



MultiView Independent Component Analysis with Delays

Ambroise Heurtebise, Pierre Ablin, Alexandre Gramfort

► To cite this version:

Ambroise Heurtebise, Pierre Ablin, Alexandre Gramfort. MultiView Independent Component Analysis with Delays. MLSP 2023 - 33rd IEEE International Workshop on Machine Learning for Signal Processing, Sep 2023, Rome, Italy. hal-04317549

HAL Id: hal-04317549

<https://hal.science/hal-04317549>

Submitted on 1 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

MultiView Independent Component Analysis with Delays

Ambroise Heurtebise¹, Pierre Ablin², and Alexandre Gramfort¹

¹Université Paris-Saclay, Inria, CEA, Palaiseau, France

²Apple, Paris, France

Abstract

Linear Independent Component Analysis (ICA) is a blind source separation technique that has been used in various domains to identify independent latent sources from observed signals. In order to obtain a higher signal-to-noise ratio, the presence of multiple views of the same sources can be used. In this work, we present MultiView Independent Component Analysis with Delays (MVICAD). This algorithm builds on the MultiView ICA model by allowing sources to be delayed versions of some shared sources: sources are shared across views up to some unknown latencies that are view- and source-specific. Using simulations, we demonstrate that MVICAD leads to better unmixing of the sources. Moreover, as ICA is often used in neuroscience, we show that latencies are age-related when applied to Cam-CAN, a large-scale magnetoencephalography (MEG) dataset. These results demonstrate that the MVICAD model can reveal rich effects on neural signals without human supervision.

1 Introduction

Independent Component Analysis (ICA) allows to separate mixed signals without knowing the mixing operator [Comon, 1994; Hyvärinen and Oja, 2000]. In linear ICA [Hyvärinen, 1999; Ablin et al., 2018], the mixing is assumed linear and the goal is to find a linear transformation of the mixed signals that maximizes the statistical independence of the resulting sources. The parameters of the transformation matrix are estimated by optimizing a cost function, such as a likelihood.

In real-world data, one can sometimes access multiple views from the same data. For example, a doctor may have an MRI scan, a CT scan, and the answers to a clinical questionnaire for a diseased patient. In the context of ICA, a naive method would be to use ICA on each view and compare the resulting sources, but this solution would not take advantage of the group structure. Thus, many multi-view ICA methods have been developed to take advantage of multiple data views and increase the signal-to-noise ratio of the obtained sources (see Section 2). As multiple views are used to obtain a better recovery of the sources, it is typically assumed that the sources are the same for all views.

MultiView Independent Component Analysis (MVICA) [Richard et al., 2020] is an algorithm that makes the latter assumption. It aims to jointly estimate the transformation matrices for multiple views while assuming that they all share the exact same sources. While MVICA guarantees the identifiability of its model and produces state-of-the-art results for several machine learning tasks, it assumes that the sources are perfectly identical and temporally aligned between views. In other words, MVICA ignores the view-specific variability of the sources, which can lead to suboptimal separation results for certain applications.

This kind of inter-view variability typically arises in neuroscience, a critical application of ICA. Indeed, ICA has been extensively used for the analysis of electroencephalography (EEG) [Makeig et al., 1995] and magnetoencephalography (MEG) signals [Vigário et al., 2000], which are non-invasive recordings of the electrical potentials and magnetic fields produced by neurons in the brain. Since these signals are a mixture of the brain’s complex neural activity, ICA has been successfully used to separate out the different sources of activity. In this context, using data from multiple subjects at once can reveal better insights about brain functional organization. Yet, neural sources can present temporal delays between different subjects due to individual differences in brain anatomy and functional connectivity [Price et al., 2017; Roberts et al., 2010]. These delays can affect the performance

of standard multi-view ICA algorithms, as they assume that the sources are uncorrelated and have the exact same temporal structure across subjects.

In this work, building on the MVICA model, our goal is to develop an ICA algorithm that is robust to temporal delays. Specifically, we propose a novel method, called MultiView Independent Component Analysis with Delays (MVICAD), that incorporates an iterative temporal alignment step in the estimation procedure. We demonstrate the effectiveness of our method using simulations and MEG data, showing that it outperforms the MVICA algorithm. Besides, we report that the estimated delays of certain sources are related to age.

2 Related works

Many ICA algorithms have been proposed to process data from multiple datasets at once. They vary by their assumptions: some methods assume that the datasets have the same dimension, while others consider that all datasets share the same sources. In this section, we describe some of these algorithms.

MVICA [Richard et al., 2020] is our reference algorithm. It is designed to identify shared independent components across multiple datasets. It assumes that the sources in each dataset are independent and that the shared sources have the same temporal patterns across datasets. Unlike some other multi-view ICA approaches, MVICA’s likelihood can be written in closed form and is then optimized with a quasi-Newton algorithm, using the gradient and an approximation of the Hessian of the likelihood. Since comparisons with other multi-view ICA methods have been made in [Richard et al., 2020], we compare our method to MVICA.

Group ICA refers to a whole category of ICA algorithms for processing multiple views of the same data. It usually consists in first applying Principal Component Analysis (PCA) on each view separately. Then all the resulting sets are assembled into one set, using either concatenation or multi-set CCA [Kettenring, 1971; Correa et al., 2010], and PCA is used again on the obtained set. Finally, an ICA algorithm is used. Concatenation can be either spatial or temporal: temporal concatenation produces individual sources and a common mixing matrix, whereas spatial concatenation gives common sources and individual mixing matrices. Since we look for shared sources, we can compare our method to Group ICA with spatial concatenation [Calhoun et al., 2009]. Adding a back-reconstruction step is also possible to recover individual sources. However, this method does not bring statistical guarantees like consistency or asymptotic efficiency, contrary to MVICA. It also gives poorer results than MVICA on several neuroimaging tasks [Richard et al., 2020]. One strength of CCA-based ICA [Varoquaux et al., 2009; Tsatsishvili et al., 2015] is that they project multiple datasets onto a common underlying space by maximizing correlations between sets, which can be used for data fusion. Some of these methods can handle datasets with different numbers of sources. However, taking temporal delays into account in Group ICA would only be possible in the assembling step (done with PCA or multi-set CCA), which is less explicit than doing it during ICA optimization as we do in this paper.

Shared and Individual Multiview Independent Component Analysis [Pandeva and Forré, 2022] can identify both shared and individual sources at the same time. This method assumes that the sources are independent within each dataset but allows for both shared and individual sources to contribute to the data. It uses a joint estimation approach to simultaneously estimate the shared and individual components and their corresponding spatial and temporal maps. However, the number of shared components needs to be specified, which can be difficult.

3 METHOD

Notation The absolute value of the determinant of a matrix W is $|W|$. The ℓ_2 norm of a vector \mathbf{s} is $\|\mathbf{s}\|$, and $\|S\|$ is the Frobenius norm of a matrix S . For a scalar valued function f and a matrix $S \in \mathbb{R}^{p \times n}$, we write $f(S) = \sum_{j=1}^p \sum_{i=1}^n f(S_{ji})$.

3.1 Model and likelihood

Given m views and p sources, we model the observed signals $X^i \in \mathbb{R}^{p \times n}$, $i = 1, \dots, m$, as a linear combination of shared but delayed sources plus noise. Here n is the number of observations (assumed to be common to each view) and p is the number of features, typically corresponding to the number




Figure 1: Sources of 2 views. We observe different delays τ_j^i for each view i and each source j . Blue curves represent the first view’s sources and orange ones represent the sources of the second view. Left: first source. Right: second source.

of channels for multivariate signals. The model is:

$$X^i = A^i(\mathcal{T}_{\boldsymbol{\tau}^i}(S) + N^i) \quad , \quad i = 1, \dots, m \quad , \quad (1)$$

where the $A^i \in \mathbb{R}^{p \times p}$ are view-specific mixing matrices, assumed to be full-rank, \mathcal{T} is a shift operator that delays sources across samples, $S \in \mathbb{R}^{p \times n}$ are the shared sources, the vector $\boldsymbol{\tau}^i \in \{-\tau_{\max}, \dots, \tau_{\max}\}^p$ contains the integer delays for the p sources for a given view and $N^i \in \mathbb{R}^{p \times n}$ are the noise matrices. In practice, if there are 2 sources and delays are $\{1, 2\}$, the shift operator will shift to the right the first source by 1 sample and the second source by 2 samples. The hyperparameter τ_{\max} thus represents the maximum shift allowed. These delays represent a temporal shift rather than a time distortion, so they correspond to what is called “constant delays” and not to “cumulative delays” in [Roberts et al., 2010].

Figure 1 aims to help understand how delays are modeled. It shows how delay can differ for every source and every view: in the figure, the first view has a first source that peaks later than the second view but has an earlier peak for the second source. Also, we observe that the parameters τ_j^i , where i is the view and j is the source, model the delay between sources and the average source, represented in a dotted grey line.

As usual in ICA, we assume that observations are i.i.d., and so are the sources. Also, we assume that the noise is Gaussian decorrelated of variance σ^2 and independent across views and from the sources. In practice, thanks to MVICA’s robustness to noise misspecification [Richard et al., 2020], estimating the noise level is not critical, so we usually set it to 1.

Importantly, our model is identifiable (see Appendix A) up to scale and permutation for the mixing matrices and up to a common delay for the estimated delays. Therefore recovering the sources is a well-posed problem.

We insist on the fact that we represent signals as matrices containing all time samples instead of representing them as vectors. We made that choice because of the use of the circular shift operator \mathcal{T} . Indeed, here we assume periodic boundary conditions, so this operator takes as inputs p signals of length n and p integer delays and rolls each signal by its corresponding delay. Thus, we prefer to consider all observations at once. Having periodic boundary conditions on real signals can sometimes produce unwanted abrupt discontinuities when signals are noisy at the edges of the time interval. This however can be mitigated by applying a windowing function pushing signals to zero at the edges. In particular, this is appropriate in neuroscience for EEG or MEG signals because sources are supposed to emerge from baseline level and go back to baseline at the end of the period of interest.

We derive a maximum-likelihood approach to recover the parameters of the model, in a very similar way to [Richard et al., 2020]. As usual in ICA, sources reconstruction is done by estimating first the unmixing matrices $W^i = (A^i)^{-1} \in \mathbb{R}^{p \times p}$. Thus, we see the likelihood as a function of $\mathbf{W} = \{W^1, \dots, W^m\}$ and $\boldsymbol{\tau} = \{\boldsymbol{\tau}^1, \dots, \boldsymbol{\tau}^m\}$. The negative log-likelihood can be written as:

$$\mathcal{L}(\mathbf{W}, \boldsymbol{\tau}) = - \sum_{i=1}^m \log |W^i| + \frac{1}{2\sigma^2} \sum_{i=1}^m \|Y^i - \bar{S}\|^2 + f(\bar{S}), \quad (2)$$

where $Y^i = \mathcal{T}_{-\tau^i}(W^i X^i)$ are the aligned estimated sources of subject i , $\bar{S} = \frac{1}{m} \sum_{i=1}^m Y^i$ are the average estimated sources and f is a smoothed version of the logarithm of the source density d by convolution with a Gaussian kernel. As explained in [Jung et al., 1997], the density d does not need to be known and can be estimated easily as soon as we know if sources are sub- or super-Gaussian. In neuroscience, we commonly assume that neural sources are super-Gaussian.

3.2 Unmixing matrix optimization

Optimizing Equation (2) with respect to all parameters at once is complex since the variables τ^i are discrete, while the W^i are continuous, so we choose to minimize it iteratively by block coordinate descent. First, \mathcal{L} is alternatively minimized with respect to each W^i . Subsequently, the minimization process takes place alternately with respect to each τ^i , and so forth.

By applying $\mathcal{T}_{\tau^i}(\cdot)$ in the squared norm and in function f , optimization of \mathcal{L} with respect to W^i boils down to minimizing:

$$\begin{aligned} \mathcal{L}^i(W^i) = & -n \log |W^i| \\ & + \frac{m-1}{2m\sigma^2} \left\| W^i X^i - \sum_{j \neq i} \frac{\mathcal{T}_{\tau^i - \tau^j}(W^j X^j)}{m-1} \right\|^2 \\ & + f \left(\frac{1}{m} W^i X^i + \frac{1}{m} \sum_{j \neq i} \mathcal{T}_{\tau^i - \tau^j}(W^j X^j) \right). \end{aligned}$$

This formula strongly looks like Equation (4) in [Richard et al., 2020]. Indeed, the only difference lies in the fact that we obtained $\mathcal{T}_{\tau^i - \tau^j}(W^j X^j)$ instead of $W^j X^j$ in both sums over j . Intuitively, in the squared norm, we compare $W^i X^i$ to the quantity $\frac{1}{m-1} \sum_{j \neq i} \mathcal{T}_{\tau^i - \tau^j}(W^j X^j)$, where the last term is the shared source estimate of all subjects but subject i , deliberately delayed by subject i 's delays. Having similar formulas allows us to derive the relative gradient expression and Hessian approximation of \mathcal{L} with respect to W^i immediately. To do so, we need to replace the estimated sources $W^j X^j$ of subject $j \neq i$ by its shifted version $\mathcal{T}_{\tau^i - \tau^j}(W^j X^j)$ in the formulas of gradient $G^i \in \mathbb{R}^{p \times p}$ and Hessian approximation $H^i \in \mathbb{R}^{p \times p \times p \times p}$ in [Richard et al., 2020]. Then, we use a quasi-Newton approach by computing a direction $D = -(H^i)^{-1} G^i$ and finding with line search a step size ρ such that $\mathcal{L}^i((I_p + \rho D)W^i) < \mathcal{L}^i(W^i)$.

3.3 Delay optimization

Equation 2 should be optimized with respect to τ^i . Because of its complex form, the term $f(\bar{S})$ is burdensome to minimize rapidly, so we chose to discard it. Indeed, minimization of the true loss function would require an exhaustive search that has a computational complexity of $o(\tau_{\max} \times n)$, while using only the squared norm part will allow us to use cross-correlation to quickly compute the error, hence a computational complexity of $o(\tau_{\max} \times \log(n))$. As the first term, $-\sum_{i=1}^m \log |W^i|$, is independent from τ^i , the only part that needs to be minimized is:

$$\mathcal{L}^{i'}(\tau^i) = \sum_{i=1}^m \|Y^i - \bar{S}\|^2 = \sum_{i=1}^m \sum_{j=1}^p \|\mathbf{y}_j^i - \bar{\mathbf{s}}_j\|^2,$$

where $\mathbf{y}_j^i \in \mathbb{R}^n$ (resp. $\bar{\mathbf{s}}_j \in \mathbb{R}^n$) is the j -th row of Y^i (resp. \bar{S}). Since $\mathbf{y}_j^i - \bar{\mathbf{s}}_j$ only depends on τ_j^i , and not on τ_k^i for $k \neq j$, we can minimize $\mathcal{L}^{i'}$ with respect to each τ_j^i separately.

Thus, optimization boils down to finding, for each source $j \in \{1, \dots, p\}$ separately, the integer τ_j^i that maximizes $\langle \mathbf{y}_j^i, \mathbf{y}_j^{-i} \rangle$, where $\mathbf{y}_j^{-i} = \frac{1}{m-1} \sum_{k \neq i} \mathbf{y}_j^k$. This step is simply done by cross-correlation. Stacking the results in a vector gives $\tau^i = \{\tau_1^i, \dots, \tau_p^i\}$ and concludes the delay optimization step of subject i .

Note that we have to set a maximum delay parameter τ_{\max} in MVICAD. This parameter represents the maximum delay possible, in absolute value, for each source and each subject, that our algorithm will test. In the neuroscience application, there are natural choices for τ_{\max} (see Section 4.2).




Figure 2: Left: true generated shared sources. Middle: sources estimated by MVICA. Right: sources estimated by MVICAD.

Furthermore, we compared on simulations our method with an exhaustive search that includes the last term of Equation 2 in delay optimization. We observed that ignoring this term does not affect performance.

4 EXPERIMENTS

All the code is written in Python and is available on our [GitHub repo](#).

4.1 Simulation study

To evaluate the performance of the MVICAD algorithm, we conducted a number of simulations allowing us to quantify the error in parameter estimation.

In synthetic experiments, we simulate data according to our model (1). We use 5 subjects, 3 sources, and 700 samples, as this number of samples is comparable to the one available for MEG data (See below).

In Figure 2, we also use a maximum delay parameter of 40, as this number corresponds to 40 ms if the sampling rate is 1000 Hz, which is the typical maximum amount of delay in the context of MEG signals used in real data experiments [Price et al., 2017]. This figure (left) shows the true simulated shared sources. In the middle, we observe that sources found by MVICA do not fit these true sources, contrary to the sources found by MVICAD (right).

4.1.1 Amari distance

As classically done for ICA algorithms, we use the Amari distance [Moreau and Macchi, 1998] to quantify the error on mixing matrices. The benefit of this metric is that it is scale and permutation invariant. Figure 3 represents the Amari distance of both MVICA and MVICAD for several levels of delay introduced in the model. We see that our algorithm outperforms MVICA in this particular setup, as soon as there is some delay. In fact, the curves seem to be approximately linear. They start from the same point when there is no delay, and when there is a delay of 40, MVICA’s Amari distance is 1.18, whereas the one of MVICAD is equal to 0.68.

Note that, for each level of delay in the x -axis, the maximum delay parameter of MVICAD is set to this level of delay precisely. Consequently, both algorithms have the same Amari distance when the delay equals 0. Indeed, when the maximum delay parameter is set to 0, our algorithm strictly corresponds to MVICA.

4.1.2 Artificial delays

Our next experiment illustrates the ability of the model to recover delays. We used 40 subjects, 5 sources, 700 samples, a maximum delay parameter of 40, and a quantity of noise that gave a signal-to-noise ratio approximately equal to 5.

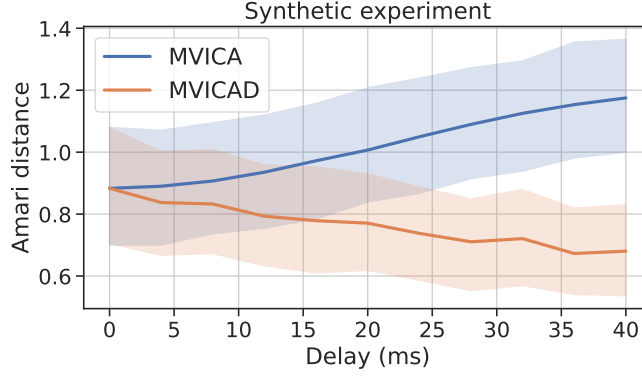


Figure 3: Amari distance of both MVICA and MVICAD for several levels of delay. We observe that our algorithm outperforms MVICA as soon as some delay is introduced in the model.

Figure 4 shows the results of this experiment. In total, there are $40 \times 5 = 200$ estimated delays. We observe that estimated delays are strongly correlated with artificial delays: the slope of the fit function is almost 1, and so does the R^2 score.

Note that not having a slope of exactly 1 can be partially explained by the presence of the noise. Searching a delay for each source of each subject means having mp delay parameters and is thus sensitive to noise. Also, it can be explained by the fact that the number of samples is rather low (700). Nevertheless, this experiment shows the ability of MVICAD to retrieve the true delays.

4.2 Real data experiments

In order to test MVICAD on real data, we use the Cam-CAN dataset. In this dataset delays usually do not exceed 60 ms for visual tasks and 20 ms for auditory tasks [Price et al., 2017]. Since the sampling rate is equal to 1000 Hz in the data, we set τ_{\max} to 60 (number of samples).

4.2.1 Cam-CAN dataset

The Cam-CAN dataset [Taylor et al., 2017] is a large and comprehensive dataset of neuroimaging, cognitive, and demographic data collected from a group of healthy adults. The dataset includes data from 661 participants, ranging in age from 18 to 88 years old, with equal numbers of participants in each 10-year age range. The participants were recruited from the general population in the Cambridgeshire area of the UK, with the aim of recruiting a sample that was representative of the local population in terms of age, sex, and education level.

This dataset is freely available to researchers and has become a resource for investigating the neurobiological mechanisms underlying healthy aging and age-related cognitive decline.

The Cam-CAN Stage 2 repository contains a subset of data from the larger Cam-CAN dataset, focused on the second wave of assessments conducted approximately two years after the initial assessment. The MEG data in the Stage 2 repository includes recordings of brain activity using a whole-head 306 channel Elekta Neuromag Vectorview system. It contains data from several cognitive tasks, including visual and auditory tasks. For each of these tasks, participants were presented a series of stimuli, and averaging the obtained signals give what is called evoked data.

In our experiments, we focus on evoked data of both tasks. We had to preprocess data and remove subjects with unusual signals, which reduced the number of subjects for the visual and auditory tasks. Also, these data last 0.7 second at a sampling rate of 1000 Hz, thus giving datasets of shape (477, 306, 701) for the visual task and (501, 306, 701) for the auditory one.

4.2.2 Delay and age correlation

The next experiment consists in studying age-delay correlation, as Cam-CAN data are often used to investigate mechanisms underlying healthy aging and age-related cognitive decline.

Figure 5 shows the scatter plot of age and delay of the 477 subjects of the visual task. This scatter plot shows a significant correlation between the age of subjects and estimated delays. To produce it,




Figure 4: Scatter plot of artificial and estimated delays. We used 40 subjects and 5 sources, and each point represents one source of one subject.




Figure 5: Scatter plot of age and delay of the 477 subjects of the visual dataset. We observe a significant correlation between aging and estimated delay.

we first applied a PCA with 10 components on the dataset, whose shape became (477, 10, 701), and then used MVICAD to get estimated delays of shape (477, 10). We report here the source whose delays are most linearly related to age (using R^2 score, and p-value of Pearson correlation). From a biological point of view, having relevant and non-relevant sources in terms of age-delay correlation makes sense because some of the 10 obtained sources represent noise, whereas others are closer to the actual neural signals affected by aging.

However, we can also investigate the model with only one delay per subject. In other words, we force estimated delays to be equal across sources in our algorithm, i.e. we only estimate one delay per subject. This method reveals some aging effects without having to visually explore delays of each source, yet it is not source-specific. The obtained results are: slope = 0.08 ; $R^2 = 0.08$; p-value = 9.69×10^{-11} . The drawback of having one delay per subject is that it produces a higher negative log-likelihood and poorer source reconstruction.

Finally, note that this experiment produces lower results for the auditory task. It was expected since auditory stimuli are known to present “cumulative delays” instead of “constant delays” [Roberts et al., 2010].

5 CONCLUSION

We have proposed an improvement of the algorithm MVICA which aims to retrieve latent sources from multiple views and can be used in various domains. To do so, we estimated delays that appear in each source of each subject and use them to correct our source reconstruction during optimization. We presented the likelihood of our model in closed form and proposed a way of minimizing it with respect to both unmixing matrices and delays. In the context of MEG signals, we showed that our estimated delays were significantly correlated to age and demonstrated on a synthetic experiment that our method outperforms MVICA in terms of Amari distance.

6 ACKNOWLEDGEMENT

Data collection and sharing for this project was provided by the Cambridge Centre for Ageing and Neuroscience (CamCAN). CamCAN funding was provided by the UK Biotechnology and Biological Sciences Research Council (grant number BB/H008217/1), together with support from the UK Medical Research Council and University of Cambridge, UK. Data used in the preparation of this work were obtained from the CamCAN repository (available at <http://www.mrc-cbu.cam.ac.uk/datasets/camcan/>). This work was supported by the ANR BRAIN (ANR-20-CHIA0016) grant.

References

- Ablin, P., Cardoso, J.-F., and Gramfort, A. (2018). Faster independent component analysis by preconditioning with hessian approximations. *IEEE Transactions on Signal Processing*, 66(15):4040–4049.
- Calhoun, V. D., Liu, J., and Adalı, T. (2009). A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. *Neuroimage*, 45(1):S163–S172.
- Comon, P. (1994). Independent component analysis, a new concept? *Signal processing*, 36(3):287–314.
- Correa, N. M., Adalı, T., Li, Y.-O., and Calhoun, V. D. (2010). Canonical correlation analysis for data fusion and group inferences. *IEEE signal processing magazine*, 27(4):39–50.
- Hyvarinen, A. (1999). Fast and robust fixed-point algorithms for independent component analysis. *IEEE transactions on Neural Networks*, 10(3):626–634.
- Hyvärinen, A. and Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural networks*, 13(4-5):411–430.
- Jung, T.-P., Humphries, C., Lee, T.-W., Makeig, S., McKeown, M., Iragui, V., and Sejnowski, T. J. (1997). Extended ICA removes artifacts from electroencephalographic recordings. *Advances in neural information processing systems*, 10.
- Kettenring, J. R. (1971). Canonical analysis of several sets of variables. *Biometrika*, 58(3):433–451.
- Makeig, S., Bell, A., Jung, T.-P., and Sejnowski, T. J. (1995). Independent component analysis of electroencephalographic data. *Advances in neural information processing systems*, 8.
- Moreau, E. and Macchi, O. (1998). Self-adaptive source separation. ii. comparison of the direct, feedback, and mixed linear network. *IEEE transactions on signal processing*, 46(1):39–50.
- Pandeva, T. and Forré, P. (2022). Multi-view independent component analysis with shared and individual sources. *arXiv preprint arXiv:2210.02083*.
- Price, D., Tyler, L. K., Neto Henriques, R., Campbell, K. L., Williams, N., Treder, M. S., Taylor, J. R., and Henson, R. (2017). Age-related delay in visual and auditory evoked responses is mediated by white-and grey-matter differences. *Nature communications*, 8(1):15671.
- Richard, H., Gresele, L., Hyvarinen, A., Thirion, B., Gramfort, A., and Ablin, P. (2020). Modeling shared responses in neuroimaging studies through multiview ICA. *Advances in Neural Information Processing Systems*, 33:19149–19162.
- Roberts, T. P., Khan, S. Y., Rey, M., Monroe, J. F., Cannon, K., Blaskey, L., Woldoff, S., Qasmieh, S., Gandal, M., Schmidt, G. L., et al. (2010). MEG detection of delayed auditory evoked responses in autism spectrum disorders: towards an imaging biomarker for autism. *Autism Research*, 3(1):8–18.
- Taylor, J. R., Williams, N., Cusack, R., Auer, T., Shafto, M. A., Dixon, M., Tyler, L. K., Henson, R. N., et al. (2017). The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: Structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *neuroimage*, 144:262–269.
- Tsatsishvili, V., Cong, F., Toivainen, P., and Ristaniemi, T. (2015). Combining PCA and multiset CCA for dimension reduction when group ICA is applied to decompose naturalistic fMRI data. In *2015 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6. IEEE.
- Varoquaux, G., Sadaghiani, S., Poline, J. B., and Thirion, B. (2009). CanICA: Model-based extraction of reproducible group-level ICA patterns from fMRI time series. *arXiv preprint arXiv:0911.4650*.
- Vigário, R., Sarela, J., Jousmiki, V., Hamalainen, M., and Oja, E. (2000). Independent component approach to the analysis of EEG and MEG recordings. *IEEE transactions on biomedical engineering*, 47(5):589–593.

A Appendix

Let us prove the identifiability of our model. We consider multiple i.i.d. r.v. $\mathbf{s}_k \in \mathbb{R}^p$, $k = 1, \dots, \tau_{\max}$, that model sources, instead of only one r.v. \mathbf{s} , as is often the case. Having multiple variables allows us to delay them with operator \mathcal{T} . For simplicity, we define $\mathbf{s} := \mathbf{s}_1$. Recall that our model is:

$$\mathbf{x}^i = A^i(\mathcal{T}_{\boldsymbol{\tau}^i}(\mathbf{s}) + \mathbf{n}^i), \quad i = 1, \dots, m,$$

with $\boldsymbol{\tau}^i \in \mathbb{R}^p$. Note that, compared to Equation (1), we used vectors instead of matrices. Assume that we also have:

$$\mathbf{x}^i = A'^i(\mathcal{T}_{\boldsymbol{\tau}'^i}(\mathbf{s}') + \mathbf{n}'^i), \quad i = 1, \dots, m,$$

for some $A'^i, \boldsymbol{\tau}'^i, \mathbf{s}'$ and \mathbf{n}'^i . Note that we also define $\mathbf{s}' := \mathbf{s}'_1$ where $\mathbf{s}'_k \in \mathbb{R}^p$, $k = 1, \dots, \tau_{\max}$, are i.i.d. r.v.

Since we assume that \mathbf{s}_k has non-Gaussian independent components (whose densities are not reduced to a point-like mass) and that the r.v. \mathbf{s}_k are i.i.d. and that \mathbf{n}^i is Gaussian decorrelated, then $\mathcal{T}_{\boldsymbol{\tau}^i}(\mathbf{s}) + \mathbf{n}^i$ also has non-Gaussian independent components. And so does $\mathcal{T}_{\boldsymbol{\tau}'^i}(\mathbf{s}') + \mathbf{n}'^i$. Following [Comon, 1994], Theorem 11, there exists a scale-permutation matrix P^i such that $A'^i = A^i P^i$. As a consequence, and since A^i is invertible, we have, for all i :

$$\begin{aligned} \mathcal{T}_{\boldsymbol{\tau}^i}(\mathbf{s}) + \mathbf{n}^i &= P^i(\mathcal{T}_{\boldsymbol{\tau}'^i}(\mathbf{s}') + \mathbf{n}'^i) \\ \Rightarrow \mathbf{s} + \mathcal{T}_{-\boldsymbol{\tau}^i}(\mathbf{n}^i) &= \mathcal{T}_{-\boldsymbol{\tau}^i}(P^i(\mathcal{T}_{\boldsymbol{\tau}'^i}(\mathbf{s}') + \mathbf{n}'^i)) \\ &= P^i(\mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}') - \mathcal{T}_{-(P^i)^\top \boldsymbol{\tau}^i}(\mathbf{n}'^i)). \end{aligned}$$

The last line comes from the equality $\mathcal{T}_{\boldsymbol{\tau}} P = P \mathcal{T}_{P^\top \boldsymbol{\tau}}$ for a scale-permutation matrix P and a vector $\boldsymbol{\tau}$.

We focus on subject 1 and subject $i \neq 1$:

$$\begin{aligned} \mathbf{s} + \mathcal{T}_{-\boldsymbol{\tau}^1}(\mathbf{n}^1) - (\mathbf{s} + \mathcal{T}_{-\boldsymbol{\tau}^i}(\mathbf{n}^i)) \\ &= P^1(\mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}') + \mathcal{T}_{-(P^1)^\top \boldsymbol{\tau}^1}(\mathbf{n}'^1)) \\ &\quad - P^i(\mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}') + \mathcal{T}_{-(P^i)^\top \boldsymbol{\tau}^i}(\mathbf{n}'^i)). \end{aligned}$$

Consequently,

$$\begin{aligned} P^1 \mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}') - P^i \mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}') \\ &= \mathcal{T}_{-\boldsymbol{\tau}^1}(\mathbf{n}^1) - \mathcal{T}_{-\boldsymbol{\tau}^i}(\mathbf{n}^i) \\ &\quad + P^i \mathcal{T}_{-(P^i)^\top \boldsymbol{\tau}^i}(\mathbf{n}'^i) - P^1 \mathcal{T}_{-(P^1)^\top \boldsymbol{\tau}^1}(\mathbf{n}'^1). \end{aligned} \tag{3}$$

Since the right-hand side of the last equality is a linear combination of Gaussian random variables, this implies that $P^1 \mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}') - P^i \mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}')$ is also Gaussian.

Let us show that there exists a vector $\boldsymbol{\tau} \in \mathbb{R}^p$ such that, for all i , $\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1 = \boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i =: \boldsymbol{\tau}$. Note that this is an equality modulo τ_{\max} . Indeed, delaying sources by $\boldsymbol{\tau}$ and by $\boldsymbol{\tau} + q\boldsymbol{\tau}_{\max}$, $q \in \mathbb{Z}$, is equivalent, according to our definition of the shift operator \mathcal{T} . For simplicity, we omit the modulo part in equalities about delays.

By contradiction, let us suppose that $\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1 \neq \boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i$. Let us call $C \subset \{1, \dots, p\}$ the set of indices such that $(\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1)_C = (\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i)_C$ and let $\bar{C} = \{1, \dots, p\} \setminus C$. By assumption, $\bar{C} \neq \emptyset$. By definition of C , $(\mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}'))_C = (\mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}'))_C$, and $(\mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}'))_{\bar{C}}$ and $(\mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}'))_{\bar{C}}$ are independent.

Thus, the left-hand side of Equation (3) only depends on three terms: $(\mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}'))_C$, $(\mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}'))_{\bar{C}}$ and $(\mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}'))_{\bar{C}}$, which are all three independent from each others.

The left-hand side of Equation (3) is thus a linear combination of independent random variables and is Gaussian. So, by Cramér's lemma, we should have that $\mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}')$ and $\mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}')$ are Gaussian too, which is absurd, given that \mathbf{s}' is assumed to be non-Gaussian. So, there exists a vector $\boldsymbol{\tau}$ such that $\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1 = \boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i =: \boldsymbol{\tau}$.

Let us define $\tilde{\mathbf{s}} := \mathcal{T}_{\boldsymbol{\tau}'^1 - (P^1)^\top \boldsymbol{\tau}^1}(\mathbf{s}') = \mathcal{T}_{\boldsymbol{\tau}'^i - (P^i)^\top \boldsymbol{\tau}^i}(\mathbf{s}')$. Since the right-hand side of Equation (3) is Gaussian, it follows that $(P^1 - P^i) \tilde{\mathbf{s}}$ is Gaussian too. Thus, the equality holds only if $P^1 = P^i$. Therefore, the matrices P^i are all equal, and there exists a scale-permutation matrix P such that $A'^i = A^i P$.

In conclusion, we proved that there exists a scale and permutation matrix P and a vector $\boldsymbol{\tau} \in \mathbb{R}^p$ such that, for all i , $A'^i = A^i P$ and $\boldsymbol{\tau}'^i = P^\top \boldsymbol{\tau}^i + \boldsymbol{\tau}$.