erol and B. Hunyadi, "Tensors for neuroimaging: A review on applications of tensors to unravel the mysteries of the brain," *Tensors Data Process.*, pp. 427–482, 2022.

[44] C. Chatzichristos et al., "Blind fMRI source unmixing via higher-order tensor decompositions," *J. Neurosci. Methods*, vol. 315, pp. 17–47, 2019.