

Source Separation of Phase-Locked Subspaces

Miguel Almeida^{1,2*}, José Bioucas-Dias¹, and Ricardo Vigário²

miguel.almeida@lx.it.pt, jose.bioucas@lx.it.pt, ricardo.vigario@hut.fi

¹ Institute of Telecommunications, Instituto Superior Técnico, Lisbon, Portugal

<http://www.ist.utl.pt/en/>

² Adaptive Informatics Research Centre

Aalto University School of Science and Technology, Finland

<http://www.aalto.fi/en/>

Abstract. We present a two-stage algorithm to perform blind source separation of sources organized in subspaces, where sources in different subspaces have zero phase synchrony and sources in the same subspace have full phase synchrony. Typical techniques such as ICA algorithms are not adequate for such signals, because phase-locked signals are not independent. We demonstrate the usefulness of this algorithm on a simulated dataset. The results show that the algorithm works very well in low-noise situations. We also discuss the necessary improvements to be made before the algorithm is able to deal with real-world signals.

Key words: phase-locking factor (PLF), synchrony, blind source separation (BSS), independent component analysis (ICA), subspaces, TDSEP

1 Introduction

Recently, synchrony phenomena have been studied with increasing frequency by the scientific community. Such phenomena have been observed in many different physical systems, including electric circuits, laser beams and human neurons [2].

Synchrony is believed to play a relevant role in the way different parts of human brain interact. For example, it is known that when humans engage in a motor task, several brain regions oscillate coherently [3, 4]. Also, several pathologies such as autism, Alzheimer and Parkinson are thought to be associated with a disruption in the synchronization profile of the brain (see [5] for a review).

To perform inference on the networks present in the brain or in other real-world systems, it is important to have access to the dynamics of the individual oscillators (which we will call “sources”). Unfortunately, in many real applications including brain electrophysiological signals (EEG and MEG), the signals from individual oscillators are not directly measurable, and one only has access to a superposition of the sources. For example, in EEG and MEG the signals measured in one sensor contain components coming from several brain regions [6]. In this case, spurious synchrony occurs, as we will show in this paper.

* Corresponding author.

Undoing this superposition is typically called a blind source separation (BSS) problem. One usually assumes that the mixing is linear and instantaneous, which is a valid approximation in brain signals [7]. However, independence of the sources is not a valid assumption, because phase-locked sources are highly dependent. In this paper we address the problem of how to separate phase-locked sources. We have previously addressed this problem with relative success [1], and in this paper we propose an improved version of that approach which yields better results, is faster and has fewer limitations. The separation algorithm we propose uses TDSEP [13] as an initialization and then uses only the phase information of the signals. The amplitude information is discarded, because signals may exhibit synchrony even when their amplitudes are uncorrelated [9].

The algorithm presented here assumes nothing specific of brain signals, and should work in any situation where phase-locked sources are mixed approximately linearly and noise levels are low.

2 Background, Notation and Algorithm

Given two oscillators with phases $\phi_j(t)$ and $\phi_k(t)$ for $t = 1, \dots, T$, the Phase Locking Factor (PLF) between those two oscillators is defined as

$$\varrho_{jk} = \left| \frac{1}{T} \sum_{t=1}^T e^{i[\phi_j(t) - \phi_k(t)]} \right| = \left| \left\langle e^{i(\phi_j - \phi_k)} \right\rangle \right|, \quad (1)$$

where $\langle \cdot \rangle$ is the time average operator. The PLF obeys $0 \leq \varrho_{jk} \leq 1$. The value $\varrho_{jk} = 1$ corresponds to two oscillators that are perfectly synchronized, i.e., that have a constant phase lag. The value $\varrho_{jk} = 0$ is attained if the two oscillators' phases are not correlated, as long as the observation period T is sufficiently long. The oscillators can be represented by real signals, since techniques such as the Hilbert Transform [10] can be used to extract the phase of real signals.¹ Typically, the PLF values are stored in a PLF matrix \mathbf{Q} such that $\mathbf{Q}(j, k) = \varrho_{jk}$.

We assume that we have a number of signals (called “sources”) that are organized in subspaces, such that the PLF between sources belonging to the same subspace is high, and the PLF between sources in different subspaces is low. Let $\mathbf{s}(t)$, for $t = 1, \dots, T$, denote the vector of true sources and $\mathbf{y}(t) = \mathbf{A}\mathbf{s}(t)$ denote the mixed signals, where \mathbf{A} is the mixing matrix, which is assumed to be square and non-singular. Our goal is to find an unmixing matrix \mathbf{W} such that the estimated sources $\mathbf{x}(t) = \mathbf{W}^T \mathbf{A}\mathbf{s}(t)$ are as close to the true sources as possible, up to permutation, scaling, and sign.

We now illustrate the effect of a linear mixture on PLF values. If the sources have very low or very high synchrony (PLF ≈ 0 or 1), their mixtures will have intermediate values of synchrony. This is illustrated in the top two rows of Fig. 1. Note that significant partial synchrony is present in all pairs of mixture signals.

¹ The Hilbert Transform, and other similar techniques, can be used only if the signals have sufficiently narrow Fourier spectra.

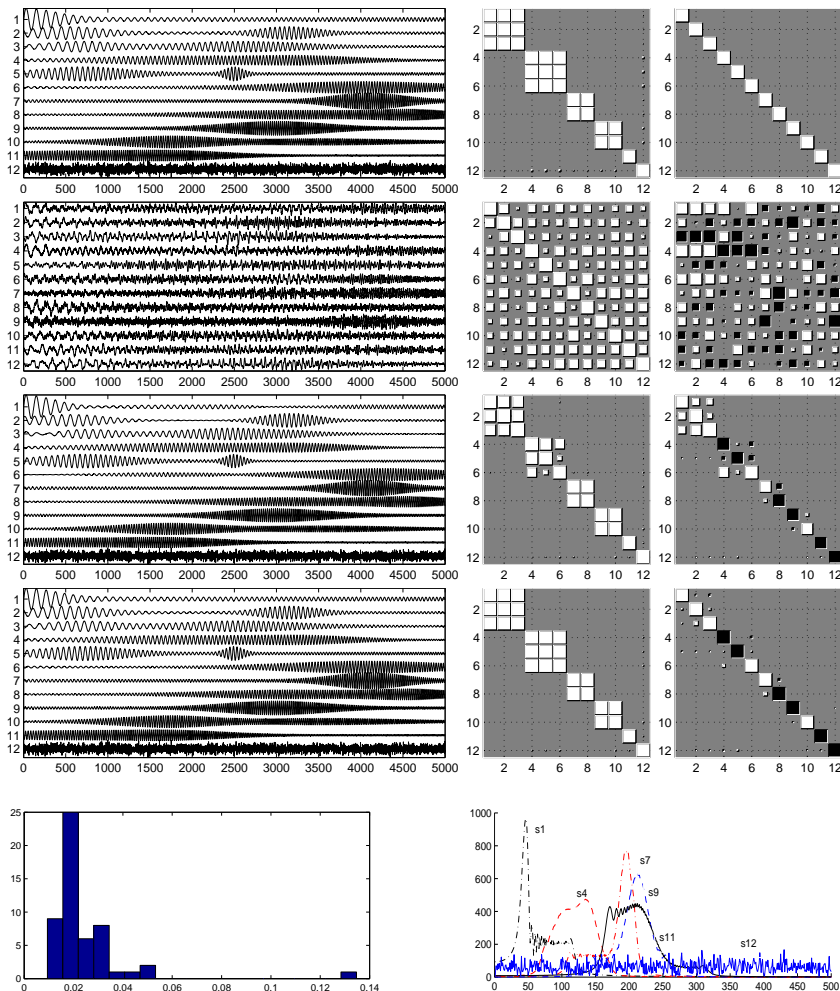


Fig. 1. The dataset used throughout this paper. (*First row*) Original sources (*left*), PLFs between them (*middle*) and the identity matrix (*right*), symbolizing that these are the true sources. (*Second row*) Mixed signals, PLFs between them and the mixing matrix \mathbf{A} . (*Third row*) Sources resulting from TDSEP (*left*). Note that the inter-subspace PLFs (*middle*) are very close to zero, but the intra-subspace PLFs are not all close to 1. Further, the intra-space separation is poor, as can be seen from inspection of the product $\mathbf{W}_{tdsep}^T \mathbf{A}$ (*right*). (*Fourth row*) Results found after the second stage of the algorithm. The estimated sources (*left*) are very similar to the original ones. This is corroborated by the PLFs between the estimated sources (*middle*) and the final unmixing matrix (*right*). In the third and fourth rows, the permutation was corrected manually. (*Bottom row*) Histogram of the Amari Performance Index for 100 runs (*left*), corresponding to 100 random mixing matrices for these sources, and Discrete Fourier Transform of the first, fourth, seventh, ninth, eleventh and twelfth sources, one from each subspace (*right*).

Since linearly mixing sources which have PLFs close to 1 or 0 yields signals which have partial synchrony, one can reason that finding an unmixing matrix such that the estimated sources have PLFs of 1 or 0 can separate such sources. We have previously used such an approach, motivated by ICA, to separate this kind of sources [1]. That approach, which we called Independent Phase Analysis (IPA), showed decent results, but was limited to near-orthogonal mixing matrices. A non-near-orthogonal mixing matrix yielded, however, poor results. This poor performance is closely related to the fact that, given two phase-locked signals, it is always possible to construct a linear combination of them which has a PLF of zero with those signals.² This implies that our problem is ill-posed, when the goal is to find an unmixing matrix such that the PLF matrix has only zeroes and ones. Some zero PLFs may, in fact, correspond to phase-locked signals.

To avoid the above referred ill-posedness, we introduce in this paper a different unmixing criterion which consists in dividing the problem into two subproblems: first, separate the subspaces from one another, even if within the subspaces some mixing remains. Second, unmix the sources within each subspace. We now discuss each of these subproblems in detail and explain why this new problem is no longer ill-posed. Since this method is an improvement of the previous approach [1], we will continue to name it Independent Phase Analysis.

2.1 Inter-subspace separation and subspace detection

The objective of the first stage is to find an unmixing matrix \mathbf{W} such that the estimated subspaces are correct, even if sources within each subspace are still mixed. We assume that signals in different subspaces have little interaction with each other, which should usually correspond to a distinct time structure. Therefore, techniques that use temporal structure to perform separation should be adequate to this first stage. We chose to use Ziehe *et al.*'s implementation of TDSEP [13] for this first subproblem, but SOBI [14] can be used instead. Although we don't know any theoretical results that support TDSEP's adequacy to this task, we have repeatedly observed that it separates subspaces quite well.

A non-trivial step is the detection of subspaces from the result of TDSEP. From our experience, TDSEP can perform the inter-subspace separation very well but cannot adequately do the intra-subspace separation. This means that PLF values within each subspace will be underestimated. Naively, one can arbitrarily define a hard PLF threshold, below which signals are considered not synchronized and above which signals are considered in full phase synchrony. However, such an approach fails to account for indirect interactions, such as "signal A phase-locked with signal B which is phase-locked with signal C". To account for such interactions, we first raise the PLF matrix to the second power.

The matrix resulting from this hard thresholding should be block-diagonal, with each block having all elements equal to 1. If this is the case, we can group the found signals into subspaces, and we further unmix each of these subspaces

² Unfortunately, the proof of this claim is too lengthy to show here.

at a time (see section 2.2). By inverting and transposing the mixing matrix estimated by TDSEP, we have a first estimate of the unmixing matrix \mathbf{W}_{tdsep} .

If the matrix resulting from the thresholding is not block-diagonal with blocks with all elements equal to 1, our algorithm considers that the subspaces were wrongly detected and fails. See section 4 for possible improvements on this.

2.2 Intra-subspace separation

In the second stage of IPA, we begin by selecting the subset of columns of the unmixing matrix \mathbf{W}_{tdsep} found by TDSEP that form the l -th subspace, which we denote by S_l . We construct a rectangular matrix $\mathbf{W}_{tdsep,l}$ from those columns. Denote the number of signals in the l -th subspace by N_l , and let $\mathbf{z}_l(t) = \mathbf{W}_{tdsep,l}^\top \mathbf{y}(t)$ be the vector of the sources of subspace S_l estimated by TDSEP.

In this second stage, our goal is to separate each of these subsets of sources. As explained above, the true sources should have a PLF of 1 with the sources in the same subspace. We should therefore unmix the N_l sources found by TDSEP such that their PLFs are as high as possible. Mathematically, this corresponds to finding a N_l by N_l matrix \mathbf{W}_l such that the estimated sources in the l -th subspace, $\mathbf{x}_l(t) = \mathbf{W}_l^\top \mathbf{W}_{tdsep,l}^\top \mathbf{y}(t) = \mathbf{W}_l^\top \mathbf{z}_l(t)$, have the highest possible PLFs.

In this second stage, for each subspace l , the objective function to be maximized is

$$J_l = (1 - \lambda) \sum_{j,k \in S_l} \varrho_{jk}^2 + \lambda \log |\det \mathbf{W}_l|, \quad (2)$$

where $\sum_{j,k \in S_l} \varrho_{jk}^2$ is a sum over all pairs of sources in subspace S_l of the PLF between those sources squared, the second term, similarly to ICA [8], penalizes unmixing matrices that are close to singular, and λ is a parameter controlling the relative weight of the two terms. The second term, already present in the previous version of IPA [1], serves the purpose of preventing the algorithm from finding solutions which trivially have $\varrho_{jk} = 1$. Each column of \mathbf{W} must be constrained to have unit norm to prevent trivial decreases of the penalty term.

With this formulation, the problem is no longer ill-posed as described above. Furthermore, we now need only optimize a subset of parameters at a time. This can drastically reduce the time needed to separate a set of sources.

The gradient of J_l relative to an entry w_{ij} of the weight matrix \mathbf{W}_l is given by (we omit many dependences on l for clarity)

$$\frac{\partial J_l}{\partial w_{ij}} = (1 - \lambda) 4\pi \sum_{k=1}^N [2\varrho_{jk}] \left\langle \sin(\Psi_{jk} - \Delta\phi_{jk}) \frac{Z_i}{X_j} \sin(\psi_i - \phi_j) \right\rangle - \lambda [\mathbf{W}_l^{-\top}]_{ij} \quad (3)$$

where ϱ_{jk} is the PLF between estimated sources j and k , $Z_i = |\tilde{z}_i|$ where \tilde{z}_i is the analytic signal of the i -th source estimated by TDSEP (obtained from the Hilbert transform [10]), $X_j = |\tilde{x}_j|$ where \tilde{x}_j is the analytic signal of the j -th estimated source, $\psi_i = \text{angle}(\tilde{z}_i)$ is the phase of the i -th source found by TDSEP, $\phi_j = \text{angle}(\tilde{x}_j)$ is the phase of the j -th estimated source, $\Delta\phi_{jk} = \phi_j - \phi_k$ is the phase difference of estimated sources j and k , $\Psi_{jk} = \text{angle}(\langle e^{i\Delta\phi_{jk}} \rangle)$ is

the average phase difference between estimated sources j and k , and $[\mathbf{W}_l^{-T}]_{ij}$ is the (i, j) element of the inverse of \mathbf{W}_l^T .

3 Results

We present results showing that this new approach provides drastic improvements on the separation quality. The optimization was done using a gradient algorithm with adaptive step sizes, running up to 600 gradient iterations until the average PLF within the subspace is greater than 0.9999. λ was hand tuned for optimal performance. However, the algorithm yields similar results for λ within a factor or 2 of the optimal one.

We simulate the noiseless instantaneous linear mixture of 12 sources depicted in the first row of Fig. 1. These sources belong to 6 clusters of sizes 3, 2 and 1. We generate 100 datasets from these sources, by generating 100 mixing matrices, each of which with elements i.i.d. from the Uniform(-1,1) distribution. We then run the algorithm once for each mixing matrix, for a total of 100 runs. Each run took about 1 minute on a modern laptop computer.

The second row of Fig. 1 shows the mixed signals which are the input to our algorithm. The third row shows the sources estimated by TDSEP. Inspection of the PLFs between these sources (shown on the subfigure in the second column, third row) shows that some of the estimated sources do not have high PLFs, and an inspection of the product $\mathbf{W}_{tdsep}^T \mathbf{A}$ reveals that the inter-subspace separation was very good, but the intra-subspace separation was poor.

The fourth row of Fig. 1 shows that by maximizing the intra-subspace PLFs we can significantly improve the separation within each subspace. This is clearly visible in the product $\mathbf{W}^T \mathbf{A}$, depicted on the third column.

We measure the performance of IPA using the Amari Performance Index (API) [12], which measures the average relative contamination in each estimated source from all other sources. The API is non-negative and decreases to zero as the separation quality increases. A histogram of the API for the 100 runs of IPA is shown in Fig. 1. The mode of this histogram corresponds to an API of 0.0175. This is very similar to the example in Fig. 1, which has an API of 0.0171.

We used a threshold of 0.1 on the squared PLF matrix for the detection of subspaces. In 18% of the runs, the resulting matrix is not block-diagonal with each block full of ones, and therefore the algorithm stops.

4 Discussion

The above results demonstrate that IPA can successfully extract mixed sources based on their phase synchrony values and that it performs considerably better than the previous version. The previous version had 71% of separations with API below 0.3 [1], while this new version has 99% below 0.06, which is a remarkable improvement. Furthermore, the new version of IPA works well with any mixing matrix, even if it is close to singular. The previous version worked well only if the mixing matrix was close to orthogonal [1].




Fig. 2. Histograms of API for 100 runs of separate datasets containing a single cluster of 2 (*far left*), 3 (*middle left*), 4 (*middle right*) and 5 (*far right*) signals.

It should be noted that a PLF value of zero does not imply that the signals have distinct frequency spectra. In fact, the subspaces used here have overlapping frequency spectra (shown in the bottom right corner of Fig. 1), and an attempt at separating the subspaces based on Fourier transforms alone would fail.

Our selection of the value 0.1 for the PLF threshold was empirical. Although all the successful separations were very good, 18% of the runs yielded cases where our (admittedly crude) subspace detection procedure failed. This suggests that an improved subspace detection procedure is the most important improvement to be made, and it is one we are actively working on. Also, although TDSEP works quite well for the inter-subspace separation for the sources shown here, there is no certainty that it will perform as well for any kind of signals.

Although this algorithm works quite well for this simulated data, several improvements must be made for it to be usable in real data. First of all, we have noted that the performance of IPA degrades when it is used with large subspaces. Subspaces of 1, 2 or 3 signals can be separated quite well, but on subspaces of 4 signals the performance begins to decrease. For 5 or more signals in a subspace, performance becomes considerably worse. To illustrate this limitation, we present in Fig. 2 histograms of 100 runs of the algorithm on datasets with 2, 3, 4 and 5 signals, all belonging to a single subspace.

Another unclear aspect is how this algorithm will perform if the true sources do not have extreme values of PLF but more intermediate values. It is possible that the second stage of the algorithm will overfit in such situations.

The major hindrance to be overcome before IPA can be applied to real signals is to make it work under noisy conditions. IPA performs quite well for noiseless mixtures, but real signals will always have some amount of noise. Preliminary results show that the algorithm can tolerate small amounts of noise (up to 1% of noise amplitude relative to the signal amplitude), but this still needs to be improved prior to its application to real world signals.

5 Conclusion

We have presented a two-stage algorithm, called Independent Phase Analysis (IPA), to separate phase-locked subspaces from linear mixtures. We have shown that this approach yields much better results than the previous version of IPA, and that it is no longer limited to near-orthogonal mixing matrices. Our results show that although TDSEP alone is not enough to adequately separate phase-

locked sources, its conjunction with a subsequent intra-subspace separation gives very good separation quality in noise-free situations. Nevertheless, improvements are necessary before this algorithm can be applied to real-world signals.

Acknowledgments. MA is funded by scholarship SFRH/BD/28834/2006 of the Portuguese Foundation for Science and Technology. This study was partially funded by the Academy of Finland through its Centres of Excellence Program 2006-2011.

References

1. Almeida, M., Vigário, R.: Source-Separation of Phase-Locked Subspaces, Proceedings of the Independent Component Analysis Conference, 2009.
2. Pikovsky, A., Rosenblum, M., Kurths, J.: Synchronization: A Universal Concept in Nonlinear Sciences. Cambridge University Press (2001).
3. Palva, J.M., Palva, S., Kaila, K.: Phase Synchrony Among Neuronal Oscillations in the Human Cortex. *Journal of Neuroscience*. 25, 3962–3972 (2005).
4. Schoffelen, J.M., Oostenveld, R., Fries, P.: Imaging the Human Motor System’s Beta-Band Synchronization During Isometric Contraction. *NeuroImage*. 41, 437–447 (2008).
5. Uhlhaas, P.J., Singer, W.: Neural Synchrony in Brain Disorders: Relevance for Cognitive Dysfunctions and Pathophysiology. *Neuron*. 52, 155–168 (2006).
6. Nunez, P.L., Srinivasan, R., Westdorp, A.F., Wijesinghe, R.S., Tucker, D.M., Silberstein, R.B., and Cadusch, P.J.: EEG Coherency I: Statistics, Reference Electrode, Volume Conduction, Laplacians, Cortical Imaging, and Interpretation at Multiple Scales. *Electroencephalography and clinical Neurophysiology*. 103, 499–515 (1997).
7. Vigário, R., Särelä, J., Jousmäki, V., Hämäläinen, M., Oja, E.: Independent Component Approach to the Analysis of EEG and MEG Recordings. *IEEE Transactions On Biomedical Engineering*. 47, 589–593 (2000).
8. Hyvärinen, A., Karhunen, J., Oja, E.: Independent Component Analysis. John Wiley & Sons (2001).
9. Rosenblum, M.G., Pikovsky, A.S., Kurths, J.: Phase Synchronization of Chaotic Oscillators. *Physical Review Letters*. 76, 1804–1807 (1996).
10. Oppenheim, A.V., Schaffer, R.W., Buck, J.R.: Discrete-Time Signal Processing. Prentice-Hall International Editions (1999).
11. Schleimer, J.H., Vigário, R.: Order in Complex Systems of Nonlinear Oscillators: Phase Locked Subspaces. In: Proceedings of the European Symposium on Neural Networks (2007).
12. Amari, S., Cichocki, A., Yang, H.H.: A New Learning Algorithm for Blind Signal Separation. *Advances in Neural Information Processing Systems*. 8, 757–763 (1996).
13. Ziehe, A., Müller, K.-R.: TDSEP - an Efficient Algorithm for Blind Separation Using Time Structure. In: Proceedings of the International Conference on Artificial Neural Networks (1998).
14. Belouchrani, A., Abed-Meraim, K., Cardoso, J.-F., Moulines, E.: A Blind Source Separation Technique Using Second Order Statistics. *IEEE Transactions on Signal Processing*. 45, 434–444 (1997).