

A Tutorial Review on Multi-subject Decomposition of EEG

René J. Huster^{1,2,3} · Liisa Raud^{1,3}

Received: 7 July 2017 / Accepted: 11 October 2017 / Published online: 23 October 2017
© Springer Science+Business Media, LLC 2017

Abstract Over the last years we saw a steady increase in the relevance of big neuroscience data sets, and with it grew the need for analysis tools capable of handling such large data sets while simultaneously extracting properties of brain activity that generalize across subjects. For functional magnetic resonance imaging, multi-subject or group-level independent component analysis provided a data-driven approach to extract intrinsic functional networks, such as the default mode network. Meanwhile, this methodological framework has been adapted for the analysis of electroencephalography (EEG) data. Here, we provide an overview of the currently available approaches for multi-subject data decomposition as applied to EEG, and highlight the characteristics of EEG that warrant special consideration. We further illustrate the importance of matching one's choice of method to the data characteristics at hand by guiding the reader through a set of simulations. In sum, algorithms for group-level decomposition of EEG provide an innovative and powerful tool to study the richness of functional brain networks in multi-subject EEG data sets.

Keywords Group ICA · EEG · Multi-subject · Group-level · Blind source separation · Decomposition

Introduction

Applications of machine-learning techniques for neuroscience data, triggered much by developments in functional magnetic resonance imaging (fMRI) methodology, have gained a strong momentum in recent years. Group-level analysis of fMRI data, for example, directly infers the latent structure of processes in resting-, task-, or disease-states from multiple subjects concurrently (e.g., Calhoun and Adali 2012). This technique has been extremely successful due to its ability to specify and disambiguate the activity of otherwise simultaneously active functional networks of the brain. Similar applications have been developed for electroencephalography (EEG) data; however, these techniques are not that widely used yet. The fast-paced changes in neural states that characterize EEG recordings, along with inter-individual variations in scalp topographies as a result of differences in neuroanatomy, create a vast space of different scenarios to which no simple and generic solution exists so far (Makeig et al. 2004). Nonetheless, in recent years a number of methods have been proposed that altogether address the majority of possible scenarios, and thorough evaluations support their validity as methods for the estimation of functional networks from EEG data.

Here, we will discuss these different scenarios and how to address them by introducing a variety of current approaches for group-level component analyses of EEG data. We will start out with an overview of currently available approaches, focusing on a description at the conceptual level rather than elaborating on mathematical specifics or the computational implementation. Thereafter, the discussion of methods will

This is one of several papers published together in Brain Topography on the “Special Issue: Multisubject decomposition of EEG—methods and applications”.

✉ René J. Huster
rene.huster@psykologi.uio.no

¹ Multimodal Imaging and Cognitive Control Lab, Department of Psychology, University of Oslo, Oslo, Norway

² Psychology Clinical Neurosciences Center, University of New Mexico, Albuquerque, USA

³ Cognitive Electrophysiology Cluster, Department of Psychology, University of Oslo, Oslo, Norway

be guided by the formulation of specific scenarios commonly encountered in EEG. To illustrate these scenarios, simulations are provided that vary EEG characteristics such as temporal variability of stimulus-evoked activity or topographical distributions of scalp potentials across subjects. This review is accompanied by code snippets to explore these scenarios and apply selected methods for group-level decomposition, thereby giving this synopsis also a tutorial character.

Aims and Methods

EEG constitutes measurements from a number of electrodes attached to a person's scalp, with each single electrode recording a mixture of activity patterns from several concurrently active brain regions. A number of algorithms have previously been applied to EEG data to recover or isolate activity patterns of underlying brain sources or networks, thereby effectively demixing the original recordings according to specified assumptions. Classic examples for the analysis of single subject data sets are EEG source reconstruction (aka inverse modeling; e.g., Michel et al. 2004), principal component analysis (PCA; e.g., Skrandies 1993), or independent component analysis (ICA; Makeig et al. 2004; Delorme et al. 2012). Especially the latter, ICA, has gained much momentum for its ability to isolate and subtract activity associated with eye artifacts from genuine brain activity. ICA assumes that the observed variables (e.g. EEG channels) are linear mixtures of independent sources (brain regions or networks), and uses higher-order statistics to demix the original recordings. ICA then provides us with a number of source estimates, where each source is characterized by a time course that is maximally independent from all other source time courses, and a component topography that specifies how strong a given source contributes to the EEG recordings at each single channel. Please refer to Onton et al. (2006) for a more detailed explanation of the generative model of ICA and its application to EEG. In comparison to ICA, PCA merely decomposes the data to linearly uncorrelated variables, thereby leaving room for non-linear dependencies between component time courses. These and other methods have been proven very useful for the analysis of EEG data. One major problem, though, is that the latent structure estimated from single subject data sets does not directly generalize across subjects. This, however, is exactly what we are most often interested in, since the majority of our statistical analyses aim at inferring properties of a population.

Two general approaches can be differentiated that address this issue. First, the clustering of components estimated from EEG data of different subjects has been applied to identify groups of similar components expressed within a sample.

Second, a class of algorithms has been brought forward that rather relies on the concatenation of single-subject data sets to large matrices from which a latent structure of sources is estimated that is representative for the sample as a whole. In the following, we will shortly describe the general ideas behind applications proposed so far, and later illustrate their individual strengths and limitations in context of specific scenarios.

Across Subject Component Clustering

For across subject component clustering, the individual components are first derived for each single subject (e.g. via algorithms implemented in common toolboxes for EEG processing, such as EEGLAB or Fieldtrip; e.g., Onton et al. 2006). These components are then grouped together across subjects based on their similarity. The general idea is to define a set of features characterizing a component, e.g. its topography, spectral profile, or activity patterns averaged across trials in the time domain (i.e. the event-related response), and to match components across subjects to groups with highly similar feature profiles. Viola et al. (2009), for example, developed an EEGLAB plug-in (COR-RMAP) that matches individual components to a template based on correlations of the components' scalp topographies. The pre-defined template topography simply mimics the common topography usually seen with eye-artefacts, i.e. high and symmetric loadings at frontal electrodes close to the eyes (e.g., FP1, FP2), with otherwise uniformly low loadings elsewhere. Wessel and Ullsperger (2011) developed this idea further by additionally taking the temporal evolution of template and test components into account. This procedure is also available as an EEGLAB plug-in (COMPASS). Furthermore, EEGLAB itself provides procedures to cluster components obtained from different subjects into groups using different algorithms (e.g. k-means cluster analysis; more details can be found here: <http://sccn.ucsd.edu/wiki/>). Using a very basic k-means clustering for example, k clusters could be initiated by randomly choosing k independent components as representative cluster means. Then, each single component is iteratively assigned to the cluster it is most similar to, consecutively updating the corresponding cluster's mean feature vector after each assignment for the following iterations. A common measure to determine similarity (or dissimilarity) is the Euclidian distance, which simply is the root of the sum of the squared differences between the feature vectors of the cluster mean and the test component. Thus, the lower this distance measure, the higher the similarity. Procedures based on these kinds of clustering have successfully been applied for the identification of artifacts (e.g., Viola et al. 2009; Wessel and Ullsperger 2011), or genuine neural processes such as fronto-medial dynamics of working memory (Onton et al. 2005). Bigdely-Shamlo et al. (2013) further

advanced component-clustering in their measure projection analysis method. Here, functional domains are identified in a template brain that represent a mapping of (a) the spatial proximity of independent components after localization, and (b) a measure quantifying the similarity of component activity patterns for each voxel in the brain. Measure projection analysis thus constitutes an elegant statistical method the integration of source-localized EEG measures across data sets within a brain imaging framework.

Clustering procedures are very powerful, yet they usually depend on informed user input, which limits their applicability in context of multi-subject EEG analyses. Since the user has to define and derive the variables that constitute the feature vector, user choices can dramatically change the outcome of the clustering procedure. It is also not uncommon that most component clusters only receive contributions from a subset of all subjects, a situation that significantly complicates and hampers statistical assessments and conclusions at the group level. For the remainder of this review, we will therefore focus on procedures that concurrently infer a common or most representative source structure from the sample as a whole, thereby naturally generalizing to the group-level.

Multi-subject Data Decomposition

A growing number of applications aims at directly inferring the latent structure of sources driving scalp EEG recordings concurrently from multi-subject data sets. The common characteristic here is that the data sets of several subjects are concatenated to build a new data structure to which a source separation algorithm is applied. Inherent to most of these techniques is a data reduction step to render the problem computationally feasible. Whereas some of these procedures work on time series data, others can be applied in the frequency domain, thus also enabling the analysis of induced or even spontaneous EEG activity. We will start by first introducing some procedures working on the EEG in the temporal domain. Then, we will address applications that work on EEG spectra, i.e. EEG as characterized by its frequency-specific content. Figure 1 provides a schematic overview of data organization and processing for each of the approaches described below.

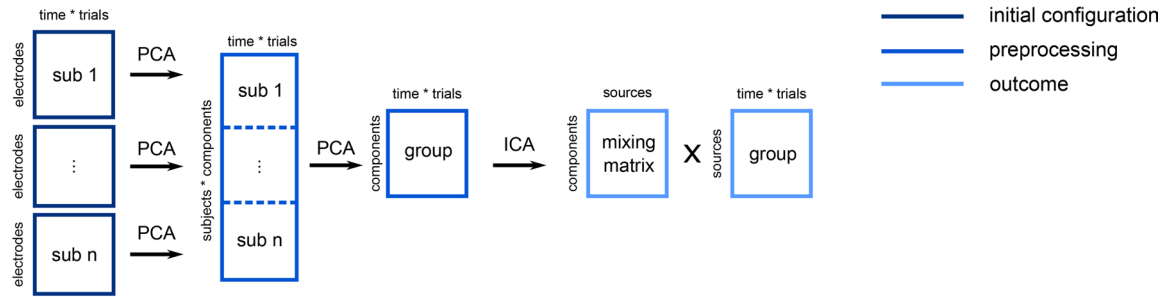
Multi-level Group ICA (mlGICA)

Eichele et al. (2011) proposed a combination of two data reduction steps, followed by ICA, for the analysis of multi-subject time-series data. Specifically, each single-subject data set containing the single-trial EEG first undergoes data reduction via PCA. Then, the selected principal


components of all subjects get concatenated vertically to form a large matrix of the size $(S \times C) \times (T \times L)$, where S and C indicate the number of subjects and components of the single-subject PCA, and T and L the number of trials and the number of data points per trial, respectively. This necessitates that the subject-specific data sets are of the same size, i.e. they contain the same number of trials, and that the order of trials with respect to the conditions they belong to is consistent across data sets as well. A second PCA is computed on this concatenated data matrix, again selecting a certain number of components, to which finally ICA is applied. Resulting group independent components represent statistically independent sources representative for the sample as a whole. It is possible to back-reconstruct independent component time-courses and spatial maps for each subject individually, since individual demixing-matrices can be constructed from the two PCA and the ICA coefficient matrices, then applying them to the individual EEG time-courses (for details see Eichele et al. 2011; Huster et al. 2015).

Note that PCAs are commonly used as processing steps prior to ICA for two reasons: first, by finding uncorrelated components, PCA simplifies the problem to be solved by ICA, because pair-wise dependencies of principal components are already reduced. Also, PCA is commonly used for data reduction by selecting only a certain number of components for further processing, e.g. the first \times components that altogether explain 90% of the variance of the data set. Starting out with 64-channel recordings of 30 participants, for example, without any data reduction a total of 1920 components would have to be processed and later analyzed and eventually interpreted. However, it is not uncommon that about 10–20 components explain about 90% of the variance of an EEG recording, and that the number of components at this percentage is similarly high for the second or group-level PCA. Thus, a common procedure is to extract the same number of components at both the single-subject and the group-level, thereby significantly reducing the computational complexity of the analysis, as well as simplifying the actual interpretation of the data. However, data reduction always entails the risk of disregarding activity patterns of potential interest that might be represented in the unselected principal components, and we will discuss this issue in a bit more detail in context of certain use-cases. Also, the combination of data organization (stacking the subjects along the spatial dimension, i.e. the dimension originally representing electrodes) causes the PCA to identify correlational patterns across subjects and trials, resulting in a bias of this method towards event-related activity (Huster et al. 2015). More precisely, induced activity that is not strictly time-locked across trials necessarily also varies across subjects, causing low correlations of corresponding activity patterns.

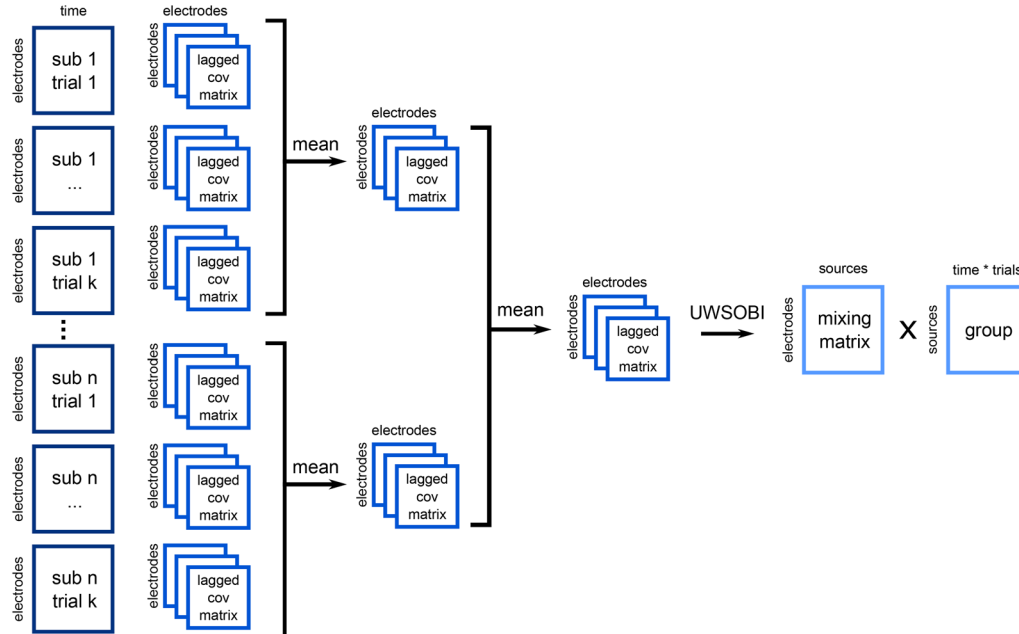
Multilevel group ICA




Temporal-concatenation group ICA



Temporal group UWSOBI



Multi-way decomposition



Spatiospectral group ICA

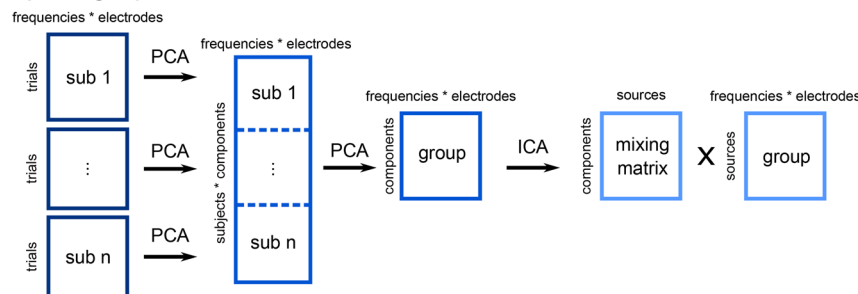


Fig. 1 Schematic depiction of basic data organization and processing steps for each of the group-level decomposition approaches discussed in this overview. Schematic workflows for multilevel group ICA and temporal-concatenation group ICA were adapted from Huster et al. (2015), and that for temporal group UWSOBI was adapted from Lio and Boulinguez (2016)

Temporal-Concatenation Group ICA (tcGICA)

Cong et al. (2013) formalized a related procedure, but instead of a spatial concatenation as in mlGICA, the data are concatenated in the temporal dimension. Hence, EEG data sets of different subjects are not stacked along the dimension representing the channels, but rather are concatenated temporally, such that the resulting matrix is of size $E \times (S \times T \times L)$. Here, E represents the number of electrodes, and S , T , and L are defined as above (S = number of subjects; T = number of trials; L = number of samples per trial). Note that T and L are allowed to vary for different subjects, meaning that there is no restriction to have the same number of trials over subjects. Generally speaking, one may as well concatenate recordings of spontaneous EEG (e.g., resting state activity) of different subjects. This composite data matrix then is subjected to ICA (with or without previous PCA). Each of the resulting group independent components is defined by a common map and the time-courses of the different subjects, which can easily be extracted or aggregated for further processing. Inherent to this procedure is the assumption that the mixing matrix, i.e. the number of sources and how these sources in the brain project to electrodes, is constant across subjects. This means that differences in the scalp distribution (or topography) of a given event-related response across subjects represent or contribute to conditions violating this assumption. Consequently, there is no representation of individual component topographies. On the other hand, since the decomposition works along the temporal dimension, this procedure is sensitive to evoked, induced, as well as spontaneous activity.

Temporal Group-UWSOBI

Lio and Boulinguez (2016) devised and evaluated another approach relying on the joint diagonalization of time-lagged covariance matrices. Specifically, for each single trial of a subject's data set a covariance matrix is computed relating each channel to all other channels; in addition, this procedure is repeated with lagged time-series by shifting the data by one time point per iteration. Then, these lagged covariance matrices are averaged across trials for every subject, and consecutively also across subjects. At last, these group-level covariance matrices (e.g., 50 in case of 50 lags) are subjected to a joint diagonalization to estimate the demixing matrix. Although this procedure does not allow for the

estimation of individual component or source topographies, it was shown that it resulted in good representation of source constellations at the group-level. Specifically, the authors extensively simulated EEG data with strongly varying source constellations (i.e. with inter-individual differences in source locations and source orientations) in different regions of the brain, and compared decompositions for group-UWSOBI and temporal-concatenation group-ICA (tc-GICA). It was found that group-UWSOBI showed better performance regarding both the quality of waveform estimation as well as the precision of source localization, regardless of the exact number of subjects used for group-level decomposition. Interestingly, the evaluation showed that the number of subjects needed for a satisfactory source reconstruction differs across the regions of signal origin in the brain, with larger groups needed in regions such as the dorsolateral prefrontal cortex or the temporo-parietal junction. It can be speculated that this may be due to the special morphology of these regions, where slight variations of cortical shape may cause signal cancellations when averaging across subjects.

Multi-way Decomposition for Time–Frequency Transformed EEG

In many cases we might not be interested in time-domain data, but rather in frequency-specific changes of activity patterns over time. The ERPWAVELAB toolbox (Mørup et al. 2007), available as an EEGLAB plug-in, has been developed with this in mind. Note that the previously introduced procedures mostly work on matrices (two-dimensional arrays, with dimensions representing channels and time points, and the data of subjects being concatenated to one larger two-dimensional array). Computing a time–frequency transform of plain time series data adds another dimension though (channel, time, and frequency per subject), so that the generalization of the previously described methods to this scenario is not trivial. However, the ERPWAVELAB toolbox not only provides the possibility to compute time–frequency transforms and derive various measures from it (such as amplitude- or coherence-based measures to compare channels or trials), but it also implements algorithms to decompose such multi-dimensional data sets. Implementations rely on the PARAFAC (Carroll and Chang 1970) and TUCKER (Tucker 1966) decompositions, which are generalizations of decomposition techniques such as PCA, ICA, or the singular value decomposition. These generalizations can well be applied to multi-way arrays (tensors) with more than two dimensions. This makes the decomposition of multi-subject time–frequency data possible by applying the algorithms to tensors of high order, e.g. covering time, frequency, channel, and subject. Although the toolbox is optimized to work on aggregate data per subject, such as the event-related spectral perturbation (ERSP; Delorme and Makeig 2004), the

general framework does allow for the inclusion of single trial data as well. Other than that, the outcome is comparable to the group ICA approaches, where the components are characterized by their topographies (scalp distribution) and activity profiles (in time and frequency). The reconstruction of individual topographies and activity patterns, as well as the comparison of groups of subjects has also already been implemented in the toolbox.

Spatiospectral Group ICA

Similar to the previous approach, Bridwell et al. (2016) developed a procedure for the analysis of EEG spectra, or more specifically, the characterization of epoched EEG by amplitudes at different frequencies. This approach again applies blind source separation to a 2-dimensional matrix, constructing it slightly differently though. The amplitude spectra of a given subject are stacked across trials, generating a $[T \times (F \times E)]$ matrix, where T corresponds to trials, F to frequencies, and E to electrodes. From here, this approach very much follows the procedure described for multi-level group ICA, and is also implemented in the EEGIFT toolbox. Thus, these subject-specific matrices are subjected to a PCA, then the selected components of all subjects are concatenated vertically (along the dimension originally representing trials), and the group-level PCA and ICA are computed consecutively. This spatio-spectral analysis therefore is conceptually closer to spatial group-level analysis as applied to fMRI data than it is to the temporal decomposition usually computed on EEG data. An observed spatio-spectral EEG map within a given trial/epoch is comparable to a spatial fMRI map for a given TR, where each map represents a linear mixture of statistically independent source maps. Thus, we receive spatio-spectral source maps alongside their temporal profiles (changes across trials). This procedure may be especially well-suited for the analysis of ongoing brain activity in the absence of task-related stimulation, since frequency content is naturally aggregated across time. Although initially Infomax ICA was applied to decompose the data, the authors now tested a variety of different blind source separation algorithms for use in this framework (Bridwell et al. 2016). COMBI (Tichavský et al. 2008), a combination of FAST ICA (EFICA) and weights-adjusted second-order blind identification (WASOBI), also proved to be a powerful candidate for the decomposition of EEG data.

The Wider Picture

Many studies use similar approaches, some of which were published even before the frameworks discussed above were presented in a more formalized fashion. For example, Kovacevic and McIntosh (2007) set up a group ICA with temporal concatenation already earlier, and used it in combination

with partial-least square analysis to assess task-dependent changes in independent components. Congedo et al. (2010) computed ICA on the grand-averaged, complex spectral Fourier matrices. Ramkumar et al. (2014) combined single-subject and group PCA (as in multilevel group ICA) to compute an ICA on time–frequency transformed data, where input variables were either based on reconstructed cortical sources or time points. Similarly, we recently used multi-level group ICA and time–frequency decomposed data using channels as input variables to derive features for single-trial classification (Huster et al. 2017). Thus, the frameworks discussed above are flexible and can easily be adapted further for more sophisticated analyses. However, the growth of this field at this moment seems to be constrained by the rather low number of software packages that offer relevant applications out of the box.

Problem Analysis and Structure

When reconstructing the neural source activity patterns via methods for the decomposition of multi-subject data, care has to be taken to choose the most appropriate method and to adapt it to the research question at hand. Thus, let us shortly consider several scenarios one might encounter. In terms of neural processes under study, at least two dimensions should be considered: first, the temporal characteristics of neural responses, i.e. whether the neural activity of interest is evoked, induced, or spontaneous; and second, the homogeneity of the structure of latent processes across the sample. In addition, we may or may not be interested in the reconstruction of single-subject source patterns in addition to group-level activity, for example to compare two groups, or to study the variability of a neural process across the whole sample.

As to the temporal characteristics of EEG, recorded neural activity could be either spontaneous, that is the EEG was measured in the absence of any task or external stimulation, or there were external and/or internal events to which the brain responded. In the latter case, evoked and induced activity is differentiated (e.g., Makeig et al. 2004). Evoked activity refers to neural activity that occurs strictly phase- or time-locked relative to an event, i.e. the temporal profile is about the same for all repetitions of that event, which for example implies that there is only little variability in onset times or peak latencies. With induced activity, on the other hand, exactly the opposite is the case: onset and peak times do vary across repetitions, causing low or non-significant correlations of the time series over corresponding trials. Note also that the detrimental effects of temporal jittering on signal averaging become more pronounced with higher frequencies. Whereas a jitter of peak amplitudes by 10 ms does not strongly affect the event-related potential of low

frequency activity such as the N2 or P3, the frequency peaks of which are in the range of upper delta (<4 Hz) and lower theta (4–8 Hz). However, such a response jittering at 50 Hz may lead to cancelations during averaging since 10 ms difference from one trial to the next corresponds to a phase inversion. As we will see later, data organization and pre-processing strongly affect the sensitivity of a method to these different neural activity profiles.

With respect to the homogeneity of the latent structure across the sample, two aspects are worth considering. On the one hand, most of our inferences and cognitive models generalize to the level of populations, inherently assuming that the number and nature of (cognitive) processes is the same for different individuals. Also, these processes are assumed to originate from the same brain regions or networks. Obviously, this is an oversimplification known not to be true. Individuals may rely on different strategies for task processing, engage brain structures in slightly different constellations, or may even vary in the number of differentiable sources or processes that contribute to performance. Not least, individual differences in brain morphometry cause differences in the topography of scalp potentials. Individuals differ with respect to the positioning of cytoarchitectural areas in the cortex, or specifically its gyri and sulci. This leads to variations in the orientation of current dipoles relative to the skull, thereby shifting scalp potential distributions. These factors translate to differences in the mixing matrices of different subjects, since the mixing matrix also codes the projection of the source activity patterns to the scalp electrodes.

At last, the research question at hand dictates whether one wants to work with the estimated multi-subject components or sources exclusively, or whether there is need to reconstruct the source activities for individuals (or groups thereof). This aspect is only partially independent from the previous one, the homogeneity of the latent structure. As explained earlier, group-level decomposition rests on the assumption that there is a common underlying structuring of the latent processes of the sample. Let's consider the two most extreme cases. On the one hand, all subjects could rely on the exact same source constellation (identical numbers

of sources, qualitatively the same processes, identical temporal profiles, etc.), such that a multi-subject decomposition and the estimated group-level components or sources would perfectly represent the sample characteristics. The other extreme would be that all subjects show a unique constellation, i.e. no two subjects would exhibit an identical source (in terms of a cognitive process, brain region, or network); in this case there simply is no common structure which could be considered representative for the sample. In most cases the truth will be somewhere in-between: in many task-contexts, such as plain speeded response tasks, early sensory or motor processes are very similar across individuals, whereas in more complex-decision making tasks, processes in-between sensory evaluation and response generation leave more room for inter-individual variability. In any case, many research questions aim to compare not only conditions, but also different groups, such as healthy controls and clinical samples. Thus, in some cases we need to be able to reconstruct and group individually-reconstructed sources to compare corresponding sources between groups quantitatively, or even to assess differences in the latent structure between groups in its entirety.

Use-Case Examples

Based on the considerations of the previous section, we will now discuss some of the most common analysis scenarios, structuring the discussion according to the neural activity patterns under study. We will further need to take the homogeneity of the latent structure across subjects into consideration. Needless to say that the previously discussed case of a perfect non-overlap of the latent structure across subjects (i.e., utterly distinct source neural activity patterns for each subject) would preclude the application of any multi-subject decomposition. However, it is fair to assume that this is a very unlikely case as long as participants process the same task (or more generally put: recordings are taken in similar mental states). Figure 2 provides an overview of the previously discussed methods and a rough guide to what scenarios each method covers. This is not meant to be an

Fig. 2 Overview of the different approaches and possible scenarios in terms of data characteristics and analysis goals

	Time-domain	Frequency-domain	Evoked	Induced	Spontaneous	Individual activity	Individual topography
Multi-level group ICA	✓		✓	(✓)		✓	✓
Temporal-concatenation group ICA	✓		✓	✓	✓	✓	
Temporal group UWSOBI	✓		✓	✓		✓	
Multi-way time-frequency decomposition		✓	✓	✓		✓	✓
Spatiospectral group ICA		✓			✓	✓	✓

exhaustive overview, but should rather serve as a starting point for the reader's considerations.

To illustrate some of the concepts, a set of simulations was generated that represent three sources with activity at 8, 20, and 40 Hz, and fronto-medial, lateral–central, and occipital topographies, respectively. A total of 50 trials were simulated with a length of 3 s at a sampling rate of 200 Hz, and a virtual baseline of -1500 ms. These sources are embedded in Gaussian noise. When discussing the different scenarios, certain parameters of these data sets are varied, as for example the variability of the individual mixing matrices (i.e. the projection of the sources to the scalp electrodes), or onset variability of sources across trials. Details on data generation as well as a quantitative analysis of the effects of certain parameters on source reconstruction can be found in Eichele et al. (2011), and Huster et al. (2015).

The main aim of this tutorial review is to guide researchers new to the field of multi-subject data decomposition in their choice of procedures and to equip them with a basic understanding of underlying principles and methods. To this means, we will focus on multi-level and temporal-concatenation group-level decompositions using ICA (mIGICA and tcGICA, respectively), which are powerful, yet relatively easy to understand and simple to modify. The code used for the purpose of this paper, including some routines for multi-subject decomposition, can be found here: <http://www.rjhuster.com/downloads>. The code is accompanied by instructions to simulate the data and run the analyses executed for this tutorial review. An extensive comparison of all methods or a discussion of the intricate details of every single approach cannot be provided within this overview, not least because their implementations do not yet exist in a unified computational framework. However, most of the techniques introduced in the previous sections are accessible via separate routines or software packages, as well as tutorials and manuals. Thus, this tutorial review should be considered a practical starting point for further exploration of this exciting field.

Evoked Activity

Let us cover the easiest case first, namely the exclusive interest in or occurrence of evoked activity with a similar latent structure across subjects. All methods covered above are well-suited to treat this case, and the major question of interest is whether we want the data to be analyzed in the time or frequency domain. But since the activity patterns across both trials and subjects show only minimal temporal offset anyways, we may as well stay in the time-domain. Then, both tcGICA and mIGICA may be worth considering, because both should be able to recover the generating sources well.

Figure 3 depicts the simulated source activity time courses, their scalp projections, as well as results of the

group-level decompositions using mIGICA and tcGICA. Please refer to the MATLAB-scripts available on our webpage to check for all details of the analyses. When running mIGICA, the inspection of the subject-specific PCAs used for whitening and data reduction indicates that on average about 56 components explain just about 90% of the variance of the data. Thus, for now we will extract 56 components at both the first PCA and the group-level PCA, consequently also arriving at 56 independent components. As can be seen, both mIGICA and tcGICA nicely reconstruct the three simulated sources in three separate components alongside their corresponding topographies. The remaining 53 (not depicted in Fig. 3, but accessible via the script), reflect pure noise as indicated by unstructured topographies and a lack of stimulus-related activity. When extracting the single-trial peak amplitudes from the reconstructed sources of each subject, we find high correlations between the reconstructed and simulated trial-by-trial source amplitude variations. Thus, data indicate that both procedures nicely reconstructed the original source activity patterns.

Induced Activity

What happens, however, if the activity patterns are not perfectly time-locked across trials and, as a consequence, also not across subjects? It can be predicted that this poses a problem for mIGICA, because the data organization inherently assumes a strong correlation of time-courses across subjects; here, induced responses will effectively decorrelate the data across subjects, causing a drop in reconstruction performance. Thus, we now repeat the simulation allowing source activity patterns to jitter up to ± 50 ms randomly from trial-to-trial. As can be seen in Fig. 4, tcGICA nicely reconstructs all three sources. Although the averaged time-courses may appear somewhat distorted, this is merely a byproduct of computing ERPs (i.e., the average across trials in the temporal domain) of induced activity patterns. This effect becomes more apparent when additionally inspecting the event-related spectral perturbations (ERSP) of the resulting components. Here, every source exhibits a clearly defined peak response. An inspection of the sources reconstructed with mIGICA, on the other hand, suggests that only the theta-source was successfully reconstructed. All other components of this decomposition are noisy beyond clear identification, yet an inspection of their corresponding ERSPs also suggests that all three simulated sources were disaggregated to different degrees, and can be found mixed in nearly all other components (Fig. 4 depicts only the first three components though). The effect that a given degree of temporal jittering over trials causes stronger deteriorations in reconstruction quality of high- as compared to low-frequency sources can easily be explained: a maximum jitter of 50 ms, for example, corresponds to a maximum phase

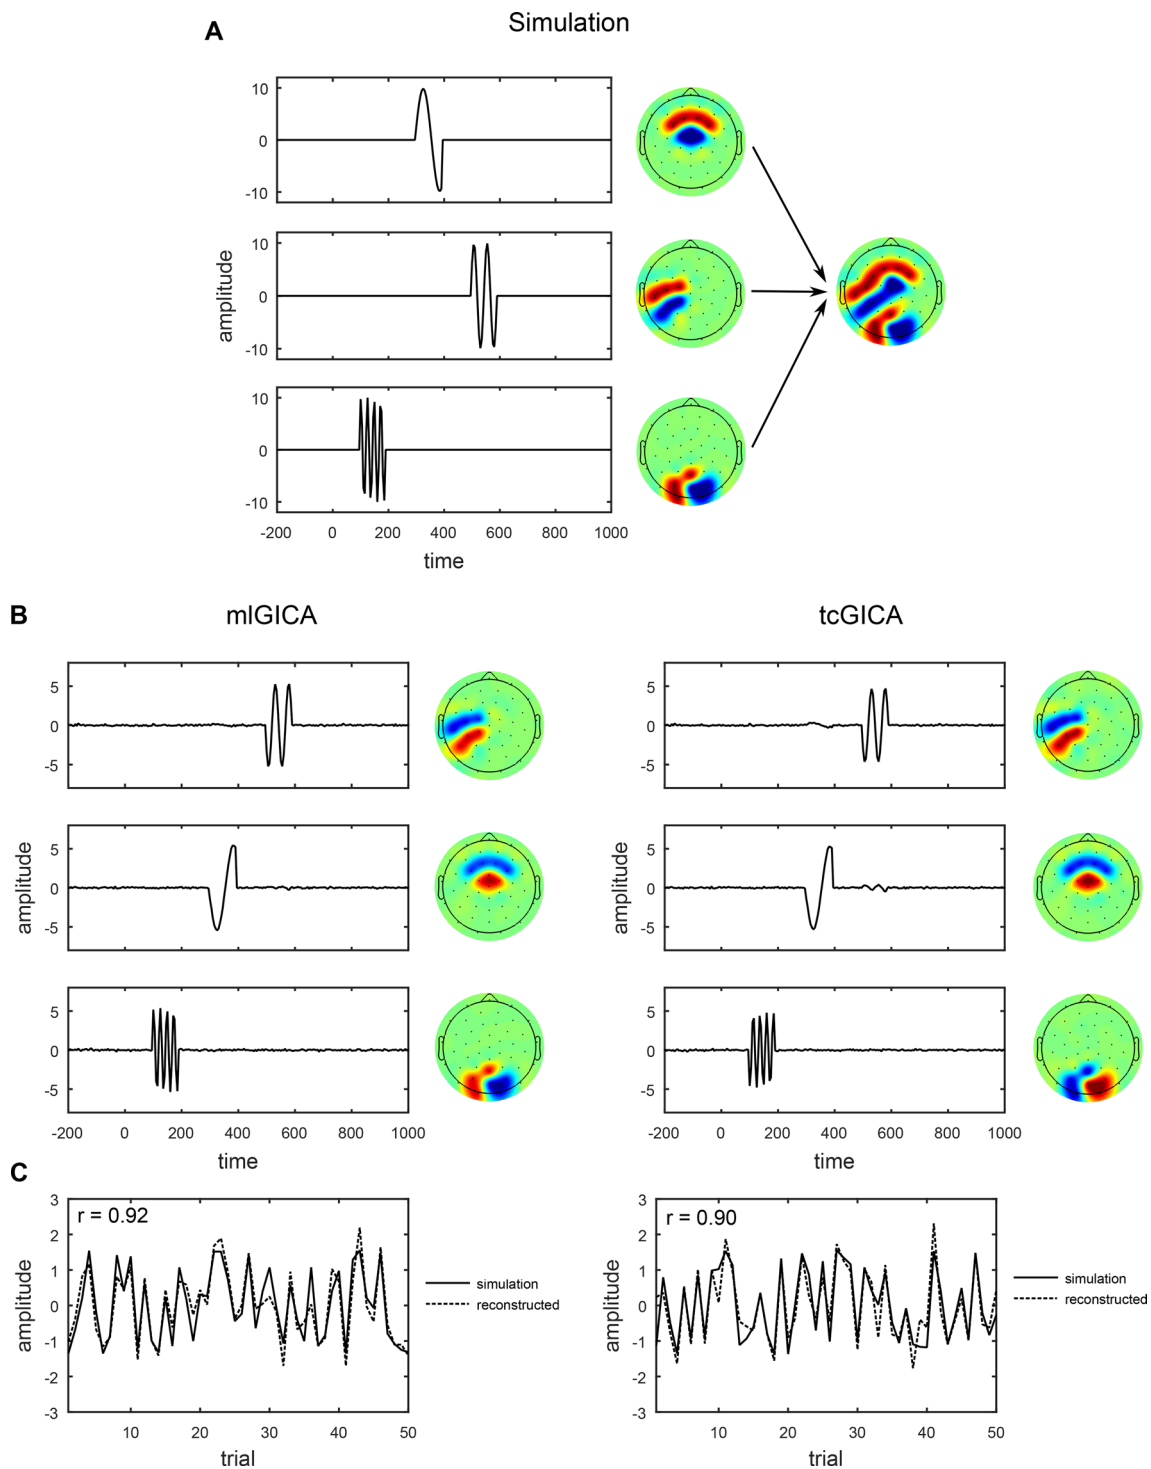


Fig. 3 Comparison of mIGICA and tcGICA for the reconstruction of sources characterized by strong time-locking and stable topographies. **a** Simulated source time courses (left), their respective topographies (middle) and their combined topography as seen for a single subject

(right). **b** Comparison of mIGICA (left) and tcGICA (right) for the reconstruction of sources characterized by strong time-locking and stable topographies. **c** Trial-by-trial variation of the original simulations and reconstructed source amplitudes

shift of 180° at 10 Hz, yet to shifts of 360° and beyond for frequencies equal or higher than 20 Hz, respectively. One might argue that the total maximum jitter as applied for these

simulations may be considered extreme for many neural response profiles in context of cognitive tasks. Note though that response times often show variation even exceeding

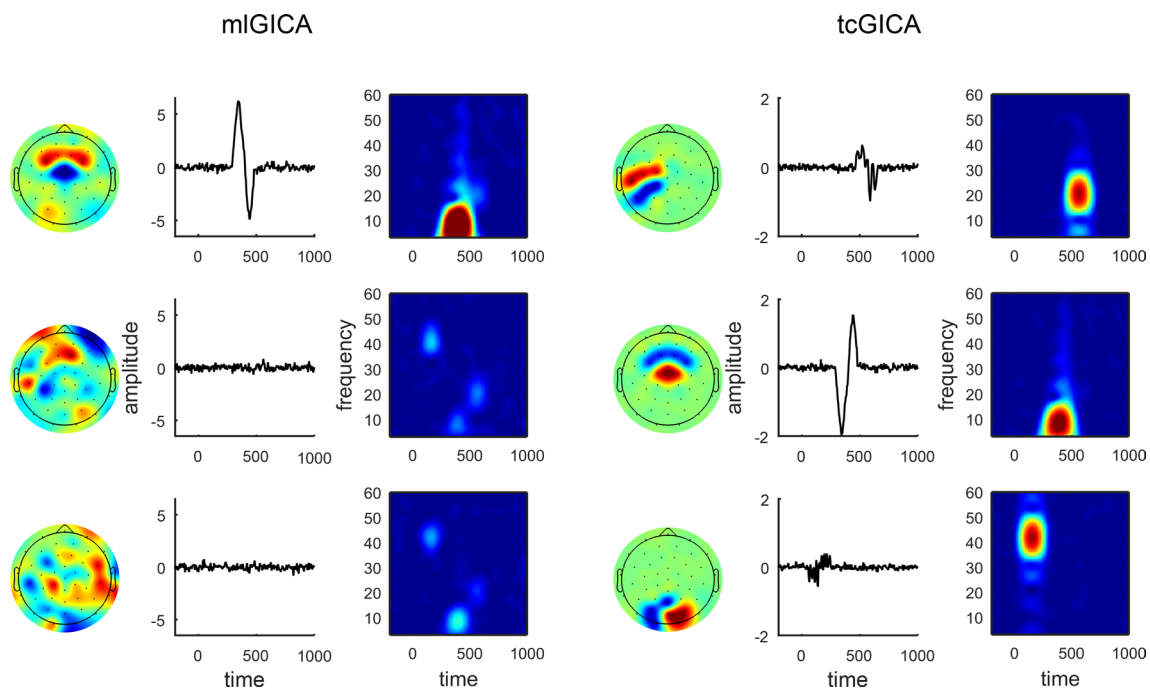


Fig. 4 Comparison of mIGICA and tcGICA for the reconstruction of sources characterized by strong temporal jittering across trial onsets (induced activity) and stable topographies. From left to right: recon-

structed topographies, event-related potentials and event-related spectral perturbations

100 ms from fastest to slowest responses. A more thorough evaluation of the effects of temporal jittering on group-level decompositions can be found in Huster et al. (2015).

Topographical Variation

Further sources of variation across subjects are differences in electrode placement, and even more severely, brain morphology. Both contribute to inter-individual differences in the scalp topography of a given ERP. In mathematical terms, this means that the mixing matrixes, i.e. the model specifying the mapping of sources to the scalp electrodes, deviate from each other across subjects. This again violates the assumption inherent to tcGICA, where both PCA and ICA are applied exclusively at the group-level. With mIGICA, however, a subject-specific PCA precedes the group-level analyses, thereby allowing for inter-individual differences in the mapping of sources to scalp electrodes.

To illustrate the resulting differences in source reconstruction, we again simulate data, this time introducing topographical variability while again focusing on perfectly time-locked activity patterns across trials. As can be seen in Fig. 5, mIGICA reconstructs the three sources in the first three independent components, with consecutive components capturing noise. In addition, for the first component (occipital gamma) the subject-specific reconstruction of topographies for five subjects are shown, elucidating the

algorithm's ability to capture and reconstruct differences in mixing matrixes across subjects. In contrast, when inspecting the first five independent components reconstructed via tcGICA, we notice that the first four all represent fronto-central theta activity, which should have been captured within a single group component, whereas the fifth represents a reconstructed beta source. When inspecting the corresponding time courses, we further see that their reconstructions strongly differ with respect to the achieved signal to noise ratio. Since tcGICA does not account for inter-individual differences in source mixing, inter-subject variability causes a deteriorated reconstructed performance for all sources, with subjects contributing differently to the different independent components.

Spontaneous Activity

How to analyze spontaneous EEG, that is EEG in the absence of external experimental stimulation? By definition, this situation is characterized by the lack of phase- or time-locking of brain activity across subjects. Consequently, the situation conceptually mimics the one of induced responses with extreme jittering of neural activity across trials (where a trial would correspond to a random period of resting EEG). The outcome will thus mirror that of our earlier inspection, just in an even more amplified way. Since mIGICA was not even designed with this

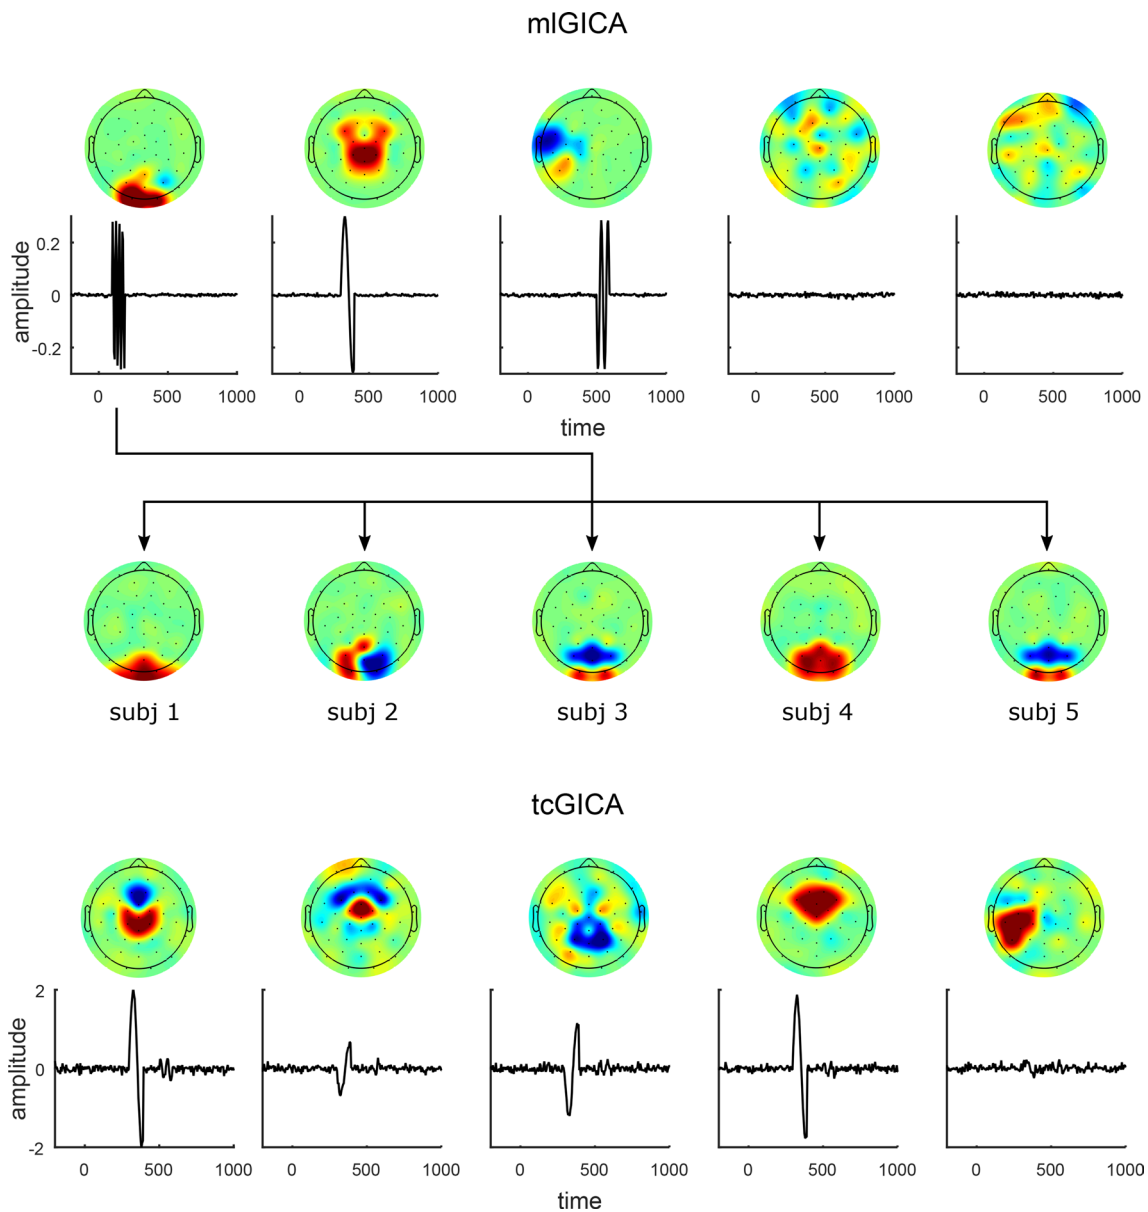


Fig. 5 Comparison of mIGICA and tcGICA for the reconstruction of time-locked sources (evoked activity), but topographies showing inter-subject variability. For both methods, the top two rows represent the reconstructed source topographies and event related poten-

tials for the first five sources. In mIGICA, the bottom row represents the individually reconstructed topographies of the first source for five individual participants. In tcGICA, the reconstruction of individual topographies is not possible

scenario in mind, it will fail to reliably reconstruct such sources, and we refrain from this rather trivial assessment here (the reader may very well simulate this case using our scripts, simply by setting the trial-wise jitter to unrealistically high values). Then again, tcGICA will provide a valid reconstruction of spontaneous sources as long as inter-subject differences of the mixing matrices are low. But as exemplified in Fig. 2, considering alternative procedures directly working with the frequency-representation of the data, such as spatio-spectral group ICA, may be preferable.

Adaptations

In addition to the variety of algorithms already available, the actual power of the framework for multi-subject data decomposition is its flexibility and adaptability. Just taking mIGICA as an example, two major issues can be identified that potentially limit the applicability of this algorithm. As discussed earlier, a PCA-based data reduction step concurrently calculated across all subjects may cause two major limitations: (1) limited power in the detection and reconstruction of sources with induced activity patterns, and (2)

a bias towards activity patterns that are not only time-locked but also large in amplitude, thereby causing large variations in scalp EEG. The latter may, for example, be the case with large ERPs such as the P300, as opposed to early sensory activity patterns that usually are smaller in amplitude and more focal, and thus also cause less systematic variability in the data. However, these scenarios can be addressed by computing time–frequency decompositions prior to group-level data decomposition, and the re-organization of the data. Single-trial time–frequency data could easily be arranged in a two-dimensional matrix where the first dimension again represents channels, whereas the second dimension stores the multiplexed time-/frequency data (e.g., $ch1t1f1$, $ch1t1f2 \dots ch1t2f1$, $ch1t2f2 \dots$). Figure 6 displays the rationale of this procedure, as well as three group independent components computed on the time–frequency transforms of simulated data with loose time-locking (the same data

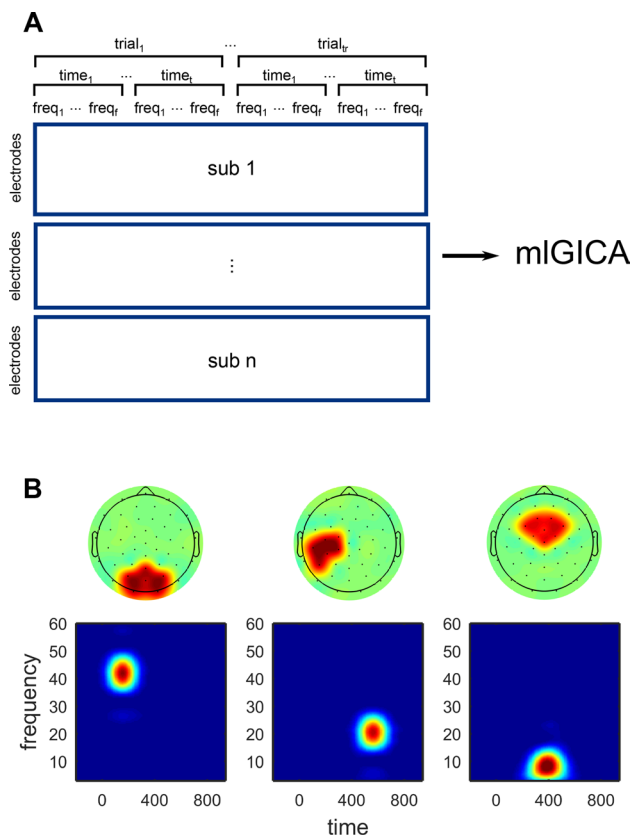


Fig. 6 Adaptation of mIGICA for the reconstruction of sources with strong temporal jittering across trial onsets (induced activity). **a** The procedure simply relies on the reorganization of the data prior to the analysis, such that the frequency-specific amplitude values get multiplexed with respect to the samples in a given trial as well as across trials (tr trial number; t time point or sample within a trial; f frequency). **b** The component activity patterns (lower row) computed on the time–frequency decompositions show that mIGICA now nicely captures and reconstructs all three sources (compare to Fig. 4, with mIGICA computed on time-domain data)

as used in our previous example; Fig. 4). As can be seen, the thereby adapted procedure now again captures all three sources very well. In addition, a channel-wise z-scoring of the data can alleviate a potential bias towards large potentials relative to rather focal and small amplitude EEG-effects by causing a relative down-scaling of the former, and an up-scaling of the smaller amplitude signals (given that the topographies do not overlap completely). Thus, all methods can be tweaked to overcome at least some of their limitations, and the power of such adaptations has hardly been explored yet (but see Huster et al. 2015; Bridwell et al. 2016; or Lio and Boulinguez 2016, for the evaluation of performance differences caused by the integration of different BSS algorithms in these frameworks).

Model Order Selection

One final issue that needs commenting relates to the specification of the model order, i.e. the number of sources to be estimated in a given analysis. In the previous examples, we simply based this decision on the variance cumulatively explained by principal components (e.g., the number of components that together explain 90% of the variance in the data). Here, we fixed the model order based on the cumulatively explained variance of the subject-specific PCAs, keeping this parameter constant for all subjects as well as the group-level analysis. It can be argued that this may result in misrepresentations or biases of the activity patterns of single subjects. Yet, one must not forget that these group-level decompositions are meant to derive source estimates concurrently representative for multiple subjects, thus inherently trading the optimized representation of the group against that of single subjects. In addition, whereas ICA minimizes statistical dependencies between components and consequently accounts for both linear as well as non-linear associations, PCA merely de-correlates variables and thus disregards non-linear dependencies. Thus, although the variance-based criterion is commonly used, it is not without its shortcomings. Other procedures to estimate the model order rely on information theoretic concepts. The minimum description-length criterion, Akaike's information criterion and Bayesian information criterion have been implemented in popular ICA software packages, such as MELODIC and GIFT, and are widely used in context of fMRI analyses (e.g., Williams 1994; Stoica and Babu 2012). It has been argued though that these procedures are sensitive to data with low signal-to-noise ratios, then potentially estimating source numbers with low reliability only. Alternatively, one may test the statistical reliability of ICA estimates using different initial values for the number of components. ICASSO (Himberg et al. 2004) is an algorithm that runs ICA several times and produces different component estimates for each run; it then clusters the components of all runs based on their

similarity. Components that can reliably be estimated across runs correspond to tight clusters (groups of components with high similarity). When comparing such cluster solutions for ICA estimates with different numbers of sources, the one with the highest reliability (the overall “tightest clusters”) should be preferred. However, as of now there is no gold standard for the estimation of the most appropriate model order, and the comparison of several measures for a given data set may be the most appropriate way to estimate the number of components to be estimated.

Conclusion

No size fits all! As our practically-guided discussion of scenarios exemplified, all current methods come with their strengths and weaknesses. Thus, care has to be taken to choose the approach best-suited to address the research question at hand and to match the present data characteristics. Where this is the case though, group-level or multi-subject decomposition of EEG concurrently solves two major problems: (1) the spatio-temporal separation of activity originating from different brain regions and representing independent neurocognitive processes; (2) the matching of recovered sources across subjects.

Similar approaches used for the analyses of fMRI data surfaced about one-and-a-half decades ago, and since then have extensively been applied to study brain states. Group-ICA of fMRI is one of the two major techniques to extract and analyze resting state networks (e.g., Calhoun and Adali 2012). The assessment of these intrinsic functional connectivity networks now promises to become a useful diagnostic tool and a predictor for treatment outcome measures (e.g., Rashid et al. 2016). With respect to EEG, still much more work is needed to mirror the widespread use of such algorithms as seen in fMRI data analyses. However, we now have a number of tools available to address the most common scenarios relevant for the analysis of both task and resting state EEG. Indeed, first applications prove their utility. Group-level decompositions are now commonly used in combination with other advanced signal analysis techniques, e.g. for (1) data fusion of simultaneously acquired EEG and fMRI (e.g., Bridwell et al. 2013), (2) feature generation for single-trial prediction (e.g., Huster et al. 2017), (3) the study of the genetic underpinnings of neural oscillations (e.g., Antonakakis et al. 2016), (4) the comparative analysis of resting state networks inferred from fMRI and EEG (Yuan et al. 2012), (5) the assessment of developmental trajectories of neurocognitive processes across the lifespan (e.g., Enriquez-Geppert and Barceló 2016; van Dinteren et al. 2017), (6) the characterization of the neural dynamics of depressive and psychotic symptoms (Bridwell et al. 2014, 2015), as well as (7) functional and effective connectivity analyses of brain

networks (Huster et al. 2014). Although this list is far from being complete, it clearly showcases the wide applicability of multi-subject EEG data decomposition. In sum, group-level decomposition is a powerful tool for analyzing EEG data, for which a strong propagation over the next decade is to be expected.

References

- Antonakakis M, Zervakis M, van Beijsterveldt CEM et al (2016) Genetic effects on source level evoked and induced oscillatory brain responses in a visual oddball task. *Biol Psychol* 114:69–80. doi:10.1016/j.biopsycho.2015.12.006
- Bigdely-Shamlo N, Mullen T, Kreutz-Delgado K, Makeig S (2013) Measure projection analysis: a probabilistic approach to EEG source comparison and multi-subject inference. *Neuroimage* 72:287–303. doi:10.1016/j.neuroimage.2013.01.040
- Bridwell DA, Wu L, Eichele T, Calhoun VD (2013) The spatio-spectral characterization of brain networks: fusing concurrent EEG spectra and fMRI maps. *Neuroimage* 69:101–111. doi:10.1016/j.neuroimage.2012.12.024
- Bridwell DA, Kiehl KA, Pearlson GD, Calhoun VD (2014) Patients with schizophrenia demonstrate reduced cortical sensitivity to auditory oddball regularities. *Schizophr Res* 158:189–194. doi:10.1016/j.schres.2014.06.037
- Bridwell DA, Steele VR, Maurer JM et al (2015) The relationship between somatic and cognitive-affective depression symptoms and error-related ERPs. *J Affect Disord* 172:89–95. doi:10.1016/j.jad.2014.09.054
- Bridwell DA, Rachakonda S, Silva RF et al (2016) Spatiospectral decomposition of multi-subject EEG: evaluating blind source separation algorithms on real and realistic simulated data. *Brain Topogr*. doi:10.1007/s10548-016-0479-1
- Calhoun VD, Adali T (2012) Multisubject independent component analysis of fMRI: a decade of intrinsic networks, default mode, and neurodiagnostic discovery. *IEEE Rev Biomed Eng* 5:60–73. doi:10.1109/RBME.2012.2211076
- Carroll JD, Chang J-J (1970) Analysis of individual differences in multidimensional scaling via an n-way generalization of “Eckart-Young” decomposition. *Psychometrika* 35:283–319. doi:10.1007/BF02310791
- Cong F, He Z, Hämäläinen J et al (2013) Validating rationale of group-level component analysis based on estimating number of sources in EEG through model order selection. *J Neurosci Methods* 212:165–172. doi:10.1016/j.jneumeth.2012.09.029
- Congedo M, John RE, De Ridder D, Prichep L (2010) Group independent component analysis of resting state EEG in large normative samples. *Int J Psychophysiol* 78:89–99. doi:10.1016/j.ijpsycho.2010.06.003
- Delorme A, Makeig S (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods* 134:9–21. doi:10.1016/j.jneumeth.2003.10.009
- Delorme A, Palmer J, Onton J et al (2012) Independent EEG sources are dipolar. *PLoS ONE* 7:e30135. doi:10.1371/journal.pone.0030135
- Eichele T, Rachakonda S, Brakedal B et al (2011) EEGIFT: group independent component analysis for event-related EEG data. *Comput Intell Neurosci*. doi:10.1155/2011/129365
- Enriquez-Geppert S, Barceló F (2016) Multisubject decomposition of event-related positivities in cognitive control:

- tackling age-related changes in reactive control. *Brain Topogr*. doi:10.1007/s10548-016-0512-4
- Himberg J, Hyvärinen A, Esposito F (2004) Validating the independent components of neuroimaging time series via clustering and visualization. *Neuroimage* 22(3):1214–1222
- Huster RJ, Plis SM, Lavalée CF et al (2014) Functional and effective connectivity of stopping. *Neuroimage* 94:120–128. doi:10.1016/j.neuroimage.2014.02.034
- Huster RJ, Plis SM, Calhoun VD (2015) Group-level component analyses of EEG: validation and evaluation. *Front Neurosci* 9:254. doi:10.3389/fnins.2015.00254
- Huster RJ, Schneider S, Lavalée CF et al (2017) Filling the void-enriching the feature space of successful stopping. *Hum Brain Mapp* 38:1333–1346. doi:10.1002/hbm.23457
- Kovacevic N, McIntosh AR (2007) Groupwise independent component decomposition of EEG data and partial least square analysis. *Neuroimage* 35:1103–1112. doi:10.1016/j.neuroimage.2007.01.016
- Lio G, Boulinguez P (2016) How does sensor-space group blind source separation face inter-individual neuroanatomical variability? Insights from a simulation study based on the PALS-B12 atlas. *Brain Topogr*. doi:10.1007/s10548-016-0497-z
- Makeig S, Debener S, Onton J, Delorme A (2004) Mining event-related brain dynamics. *Trends Cogn Sci (Regul Ed)* 8:204–210. doi:10.1016/j.tics.2004.03.008
- Michel CM, Murray MM, Lantz G et al (2004) EEG source imaging. *Clin Neurophysiol* 115:2195–2222. doi:10.1016/j.clinph.2004.06.001
- Mørup M, Hansen LK, Arnfred SM (2007) ERPWAVELAB a toolbox for multi-channel analysis of time-frequency transformed event related potentials. *J Neurosci Methods* 161:361–368. doi:10.1016/j.jneumeth.2006.11.008
- Onton J, Delorme A, Makeig S (2005) Frontal midline EEG dynamics during working memory. *Neuroimage* 27:341–356. doi:10.1016/j.neuroimage.2005.04.014
- Onton J, Westerfield M, Townsend J, Makeig S (2006) Imaging human EEG dynamics using independent component analysis. *Neurosci Biobehav Rev* 30:808–822. doi:10.1016/j.neubiorev.2006.06.007
- Ramkumar P, Parkkonen L, Hyvärinen A (2014) Group-level spatial independent component analysis of Fourier envelopes of resting-state MEG data. *Neuroimage* 86:480–491. doi:10.1016/j.neuroimage.2013.10.032
- Rashid B, Arbabshirani MR, Damaraju E et al (2016) Classification of schizophrenia and bipolar patients using static and dynamic resting-state fMRI brain connectivity. *Neuroimage* 134:645–657. doi:10.1016/j.neuroimage.2016.04.051
- Skrandies W (1993) EEG/EP: new techniques. *Brain Topogr* 5:347–350
- Stoica P, Babu P (2012) On the proper forms of BIC for model order selection. *IEEE Trans Signal Process* 60:4956–4961
- Tichavský P, Koldovský Z, Yeredor A et al (2008) A hybrid technique for blind separation of non-gaussian and time-correlated sources using a multicomponent approach. *IEEE Trans Neural Netw* 19:421–430. doi:10.1109/TNN.2007.908648
- Tucker LR (1966) Some mathematical notes on three-mode factor analysis. *Psychometrika* 31:279–311. doi:10.1007/BF02289464
- van Dinteren R, Huster RJ, Jongsma MLA et al (2017) Differences in cortical sources of the event-related P3 potential between young and old participants indicate frontal compensation. *Brain Topogr*. doi:10.1007/s10548-016-0542-y
- Viola FC, Thorne J, Edmonds B et al (2009) Semi-automatic identification of independent components representing EEG artifact. *Clin Neurophysiol* 120:868–877. doi:10.1016/j.clinph.2009.01.015
- Wessel JR, Ullsperger M (2011) Selection of independent components representing event-related brain potentials: a data-driven approach for greater objectivity. *Neuroimage* 54:2105–2115. doi:10.1016/j.neuroimage.2010.10.033
- Williams DB (1994) Counting the degrees of freedom when using AIC and MDL to detect signals. *IEEE Trans Signal Process* 42:3282–3284
- Yuan H, Zotev V, Phillips R et al (2012) Spatiotemporal dynamics of the brain at rest—exploring EEG microstates as electrophysiological signatures of BOLD resting state networks. *Neuroimage* 60:2062–2072. doi:10.1016/j.neuroimage.2012.02.031