Performance assessment of IVA-based subgroup identification methods

and their application in experimental fUS data

C.W.H. Bot

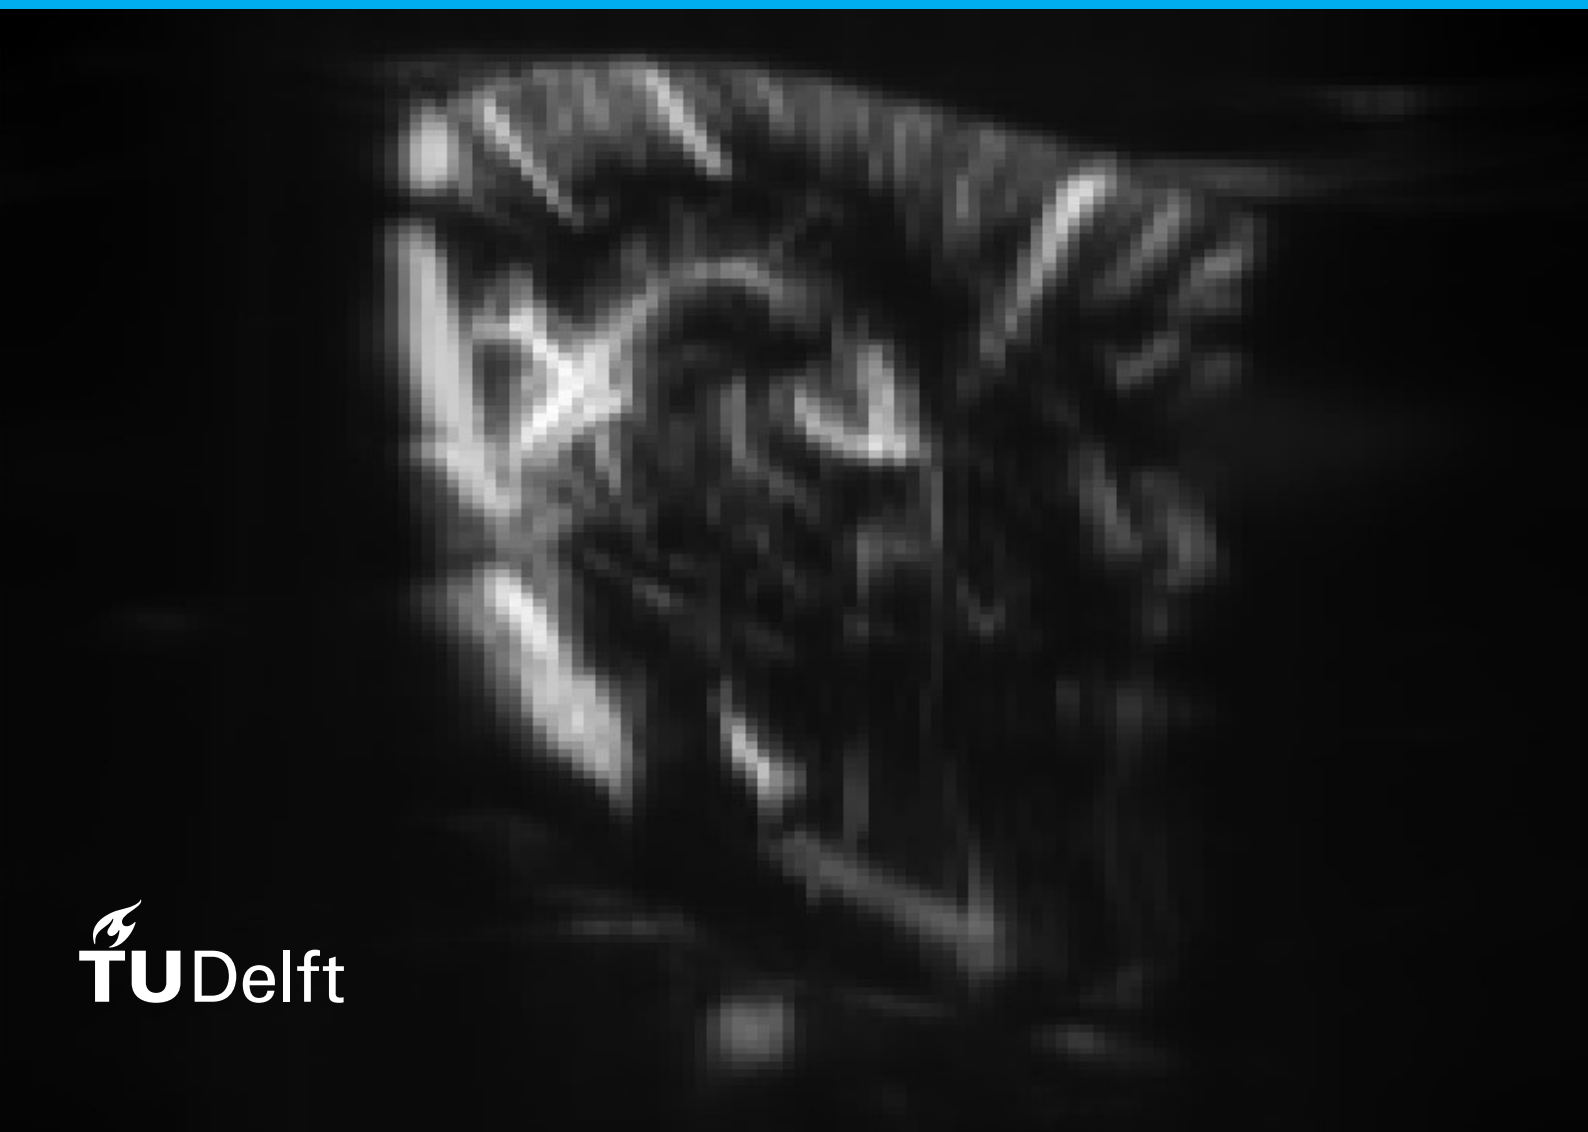

TU Delft

Performance assessment of IVA-based subgroup identification methods

and their application in experimental fUS data

by

C.W.H. Bot

To obtain the degree of Master of Science in Electrical Engineering,
Track: Signal processing systems,
at the Delft University of Technology.

Author:	Christiaan Bot	4927133
Supervisor:	Dr. B. Hunyadi	TU Delft
Chair:	Dr. B. Hunyadi	TU Delft
Committee:	Dr.ir. R.F. Remis	TU Delft

16 Aug. 2024

TU Delft

Abstract

Using independent vector analysis (IVA) to analyze and find subgroups in functional magnetic resonance imaging (fMRI) and functional ultrasound (fUS) data requires a lot of manual labour. Recently methods like subgroup identification using IVA (SI-IVA) and IVA for common subspace identification (IVA-CS) have tried to reduce this labour through automation. However, both methods did not test for accuracy. This thesis shows through simulations that these proposed methods are not accurate or robust enough to be trusted and that spectral clustering is a better alternative for automatic subgroup identification. Spectral clustering is then incorporated into the analysis of experimental fUS data of two groups of mice to try and identify these automatically. In this analysis, adaptive constrained IVA (acIVA) was used to incorporate references, further improving the interpretability of the results as components are directly linked to prior constraints. However, applying subgroup analysis showed that the mice could not be clustered based on their response to the stimuli. Still, spectral clustering is more accurate in the simulations making it a promising alternative for automatic subgroup identification. Furthermore, combining spectral clustering with acIVA makes the results more interpretable due to constrained components not being subject to permutation ambiguity.

Contents

1 Introduction

One of the best-kept secrets is how the brain works, bit by bit this mystery is being chipped away at through new functional imaging techniques like functional ultrasound (fUS), giving more insights as to what happens inside the brain. This data has to be processed and interpreted to reach further conclusions, requiring improvements to existing analysis methods to get there. This thesis will assess some of these methods and thereafter examine experimental fUS data with those methods.

In the field of functional imaging, functional magnetic resonance imaging (fMRI) has been the standard when investigating the brain for more than two decades [1]. This is because the fMRI measures blood-oxygen-level-dependent (BOLD) signals. This BOLD signal is of interest as it is linked to neural connectivity through neurovascular coupling. While fMRI is often used, there are spatial and temporal limitations when measuring BOLD signals. The spatial resolution for 3 T experiments is 2-3 mm and the measurement frequency is below 0.5 Hz [2]. This makes it hard to accurately analyze small subjects and higher frequency signals. In 2011 a new functional imaging method called fUS was introduced, which shows improvements in spatial and temporal resolution compared to fMRI [3]. Instead of measuring BOLD signals, fUS measures the cerebral blood volume (CBV). CBV is also linked to neural connectivity through neurovascular coupling, making fUS an interesting alternative to generate more insights about the functional connectivity for stimulus-related responses [4].

To find differences between the subjects that are imaged it is first necessary to extract the underlying sources that generate the recorded fUS data. Extracting these underlying sources can be done in two ways for fMRI and fUS [5]. The first focuses on a model-based analysis that is heavily reliant on good a priori assumptions. The second method is a more data-driven approach based on blind source separation (BSS), this is a more explorative approach to analyzing the data, where fewer assumptions are made. A classic BSS approach to capture components in fMRI is through independent component analysis(ICA) [1], which tries to separate the data into components (sources) that are independent of one another. This allows ICA to reveal regions of interest and associated time courses in a data-driven manner [5]. This thesis will focus on an extension of the latter method called adaptive constrained independent vector analysis (acIVA) which remains data-driven but also allows for joint analysis of multiple subjects and allows for including prior information through constraints.

However, after the components have been extracted it is still necessary to manually identify and compare each component to infer differences between groups of subjects. This takes a large amount of manual labour [6]. To reduce this labour, two methods called: subgroup identification using IVA (SI-IVA) and IVA for common subspace identification (IVA-CS), have been proposed, these methods try to automatically highlight subgroups within the extracted sources through automatic subgroup identification. These methods successfully found subgroups within the datasets to which they were applied, however, there are no indicators given as to how accurate these methods are. This thesis will perform an analysis of these subgroup identification methods and compare them with other clustering algorithms to find a robust and reliable subgroup identification method. After this performance analysis, the best method among those tested will be used with experimental fUS data where two groups of mice are subject to the same stimuli. By applying the selected subgroup identification method it is the question to see if these groups can be extracted from the data.

2

Background

Before the fUS signals are decomposed into sources and used to find subgroups, a short introduction to fUS is included to explain what the imaging technique is, what is being measured and how this relates to neural activity and the stimuli.

2.1. Functional Ultrasound

fUS is a relatively new imaging technique introduced in 2011 [3]. fUS uses a probe that pulses ultrasonic waves into a targeted area, and these ultrasonic waves create backscattered echoes when interacting with tissue. Capturing these echoes is what allows fUS to create an image. When this ultrasonic wave touches tissue that moves (such as blood) a Doppler shift happens to the frequency of the echo relative to the movement of the tissue. By capturing these shifts fUS measures the CBV changes in a pixel. Like BOLD, CBV does not directly measure neural activity but can highlight regions where neural activity takes place through neurovascular coupling, as the CBV changes highlight where more or less blood is flowing at each moment signifying neural activity.

To accurately measure this signal, plane-wave illumination is used where multiple tilted waves are emitted at different angles. By combining these measurements fUS is capable of making an image every millisecond. These are still subject to large amounts of noise and outliers, therefore, after acquiring a number of these images a compound image is computed to create a good contrast and resolution and reduce the effects of noise. Since the measurements include movements from any tissue within the target area, signals not related to the CBV are also captured. To capture only blood movements, a high-pass filter is used to remove signals that move slower than blood [7]. This results in the eventual power Doppler image (PDI) to be described by an average of the high-pass filtered single images [7]. Resulting in the intensity I of the pixel at location (x, z) in the PDI to be given by:

$$I(x,z) = \frac{\sum_{i=1}^{N} |s(x,z,t_i)|^2}{N}, \tag{2.1}$$

where N is the number of images used for the compound image, and $s(x, z, t_i)$ is the value at pixel (x, z) from image t_i. N can be freely chosen but is generally around 200, this results in each compound image consisting of 200 images that are measured at a frequency of 1 KHz, meaning that the final PDIs are measured at a frequency of around 4 to 5 Hz as it depends on N. The spatial resolution that was achieved in the introduction of fUS was $100\mu m \times 100\mu m \times 200\mu m$. Showing that fUS has a better spatiotemporal resolution than fMRI which had a spatial resolution in the millimetres and a temporal resolution in the seconds.

Since fUS measures a different signal compared to fMRI the methods may not be comparable, therefore, research has been performed to see how both methods compare and differ in terms of extracted networks. For example, in [8] fUS is compared to fMRI when it comes to extracting resting state networks (RSN) of mice. It was found that applying the same pipeline on both methods results in the same RSNs being extracted, the major difference between the methods was that fUS gave rise to bilateral components that were unilateral in fMRI. This means that fMRI extracts the RSNs such that both brain halves are independent, and fUS extracts RSNs where both halves are dependent. Nevertheless, it was deemed that fUS is a suitable alternative for functional imaging, and methods used for fMRI are also applicable to fUS.

2.2. Hemodynamic Response Function

It should be of note that the captured signals (which are a hemodynamic response) will not be of the same shape as the stimulus, due to the brain having to process and respond to what is happening when an input is supplied. The processing of this input is modelled through a linear time-invariant system where the output is generated through a transfer function. Within functional imaging, this transfer function is usually denoted as the hemodynamic response function (HRF). This means that the HRF transforms the input through a convolution to what is eventually measured. A general model for what an HRF looks like under physiological constraints in fUS is shown below [9]:

$$h(t,\theta) = \theta_1 \left(\Gamma(\theta_2)^{-1} \theta_3^{\theta_2} t^{\theta_2 - 1} e^{-\theta_3 t} \right), \tag{2.2}$$

where Γ is the gamma function, and θ is a parameter set that controls the shape of the HRF, with θ_1 controlling the height, and θ_2 and θ_3 controlling the delay and dispersion of the HRF. An example of what these look like under various settings for θ_2 and θ_3 is shown in Figure 2.1.

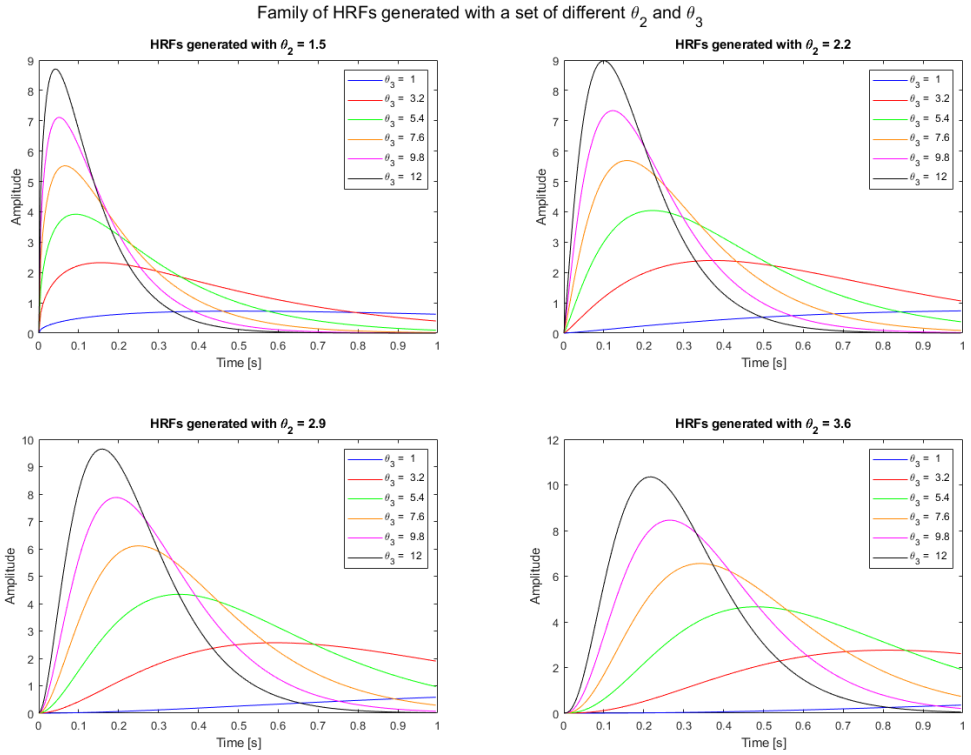

Figure 2.1: A family of HRFs under different conditions for θ_2 and θ_3. It can be seen that changing only a single parameter will affect both the delay and dispersion of the HRF.

The θ parameters of the HRF can vary depending on the input, meaning that when two different types of stimuli are provided, it also is necessary to find the parameters for the HRF for both stimuli.

3

Methodology

This chapter introduces the data model and how it is possible to recover the individual sources that make up the data through ICA. This is followed by an introduction of independent vector analysis (IVA) and its extension acIVA and how these improve upon the usage of ICA. The chapter ends by introducing the methods that perform automatic subgroup identification and an alternative method for subgroup identification that will also be included in the performance assessment.

3.1. The data model and ICA

The brain can be considered to consist of regions that perform specific tasks, these are sources that need to be extracted. However, when imaging the brain the signals measured are a combination of all these underlying sources, obfuscating which region is performing which task. The goal is to extract these underlying sources to see how every subject responds to specific stimuli. By assuming that these tasks are independent it is possible to perform ICA and extract independent sources within the brain. In this section, the fUS data is first expressed mathematically, which is followed by how ICA manages to extract these underlying sources.

Expressing the data mathematically is done as follows. For all K mice a fUS dataset is recorded that consists of V voxels measured at T samples, this data can be interpreted as being generated by N independent sources that are instantaneously mixed to obtain the observation $\mathbf{x}^{[k]}(t) \in \mathbb{R}^V$ measured at time t for subject k. This can mathematically be expressed as:

$$\mathbf{x}^{[k]}(t) = \mathbf{A}^{[k]}\mathbf{s}^{[k]}(t), \tag{3.1}$$

where $\mathbf{X}^{[k]} = \left[\mathbf{x}^{[k]}(1), \mathbf{x}^{[k]}(2), \cdots, \mathbf{x}^{[k]}(T)\right] \in \mathbb{R}^{V\times T}$, $\mathbf{S}^{[k]} = \left[\mathbf{s}^{[k]}(1), \mathbf{s}^{[k]}(2), \cdots, \mathbf{s}^{[k]}(T)\right] \in \mathbb{R}^{N\times T}$, and $\mathbf{A}^{[k]} \in \mathbb{R}^{V\times N}$ are the recorded data, the source matrix, and the spatial mixing matrix respectively. This allows for Equation 3.1 to be written in matrix form as:

$$\mathbf{X}^{[k]} = \mathbf{A}^{[k]}\mathbf{S}^{[k]}, \tag{3.2}$$

which is the data model for the fUS data. Finding $\mathbf{S}^{[k]}$ in this situation is defined as a BSS problem, which is where only $\mathbf{X}^{[k]}$ is known with the goal being to estimate $\mathbf{S}^{[k]}$. One popular method of solving the BSS problem is through the use of ICA which approximates $\mathbf{S}^{[k]}$ under the assumption that the sources are independent and that the mixing matrix does not change over time [10]. Which fits the data model described in Equation 3.2.

ICA finds these sources by finding an unmixing matrix $\mathbf{W}^{[k]}$ which is ideally the inverse of $\mathbf{A}^{[k]}$, such that when the data is multiplied with the unmixing matrix an estimate for the sources is returned as seen in the following equation:

$$\mathbf{W}^{[k]}\mathbf{X}^{[k]} = {\mathbf{A}^{[k]}}^{-1}\mathbf{A}^{[k]}\mathbf{S}^{[k]} = \hat{\mathbf{S}}^{[k]}, \tag{3.3}$$

with $\hat{\mathbf{S}}^{[k]}$ being the estimated source matrix, while subject to containing independent sources.

Calculating the unmixing matrix under the assumption of independent sources can be done in various ways, one way is through the minimization of mutual information. Minimizing the mutual information between two variables results in both variables not explaining one another making them independent. This is the approach of the InfoMax algorithm [11]. Another popular approach is that of fastICA which maximizes

negative entropy, with the idea that all source probability density functions should be as far from normally distributed as possible to achieve independence between the sources [12]. Once the unmixing matrix has been found and the independent time courses have been extracted it is possible to extract the spatial mixing matrix $\mathbf{A}^{[k]}$ through back-reconstruction [5].

In this way ICA can be applied to extract independent time series, usually called temporal ICA (tICA) [5]. However, ICA can also be applied to the transpose of the data matrix, resulting in the source matrix containing N spatially independent voxel maps. These voxel maps are spatially independent regions within a subject, this method is usually called spatial ICA (sICA). In this thesis, the temporal response to the stimulus is of interest, therefore the focus is placed on extracting independent time series through temporal analysis as these can be compared with the stimulus to see how the subjects respond to the stimulus.

Performing ICA on each subject separately will result in K independently estimated source matrices that are subject to permutation ambiguity. Theoretically, these can be identified manually but doing this by hand would require a lot of manual labour when the number of subjects and sources increase [6]. One way to solve this permutation ambiguity is by automatically organising the components with self-organized clustering as is done in self-organized ICA (sogICA) [13]. This method still performs ICA on each subject separately, but sogICA performs post-processing to automatically find similar sources across subjects and groups them according to similarity. While this method is capable of grouping components, it is not maximally using the available information. It is expected that the subjects behave similarly within a group and this can be leveraged by extending ICA to simultaneously analyze all subjects. Examining the subjects simultaneously instead of independently turns the BSS problem into a joint BSS (JBSS) problem, where dependencies across subjects are taken into account. By turning this into a JBSS problem this shared information can be used and allows for the grouping of the estimated sources during the component extraction without the need for post-processing. Performing this joint analysis requires changes to ICA which will be discussed in the next section.

3.2. Independent vector analysis

Solving JBSS can be done in various ways, one approach is to concatenate the datasets of all subjects either spatially or temporally and perform a single ICA on the resulting dataset [14]. This is called group ICA (GICA) which is a popular method for making group inferences in fMRI data [5]. However, by creating this single large dataset GICA assumes that all subjects share a common subspace, which may not be the case when two groups are analyzed. It is also possible to define the data as a tensor, taking the subjects as a third dimension, allowing for tensor ICA [15], where the data is a trilinear combination of three outer products that represent the group spatial maps, group time courses and subject loadings. These subject loadings represent how much each subject is represented by the same outer product of a group spatial map with a group time course. This method also assumes a common subspace for all subjects which again might not be the case. Therefore, another method is of interest when looking for subgroups, this method is called IVA. IVA does not assume a common subspace like GICA, instead promoting a common subspace while leaving room for individual variability [16], [17]. This makes it possible to capture differences within a single dataset containing both groups while also promoting the extraction of components common to all subjects, how IVA achieves this is further explained in this section.

For IVA to extract these common components, it is first necessary to go back to ICA. Where instead of doing ICA for each subject separately IVA concatenates the data matrices and assumes that the spatial mixing matrix for each subject can be unique, unlike GICA which assumes that the mixing matrix is the same for all subjects. This results in the following equation:

$$\begin{bmatrix} \mathbf{X}^{[1]} \\ \vdots \\ \mathbf{X}^{[K]} \end{bmatrix} = \begin{bmatrix} \mathbf{A}^{[1]} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}^{[K]} \end{bmatrix} \begin{bmatrix} \mathbf{S}^{[1]} \\ \vdots \\ \mathbf{S}^{[K]} \end{bmatrix} \Longleftrightarrow \mathbf{X} = \mathbf{AS}, \tag{3.4}$$

where $\mathbf{A}$ is now a block diagonal matrix, and the components within each $\mathbf{S}^{[k]}$ are still independent. To promote similarity across subjects, IVA allows for dependence across $\mathbf{S}^{[k]}$ by adding a construct called the source component vector (SCV), which identifies similar components across subjects and groups them during optimization. Let $\mathbf{s}^{[k]}(t) = \left[s_1^{[k]}(t), \cdots, s_n^{[k]}(t), \cdots, s_N^{[k]}(t) \right]^T$ then the SCV for source n is defined as:

$$\mathbf{s}_n(t) = \left[s_n^{[1]}(t), \cdots, s_n^{[K]}(t) \right]^T, \tag{3.5}$$

taking the T SCVs corresponding to a component n and stacking them into a matrix as follows:

$$\mathbf{S}_n = [\mathbf{s}_n(1), \mathbf{s}_n(2), \cdots, \mathbf{s}_n(T)], \tag{3.6}$$

results in a source component matrix (SCM). Optimizing the dependence within the SCVs will result in similar sources being grouped in the same SCV. This dependence is enforced by maximizing the mutual information within each SCV. This way the n-th SCV will compare the n-th source across all subjects, and if all subjects share this source it will be visible through the SCV. Below is a visualisation of IVA with the SCMs Figure 3.1 taken from [18].

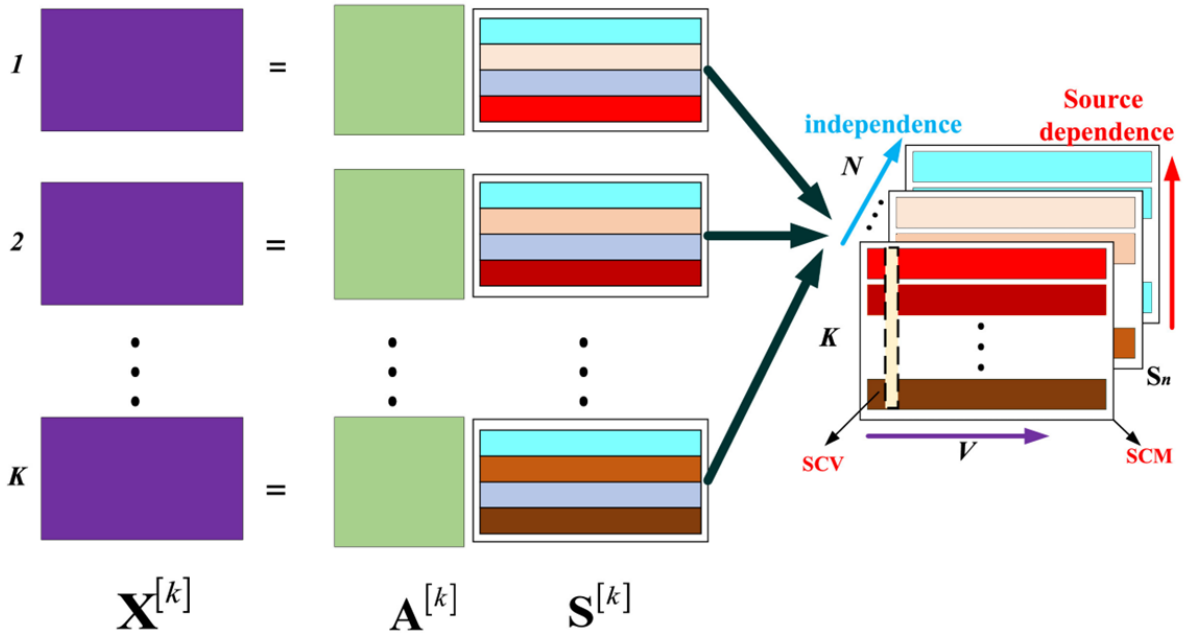

Figure 3.1: Visual representation of the IVA process, with each colour being an independent source. The SCMs try to maximize the dependence within themselves, grouping similar sources across subjects.

To get the unmixing matrix $\mathbf{W}$ a cost function is minimized. Taking the approach of minimization of mutual information to get independent sources and applying it to the estimated sources $[\mathbf{s}_1, \mathbf{s}_2, \cdots, \mathbf{s}_N]$ gives the following cost function:

$$\mathscr{I}_{IVA}(\mathbf{W}) = I[\mathbf{s}_1, \mathbf{s}_2, \cdots, \mathbf{s}_N] = \sum_{n=1}^{N} H[\mathbf{s}_n] - H[\mathbf{s}_1, \mathbf{s}_2, \cdots, \mathbf{s}_N] \tag{3.7a}$$

$$= \sum_{n=1}^{N} H[\mathbf{s}_n] - H\left[\mathbf{W}^{[1]}\mathbf{x}^{[1]}, \mathbf{W}^{[2]}\mathbf{x}^{[2]}, \cdots, \mathbf{W}^{[K]}\mathbf{x}^{[K]}\right] \tag{3.7b}$$

$$= \sum_{n=1}^{N} H[\mathbf{s}_n] - \sum_{k=1}^{K} \log\left|\det\left(\mathbf{W}^{[k]}\right)\right| - C, \tag{3.7c}$$

where I is the mutual information, $H(\mathbf{s}_n)$ the differential entropy, and C the constant $H\left[\mathbf{x}^{[1]}, \mathbf{x}^{[2]}, \cdots, \mathbf{x}^{[K]}\right]$, where Equation 3.7c is obtained through the entropy of the linear transformation $\mathbf{Wx}$ also being defined as $\log|\det(\mathbf{W})| + H[\mathbf{x}]$. Furthermore, the differential entropy $H[\mathbf{s}_n]$ can be rewritten into $\sum_{k=1}^{K} H\left[s_n^{[k]}\right] - I(\mathbf{s}_n)$, with $I(\mathbf{s}_n)$ being the mutual information within the n-th SCV. Substituting this into Equation 3.7c gives the following cost function for IVA:

$$\mathscr{I}_{IVA}(\mathbf{W}) = \sum_{n=1}^{N}\left[\sum_{k=1}^{K} H\left[s_n^{[k]}\right] - I(\mathbf{s}_n)\right] - \sum_{k=1}^{K} \log\left|\det\left(\mathbf{W}^{[k]}\right)\right| - C, \tag{3.8}$$

which shows that minimizing the cost function simultaneously optimizes for maximum mutual information within each SCV while also keeping the components independent within a subject.

This shows that IVA does not enforce a common subspace for each component, instead trying to promote it while allowing for subject variability within each component. This allows IVA to capture both groups without them interfering with each other. This does come at a computational cost due to a larger number of parameters that must be estimated with $\mathbf{W}$. Furthermore, while the components are grouped according to similarity, it is still required to manually identify which component is related to the stimuli. To do this automatically acIVA is used which will be discussed in the following subsection.

Adaptive constrained IVA

During the experiment one or two stimuli are provided, making the components related to these stimuli of special interest. Unfortunately, it is necessary to identify the components related to the stimuli through some form of post-processing as the components are subject to permutation ambiguity each time IVA is run again. This ambiguity can be removed by specifically targeting the extraction of the components related to the stimuli. This can be done through adaptive constrained IVA acIVA, which targets references (constraints) during the IVA analysis [19].

Let L be the number of references that are expected in the data, then $\mathbf{d}_l$ is the l-th reference signal. To target these during extraction a constraint is added to the cost function in Equation 3.8, resulting in the new IVA cost function being:

$$\mathscr{J}_{\text{acIVA}} = \mathscr{J}_{\text{IVA}} - \sum_{l=1}^{L} \frac{1}{2\gamma_l} \times \sum_{k=1}^{K} \left\{ \left[\max\left\{0, \mu_l^{[k]} + \gamma_l g\left(\hat{\mathbf{s}}_l^{[k]}, \mathbf{d}_l\right)\right\}\right]^2 - \left(\mu_l^{[k]}\right)^2 \right\}, \tag{3.9}$$

where γ_l is the penalty parameter, $\mu_l^{[k]}$ the Lagrange multiplier, and $g(\hat{\mathbf{s}}_l^{[k]}, \mathbf{d}_l)$ is the inequality constraint which is defined as:

$$g\left(\hat{\mathbf{s}}_l^{[k]}, \mathbf{d}_l\right) = \rho_l^{[k]} - \epsilon\left(\hat{\mathbf{s}}_l^{[k]}, \mathbf{d}_l\right) \leq 0, \tag{3.10}$$

with ρ_l being the constraint parameter, and $\epsilon(\hat{\mathbf{s}}_l^{[k]}, \mathbf{d}_l)$ is some function that measures similarity between the reference and estimated component. In this thesis, ϵ is chosen to be the absolute correlation coefficient. Which results in the cost function being that of constrained IVA [20].

In constrained IVA the selection of ρ_l is of large importance as it heavily impacts the estimated sources. Selecting a large value for ρ_l results in the estimated source being forced to be the reference, leaving no room for variability, inversely setting ρ_l to be too low could cause the estimated source to not be related to the reference at all. By making this ρ_l adaptive, constrained IVA becomes acIVA. acIVA uses a set of possible values (ρ_i) with $i \in \mathbb{N}$ that is given for $\rho_l^{[k]}$ in ascending order, during gradient descent for each step the highest value of ρ is found that satisfies Equation 3.10, this is then taken to be $\rho_l^{[k]}$ from which the Lagrange multiplier is computed and a new gradient is found. This means that the cost function in Equation 3.9 is used for the L components linked to a reference, and the cost function in Equation 3.8 is used for $N-L$ components not linked to a reference.

Making the constraint parameter adaptive allows IVA to be guided to the solution without incorrectly enforcing the reference. On top of that it can be seen from Equation 3.9, that acIVA improves interpretability since the first L components are not subject to permutation ambiguity, instead being linked to the first L references. For example, adding a single reference related to a puff stimulus would mean that only the first component has to be checked to see if the subjects react to the puff stimulus. By checking the first SCM it can then be directly seen which subjects react to the first reference.

3.3. Subgroup identificaiton

Analyzing the SCMs can now give insights into how subjects react to a stimulus. By applying clustering it may even be possible to detect specific patterns in how subjects react, in turn allowing for clustering of these subjects. This is what the two subgroup identification methods using IVA have proposed. These methods are called: IVA for common subspace analysis (IVA-CS) and subgroup identification using IVA (SI-IVA) [21], [22]. Both subgroup identification methods rely on the covariance matrix of the SCMs to perform clustering. Therefore, the covariance matrix is introduced first, this is followed by a discussion of the proposed methods and one of the other clustering methods that has been tested.

The covariance matrix of each SCM is of interest as it can highlight which subjects share which component. One can compute the covariance matrix $\mathbf{C}_n$ for SCM $\mathbf{S}_n$ as follows:

$$\mathbf{C}_n = (1/T)\mathbf{S}_n\mathbf{S}_n^T, \tag{3.11}$$

where the covariance matrix is equal to the correlation matrix when the SCM is standardized beforehand. This way the (i, j)-th element of $\mathbf{C}_n$ highlights the correlation between subjects i and j for source n. A visual representation of this is shown in Figure 3.2. Note that the correlation coefficient is extended to be the absolute value of the correlation coefficient, this is done to remove the effects of sign ambiguity from IVA which can influence the clustering.

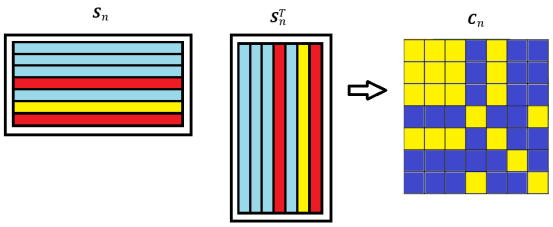

Figure 3.2: The covariance matrix $\mathbf{C}_n$ will show which subjects are correlated with one another. Here the colours within the SCM $\mathbf{S}_n$ highlight the independent sources. Resulting in the covariance matrix having high correlations (yellow) at position (i, j) when i and j have the same source and low correlations (blue) when i and j have different sources. For example, subjects 1, 2, 3, and 5 have the blue component within $\mathbf{S}_n$, resulting in $\mathbf{C}_n$ having columns 1, 2, 3, and 5 being similar.

These covariance matrices are nice for clustering due to the covariance matrix quickly showing which subjects share a component and which do not. IVA-CS and SI-IVA also use these covariance matrices to identify subgroups, by either interpreting it as a weighted adjacency matrix where a subgroup is a cluster in the graph or taking each column (or row) as a data vector and trying to group these through clustering methods like k-means. Both methods will be discussed in the following subsections, starting with SI-IVA.

3.3.1. Subgroup identification using IVA

SI-IVA clusters the subjects by performing k-means clustering for each covariance matrix, however, this means that the number of subgroups M within the covariance matrix has to be determined beforehand. Selecting the number of subgroups for k-means is usually done through heuristic methods like the elbow or silhouette methods. However, instead of employing these heuristics, SI-IVA proposes a new method of finding M using the Gershgorin Disc theorem [22], [23].

The Gershgorin disc theorem bounds eigenvalues of a square matrix inside discs, these discs are defined as follows: The radius $R_i = \sum_{j \neq i} \left| c_n^{[i,j]} \right|$ of disc i is the sum of the absolute values of the non-diagonal entries of the i-th row of $\mathbf{C}_n$, with the disc being centered at $c_n^{[i,i]}$. Resulting in the gershgorin discs being defined as $\left\{ z \in \mathbb{R} : |z - c_n^{[i,j]}| \leq R_i \right\}$. Each Gershgorin disc contains one eigenvalue of $\mathbf{C}_n$. Applying this to the covariance matrix gives that all discs are centered at $(1, 0)$ since the diagonal is all ones and all eigenvalues should be within the largest disc. Let $R_{\min} = \min_i R_i$ be the radius of the smallest disc of $\mathbf{C}_n$, then according to SI-IVA, the number of eigenvalues outside of the smallest disc is the number of subgroups in $\mathbf{C}_n$. The reasoning behind finding the subgroups this way is that the smallest disc represents an energy threshold and the eigenvalues that surpass this threshold represent clusters whose energy exceeds the threshold of this smallest disc. A visual representation of this process is shown in Figure 3.3 taken from [22].

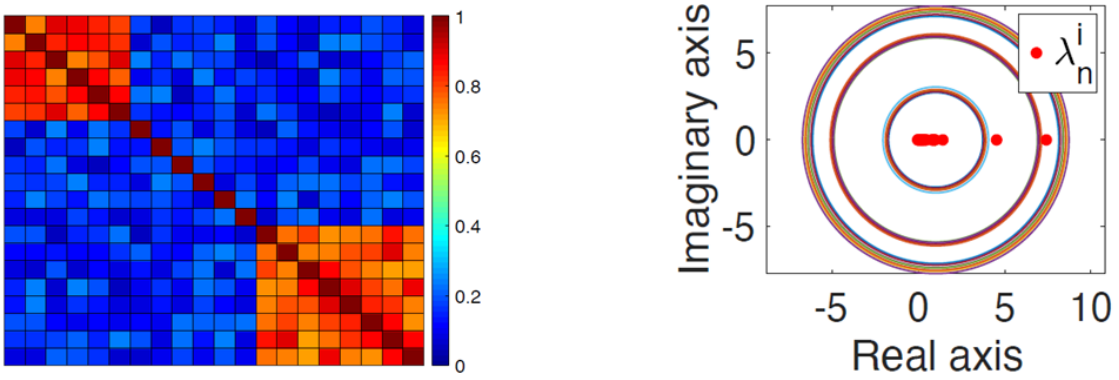

Figure 3.3: SI-IVA applied to the covariance matrix shown on the left which contains two subgroups. On the right, the Gershgorin discs are drawn and the eigenvalues are plotted, showing that two eigenvalues cross the threshold which agrees with the shown covariance matrix. These images are taken from [22]

While finding the number of subgroups SI-IVA stores the eigenvectors $\mathbf{u}_m$ of the eigenvalues that cross the threshold. This is done because they are used for clustering. SI-IVA assumes that each eigenvector represents a single subgroup, meaning that by segmenting each eigenvector into two groups using k-means will return the subgroups accordingly.

The authors give no bounds on the accuracy of SI-IVA or when it will work. However, in this thesis, a theoretical analysis was performed showing that SI-IVA will always identify there being one subgroup, even in the case when there is none. This analysis can be found in Appendix A. This analysis does not show if SI-IVA can correctly estimate the subgroups in the case that there are one or more subgroups in $\mathbf{C}_n$.

3.3.2. IVA for common subspace identification

IVA-CS uses the covariance matrix to identify subgroups but does not use k-means, instead opting for modularity maximization of the covariance matrix. Modularity is a graph theory measure that helps identify communities within a network [24]. In this case, the covariance matrix is interpreted as a weighted adjacency matrix where each subject is now a vertex and their correlations are the connecting edges. The modularity Q is given by summing over $\mathbf{C}_{ij} - k_i k_j / 2m$ for all pairs of subjects i, j that are considered to be in the same subgroup, where k_i is the degree of vertex i, and m is the total degree of the covariance matrix.

This can be expressed mathematically through the modularity matrix $\mathbf{B}$ whose elements are defined as:

$$\mathbf{B}_{ij} = \mathbf{C}_{ij} - k_i k_j / 2m, \tag{3.12}$$

which gives that the modularity Q can be defined as:

$$Q = \mathbf{s}^T \mathbf{B} \mathbf{s}, \tag{3.13}$$

where $\mathbf{s}$ is a column vector of 1s or -1s, where the elements indicate to which group a subject belongs such that if $s_i = 1$ it means that subject i belongs to group 1 and $s_i = -1$ if subject i belongs to group 2. This shows that if two unrelated subjects are grouped a negative value is added to the modularity due to them having a low value for $C_{i,j}$. Therefore, the goal of modularity maximization is to find the correct groupings to maximize the modularity Q.

It is possible to express $\mathbf{s}$ as the linear combination of the normalized eigenvectors $\mathbf{u}_i$ of $\mathbf{B}$, such that $\mathbf{s} = \sum_{i=1}^{N} a_i \mathbf{u}_i$ and $a_i = \mathbf{u}_i^T \cdot \mathbf{s}$. Plugging this into Equation 3.13 gives:

$$Q = \sum_{i=1}^{N} (\mathbf{u}_i^T \cdot \mathbf{s})^2 \lambda_i, \tag{3.14}$$

with λ_i being the eigenvalue related to eigenvector $\mathbf{u}_i$. It can be seen, that by sorting the eigenvalues from largest to smallest $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$ it will follow that $\mathbf{s}$ should be chosen to be the principal eigenvector λ_1 resulting in a maximization of $\mathbf{Q}$. However, the principal eigenvector does not consist of only ± 1 so to allow for groupings all non-negative values in $\mathbf{s}$ are set to $+1$ and all negative values are set to -1. To see if this is the best segmentation of $\mathbf{C}_n$, all subjects will change their group once to see if the modularity is higher by doing

so. This can also result in the modularity showing that all subjects belong to the same group as long as Q is at its maximum when $\mathbf{s} = \underline{1}$. If a segmentation is done, it is possible to apply modularity on the segmented matrices again, this can be done recursively until the modularity is maximized. Showing that modularity does not need to know the number of subgroups beforehand.

3.3.3. Spectral clustering

With both SI-IVA and IVA-CS discussed, the methods from literature have been covered. However, it can be seen that the subgroup identification is mainly a clustering or community detection problem. Therefore, other existing clustering methods can also be used for subgroup identification. The tested clustering methods were kept simple with k-means on the complete covariance matrix and spectral clustering. Spectral clustering will shortly be described in this subsection.

Spectral clustering makes use of a lower-level representation of the data to effectively cluster the data with k-means. By representing the data through a limited set of eigenvectors spectral clustering allows itself to cluster datasets where k-means would fail [25]. The spectral clustering process goes as follows: the i-th column $\mathbf{c}_i$ of the covariance matrix will be considered a data point for subject i, with which an affinity matrix $\mathbf{A}$ can be constructed whose elements are defined as:

$$A_{ij} = \begin{cases} \exp(-\left|\left|\mathbf{c}_i - \mathbf{c}_j\right|\right|^2 / 2\sigma^2) & \text{if } i \neq j \\ 0 & \text{if } i = j \end{cases}, \tag{3.15}$$

where σ is a scaling parameter of the affinity drop-off. Let $\mathbf{D}$ be a diagonal matrix where the (i,i)-th element is the sum of $\mathbf{A}$'s i-th row, then construct a matrix $\mathbf{L}$ as follows:

$$\mathbf{L} = \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}, \tag{3.16}$$

now take the k largest eigenvectors (with k being the number of subgroups) $\mathbf{u}_1, \mathbf{u}_2, \cdots, \mathbf{u}_k$ of $\mathbf{L}$ and place them in a new matrix $\mathbf{X} = [\mathbf{u}_1, \mathbf{u}_2, \cdots, \mathbf{u}_k] \in \mathbb{R}^{N \times k}$. After normalizing the rows of this matrix k-means will be performed where now the points for each subject are the rows of this $\mathbf{X}$ matrix.

In the ideal case, the affinity is 1 between subjects part of the same subgroup and 0 when they are not part of the same subgroup, this will result in $\mathbf{A}$ being block diagonal which in turn makes $\mathbf{L}$ block diagonal. Taking the k largest eigenvectors of this block diagonal matrix consisting of k blocks will result in $\mathbf{X}$ also being block diagonal. For example, in the case of there being 3 subgroups, $\mathbf{L}$ will consist of 3 blocks and each eigenvector contains non-zero entries for the respective block. Stacking this into $\mathbf{X}$ and normalizing the rows results in the following:

$$\mathbf{X} = \begin{bmatrix} \underline{1} & \underline{0} & \underline{0} \\ \underline{0} & \underline{1} & \underline{0} \\ \underline{0} & \underline{0} & \underline{1} \end{bmatrix}, \tag{3.17}$$

with $\underline{0}$ being the zero vector and $\underline{1}$ being the ones vector. Performing k-means on this new matrix will cluster each subgroup due to each subject only having a non-zero entry for their respective subgroup. In real data, the matrix is not block diagonal but the affinity matrix should still signify the subgroups, in turn meaning that the eigenvectors also represent the subgroups. As with SI-IVA, it is still necessary to determine the number of subgroups. For this, the elbow method is applied on the eigenvalues of the covariance matrix. This is done as the eigenvalues give a clear cutoff point as to how many eigenvectors are required to explain the covariance matrix, which should be close to the number of subgroups.

4

Simulation results

From the methods discussed no performance guarantees were given on how well they would be able to correctly identify subgroups from the SCMs. Therefore, before going to real data and letting one of these methods automatically identify subgroups it is first necessary to see which method is most robust and accurate. To do this a large number of covariance matrices will be generated on which these methods will be evaluated. This chapter will highlight this process, starting with how the covariance matrices are simulated, followed by the performance results of these methods, after which a discussion will follow that explores and explains these results.

4.1. Simulating covariance matrices

To test the accuracy of these methods it is necessary to simulate a large number of random covariance matrices. A covariance matrix can take a large number of appearances depending on a number of variables that are linked to the data, these are: The number of subjects K, the number of subgroups M, the minimum number of subjects in a subgroup n, the distribution of correlations between the subjects in the same subgroup $f_{\text{subgroup}}(x)$, and the distribution of correlations between subjects that are not part of the same subgroup $f_{\text{noise}}(x)$. Using these parameters, a random covariance matrix can be simulated.

The following pipeline generates the covariance matrices: First, each subject is randomly assigned to be part of a subgroup or part of none, resulting in $M+1$ sets of subjects. Where M sets have the subjects that correlate within their subgroup, and the last set is the group of subjects with no correlations with any other subject. With these subgroups defined, it is possible to simulate a covariance matrix $\mathbf{C}$, this is done by first generating a lower triangular matrix whose elements are given by:

$$\mathbf{C}_{\text{tril},ij} = \begin{cases} f_{\text{subgroup}}(x) & \text{if } i < j \text{ and } i,j \text{ are in the same subgroup} \\ f_{\text{noise}}(x) & \text{if } i < j \text{ and } i,j \text{ do not share a subgroup} \\ 0 & \text{otherwise} \end{cases}, \tag{4.1}$$

which gives a simulated covariance matrix after the following:

$$\mathbf{C} = \mathbf{C}_{\text{tril}} + \mathbf{C}_{\text{tril}}^T + I, \tag{4.2}$$

with I being the identity matrix of size $K \times K$. The transpose ensures the simulated covariance matrix is mirrored along the diagonal, as the covariance matrix is also always mirrored along the diagonal.

The distributions of the correlations in and outside of subgroups can be freely chosen, examples would be that of a truncated normal distribution or a uniform distribution. For these simulations, the uniform distribution is used. An example of the random matrix generation is shown in Figure 4.1, where two matrices are generated with the same hyperparameters, from this figure it can also be seen how the larger yellow squared are the subgroups that need to be identified.

4.2. Results

To test the accuracy, 10000 matrices are randomly generated given a set of hyperparameters. To highlight how performance changes when changing a parameter 6 cases will be examined. One set of hyperparameters is

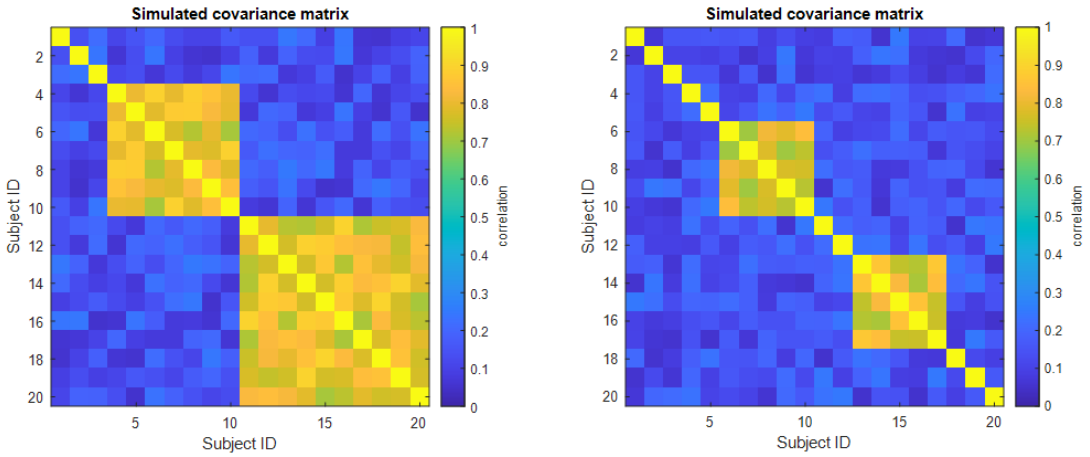

Figure 4.1: An example of two simulated covariance matrices with the same simulation parameters: $K = 20$ subjects, $L = 2$ subgroups, $m = 5$ minimum size of each subgroup, $f_{\text{subgroup}}(x) = U[0.7, 0.9]$, $f_{\text{noise}}(x) = U[0.05, 0.25]$. The two matrices have different subgroups due to the random initialization. The left matrix has a subgroup from subjects 4 to 10 and a subgroup from subjects 11 to 20. The right matrix has a subgroup from subjects 6 to 10 and 13 to 17.

the baseline, the remaining 5 cases will have only one parameter changed with respect to this baseline. The chosen parameters for each case are shown in Table 4.1.

Table 4.1: Table containing the hyperparameters used in each experiment, with the first case being the baseline and the second till the sixth cases having one hyperparameter changed with respect to the baseline.

	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6
Number of subjects K	30	20	30	30	30	30
Number of subgroups M	2	2	1	2	2	2
Minimum size n of subgroup	5	5	5	8	5	5
$f_{\text{subgroup}}(x)$	$U[0.7, 0.9]$	$U[0.7, 0.9]$	$U[0.7, 0.9]$	$U[0.7, 0.9]$	$U[0.5, 0.6]$	$U[0.7, 0.9]$
$f_{\text{noise}}(x)$	$U[0.05, 0.25]$	$U[0.05, 0.25]$	$U[0.05, 0.25]$	$U[0.05, 0.25]$	$U[0.05, 0.25]$	$U[0.1, 0.4]$

The results of these tests are shown in Table 4.2.

Table 4.2: The accuracy of the tested clustering methods in the 6 cases as shown in Table 4.1. Spectral clustering performs best from these tested methods, noteworthy is SI-IVAs bad performance which can be linked to the number of subgroups being above 1.

	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6
SI-IVA	2.72%	8.79%	100%	5.8%	0.3%	0.08%
IVA-CS	74.5%	92.38%	98.95%	93.09%	65.8%	52.01%
Spectral clustering	99.73%	99.88%	99.96%	100%	99.35%	99.27%
K-means	98.68%	98.37%	99.95%	99.27%	91.37%	97.06%

These results show spectral clustering to be a robust and reliable method for this range of parameters. More cases are tested to verify the robustness of spectral clustering in a wider array of scenarios. These parameters are shown in Table 4.3 and the corresponding results in Table 4.4.

4.3. Discussion

From the results, it is clear that spectral clustering performed best out of the compared methods but it can not guarantee high accuracy for every type of scenario (see cases 8 to 10). The methods from the literature generally perform worse and rarely on par with spectral clustering, even k-means outperforms the proposed methods.

Table 4.3: The extended cases tested to assess the robustness of spectral clustering that vary more from the baseline than previously tested,

	Case 7	Case 8	Case 9	Case 10
Number of subjects K	10	40	20	40
Number of subgroups M	1	2	3	3
Minimum size n of subgroup	3	8	3	5
$f_{\text{subgroup}}(x)$	$U[0.6, 0.9]$	$U[0.6, 0.9]$	$U[0.7, 0.9]$	$U[0.5, 0.9]$
$f_{\text{noise}}(x)$	$U[0.05, 0.4]$	$U[0.05, 0.5]$	$U[0.05, 0.4]$	$U[0.05, 0.4]$

Table 4.4: Results for the extra cases tested, it is found that SI-IVA is accurate when there is only one subgroup. Furthermore, spectral clustering shows signs of robustness but many small clusters can lower its accuracy.

	Case 7	Case 8	Case 9	Case 10
SI-IVA	96.92%	0.03%	0%	0%
IVA-CS	96.9%	29.47%	20.31%	5.8%
Spectral clustering	99.05%	98.83%	87%	90.21%
K-means	88.17%	95.51%	49.45%	73.93%

The bad performance of SI-IVA can largely be attributed to the wrong prediction of the true number of subgroups which in turn causes the clustering to fail. However, this is not the only reason SI-IVA fails, performing clustering on the eigenvectors one by one also creates issues, as this can cause the two subgroups to become a single large one. As is done in spectral clustering, taking more eigenvectors and performing k-means simultaneously may solve this issue. The results further suggest that SI-IVA is only capable of determining the correct number of subgroups in a situation where there is only one subgroup. To confirm this further numerical experiments were performed to see when SI-IVA correctly determines the correct number of subgroups. For this experiment, 1000 matrices are generated with the same distribution as in case 1 but the number of subjects and subgroups are being changed. This gives the results shown in Table 4.5.

Table 4.5: Results showing the accuracy of SI-IVA when it comes to determining the number of subgroups in a covariance matrix. It can be seen that for more than 1 subgroup SI-IVA tends to underestimate the number of subgroups unless the subgroups are relatively large and cover a large part of the covariance matrix. The average estimated subgroups also highlight that SI-IVA estimates there to be at least one subgroup even when there are zero.

Number of subjects	Minimum subjects per subgroup	True number of subgroups	Accuracy	Average estimated subgroups
10	5	0	0 %	1
10	5	1	100 %	1
20	5	0	0 %	1
20	5	1	100 %	1
20	5	2	88.9 %	1.889
20	5	3	6.5 %	1.6110
30	5	0	0 %	1
30	5	1	100 %	1
30	5	2	30.8 %	1.308
30	5	3	0.1 %	1.322
30	5	4	0 %	1.105

These results numerically confirm that SI-IVA estimates the number of subgroups to be 1 no matter the covariance matrix. These results also suggest that SI-IVA can be accurate when the subgroups are relatively large and the number of subgroups small. Since the accuracy is higher when the covariance matrix is smaller. SI-IVA is not accurate when the subgroups are small since in that case the second-largest subgroup does not have an eigenvalue that passes the threshold set by SI-IVA. To pass this threshold it is required that this second subgroup has a large enough size to pass this energy threshold. This experiment also shows that the

clustering approach from SI-IVA fails, as case 1 is repeated in this experiment. Table 4.5 shows that 30% of the time SI-IVA has the correct number of subgroups estimated in case 1. This does not align with the 3% accuracy seen in Table 4.2, this discrepancy is explained by the bad-performing clustering. Showing that SI-IVA is not a robust approach for subgroup identification.

The performance of IVA-CS is better compared to SI-IVA, however, it is not great. In Figure 4.2, an example is shown where IVA-CS is incorrect, with the left showing the original covariance matrix and the right showing the identified subgroups from IVA-CS through alpha maps.

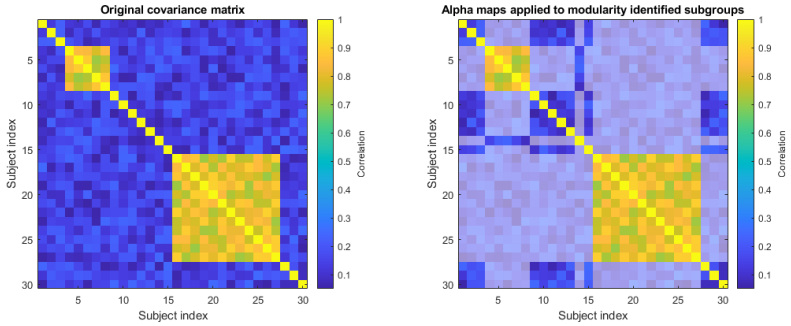

Figure 4.2: An example covariance matrix from case 1 where subgroups found by IVA-CS are highlighted in the right image, the method correctly finds the larger subgroup, yet it includes subject 14 with the smaller subgroup while this subject does not correlate with said subgroup

From Figure 4.2, it can be seen that modularity incorrectly includes a subject into a smaller subgroup. Analyzing the principal eigenvector in this case and setting it to ± 1 would give the correct subgroups, but changing the sign of subject 14 results in higher modularity, therefore it is included with the subgroup ranging from 4 to 8. This is not the only error that IVA-CS makes, it also happens that modularity partitions a correct subgroup into smaller ones due to there being a higher modularity possible if this large subgroup is segmented further. The inverse can also happen where the method stops too early as partitioning a small subgroup is not seen to be effective under modularity maximization. These cases make modularity and therefore IVA-CS unfit for automatic detection of subgroups since these errors happen at both extremes, which makes it hard to resolve this with further extensions. It is also questionable to use modularity to find subgroups in this case, this is because modularity is originally intended for connected unweighted graphs, yet here IVA-CS assumes the covariance matrix to be a fully connected weighted graph.

Going back to the original results shows that spectral clustering and k-means are better suited for this subgroup detection when compared to the methods that have been proposed in literature. This does not mean that they are perfect, especially case 9 proved difficult for all methods, in Figure 4.3 a situation is shown where both k-means and spectral clustering fail in identifying the correct subgroups in case 9.

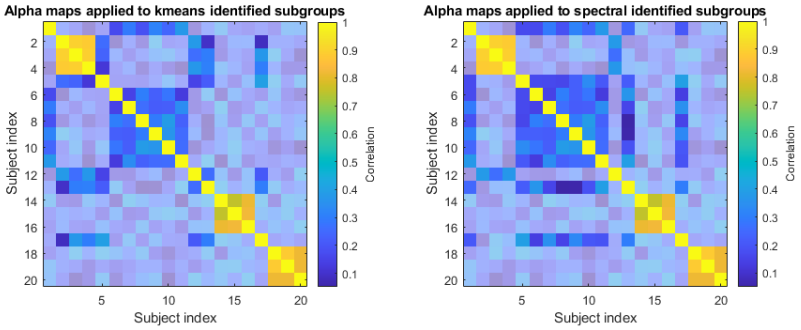

Figure 4.3: K-means and spectral clustering performed on the same covariance matrix in case 9. On the left, the subgroups from k-means are highlighted showing that subjects 5, 13, 14, and 17 are included with the group from 2 to 4. On the right are the subgroups found by spectral clustering, showing only subject 13 to be included with this small subgroup.

The results with case 9 can possibly be ignored due to the relatively small subgroup sizes. However, these results also highlight the rare errors of these methods in the other cases and are therefore further analyzed. The error seen in Figure 4.3, is that extra subjects outside of a subgroup are added to a subgroup. This addition of these subjects outside of the subgroup happens due to random initialization of the centroids of k-means. It is prone to fail when two centroids are initialized close to the origin causing the uncorrelated subjects to be partitioned instead of the subgroups. Even when using spectral clustering this appears to happen in rare cases as is seen in Figure 4.3, but this cause of error is more common in k-means which is why it performs worse when compared to spectral clustering. Since spectral clustering shows a better accuracy and is more robust towards the previously discussed errors it is the obvious choice to use for automatic subgroup identification in the rest of this thesis.

5

Experimental results

The simulation results highlight that spectral clustering is the preferred method for automatic subgroup identification, therefore, this method will be used when examining experimental fUS data. This chapter will explain the fUS data used, the pre-processing of this data, how the references are selected for acIVA, and eventually the results from applying acIVA and spectral clustering.

5.1. The data

There are two fUS datasets used in this thesis, both are recorded by the CUBE research group from Erasmus. These datasets consist of two groups of genetically different mice nine in total. These mice are subject to the same stimuli while being recorded. In the first dataset, the mice are subject to two stimuli. The first stimulus is an LED light flashing into the mouse's left eye for a short moment. The second stimulus in this experiment is a short air puff that is applied to the left eye a moment after the LED has been turned off. This pattern is repeated 20 times during a recording, and in Figure 5.1 this pattern is visually shown.

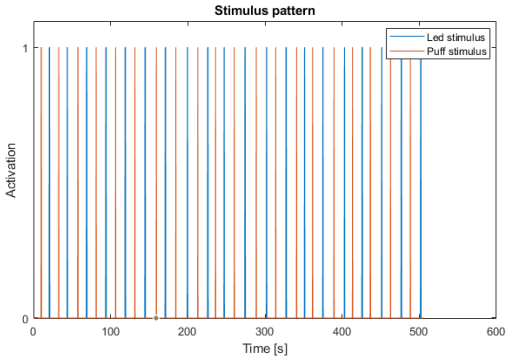

Figure 5.1: The stimulus pattern applied in the first experiment, with the LED and puff stimulus alternating.

For the second experiment, only the puff stimulus is applied, however, the recordings are made over several days, allowing the mice to learn the pattern. By analyzing these datasets it is possible to see if the mice react differently with respect to the stimuli and if they react differently over time due to them learning the stimulus pattern. The results will be related to the first experiment unless stated otherwise.

During recording 4 slices are recorded in a sagittal plane. In Figure 5.2 the sagittal plane is shown, with an image taken from [26]. Both halves of the brain will therefore have 2 slices imaged.

Since each slice is sagittal it results in the PDIs being stored as $(X \times Z)$, where x and z define the location in the X and Z dimension, on top of this the 4 slices can be considered to be in the Y dimension. Since each recording consists of 2023 compounded PDIs it gives that the data can be described as $\mathbf{X}^{[k]} \in \mathbb{R}^{X \times Z \times Y \times T}$ resulting in $\mathbf{X}^{[k]}$ being a 4-D array, with T being the number of samples (compounded images).

For interpretability, a background map is created for each mice, to help with the identification of the functional regions of each mouse. This background map will be created using an averaging of the data over

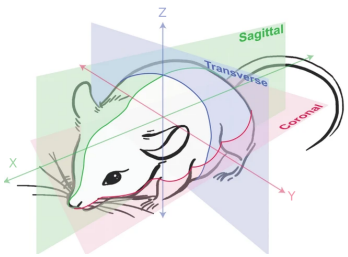

Figure 5.2: The three planes used in anatomy. In this thesis, the dataset used consists of images in the sagittal plane, meaning that the mice can be split into a left and right portion. This image is taken from [26].

time to highlight regions with high intensity. This background map, which is created for each slice, is a 3-D array $\mathbf{Z}^{[k]} \in \mathbb{R}^{X \times Z \times Y}$ which has entries defined as follows:

$$\mathbf{Z}^{[k]}_{x,z,y} = \frac{\sum_{t=1}^{T} \log\left(\mathbf{X}_{x,z,y,t}\right)}{T}, \tag{5.1}$$

such that each voxel is the mean of that specific voxel over the length of the recording. An example of such a background map is shown in Figure 5.3, where the white indicates areas with a high average intensity. Slice 1 is the leftmost slice and slice 4 is the rightmost slice.

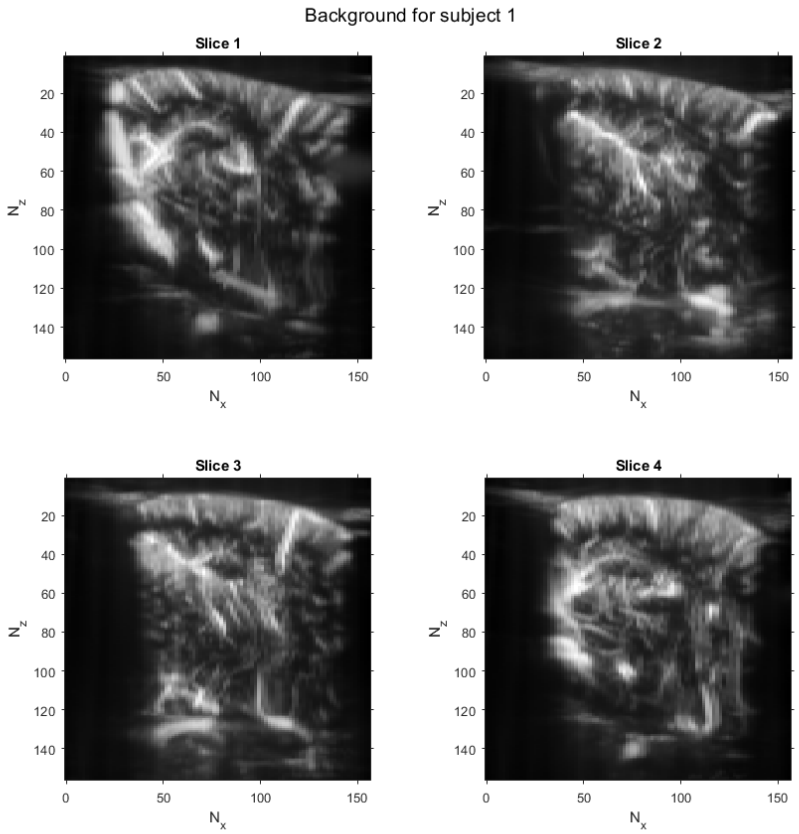

Figure 5.3: The background maps for subject 1. The white areas show vessels and functional areas where CBV changes are more common. Slice 1 is the leftmost slice and slice 4 is the rightmost slice.

5.2. Pre-processing

The raw PDI data could theoretically be directly used with the acIVA algorithm, however, the compounded images are still subject to noise and these images capture more than just the brain which can interfere with and slow down the processing. Therefore, to speed up the process and reduce the effect of noise certain pre-processing steps are applied to the PDI data.

The first pre-processing step is the reduction of the data that is processed. As can be seen in Figure 5.3 there are large black bars on both sides of the image, these are areas that are outside of the brain and therefore are not of interest. This means that especially the X dimension of the data can be reduced, as there are no black bars at the top or bottom. Cutting of the black bars results in X being reduced by a quarter for each slice, the Z dimension is left unchanged.

After the data reduction, the data is smoothed in both the time and spatial dimensions, this is done to reduce the effect of noisy spikes in the data and make the active regions require that the surrounding pixels share a similar time course. This makes it harder for single pixels with outliers to become part of active regions as the smoothing causes the surrounding pixels to dampen the outlier. This way peaks related to the stimuli are not removed since the surrounding regions will behave similarly. This smoothing is done with a Gaussian kernel of size 5 x 5 x 5 and performed on each slice separately, meaning that the kernel is in $(X \times Z \times T)$. The Y dimension is not taken into account for the smoothing because the slices are too far apart to perform smoothing.

This concludes the major data cleaning steps are performed, but to perform acIVA it is necessary to have a matrix instead of a 4-D array. To achieve this the spatial dimensions have to become one, this is done by first vectorizing each slice, resulting in the data being in the form $\mathbf{X}^{[k]} \in \mathbb{R}^{XZ \times Y \times T}$. This is followed by concatenating the vectorized slices under one another resulting in the final data format being $\mathbf{X}^{[k]} \in \mathbb{R}^{XZY \times T}$. This makes it possible to perform acIVA and get an unmixing matrix and source matrix back.

Since XZY is a large number processing this with acIVA will take an enormous amount since acIVA will try and extract XZY components. This is not only computationally very expensive but also unlikely to be necessary as it is expected that the number of underlying sources (components) is no more than 50. Therefore, dimensionality reduction is applied to reduce the necessary computations and only extract the expected amount of underlying sources. This dimensionality reduction is done through PCA which reduces the spatial dimension of the data from XZY to C which can be freely chosen. C is the number of components kept, such that these C components explain as much of the original data's variance as possible. Selecting C is of importance as it directly describes how many underlying sources make up the data, selecting it too low and sources of interest may be lost, selecting it too large means extra components may be kept that are just noise interfering with the results. Various information-theoretic criteria exist that can help identify what the number of components could be, these are methods like the Akaike information criterion [27], or Bayesian information criterion [28] that have been used in similar studies. However, in this thesis, the elbow method has already been implemented for spectral clustering and will therefore be used again. Using the elbow method it is found that 63 components should be kept, however, this is a relatively large number when compared with similar studies of mice with fMRI that find 15 to 30 components [29], [30]. Therefore, the number of components kept is set to 50 as a compromise. An investigation is done later on in subsection 5.4.2 to see what happens when the number of components is further reduced to what is usually seen in literature. To complete the pre-processing, after PCA the data is normalized such that acIVA can be performed.

5.3. Reference selection

It is still necessary to create the references for acIVA. With these experiments, there are only 1 or 2 stimuli provided meaning that only 1 or 2 references have to be supplied to acIVA. As discussed in chapter 2, it is not possible to directly use the stimulus itself as a reference, since it is transformed through the HRF. Therefore, two HRFs have to be estimated, one for each stimulus.

To accomplish this a set of HRFs is generated creating potential references when convolved with the stimuli. These potential references are all correlated with the time course in each pixel of the data. The best reference is found by looking for the reference that has the highest correlation with one of the pixels. An example of the reference found when looking at the LED stimulus is shown in Figure 5.4, with both the estimated reference compared to the time course in the voxels and a heat map showing which areas are also correlated with this stimulus. From these areas it can be seen that the top left of the image has an active region corresponding to the visual cortex, showing that this region responds to the LED stimuli. Doing this for all subjects and averaging their results provides the following θ parameters for the HRFs as shown in Table 5.1

for both the LED and puff stimulus.

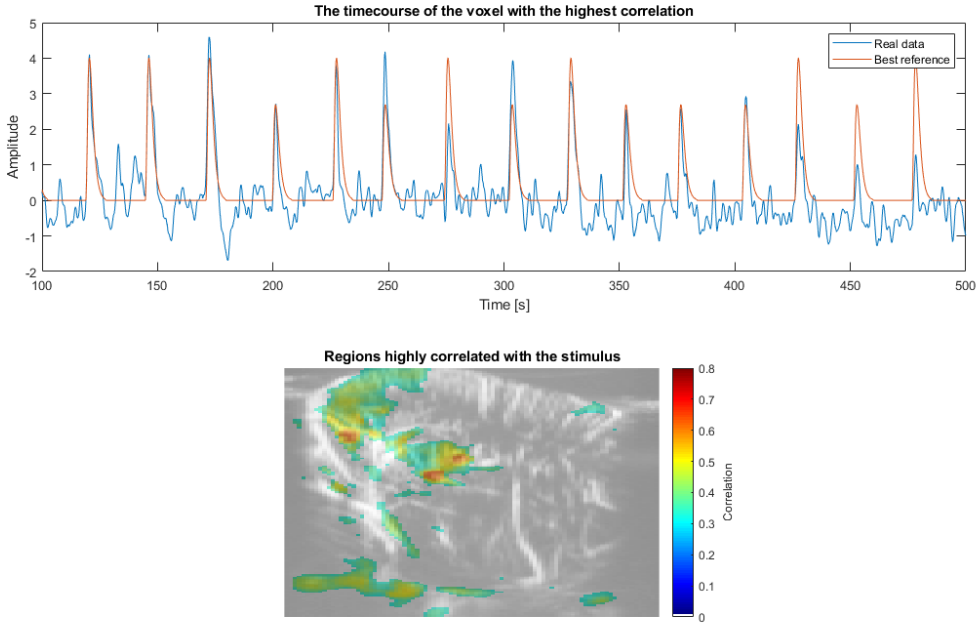

Figure 5.4: A visual representation of the correlation analysis is shown for a single subject. The top shows the identified LED reference together with the time course which it has the highest correlation with. The bottom shows other areas that are also correlated with the found stimulus. The top left is where the visual cortex is, showing that this region responds to the LED stimulus

Table 5.1: The θ parameters found for both stimuli through these experiments.

	LED	Puff
θ_1	4	2
θ_2	2.2	3.6
θ_3	1	1

5.4. Results

With the pre-processing and reference estimation done the subgroup identification can be done after performing acIVA. Since acIVA is relatively new it will also be compared to IVA in the following subsection. This will be followed by multiple minor analyses where parameters like PCA dimensions and data formatting will be discussed within this pipeline. It should again be noted that the results shown are that for the experiment with both the LED and puff stimuli unless mentioned otherwise.

5.4.1. IVA and acIVA results

Since acIVA uses a reference it may potentially be more prone to be biased when compared with IVA which does not use a reference to decompose the data. By performing both methods on the same dataset it is possible for their differences and similarities to be examined.

The conditions under which these methods are run are those previously mentioned, the number of components the data is reduced to through PCA is 50. All slices are analyzed simultaneously and for acIVA the references are constructed using the HRF parameters found in Table 5.1. Due to the random initializations of the unmixing matrices in both methods, it can be hard to compare the two with only a single run, to promote the robustness of the decompositions a robustness metric called cross intersymbol interference (cross-ISI) is used [31]. Through cross-ISI it becomes possible to find a more robust unmixing matrix by running IVA a number of times and finding the unmixing matrix that is closest to all other unmixing matrices. In this thesis, 10 unmixing matrices are generated after which the most robust is selected through cross-ISI.

Running both methods and constructing the covariance matrices will show differences and potential subgroups already. Since the covariance matrices are not sorted after IVA it is chosen to sort them by highest mean correlation to allow for easier comparisons. The resulting covariance matrices are shown for IVA and acIVA in Figure 5.5 and Figure 5.6 respectively.

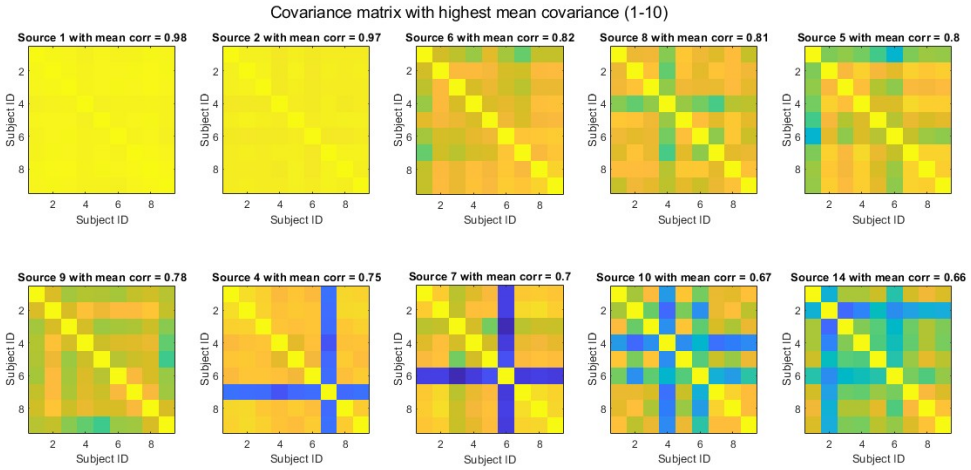

Figure 5.5: The components with the highest mean correlation between subjects as found by IVA. These matrices show that the first 10 components are quite similar across the subjects. From manual analysis it is found that component 4 is related to the LED and component 14 is related to the puff.

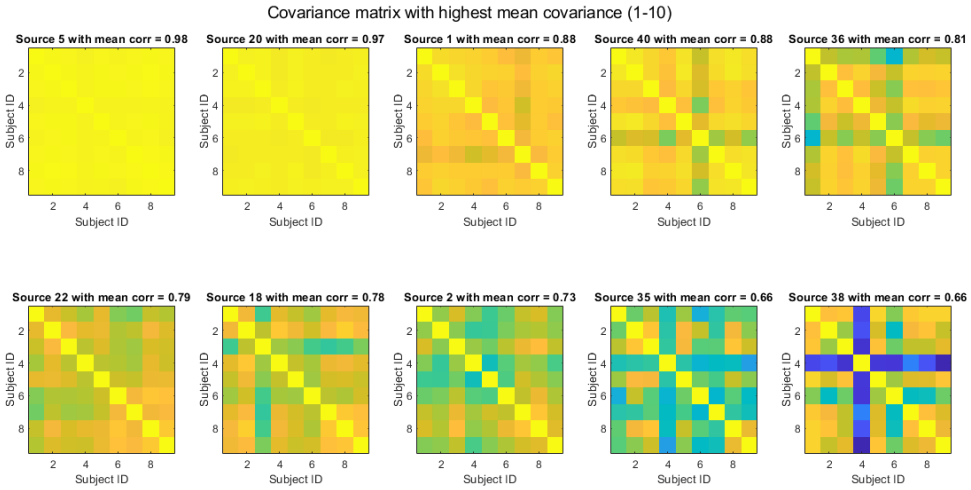

Figure 5.6: The components with the highest mean correlation between subjects as found by acIVA. Less components are extracted that are common across all subjects. Source 1 and source 2 are those found using the LED and puff reference respectively, showing that these time courses are found in all subjects, but weaker with respect to the puff.

When comparing the results in Figure 5.5 with the results in Figure 5.6, it should be noted that the component numbers can not be compared due to the random initializations creating permutation ambiguity. However, when looking at the results it can be seen that both methods capture the same highly correlated signals. From further investigation, it can be seen that these are the same low-frequency signals unrelated to the stimuli and therefore of no further interest.

When looking for the signals related to the stimuli it is not hard to find them using acIVA as the first 2 components are those related to the 2 references supplied. This means that components 1 and 2 are related to the LED and puff stimuli respectively as long as these exist within the data. For IVA these have to be manually extracted, manual analysis reveals components 4 and 14 to be related to the LED and puff respectively in

this specific case (it can be different when run again). For both methods, the components related to the stimuli show no subgroups according to spectral clustering. Analyzing the time courses and spatial maps of these components also shows that both methods extract a similar signal. An example of this is shown in the time courses extracted for the LED signal for both methods as shown in Figure 5.7 and Figure 5.8. The only difference is that IVA leaves out one subject and a low-frequency signal is added on top of the LED signal.

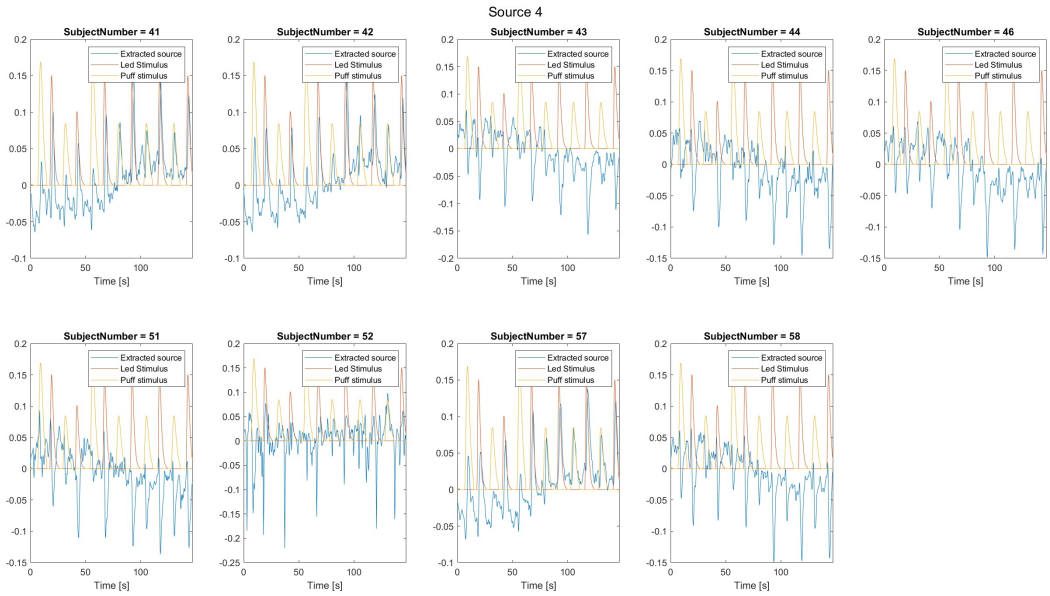

Figure 5.7: The time course of the fourth source extracted by IVA for each subject. Showing that all subjects correlate except for subject 52 (7) as can also be seen in the covariance matrix. The sign ambiguity of IVA can also be seen here.

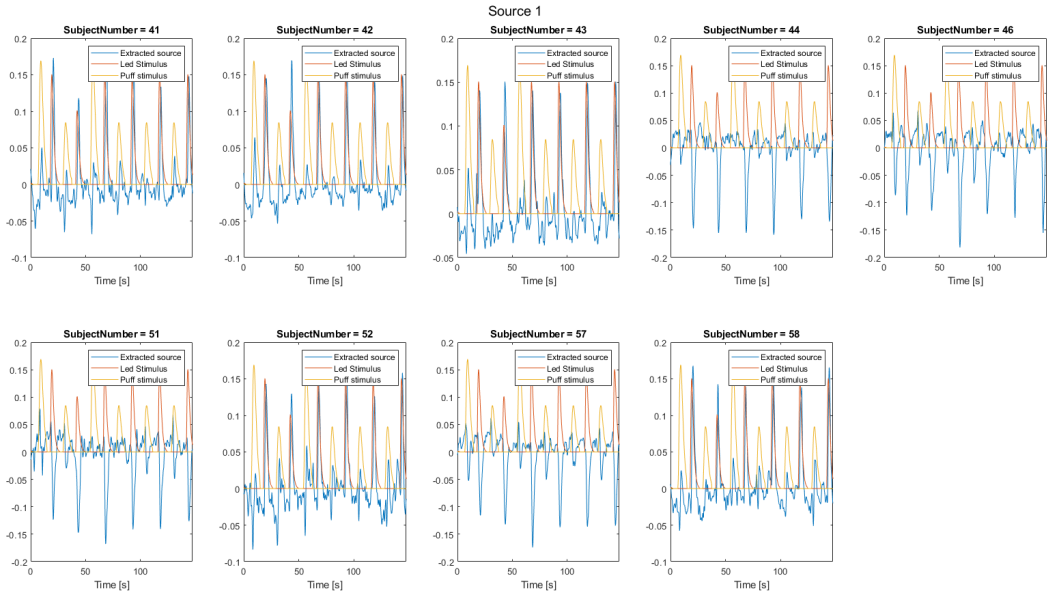

Figure 5.8: The time course of the first source extracted by acIVA for each subject. This is the component that uses the LED reference during IVA. In this case acIVA shows that all subjects react to the LED stimulus even subject 52.

Examining the time courses is only one part of the story, as the spatial maps corresponding to these time courses reveal where this activity pattern takes place. Examining the spatial maps for the LED component of both methods reveals that both methods agree that activation takes place in the visual cortex. To highlight this activation, heat maps are created showing where activation takes place. These heat maps are created by first performing z-score standardization on the spatial map and only showing the pixels with an intensity $|I| > 2$. Doing this for both methods and taking the average of these activation maps results in Figure 5.9. This figure shows that both methods extract similar regions, with the differences being likely due to random initializations.

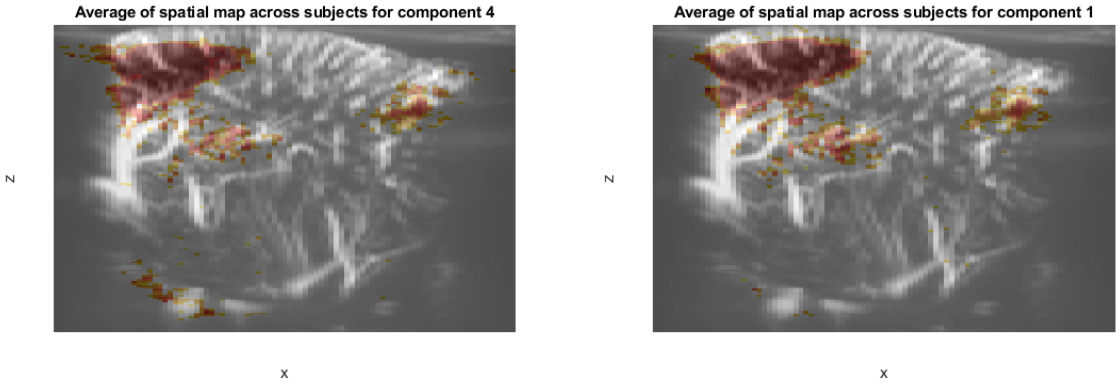

Figure 5.9: The spatial maps for slice 1 extracted for the LED component in IVA (left) and acIVA (right). Both methods extract the visual cortex in the top left for all subjects, with only minor differences between the spatial maps. These differences can be attributed to the random initializations.

These spatial maps and time courses show that both methods extract the same signals related to the stimuli, making acIVA a suitable alternative. Furthermore, acIVA is capable of directly returning the components related to the stimuli due to the references removing the need for manual identification for some of the components. For these reasons, acIVA will also be used for the remaining analyses.

The covariance matrices related to the stimuli showed no sign of subgroups according to spectral clustering, however, it is still possible to apply spectral clustering on the remaining covariance matrices found by acIVA. Doing this results in one component showing interesting subgroups. This is component 19 and the covariance matrix with the found subgroups is shown in Figure 5.10, with the subgroups highlighted by the red squares.

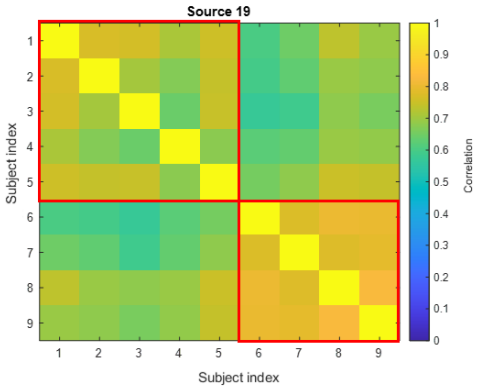

Figure 5.10: Component 19 extracted by acIVA which has a subgroup from subject 1 to 5 and subject 6 to 9 according to spectral clustering. It can also be seen that the subgroups correlate relatively well with one another, indicating that the extracted component is only slightly different when comparing across the subgroups

Examining the time courses of this component shows that this component consists of both the LED and

puff stimulus simultaneously, as shown in Figure 5.11. Investigating the time courses shows that one subgroup reacts half a second later to the stimuli when compared with the other subgroup, indicating that this appears to be a good subgroup identification by spectral clustering. However, these subgroups do not categorize the mice according to their real subgroups.

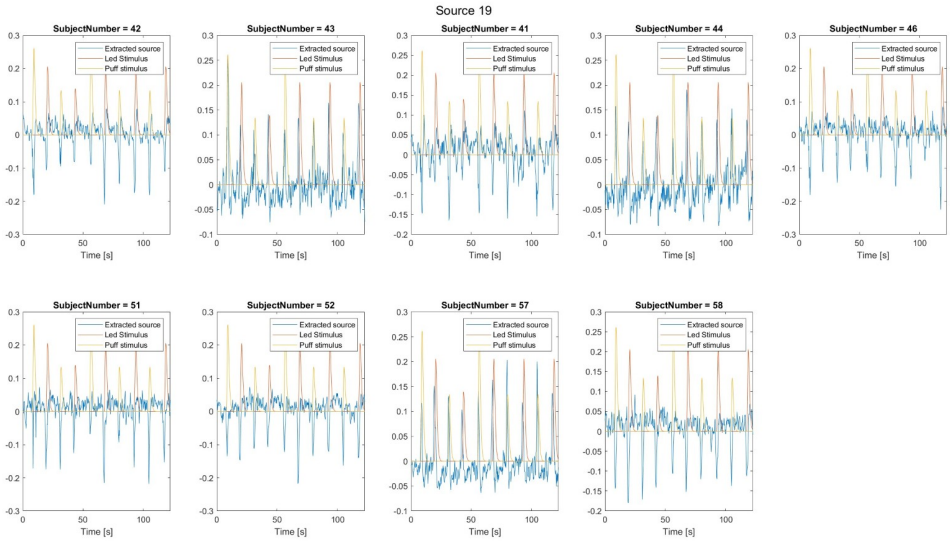

Figure 5.11: The time course for each subject for component 19. Showing that all subjects react to both stimuli. The subgroups found by spectral clustering indicate that the top row would be a subgroup and the bottom row would be a subgroup. The top row reacts half a second earlier to the stimuli when compared to the bottom row.

Performing this analysis for the second experiment where the mice can learn about the puff stimulus also does not reveal the subgroups through spectral clustering, nor does it reveal subgroups related to learning. In this experiment, there are only 8 mice analyzed since one mouse missed a slice, but of each mouse, there is a recording where they encounter the stimulus for the first time and one where they have had time to learn about the stimulus. This makes it a total of 16 recordings that are analyzed simultaneously, with the first 8 before training and the last 8 after training. The resulting covariance matrices with the highest mean correlation are shown in Figure 5.12, with component 1 being the component related to the puff stimulus.

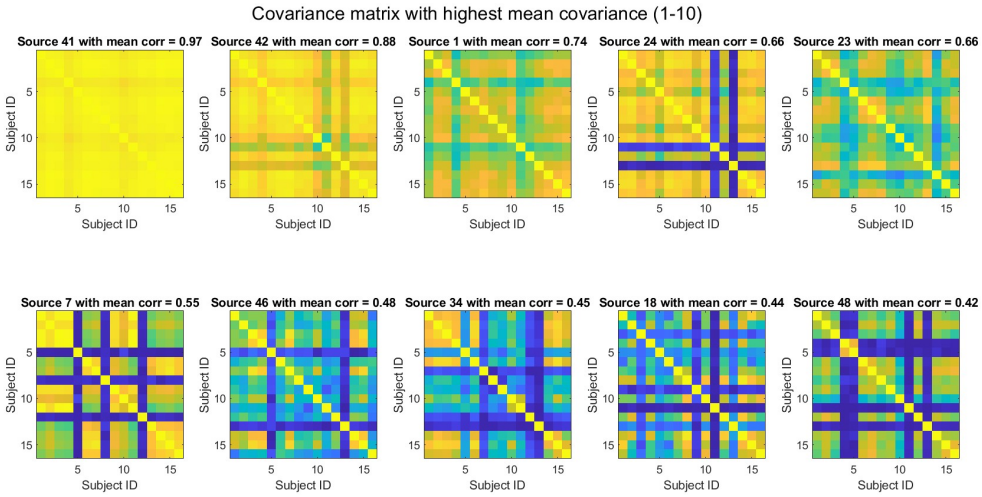

Figure 5.12: The covariance matrices from the second experiment where the mice have time to learn. The first 8 id's represent the mice before training and the last 8 id's represent the same mice but after training. Source 1 represents the puff component, which shows no subgroup over time, as the correlations are high between the first 8 recordings and the last 8 recordings.

5.4.2. PCA dimensions

As discussed earlier, the number of components that are usually used during data reduction is half of what is estimated here. This section will investigate the differences when the number of components is chosen to what is seen in similar studies. The number of components usually found within literature is 20 to 30 components [29], [30], therefore, both 20 and 30 components were tested, but only the case with 30 components will be discussed due to the results behaving similarly in both cases.

Examining the covariance matrices when only 30 components are extracted already shows a clear difference, these covariance matrices are shown in Figure 5.13. Since acIVA is used, it can be seen directly if the stimuli-related components are extracted. Looking at component 1, related to the LED stimulus, it is seen that a signal common to all subjects is extracted, examining the time courses confirms that this component is also related to the LED stimulus. Looking at component 2, which is related to the puff stimulus, it is seen that no common signal is extracted. Examining the time courses of this component indicates that the puff component is lost due to the lower order PCA. Furthermore, the component previously highlighting potential subgroups through a combination of both stimuli is also lost, nor did any new covariance matrices show up that highlight the actual subgroups through spectral clustering.

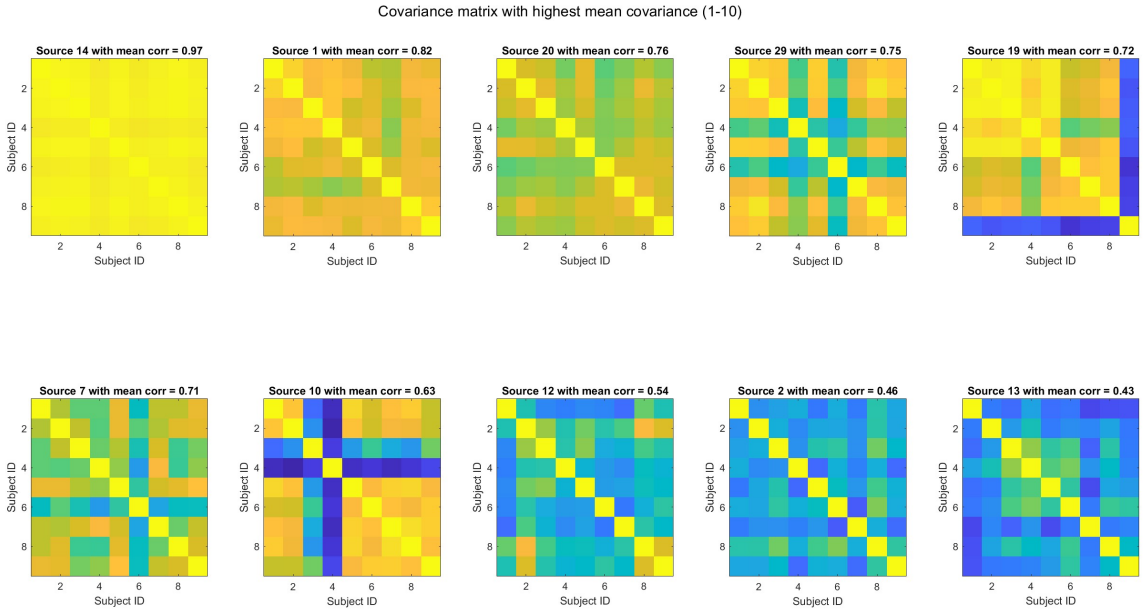

Figure 5.13: The ten components with the highest mean correlation across the subjects when performing acIVA with 30 components. Component 1 related to the LED is still retrieved and is shared across the subjects as seen with 50 components. The second component related to the puff is not shared between the subjects meaning that for at least some subjects the puff component is removed with the lower order PCA.

There might be various reasons as to why this higher-order PCA decomposition is required in this specific analysis. The first reason could be due to multiple slices being analyzed simultaneously, as opposed to analyzing each slice separately which is done in the works using 20 to 30 components. By including all slices more components may be necessary since some slices may contain sources specific to that slice. Furthermore, this order of 20 to 30 components is used in the context when group PCA is applied as an added processing step. This is performed after an individual PCA order reduction has been performed where each subject keeps more than 30 components. This is followed by concatenating all subjects and performing a single group PCA with an order of 20 or 30. This is done to further reduce the computational complexity in datasets with a large number of subjects [32], [33]. However, since the number of subjects is small and therefore the number of total components also being small, it was deemed unnecessary to use group PCA. Using group PCA to get to 20 components would result in a fairer comparison but has not been implemented and tested due to limited time. However, it is also hypothesized that the multi-slice analyses might be the cause for the higher order PCA, which will be further investigated in the following subsection.

5.4.3. Single - and Multi-Slice Analysis

Analyzing four slices instead of a single slice at a time, results in a reduced number of components to estimate. This reduced number of components improves efficiency due to fewer computations and fewer results to analyze. However, it may be that this analysis oversimplifies the data since this method assumes that the signals measured aren't shifted in time over the slices and instead happen simultaneously. To see if this is the case a single-slice analysis is performed to compare the results. The PCA order discussion will also be continued in this subsection as the multi-slice analysis may be the reason for the earlier seen discrepancy in order estimation.

Going from the multi-slice analysis to the single-slice analysis means that the usual pre-processing steps are applied except for the concatenation of the slices. Now each slice y will be considered a separate data matrix $\mathbf{X}_y^{[k]} \in \mathbb{R}^{XZ \times T}$, with only the y-th slice of each subject being supplied to the acIVA algorithm.

In the first tests, the PCA order was kept at 50 for each dataset. After performing acIVA it can be seen that the stimuli components are more significant in the left-most slices while being weaker in the right-most slices. This is as expected as the stimuli are only being applied to the left eye. However, this is also no new information as the spatial maps of the multi-slice analysis also show this weaker activation in the right slices within the LED and puff component. The subgroup analysis also reveals no new subgroup patterns in any stimuli-related component in any of the slices. The covariance matrices with the highest mean are shown for the two outer slices in Figure 5.14 and Figure 5.15.

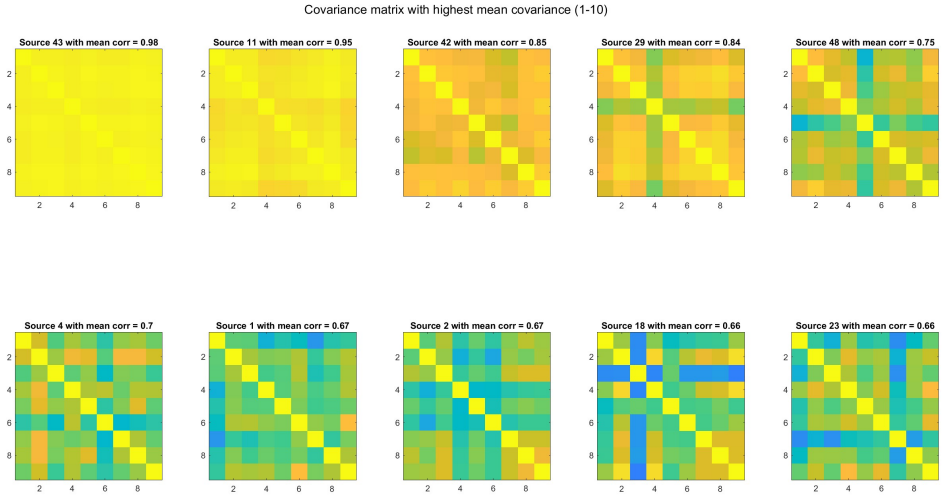

Figure 5.14: Covariance matrices for slice 1. The LED and puff components are extracted similarly as in the multi-slice analysis, with only minor deviations in the correlations.

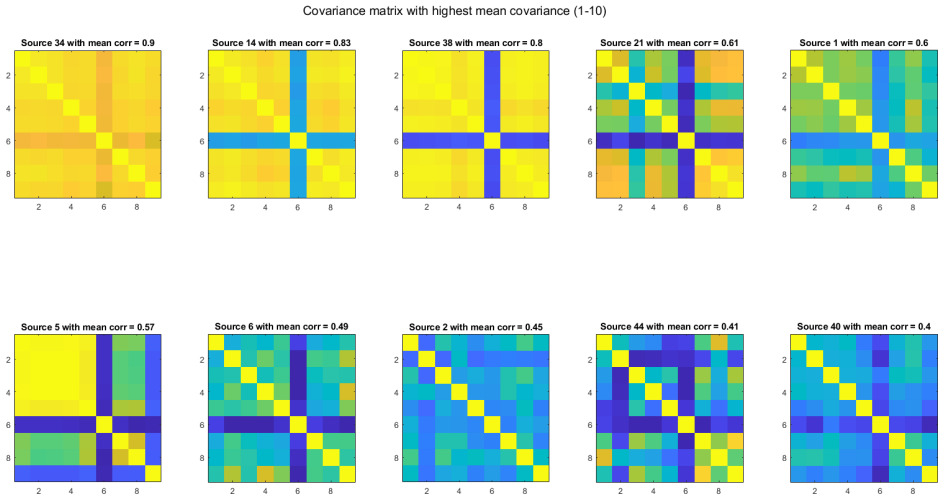

Figure 5.15: Covariance matrices for slice 4. The LED and puff components are extracted but weaker, showing that the stimuli are only applied to the left eye. This also highlights the unilateral response to the stimuli, with only one brain half reacting significantly to the stimuli.

To see if there is a delay over the slices, the peaks are extracted from the stimuli-related time courses. Doing this shows no delay between the slices, the peaks are also not shifted when compared with the multi-slice analysis. This does not mean that there is no delay between the slices, as it could also be that the delay is smaller than the measurement frequency. Comparing the spatial maps related to the stimuli also highlights no major differences between the two forms of analysis, as an example the average spatial map for slice 1 related to the LED is shown in Figure 5.16, showing the same activation regions as seen in Figure 5.9.

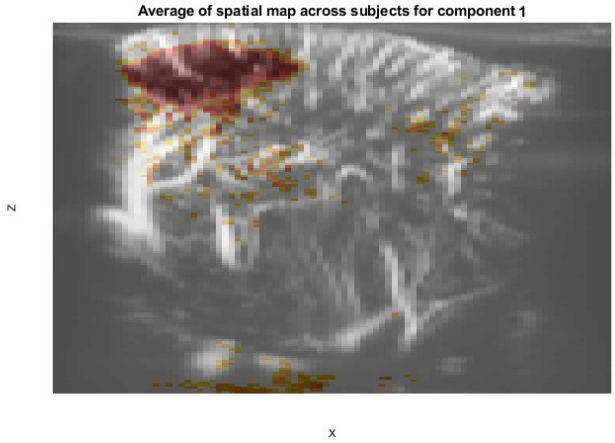

Figure 5.16: Average spatial map for the LED component in the first slice with single-slice analysis. Again the visual cortex is highlighted in the top left as expected and is similar to what is extracted with multi-slice analysis.

Gathering these results is also more time consuming due to there being four acIVAs being executed with the same input data size as is done in the multi-slice analysis. This also gives four times the number of components to look through, these components are slightly different from those seen in the multi-slice analysis, this could be an improved accuracy or an effect of random initialization conditions. Therefore, it is hard to argue whether the extracted components in the single-slice analysis are better. This makes the multi-slice analysis a faster and more interpretable method due to there being fewer components in total, this may come at a cost of accuracy but that requires further investigation.

Reducing the number of components to 30 in the single-slice analysis results in similar results as in the multi-slice analysis with 30 components. In the leftmost slices, the LED component is extracted and the puff

component is extracted in some but not all mice. For the right slices, the same pattern happens with weaker correlations. This indicates that the higher-order PCA may be required with fUS experiments since valuable information regarding the stimuli is lost when keeping only 30 components.

6

Conclusions

Methods like SI-IVA and IVA-CS perform automatic subgroup identification using IVA. These methods lacked a performance assessment making it unclear when their results can be trusted. In this thesis, a performance assessment of these methods was performed by simulating a large number of covariance matrices and checking the accuracy of these methods. The simulations show that spectral clustering is the preferred clustering algorithm over those used in IVA-CS and SI-IVA.

Applying spectral clustering to a stimulus-related fUS dataset with two groups of mice did not reveal the true subgroups in the components related to the stimulus. However, using acIVA instead of IVA helped with interpretability as it revealed the components related to the stimulus without the need for manual identification.

It was also found that the number of components kept when performing data reduction in these studies might need to be higher, as the lower-order data reduction resulted in the loss of components related to the stimuli.

While the true subgroups were not revealed by spectral clustering, it was shown through the simulations that this method is a more robust and reliable alternative to SI-IVA and IVA-CS when looking for subgroups in fUS data.

6.1. Future research

While this thesis answered some questions, it also generated new ones that could be used as a possible direction for future work. These are as follows:

Using spectral clustering as a subgroup classification algorithm was shown to be a better approach, however, this is not the only clustering method that exists and improvements are still possible. By comparing more clustering methods it may be possible to further increase the accuracy over a larger number of scenarios, resulting in more trust in the results when applied to real data.

Analyzing the subgroup identification methods only through simulations means that we can only get so close to real life. Having a dataset of covariance matrices from real experiments that are labelled by an expert could give a better insight into the accuracy of these methods when applied to real data.

This thesis made use of a temporal IVA, resulting in the components being independent time courses. These were used to try and identify subgroups, however, spatial IVA is also possible. Therefore, it may be of interest to also try spatial IVA and see if the independent spatial maps may result in the correct identification of the subgroups.

Using acIVA resulted in increased interpretability through the stimuli being used as references. Adding more references would result in even fewer components that would need to be manually identified. Therefore, an investigation into standard underlying sources could be of interest, as these can always be used no matter the dataset, improving the interpretability of IVA results.

In this thesis, the number of components kept when applying PCA is different from what is common in literature. It is still unclear as to what causes this discrepancy. Further investigation will highlight what the actual number of underlying sources needs to be for future research.

List of abbreviations

acIVA adaptive constrained independent vector analysis

BOLD blood-oxygen-level-dependent

BSS blind source separation

CBV cereberal blood volume

cross-ISI cross intersymbol interference

fMRI functional magnetic resonance imaging

fUS functional ultrasound

GICA group independent component analysis

HRF hemodynamic response funciton

ICA independent component analysis

IVA independent vector analysis

JBSS joint blind source separation

PCA principal component analysis

PDI Power Doppler Imaging

SCM source component matrix

SCV source component vector

sICA spatial independent component analysis

sogICA self-organized independent component analysis

tICA temporal independent component analysis

References

[1] M. J. McKeown, S. Makeig, G. G. Brown, *et al.*, "Analysis of fmri data by blind separation into independent spatial components," *Human brain mapping*, vol. 6, no. 3, pp. 160–188, 1998.

[2] M. J. Courtemanche, C. J. Sparrey, X. Song, A. MacKay, and R. C. D'Arcy, "Detecting white matter activity using conventional 3 tesla fmri: An evaluation of standard field strength and hemodynamic response function," *Neuroimage*, vol. 169, pp. 145–150, 2018.

[3] E. Macé, G. Montaldo, I. Cohen, M. Baulac, M. Fink, and M. Tanter, "Functional ultrasound imaging of the brain," *Nature methods*, vol. 8, no. 8, pp. 662–664, 2011.

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

[5] V. D. Calhoun, J. Liu, and T. Adalı, "A review of group ica for fmri data and ica for joint inference of imaging, genetic, and erp data," *Neuroimage*, vol. 45, no. 1, S163–S172, 2009.

[6] L. Griffanti, G. Douaud, J. Bijsterbosch, *et al.*, "Hand classification of fmri ica noise components," *Neuroimage*, vol. 154, pp. 188–205, 2017.

[7] E. Mace, G. Montaldo, B.-F. Osmanski, I. Cohen, M. Fink, and M. Tanter, "Functional ultrasound imaging of the brain: Theory and basic principles," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 60, no. 3, pp. 492–506, 2013.

[8] K. Hikishima, T. Tsurugizawa, K. Kasahara, R. Takagi, K. Yoshinaka, and N. Nitta, "Brain-wide mapping of resting-state networks in mice using high-frame rate functional ultrasound," *NeuroImage*, vol. 279, p. 120 297, 2023.

[9] 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, 2023.

[10] D. Langlois, S. Chartier, and D. Gosselin, "An introduction to independent component analysis: Infomax and fastica algorithms," *Tutorials in Quantitative Methods for Psychology*, vol. 6, no. 1, pp. 31–38, 2010.

[11] S.-i. Amari, A. Cichocki, and H. Yang, "A new learning algorithm for blind signal separation," *Advances in neural information processing systems*, vol. 8, 1995.

[12] A. Hyvarinen, "Fast ica for noisy data using gaussian moments," in *1999 IEEE international symposium on circuits and systems (ISCAS)*, IEEE, vol. 5, 1999, pp. 57–61.

[13] F. Esposito, T. Scarabino, A. Hyvarinen, *et al.*, "Independent component analysis of fmri group studies by self-organizing clustering," *Neuroimage*, vol. 25, no. 1, pp. 193–205, 2005.

[14] V. D. Calhoun, T. Adali, G. D. Pearlson, and J. J. Pekar, "A method for making group inferences from functional mri data using independent component analysis," *Human brain mapping*, vol. 14, no. 3, pp. 140–151, 2001.

[15] C. F. Beckmann and S. M. Smith, "Tensorial extensions of independent component analysis for multisubject fmri analysis," *Neuroimage*, vol. 25, no. 1, pp. 294–311, 2005.

[16] T. Kim, T. Eltoft, and T.-W. Lee, "Independent vector analysis: An extension of ica to multivariate components," in *International conference on independent component analysis and signal separation*, Springer, 2006, pp. 165–172.

[17] J.-H. Lee, T.-W. Lee, F. A. Jolesz, and S.-S. Yoo, "Independent vector analysis (iva): Multivariate approach for fmri group study," *Neuroimage*, vol. 40, no. 1, pp. 86–109, 2008.

[18] Z. Luo, "Independent vector analysis: Model, applications, challenges," *Pattern Recognition*, vol. 138, p. 109 376, 2023.

[19] S. Bhinge, Q. Long, V. D. Calhoun, and T. Adalı, "Adaptive constrained independent vector analysis: An effective solution for analysis of large-scale medical imaging data," *IEEE journal of selected topics in signal processing*, vol. 14, no. 6, pp. 1255–1264, 2020.

[20] S. Bhinge, Q. Long, Y. Levin-Schwartz, Z. Boukouvalas, V. D. Calhoun, and T. Adalı, "Non-orthogonal constrained independent vector analysis: Application to data fusion," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2017, pp. 2666–2670.

[21] Q. Long, S. Bhinge, V. D. Calhoun, and T. Adali, "Independent vector analysis for common subspace analysis: Application to multi-subject fmri data yields meaningful subgroups of schizophrenia," *NeuroImage*, vol. 216, p. 116 872, 2020.

[22] H. Yang, M. A. B. S. Akhonda, F. Ghayem, Q. Long, V. Calhoun, and T. Adali, "Independent vector analysis based subgroup identification from multisubject fmri data," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2022, pp. 1471–1475.

[23] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge university press, 2012.

[24] M. E. Newman, "Modularity and community structure in networks," *Proceedings of the national academy of sciences*, vol. 103, no. 23, pp. 8577–8582, 2006.

[25] A. Ng, M. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," *Advances in neural information processing systems*, vol. 14, 2001.

[26] P. Liu and D. Liu, "Commonly used regional exposure," in *Liu's Principles and Practice of Laboratory Mouse Operations: A Surgical Atlas*, Springer, 2023, pp. 41–105.

[27] H. Akaike, "Information theory and an extension of the maximum likelihood principle," in *Selected papers of hirotugu akaike*, Springer, 1998, pp. 199–213.

[28] G. Schwarz, "Estimating the dimension of a model," *The annals of statistics*, pp. 461–464, 1978.

[29] J. Grandjean, C. Canella, C. Anckaerts, *et al.*, "Common functional networks in the mouse brain revealed by multi-centre resting-state fmri analysis," *Neuroimage*, vol. 205, p. 116 278, 2020.

[30] E. Jonckers, J. Van Audekerke, G. De Visscher, A. Van der Linden, and M. Verhoye, "Functional connectivity fmri of the rodent brain: Comparison of functional connectivity networks in rat and mouse," *PloS one*, vol. 6, no. 4, e18876, 2011.

[31] Q. Long, C. Jia, Z. Boukouvalas, B. Gabrielson, D. Emge, and T. Adali, "Consistent run selection for independent component analysis: Application to fmri analysis," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2018, pp. 2581–2585.

[32] Y. Du and Y. Fan, "Group information guided ica for fmri data analysis," *Neuroimage*, vol. 69, pp. 157–197, 2013.

[33] S. M. Smith, A. Hyvärinen, G. Varoquaux, K. L. Miller, and C. F. Beckmann, "Group-pca for very large fmri datasets," *Neuroimage*, vol. 101, pp. 738–749, 2014.

A

Analysis of SI-IVA

An analysis of SI-IVA is performed in this chapter for when there is no subgroup in the covariance matrix that SI-IVA is applied to. This could mean that the extracted component exists in all subjects or that the extracted component does not correlate with any of the subjects. For cases with subgroups in the covariance matrix, no bounds are provided due to this requiring bounds for the second-highest eigenvalue or smaller.

SI-IVA detecting subgroups

SI-IVA performs subgroup identification by using the Gershgorin disc theorem on the covariance matrix $\mathbf{C}$. The Gershgorin disc theorem states the following: For a complex matrix $\mathbf{A} \in \mathbb{C}^{n\times n}$ with entries a_{ij}. Let $R_i = \sum_{j\neq i} |a_{ij}|$ be the absolute sum of the non-diagonal entries of row i of $\mathbf{A}$ which will be used to define a disc. Since in this analysis, the covariance matrix consists of purely positive values the radius can be described as: $R_i = \sum_{j\neq i} c_{ij}$. With this, a Gershgorin disc can be defined as a disc with a radius R_i centered at the point c_{ii}, in this case the diagonal of $\mathbf{C}$ contains all ones which places the center of each disc at one. These discs give bounds as to where the eigenvalues of $\mathbf{A}$ are located. SI-IVA uses these discs in combination with the eigenvalues to determine the number of subgroups in $\mathbf{C}$. Let $R_{\min} = \min_i R_i$ be the radius of the smallest disc of $\mathbf{C}$, then according to SI-IVA, the number of eigenvalues outside of the smallest disc is the number of subgroups. The pseudocode that performs subgroup detection in SI-IVA is shown below.

Algorithm 1 SI-IVA algorithm

```
1: procedure SI-IVA(C)
2:     m = 0
3:     λ_i ←sort( eig(C))
4:     R_i ← Σ_{j≠i} |c_ij|
5:     if λ_i > R_min + 1 then
6:         m ← m + 1
7:     end if
8: end procedure
```

Analyzing eigenvalues of the covariance matrix

Now that SI-IVA is explained it is possible to examine when it is correct for the case of no subgroups, this is achieved by examining the eigenvalues of $\mathbf{C}$. For this examination the eigenvalues of $\mathbf{C}$ are sorted in descending order:

$$\lambda_{\max} = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{n-1} \geq \lambda_n = \lambda_{\min}, \tag{A.1}$$

since the number of subgroups m should be 0, it follows that all eigenvalues should be smaller than $R_{\min} + 1$, therefore it is sufficient to check if the following holds:

$$\lambda_{\max} \leq R_{\min} + 1 = \min_i \sum_{j=1}^{n} c_{ij}, \tag{A.2}$$

since if this statement holds it is given that m will not increase and the number of subgroups stays 0, inversely if $\lambda_{\max} > R_{\min} + 1$ while there are no subgroups in $\mathbf{C}$ a proof is given that this method incorrectly determines that there are subgroups while there are none.

Start by reconsidering what the bound is that an eigenvalue needs to pass to increase the number of subgroups, this being $R_{\min} + 1$. It can be seen that due to the off-diagonal entries all being positive and all diagonal entries being 1 that this threshold is also equal to:

$$R_{\min} + 1 = \min_i \sum_{j=1}^{n} c_{ij}, \tag{A.3}$$

since $R_{\min} = \min_i \sum_{j=1,\ j\neq i}^{n} c_{ij}$ and $c_{ii} = 1$. Meaning that $R_{\min} + 1$ is the minimum row sum.

Now it is possible to compare the bounds of the maximum eigenvalue with this threshold of the minimum row sum. One bound for the eigenvalues of a Hermitian matrix (which the covariance matrix is) is the Rayleigh quotient. The Rayleigh quotient says the following about the eigenvalues of $\mathbf{C}$:

$$\lambda_{\min} \leq \mathbf{x}^* \mathbf{C} \mathbf{x} \leq \lambda_{\max}, \tag{A.4}$$

for any unit vector $\mathbf{x} \in \mathbb{C}^n$ with $\mathbf{x}^*$ being the conjugate of $\mathbf{x}$. Equality holds on the right side if $\mathbf{x}$ is the principal eigenvector such that $\mathbf{C}\mathbf{x} = \lambda_{\max}\mathbf{x}$. This makes it that $\lambda_{\max} > \mathbf{x}^*\mathbf{C}\mathbf{x}$ for any unit vector $\mathbf{x}$ that is not the principal eigenvector of $\mathbf{C}$. Take $\mathbf{x} = [1, 1, \cdots, 1, 1]^T * \frac{1}{\sqrt{n}}$, then $\mathbf{x}^*\mathbf{C}\mathbf{x} = \frac{1}{n}\sum_i \sum_j c_{ij}$ which is the average of $\mathbf{C}$ multiplied with n, comparing this with the minimum row sum gives the following:

$$\lambda_{\max} \geq \frac{1}{n} \sum_i \sum_j c_{ij} \geq \min_i \sum_{j=1}^{n} c_{ij}, \tag{A.5}$$

the equality between the minimum row sum and average can be removed under the assumption that one of the rows of $\mathbf{C}$ sums to a different value than any other row in $\mathbf{C}$ resulting in the average of the row sums not being equal to the minimum row sum. Under said assumption the statement can be turned into a pure inequality as follows:

$$\lambda_{\max} \geq \frac{1}{n} \sum_i \sum_j c_{ij} > \min_i \sum_{j=1}^{n} c_{ij}, \tag{A.6}$$

and with the minimum row sum being equal to the threshold of $R_{\min} + 1$ results in Equation A.6 being rewritten to:

$$\lambda_{\max} > \min_i \sum_{j=1}^{n} c_{ij} = R_{\min} + 1, \tag{A.7}$$

with Equation A.7 it becomes clear that SI-IVA will always detect one subgroup under the condition that the average row sum is equal to the minimum row sum. This is independent of the number of subgroups in the covariance matrix, meaning that in covariance matrices where all subjects are part of the subgroup, SI-IVA will detect a subgroup and will try to partition this covariance matrix through k-means.

Extending this analysis to check how SI-IVA performs when there are more subgroups would require bounds for the second largest eigenvalue or third largest eigenvalue. To the best of my knowledge, there are none strong enough to perform a good theoretical analysis in cases of 2 or 3 subgroups, but testing with simulations highlights that SI-IVA tends to underestimate when $m > 1$ and the number of subjects in the subgroups are small compared to the total number of subjects.