Master thesis for the CogMaster at ENS-PSL

Subspace-based methods for extraction of meaningful spatio-temporal data from electro-magnetic imaging

Grégoire Yasmine[1]

Supervisor : Alain de Cheveigné[2]

[1]*Master's Student, Département d'Études Cognitives, École Normale Supérieure de Paris*

[2]*Senior scientist, Laboratoire des Systèmes Perceptifs, École Normale Supérieure de Paris*

May, 2025

Abstract

Whole-brain electromagnetic recordings such as electroencephalography (EEG) and magnetoencephalography (MEG) are widely used in both clinical and research settings to investigate brain dynamics. While these techniques provide excellent temporal resolution, accurately localizing the spatial origin of brain activity requires solving the ill-posed electromagnetic inverse problem. In this work, I present subspace-based source localization as an efficient and versatile class of methods that leverage linear analysis techniques to extract the signal of interest prior to localization. I focus in particular on the well-known MUSIC algorithm, which uses principal component analysis (PCA) to estimate the signal subspace. Empirical results demonstrate that although MUSIC is not the least-squares optimal solution to the dipole localization problem, it can be more robust to noise than traditional dipole fitting approaches. I also explore the use of alternative linear methods for subspace extraction, which broadens the applicability of the approach. In the multiple dipole case, I compare MUSIC to recursively applied dipole fitting and show that subspace methods can offer performance advantages. I further examine whether Independent Component Analysis (ICA) can improve source unmixing within the subspace, offering an alternative to standard recursive strategies such as RAP-MUSIC. Finally, I provide theoretical connections between subspace-based methods and classical dipole fitting, clarifying their similarities and differences within a unified framework.

Keywords: *inverse problem, magneto-encephalography, source localization, subspace-based methods, blind-source separation*

Declaration of originality :

This work originated from a method initially referred to as zero-set imaging. It was later discovered that this approach is closely related, if not equivalent, to the well-established MUSIC algorithm. As such, some of the developments presented in this thesis may overlap with existing work. However, to the best of my knowledge, the detailed empirical and theoretical comparison between MUSIC and dipole fitting conducted here—both in depth and scope—is novel.

Laboratoire des Systèmes Perceptifs

ENS DEC DÉPARTEMENT D'ÉTUDES COGNITIVES

Pre-Registration

Source localisation in magnetoencephalography and encephalography using zero-set imaging

Grégoire Degobert-Yasmine, supervised by Alain de Cheveigné

Edited February 6, 2025

1 Introduction

1. Background

Improving the spatial resolution of magneto-encephalography (MEG) and electro-encephalography (EEG) is a promising avenue for expanding these recording methods to new clinical and laboratory applications, allowing for a more precise analysis of the spatio-temporal dynamics of brain activity. The source localization problem consists in determining the position of an electromagnetic source inside the skull based on the activity recorded at the surface. This is a severely ill-posed problem, and various approaches aim to obtain a unique solution by applying regularization techniques, such as beamforming, dipole fitting, and minimum-norm current estimation.

2. Method of interest

Zero-Set Imaging (ZSI) is a novel approach that leverages linear analysis to identify both source activity and localization. Linear analysis methods, such as Principal Component Analysis (PCA), enhance source activity through linear transformations of the recorded activity vector. A key feature of these methods is the derivation of a large set of orthogonal transformations—referred to as null filters—that effectively cancel source activity. These null filters are expected to exhibit zero gain for signals originating from the true source location. Through forward modeling, a rectangular gain matrix can be constructed to map activity from the source space to the sensor space. The zero-set of the null filter corresponds to the set of locations that yield zero gain when mapped to sensor activity and projected onto the null-filter direction. The true source location is necessarily at the intersection of these zero-sets. This intersection can be determined by minimizing the location gain over the source space, providing a novel and potentially robust approach to source localization.

3. Research Question

This project aims to evaluate the robustness of the method with respect to noise and forward model accuracy, as well as to compare it with classical approaches.

4. Hypothesis We hypothesize that this method will provide perfect localization estimates under favorable conditions (high signal-to-noise ratio (SNR), accurate forward model). However, we expect noise and forward model approximations to affect the precision of the approach. Its robustness, relative to other methods, remains unknown and will be investigated. Putative estimators of the model uncertainty will also be evaluated.

2 Experimental design

5. Study type Benchmarking, numerical experiment

6. Study Design The robustness of ZSI and three other techniques (beamforming, dipole fitting, minimum-norm estimation) will be empirically evaluated on randomly generated synthetic data with known point-like sources. The capabilities of Zero-Set Imaging on real data will be assessed using high-quality benchmarking datasets available in the literature.

3 Data

7. Synthetic Data Each technique will be evaluated across a wide range of simulated conditions, including varying SNR levels, number of sources, source proximity, and forward model accuracy. Source activity will be modeled as an Ornstein–Uhlenbeck process at a specific location. Activity at both source and silent locations will be corrupted by white noise according to the expected SNR. Additional signal types (Wiener process, deterministic signals, filtered noise) may later be considered to better approximate real conditions. Multiple trials will be conducted for each method to mitigate noise-driven accuracy variability. The sensor-space projection will be obtained by applying linear transformations derived from a precise forward model. Various levels of forward model accuracy (source space resolution, head geometry approximations) will then be used for source localization.

8. Sample Size For each fixed set of experimental conditions, at least 20 trials per source location and five different source locations will be considered.

9. Real Data We will use the WWU DUNEuro reference dataset for combined EEG/MEG source analysis, as well as similar data on visually evoked potentials from the same authors. These datasets provide high-quality, trial-averaged activity and anatomical data that enable precise forward modeling. While ground truth source locations are not accessible in real data, we will compare ZSI predictions with those from state-of-the-art localization methods mentioned earlier. Biases arising from forward model inaccuracies can also be quantified using this dataset.

4 Analysis

10. Effective SNR The effective SNR is defined as the ratio of the measured signal variance to the measured noise variance in the sensor space. This metric may differ from the expected SNR used for data generation due to random variability and signal amplification at source locations. The validity of a putative estimator for this metric—the ratio between the cumulative variance of the top signal components and that of the null-filter components—will be evaluated.

11. Accuracy Metrics The accuracy of source localization will be assessed using multiple metrics:

- Pairwise Euclidean distance between the estimated sources and their closest true sources, measured either on the flattened cortical surface or in the brain volume.
- Average index of the true source locations, when sources are ranked by increasing loss.

12. Uncertainty An estimation of uncertainty, based on the depth of the loss minima, will be considered. In particular, we will investigate whether this metric serves as a reliable indicator of localization accuracy.

Laboratoire des Systèmes Perceptifs

ENS DEC DÉPARTEMENT D'ÉTUDES COGNITIVES

Pre-Registration

Source localisation in magnetoencephalography and encephalography using zero-set imaging

Grégoire Degobert-Yasmine, supervised by Alain de Cheveigné

Edited February 6, 2025

1 Introduction

1. Background

Improving the spatial resolution of magneto-encephalography (MEG) and electro-encephalography (EEG) is a promising avenue for expanding these recording methods to new clinical and laboratory applications, allowing for a more precise analysis of the spatio-temporal dynamics of brain activity. The source localization problem consists in determining the position of an electromagnetic source inside the skull based on the activity recorded at the surface. This is a severely ill-posed problem, and various approaches aim to obtain a unique solution by applying regularization techniques, such as beamforming, dipole fitting, and minimum-norm current estimation.

2. Method of interest

Zero-Set Imaging (ZSI) is a novel approach that leverages linear analysis to identify both source activity and localization. Linear analysis methods, such as Principal Component Analysis (PCA), enhance source activity through linear transformations of the recorded activity vector. A key feature of these methods is the derivation of a large set of orthogonal transformations—referred to as null filters—that effectively cancel source activity. These null filters are expected to exhibit zero gain for signals originating from the true source location. Through forward modeling, a rectangular gain matrix can be constructed to map activity from the source space to the sensor space. The zero-set of the null filter corresponds to the set of locations that yield zero gain when mapped to sensor activity and projected onto the null-filter direction. The true source location is necessarily at the intersection of these zero-sets. This intersection can be determined by minimizing the location gain over the source space, providing a novel and potentially robust approach to source localization.

3. Research Question

This project aims to evaluate the robustness of the method with respect to noise and forward model accuracy, as well as to compare it with classical approaches.

4. Hypothesis We hypothesize that this method will provide perfect localization estimates under favorable conditions (high signal-to-noise ratio (SNR), accurate forward model). However, we expect noise and forward model approximations to affect the precision of the approach. Its robustness, relative to other methods, remains unknown and will be investigated. Putative estimators of the model uncertainty will also be evaluated.

2 Experimental design

5. Study type Benchmarking, numerical experiment

6. Study Design The robustness of ZSI and three other techniques (beamforming, dipole fitting, minimum-norm estimation) will be empirically evaluated on randomly generated synthetic data with known point-like sources. The capabilities of Zero-Set Imaging on real data will be assessed using high-quality benchmarking datasets available in the literature.

3 Data

7. Synthetic Data Each technique will be evaluated across a wide range of simulated conditions, including varying SNR levels, number of sources, source proximity, and forward model accuracy. Source activity will be modeled as an Ornstein–Uhlenbeck process at a specific location. Activity at both source and silent locations will be corrupted by white noise according to the expected SNR. Additional signal types (Wiener process, deterministic signals, filtered noise) may later be considered to better approximate real conditions. Multiple trials will be conducted for each method to mitigate noise-driven accuracy variability. The sensor-space projection will be obtained by applying linear transformations derived from a precise forward model. Various levels of forward model accuracy (source space resolution, head geometry approximations) will then be used for source localization.

8. Sample Size For each fixed set of experimental conditions, at least 20 trials per source location and five different source locations will be considered.

9. Real Data We will use the WWU DUNEuro reference dataset for combined EEG/MEG source analysis, as well as similar data on visually evoked potentials from the same authors. These datasets provide high-quality, trial-averaged activity and anatomical data that enable precise forward modeling. While ground truth source locations are not accessible in real data, we will compare ZSI predictions with those from state-of-the-art localization methods mentioned earlier. Biases arising from forward model inaccuracies can also be quantified using this dataset.

4 Analysis

10. Effective SNR The effective SNR is defined as the ratio of the measured signal variance to the measured noise variance in the sensor space. This metric may differ from the expected SNR used for data generation due to random variability and signal amplification at source locations. The validity of a putative estimator for this metric—the ratio between the cumulative variance of the top signal components and that of the null-filter components—will be evaluated.

11. Accuracy Metrics The accuracy of source localization will be assessed using multiple metrics:

- Pairwise Euclidean distance between the estimated sources and their closest true sources, measured either on the flattened cortical surface or in the brain volume.
- Average index of the true source locations, when sources are ranked by increasing loss.

12. Uncertainty An estimation of uncertainty, based on the depth of the loss minima, will be considered. In particular, we will investigate whether this metric serves as a reliable indicator of localization accuracy.

1. Introduction

1.1. The inverse problem in electromagnetic brain recording

Pushing the spatial resolution limit Non-invasive brain recording and imaging techniques are essential tools for conducting ethical neuroscience research on human subjects. Among these, metabolism-based imaging methods such as functional magnetic resonance imaging (fMRI) are renowned for their excellent spatial resolution—approximately 1 mm. However, their limited temporal resolution (around 5 seconds) poses a significant challenge for studying fast-paced brain dynamics [8]. In contrast, electroencephalography (EEG) and magnetoencephalography (MEG) record fluctuations in electric potentials and magnetic fields generated by parallel ion currents resulting from the simultaneous firing of hundreds of neurons. These techniques offer much finer temporal resolution—on the order of 100 milliseconds—making them particularly well-suited for capturing rapid neural activity [1], [2]. By aligning recordings to specific stimuli and averaging across trials—a process known as *stimulus-locked averaging*—researchers can isolate event-related activity with an improved signal-to-noise ratio, gaining valuable insights into brain dynamics associated with particular tasks.

At this stage, however, the signal takes the form of a multidimensional time series reflecting local electromagnetic activity recorded at the scalp surface. Basic physical intuition suggests that sensors closer to active brain regions will detect stronger evoked responses, offering a rough proxy for the source location of neural activity. Nevertheless, even an expert cannot accurately determine the location of brain signals based solely on raw sensor data (see Figure 1 for an extrapolated map of sensor recordings on the scalp). This limitation underscores the need for advanced methods that enhance the spatial resolution of electromagnetic brain recordings.

The necessity of a forward model To obtain precise images of brain activity from an array of sensor recordings, one must solve the inverse problem of electromagnetism [20]. As with any inverse problem—i.e., one concerned with inferring unknown causes from observed effects—a prerequisite is the construction of a forward model that maps these causes to their observable consequences. Specifically, the forward model should describe how electromagnetic activity originating from any potential source within the brain projects onto the signals measured at the scalp.

Given the vast number of neurons in the brain, the space of possible sources must be reduced to render the problem tractable. It is widely accepted that macro-columns of aligned cortical pyramidal cells constitute the primary generators of measurable brain electromagnetic activity. These putative sources are typically modeled as current dipoles, characterized by their spatial position, orientation, and intensity [12]. The electromagnetic

field produced by such a dipole is fully described by the quasi-static form of Maxwell's equations [14], which are linear with respect to the dipole's intensity. Consequently, for a system with J sensors and a dipole defined by a set of 5 parameters—3 for position and 2 for orientation—denoted θ, the forward model yields a J-dimensional vector $g(\theta)$.

Contemporary software tools, leveraging anatomical information derived from non-functional brain imaging, can construct highly accurate forward models [11]. These tools typically discretize the cortical surface into a dense grid of locations and compute the mapping from dipole moments (defined by orientation and intensity) at each location to the recorded sensor data. To accommodate arbitrary dipole orientations, the forward model at each location is generally composed of three vectors, corresponding to the gain associated with each orthogonal direction in Cartesian space.

Due to the linearity of the electromagnetic field with respect to dipole moment, the forward vector for any arbitrary orientation can be expressed as a linear combination of the forward vectors corresponding to three orthogonal unit dipoles. A common simplification assumes that all dipole orientations are normal to the cortical surface—a biologically reasonable approximation given the organization of pyramidal cell columns.

While later sections will consider the case of unconstrained dipole orientations, I adopt the normal-orientation assumption here for the sake of algebraic clarity. Under this hypothesis, the forward model can be represented by a gain matrix G, where each row maps the activity at a given cortical location to each sensor in the recording device. Letting $x(t) \in \mathbb{R}^S$ denote a column vector representing dipole activity at each of the S source locations at time t, the observed sensor data $m(t)$ is obtained via the forward equation:

$$m(t) = Gx(t) \tag{1}$$

The ill-posedness of the inverse problem Now that we have established an expression linking source space activity to sensor space measurements, the inverse problem can be formulated in algebraic terms. In the forward problem, the goal is to compute $m(t)$ given the source activity $x(t)$. In contrast, the inverse problem aims to estimate the unknown brain activity time course $x(t)$ from observed sensor recordings $m(t)$.

The possible positions of dipoles within the cortex are typically discretized at a resolution on the order of 1 mm, resulting in tens of thousands of candidate source locations. Meanwhile, only a few hundred magnetic or electric field sensors are available at the scalp surface, rendering the problem fundamentally ill-posed: multiple distinct configurations of brain activity can produce the exact same recorded signal. This ill-posedness is further exacerbated by the existence of magnetically and/or electrically silent sources [6].

Here, the term "configuration" refers to any combination of dipoles with arbitrary positions and moments. However, not all such configurations are physiologically plausible or equally interpretable from a neuroscientific perspective. Several approaches have been proposed to model brain activity in a manner that is both interpretable and biologically

plausible.

In Bayesian terms, one may argue that incorporating prior knowledge or hypotheses about brain function allows for the identification of more plausible configurations, thereby enabling a unique and optimal solution under specific assumptions. This procedure is known in data science as regularization of the problem. Part of the regularization already occurs in the forward modeling step, where the source space is reduced to a finite set of discrete positions and, in some cases, constrained dipole orientations. Additional assumptions about the nature of brain activity can be introduced at the inverse stage, leading to different solutions depending on the chosen priors. In the following subsection, I review the foundational principles behind several historical approaches.

1.2. Major approaches to the inverse problem

Solutions to the inverse problem can be broadly divided into two categories: parametric approaches, which aim to explain the data using a small number of dipoles whose parameters are estimated; and imaging approaches, which estimate a distribution of currents over the full cortical surface. This thesis focuses on a specific class of parametric methods and therefore does not cover imaging approaches, although methods such as eLORETA [21] have demonstrated their effectiveness. Instead, I present the two principal families of parametric solutions: dipole fitting and beamforming.

Equivalent current dipole fitting (ECD) A natural approach to modeling brain activity is to fix a number n of point-like current dipoles and fit their positions and moments to best explain the observed data. This can be formalized in Bayesian terms. The model equation extends the forward equation (1) as follows:

$$m(t) = \sum_{i=1}^{n} g(\theta_i) x_i(t) + \varepsilon(t) = G(\theta)x(t) + \varepsilon(t) \tag{2}$$

Here, $\theta(t) = \{\theta_i(t)\}_{i=1}^{n}$ denotes the set of dipole positions and orientations to be estimated at each time step, $G(\theta)$ is the gain matrix with rows corresponding to sensor gains from each dipole, and $x(t)$ is the vector of dipole intensities. The term $\varepsilon(t)$ models residual noise, accounting for discrepancies between the model and the observed data.

One may assume either that both dipole intensities and positions vary over time (the *moving dipole* model), in which case the parameters are fitted independently for each time point, or that dipole positions are fixed over time (the fixed dipole model). This section focuses on the moving dipole case; the fixed dipole hypothesis will be revisited in the final section.

Since the noise term introduces stochasticity, parameter estimation relies on Bayesian inference. A common Bayesian strategy is to solve the maximum likelihood problem:

$$\max_{\theta, x} \ \mathbb{P}\left(m(t) \mid \theta, x\right) \tag{3}$$

Assuming the noise is spherical Gaussian white noise, the maximum likelihood estimate corresponds to the solution of a nonlinear least squares problem:

$$\min_{\theta, x} \ \|m(t) - G(\theta)x(t)\|_2^2 \tag{4}$$

In practice, aside from external sources and artifacts—which can often be mitigated by hardware design [6] or preprocessing methods [9], [19]—the dominant source of noise is intrinsic "resting-state" brain activity. This noise is correlated, both due to underlying neural dynamics and the non-orthogonal projections of brain activity onto sensor space. For the least squares solution to remain valid, one must either weight the residuals appropriately [4] or whiten the data [24].

To avoid the computational cost of brute-force optimization over all parameters, the problem is generally approached via variable separation. For any fixed spatial configuration θ, the optimal dipole activity vector $x^*_\theta(t)$ can be computed via linear regression. This yields the closed-form solution given by the Moore–Penrose pseudoinverse $G^+(\theta)$ of the gain matrix [14]:

$$x^*_\theta(t) = G^+(\theta)\, m(t) = \left[G^\top(\theta)G(\theta)\right]^{-1} G^\top(\theta)m(t) \tag{5}$$

For further simplification, when dipole orientations are not constrained, the estimated activity $x(t)$ is typically a $T \times 3$ matrix, where each row encodes the dipole moment components along the Cartesian axes. In this case, the matrix G is independent of orientation and contains the leadfields corresponding to each unit dipole direction. The optimization in Equation (4) thus reduces to a nonlinear search over spatial parameters, often solved using derivative-free iterative algorithms. While such methods may converge to local optima, modern computational resources permit exhaustive search over finely discretized source spaces—up to tens of thousands of candidate positions.

A key limitation of dipole fitting methods is that the number of dipoles must be specified a priori. Modern approaches address this through a sequential fitting procedure [23]: a first dipole is fitted, its contribution subtracted from the data, and a second dipole is then fitted to the residual signal. While increasing the number of dipoles generally improves the model's fit by increasing the rank of the data explained, it also raises the risk of overfitting, potentially introducing spurious dipoles that model noise rather than genuine neural activity. Bayesian criteria and expert knowledge about plausible activation patterns can help guide this choice, although the statistical reliability of such decisions remains a matter of debate.

Beamforming Another important class of approaches relies on spatial filtering to isolate signals originating from a specific location of interest. Such a filter is called a beamformer. An ideal beamformer would achieve unit gain at the target location and zero gain elsewhere. However, this ideal response is unattainable due to the limited degrees of freedom imposed by the number of sensors.

A classical solution to approximate this goal is the linearly constrained minimum variance (LCMV) beamformer. This technique minimizes the output power of the spatial filter—suppressing signals from all directions—while enforcing a unit-gain constraint at the target location. Formally, the LCMV beamformer for a given spatial position α is obtained by solving the constrained optimization problem:

$$\min_{W} \|Wm(t)\| \qquad \text{subject to} \qquad WG(\alpha) = I, \tag{6}$$

where $m(t) \in \mathbb{R}^J$ is the vector of sensor measurements at time t, $G(\alpha)$ is the leadfield matrix at source location α, W is the spatial filter (i.e., the beamformer weights), and I is the identity matrix, which is scalar in the fixed-orientation case and 3×3 in the free-orientation case.

Let $C = \mathbb{E}[m(t)m(t)^\top]$ denote the sensor covariance matrix. Solving the problem using Lagrange multipliers yields:

$$W = \left(G(\alpha)^\top C^{-1} G(\alpha)\right)^{-1} G(\alpha)^\top C^{-1} \tag{7}$$

This expression ensures unit gain at location α—passing the signal from that location without attenuation—while minimizing contributions from all other directions.

In the free-orientation case, $G(\alpha) \in \mathbb{R}^{J\times 3}$ models dipoles oriented along three orthogonal directions (typically the Cartesian axes), and $W \in \mathbb{R}^{3\times J}$ contains three corresponding spatial filters. In the fixed-orientation case, each location is assigned a single orientation (e.g., normal to the cortical surface), so $G(\alpha) \in \mathbb{R}^J$ is a column vector and the constraint becomes scalar.

The beamformer output at location α is then given by:

$$s(t;\alpha) = Wm(t), \tag{8}$$

which provides an estimate of the source time series at that location. The norm of $s(t;\alpha)$, either at each time point or over the full time course, can be used to localize sources in space.

We have now established that parametric approaches can follow two distinct strategies: either fitting a projection that explains the largest portion of the data (dipole fitting), or scanning the source space to find the position that exhibits the highest activity while suppressing contributions from all other locations (beamforming). Both approaches share the assumption that sources correspond to local maxima of data-related power, and both allow the derivation of explicit measures of goodness-of-fit at a given location.

In fact, the operations applied to the leadfields to evaluate source activity (see Equations 7 and 5) are structurally similar. However, while dipole fitting typically relies on the noise covariance matrix for pre-whitening, beamforming uses the data covariance matrix to suppress activity in unrelated regions. This reflects the beamformer's underlying assumption that brain signals at different locations are geometrically uncorrelated. As a result, beamforming

has been reported to perform poorly when attempting to separate sources with correlated leadfield geometries [17].

1.3. Subspace-based methods

Motivation This internship project is based on the idea that, through linear decomposition of high-dimensional data, one can obtain both a vector basis for the signal subspace and a basis for its orthogonal complement. The vectors that project the data onto the orthogonal subspace can be considered as null filters for the source of interest. In other words, signals originating from the source of interest are completely extinguished by these null filters. The source generating the signal is said to belong to the zero-set of locations for any projection into the orthogonal subspace. With the help of a forward model, one can therefore estimate the source location through a simple cost minimization.

This approach, which we initially referred to as zero-set imaging, is in fact already implemented in the MUSIC algorithm and its variants. While we do not claim to introduce a completely novel methodology, we frame the MUSIC algorithm as part of a broader class of subspace-based methods. In this subsection, we outline the general principles of such methods, with particular emphasis on the specific case of the MUSIC algorithm.

Linear analysis Linear analysis refers to a class of techniques that generate informative projections of observed data. Based on linear algebra and statistical principles, these methods yield powerful data representations and are widely used in signal processing and data science. In signal processing, when observed data are assumed to result from linear combinations of underlying sources and additive noise, linear analysis enables the computation of unmixing transformations to recover the original sources—an instance of the blind source separation problem. In EEG and MEG applications, this linearity assumption is often reasonable, making linear analysis a compelling approach for extracting meaningful signals from raw recordings.

Notable examples include principal component analysis (PCA), which identifies orthogonal directions of maximal variance, and independent component analysis (ICA), which seeks statistically independent components in the data. These methods rely on different assumptions about the mixing process and offer distinct strengths and limitations.

For example, PCA is effective at identifying low-rank signals embedded in Gaussian white noise, provided the signal strength exceeds the largest noise eigenvalue. However, due to its orthogonality constraint, geometrically correlated sources are not represented by a single principal component, but rather distributed across the signal subspace.

In contrast, ICA aims to maximize the non-Gaussianity of the signal projections in order to recover statistically independent sources. According to the central limit theorem, mixtures of independent variables tend to appear more Gaussian; hence, maximizing non-Gaussianity—often through higher-order statistical moments—facilitates the separation

of independent components. Although ICA is designed for noise-free environments, it can be combined with PCA-based pre-denoising and has demonstrated strong performance in separating components that may correspond to distinct, spatially localized brain sources.

While PCA and ICA are foundational examples that will be examined in this report, it is important to note that other linear subspace extraction methods are especially well suited to the problem of isolating brain source activity from MEG and EEG data. These methods vary in their assumptions about the mixing process, the statistical criteria they rely on, and the application scenarios in which they are most effective. In particular, we will present the joint diagonalization method, a linear analysis technique derived from PCA, which incorporates additional hypotheses about the signal of interest to enhance source extraction. We will also introduce ICA-based subspace localization as a means to address source mixing in linear analysis of data containing multiple simultaneous sources.

Subspace-based localization The linear analysis framework enables the decomposition of sensor-level data into two orthogonal subspaces: the signal subspace, which emphasizes components of interest (e.g., neural activity) while attenuating irrelevant or noisy components, and its orthogonal complement, which ideally suppresses the signal of interest. Projectors onto the signal subspace thus act as spatial filters that enhance activity originating from specific locations in the brain, whereas the orthogonal complement projectors serve as spatial null filters that cancel out contributions from sources at those locations.

Assume that the linear analysis correctly identifies a low-rank signal subspace, denoted by $\mathcal{P}$. Let P be the associated orthonormal projection matrix, whose k rows span the subspace. The orthogonal complement of this projection is denoted by $Z = P^{\perp}$, which projects data onto the subspace $\mathcal{Z}$ orthogonal to the signal. This subspace is generally referred to as the noise subspace, although noise may also be present within the signal subspace. For simplicity, we consider a constrained source model in which the dipole orientation is fixed to be orthogonal to the cortical surface. In this case, a source is fully characterized by its spatial location α^* and amplitude parameter x.

Under this assumption, each row of the orthogonal projector defines a zero-set condition for the true source's associated leadfield vector g^*, expressed as:

$$\forall j \in \{1, \dots, J-k\} \qquad \sum_{l=1}^{n} Z_{jl} g_l^* = 0 \tag{9}$$

or more compactly using the L_2 norm:

$$\|Zg^*\|_2 = 0 \tag{10}$$

This implies that the gain vector corresponding to the true source location satisfies a system of homogeneous equations, each defining a hyperplane in the signal space. Ideally, the correct source location lies at the intersection of these hyperplanes (zero-sets). However, a unique solution is not guaranteed: any gain vector lying entirely within the signal subspace

will satisfy this condition, resulting in a low-dimensional solution space and potentially multiple candidate locations. Nonetheless, if the linear analysis has successfully isolated the true signal, the corresponding source position α^* will lie within this intersection.

In practice, anatomical constraints—namely, that sources must lie within the brain volume—significantly reduce the likelihood of spurious solutions. However, the higher the dimensionality of the signal subspace, the greater the chance that unrelated positions may also satisfy the zero-set conditions. Furthermore, due to numerical errors and the imperfect nature of subspace estimation under noisy conditions, exact zero solutions are rarely observed. As a result, source localization is typically reformulated as a minimization problem over a discrete set of candidate positions. The following loss function is evaluated at each candidate location i:

$$\ell(i) = \frac{1}{\|g_i\|_2^2} \sum_{z \in \mathcal{Z}} \langle z, g_i \rangle^2 = \frac{\|Z g_i\|_2^2}{\|g_i\|_2^2} \tag{11}$$

Here, g_i denotes the gain vector at location i, and the sum is taken over the orthonormal basis vectors of the orthogonal complement $\mathcal{Z}$. Minimizing this loss yields the position with the smallest projection onto the noise subspace—that is, the most likely origin of the signal. If the signal subspace is multidimensional, the signal is likely generated by multiple distinct sources, and several local minima will typically be observed. A brute-force search over S candidate source locations has a computational complexity of $O(S)$, which remains tractable even for high-resolution source spaces exceeding 10,000 locations. The full cost map can also be visualized; sharp and deep local minima provide strong evidence for likely active source locations.

MUSIC, related methods, and beyond The idea of estimating a low-dimensional subspace that captures the signal of interest and identifying source locations that lie within this subspace has a long history. As with many MEG source localization techniques, the MUltiple SIgnal Classification (MUSIC) algorithm originates from the field of array signal processing, where it was originally developed to estimate the number and direction of arrival of signals impinging on a sensor array [3]. In the MEG context, MUSIC was adapted as an effective method for localizing multiple active sources by identifying locations whose leadfields lie within the estimated signal subspace. The signal subspace is typically estimated via a singular value decomposition (SVD) of the data covariance matrix. Source locations are then identified by minimizing a cost function such as Equation (11), which quantifies the orthogonality between candidate leadfields and the noise subspace [5].

Building on this framework, Sekihara et al. [15] proposed a method that projects a linearly constrained minimum variance (LCMV) beamformer onto the signal subspace. This subspace-projected beamformer exhibits improved robustness to noise and source correlation compared to the standard LCMV approach. A key refinement of MUSIC was introduced with the Recursively Applied MUSIC (RAP-MUSIC) algorithm [10], designed

to better handle multiple simultaneously active sources. While classical MUSIC relies on identifying local minima in the source space—a process that can be unstable in the presence of source correlations—RAP-MUSIC proceeds iteratively. After identifying a source, it projects the signal subspace onto the orthogonal complement of the subspace spanned by the corresponding leadfield. This recursive out-projection aims to isolate the contributions of the remaining sources in subsequent iterations.

However, RAP-MUSIC's recursive projection scheme is susceptible to error propagation, particularly in the presence of noise or forward model inaccuracies. To address this, Mäkelä et al. [25] introduced Truncated RAP-MUSIC (TRAP-MUSIC), which improves robustness by not only projecting out the identified source subspace but also reducing the dimensionality of the remaining signal subspace. Specifically, after projection, only the dominant remaining components—those associated with the largest singular values—are retained. This truncation step mitigates the accumulation of modeling errors and prevents weak or noise-aligned components from influencing subsequent iterations.

As in dipole fitting, pre-whitening has been proposed to improve PCA-based signal subspace extraction, since PCA is Bayes-optimal only under the assumption that noise is white [7]. Additional enhancements, such as frequency filtering to isolate specific frequency bands, have also been explored to improve signal extraction [13].

This thesis combines a comparison of subspace-based methods with the standard dipole fitting paradigm and an exploration of non-standard subspace extraction methods for source localization. These ideas are examined through a comprehensive study using realistic synthetic MEG data. I compare the standard MUSIC algorithm with classical dipole fitting in the single-source setting and show that, while both approaches exhibit comparable performance, MUSIC is generally more robust to low noise levels. I also evaluate each method's sensitivity to forward model inaccuracies. The results suggest that PCA-derived methods incorporating signal-specific hypotheses can yield alternative subspaces and may outperform standard MUSIC under certain conditions.

In the multiple-source setting, I show that MUSIC is consistently more robust than dipole fitting, particularly in scenarios involving spatially correlated sources. To address the challenge of source mixing in this setting, I explore whether independent component analysis (ICA) can replace recursive strategies such as RAP-MUSIC to disentangle the sources' time courses.

The results indicate that linear analysis methods specifically designed to enhance the signal-to-noise ratio can lead to more accurate subspace estimation and, consequently, higher localization precision.

I then address the multiple-dipole case by comparing an ICA-based source separation strategy with the standard RAP-MUSIC approach. Finally, I provide a theoretical analysis comparing subspace-based localization with dipole fitting, highlighting both their method-

ological similarities and their fundamental differences.

2. Methods

2.1. Data generation

Source positions To compare the accuracy of various source localization methods, I generated realistic data mimicking sensory-evoked fields. For single-source localization, I focused on sources located in the auditory cortex of one hemisphere. For multiple-source localization, source positions were selected either within the same auditory cortex—resulting in high spatial correlation between sources—or in distinct areas (e.g., left auditory cortex and right visual cortex), resulting in lower correlation. To ensure meaningful comparisons, configurations that returned too low a correlation in the high-correlation condition (and vice versa) were excluded.

I used a realistic source model by discretizing the cortical surface based on MRI-derived anatomical data from the MNE-Python software [22]. For auditory-related responses, source locations were selected arbitrarily from a subset of 164 vertices in the right auditory cortex. For visual-related responses, a similar process was applied using a subset of 58 vertices in the right visual cortex. Except when evaluating the free-orientation case, dipole orientations were constrained to be normal to the cortical surface at the given location.

Forward modelling Using MNE-Python and the anatomical and registration data described above, I computed a forward model for a subset of 7432 cortical vertices, based on the boundary element method (BEM) [16]. This resulted in a leadfield tensor of shape $[3, 7432, 203]$, representing the gain vectors for each of the three orthogonal dipole moment directions at each vertex, mapping to the 203 MEG gradiometer channels. For simplicity, I used only gradiometers, as they typically offer higher spatial sensitivity than magnetometers and avoid complications related to unit scaling. MNE-Python also allows for simplification to a normal-orientation forward model, yielding a leadfield matrix with a single projection per source location.

Noise-free data To generate noise-free data, I first constructed a time course of activity for each active source. These typically consist of a stereotyped sine-exponential bump peaking at an arbitrary latency. The peak time could be identical across sources—resulting in temporally correlated time courses—or distinct. Each dipole signal was then projected into sensor space using the corresponding leadfield vector (in the normal-orientation case) or a moment-weighted combination of the three leadfield vectors (in the free-orientation case).

Additive noise Realistic sensor noise was added using the noise covariance matrix provided in the MNE sample dataset, which was computed from resting-state measurements and

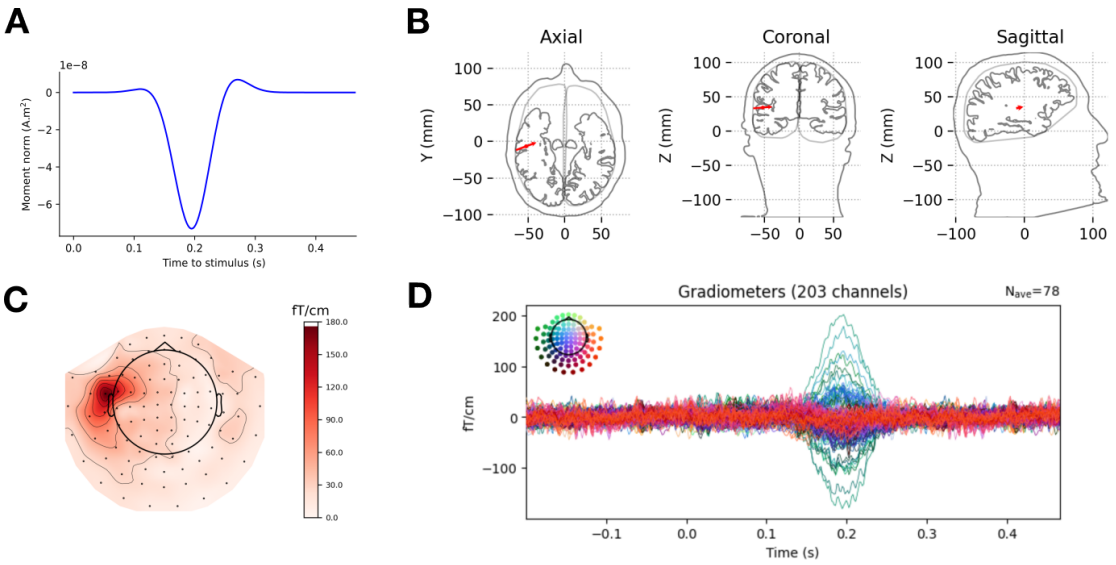

Figure 1. Generated data example. **A** Time course of the simulated dipole moment norm, peaking around 0.2 s post-stimulus, representing the temporal profile of the synthetic source. **B** Ground-truth source location used for data simulation, shown in MNI space across axial, coronal, and sagittal slices. **C** Topographic map of the simulated magnetic field gradients (fT/cm) at the peak response, projected onto the sensor array and scalp. **D** Simulated MEG gradiometer signals from 203 channels, showing trial-averaged evoked responses ($N = 78$). The inset displays the sensor layout with color-coded traces.

thus reflects a realistic sensor noise correlation structure. Noise samples were generated as independent realizations from a centered multivariate Gaussian distribution with this covariance. To further increase realism, we filtered the generated noise using a finite impulse response (FIR) filter, with parameters fitted to resting-state data.

Epoching and averaging In noisy simulations, the signal-to-noise ratio was controlled by repeating the signal of each dipole periodically, with timing locked to a symbolic, periodic stimulus onset. The dipole signals were repeated without variation across trials. While the signal magnitude remains constant, the noise magnitude decreases with the square root of the number of averaged trials. Averaging was performed using the MNE-Python environment.

2.2. Linear analysis

Pre-whitening Many data analysis methods rely on the assumption that the signal is contaminated by Gaussian white noise, which justifies the need for pre-whitening. In this study, pre-whitening was performed using the exact same noise covariance matrix that was used to generate the data, rather than an empirical estimate, in order to avoid introducing additional bias during processing. The eigenvalue decomposition of the noise covariance matrix $C = U\Lambda U^\top$ was computed using NumPy. Given $x \in \mathbb{R}^J$, a single data vector, pre-whitening was applied via the transformation

$$x_{\text{whit}} = U\Lambda^{-\frac{1}{2}}U^{\top}x,$$

such that if x contains only noise, the transformed data has unit expected covariance. While this transformation can lead to some distortion of the signal, it improves the performance of subsequent linear analysis methods.

Principal Component Analysis PCA was performed using NumPy's standard singular value decomposition algorithm, which is more numerically stable than computing the eigenvalue decomposition of the covariance matrix, especially when the number of samples is small or the data is rank-deficient (i.e., fewer samples than dimensions). The data was always pre-whitened before applying PCA, ensuring optimal recovery of low-rank signals corrupted by additive spherical noise.

Independent Component Analysis Independent Component Analysis was carried out using the FastICA implementation from the `scikit-learn` library with standard parameters. Instead of performing full whitening internally, we used the same pre-whitening transformation as described above, based on the known noise covariance matrix. FastICA aims to decompose the signal into statistically independent components. Based on the intuition that the sum of independent, identically distributed variables is more Gaussian than the individual variables themselves, FastICA seeks to maximize the non-Gaussianity of the components. Specifically, it maximizes the expected value of the logarithm of the hyperbolic cosine function applied to the components, which serves as a proxy for the negentropy.

Joint Decorrelation Joint Decorrelation was used as a versatile technique for isolating components under stronger assumptions than PCA. Assuming that a transformation in the temporal domain L improves the signal-to-noise ratio, the method proceeds by iteratively identifying orthogonal spatial projections p that maximize the ratio:

$$\frac{\|pML\|_2}{\|pM\|_2}, \tag{12}$$

which reduces to solving a generalized eigenvalue problem. In this study, we applied Joint Decorrelation using different forms of L, including stimulus-locked averaging and frequency filtering followed by averaging.

2.3. Source localization

Forward model In general, the forward model used to solve the inverse problem is identical to the one that generated the data, in order to eliminate modelling-related errors. To assess robustness to modelling inaccuracies, I also used alternative forward models: either a three-compartment spherical model with the same vertex positions provided by anatomical data, or a reduced-resolution source model.

Normally-oriented dipole hypothesis After extracting a signal subspace, the orthogonal complement of the signal-enhancing projection matrix was estimated using `SciPy`. The subspace-based loss function defined in Equation (11) was then computed for each vertex in the forward model, allowing for direct search minimization across candidate source locations.

Free-orientation dipole While the normal-orientation hypothesis simplifies the source localization problem, several parametric methods—such as equivalent current dipole fitting and beamforming—can provide estimates of both the position and orientation of active dipoles [14]. Relaxing the orientation constraint can be valuable in cases where the source of electromagnetic activity is not restricted to cortical pyramidal cells, when attempting to localize deep sources, or when using a volumetric source model. Furthermore, loosening this constraint has been shown to improve robustness to forward model inaccuracies such as limited resolution and head misalignment [18].

To accommodate such scenarios, I developed an efficient algorithm that first estimates the optimal dipole orientation for each candidate location before minimizing the subspace-based loss using the selected orientation. The algorithm still assumes that dipole orientation remains fixed over time and is therefore not suitable for modeling rotating dipoles.

Let $e = [e_x, e_y, e_z]$ be the unit-norm orientation vector of the dipole, and let $G(i) = [g_{ix}, g_{iy}, g_{iz}]$ be the leadfield matrix at position i, with columns corresponding to the three Cartesian orientations. The loss to be minimized now includes the orientation dependence:

$$\ell(i, e_x, e_y, e_z) = \frac{1}{\|e_x g_{ix} + e_y g_{iy} + e_z g_{iz}\|_2^2} \sum_{z \in \mathcal{Z}} \left\langle z, e_x g_{ix} + e_y g_{iy} + e_z g_{iz} \right\rangle^2 \tag{13}$$

or, more compactly using matrix notation:

$$\ell(i, e) = \frac{\|ZG(i)e\|_2^2}{\|G(i)e\|_2^2} \tag{14}$$

As in Equation (11), the normalization accounts for systematic variations in gain due to position and orientation. In particular, radial dipole activity is more easily detected by MEG sensors, which is corrected for by the denominator.

Minimizing this loss is a mixed optimization problem, involving a discrete search over positions and a continuous search over orientation vectors. However, for a fixed position i, the orientation minimization problem

$$\min_{e} \quad \frac{\|ZG(i)e\|_2^2}{\|G(i)e\|_2^2} \qquad \text{subject to} \qquad \|e\|_2 = 1 \tag{15}$$

reduces to a generalized eigenvalue problem for 3×3 covariance matrices. The optimal orientation e^* is given by the eigenvector associated with the smallest generalized eigenvalue, which corresponds to the minimum of the loss ratio. This problem can be efficiently solved using NumPy's linear algebra tools, yielding both the optimal orientation and the associated

loss for each candidate position. An exhaustive search over positions can then be performed to identify the most likely source location.

2.4. Multiple source handling

RAP-MUSIC Although linear analysis methods enable extraction of the signal subspace, in the multiple dipole case the activity time courses of individual sources are generally mixed within the signal components. RAP-MUSIC was used to handle the multiple-source case in a sequential fashion. After localizing each source, the gain vector g_k corresponding to the newly identified dipole is projected out of the signal subspace using:

$$P_k^{\perp} = I - \frac{g_k g_k^{\top}}{g_k^{\top} g_k}$$

To prevent re-localizing the same source, the leadfield matrix G used in subsequent iterations is modified by applying the projection:

$$G' = P_k^{\perp} G$$

This operation orthogonalizes each candidate leadfield vector with respect to g_k, ensuring that the MUSIC cost function no longer responds to components already explained by the previously identified dipole. Since MUSIC evaluates the alignment between the signal subspace and each candidate source topography, this step ensures that only novel, unaccounted-for components are tested for alignment.

ICA-based subspace decomposition As an alternative to RAP-MUSIC, we developed an approach that leverages Independent Component Analysis (ICA) to directly extract direction vectors, each ideally corresponding to a distinct source. This is feasible because ICA is specifically designed to separate independent sources from observed mixtures. After identifying the ICA components, each was treated individually: instead of projecting onto the full signal subspace, we localized sources associated with each ICA component separately.

Recursively applied dipole fitting (RAP-ECD) Similar in spirit to RAP-MUSIC, Recursively Applied Dipole Fitting (RAP-ECD) extends the standard equivalent current dipole (ECD) fitting approach to the multiple-source case through a sequential strategy. After each dipole is localized via least squares minimization, its corresponding gain vector g_k is projected out of the forward model to prevent it from being re-selected. This is done using the same outprojection operator defined in Equation (2.4.0.1) on the leadfield before continuing the analysis. By recursively applying this projection after each step, the fitting algorithm ensures that subsequent dipoles explain new variance in the data, rather than re-explaining already-modeled components.

2.5. Single dipole

Sensitivity to noise I investigated the robustness of subspace-based methods to noise in the single-dipole case. As a starting point, I compared the standard PCA-based MUSIC algorithm with the time-fixed dipole fitting approach. The source was placed arbitrarily in the left auditory cortex, and the emitted signal consisted of a stereotyped exponential bump peaking around 200 ms post-stimulus, as illustrated in Figure 1. The number of stimulus repetitions N_{ave} was varied between 5 and 250. Since signal-to-noise ratio improves with the square root of the number of averages, N_{ave} served as a proxy for SNR.

To reduce stochastic variability and ensure statistical power, I repeated the simulation 100 times for each value of N_{ave}, with the source location randomly sampled within the auditory cortex. Both MUSIC and ECD localization methods were applied to each generated dataset using the exact forward model and noise covariance matrix employed in the simulation, ensuring no modelling errors were introduced (see Methods). Under these ideal conditions, both methods should be theoretically capable of perfect localization in the absence of noise.

As shown in Figure 2, both methods perform similarly across a wide range of N_{ave} values, transitioning from high precision at large N_{ave} to total failure in very low-SNR

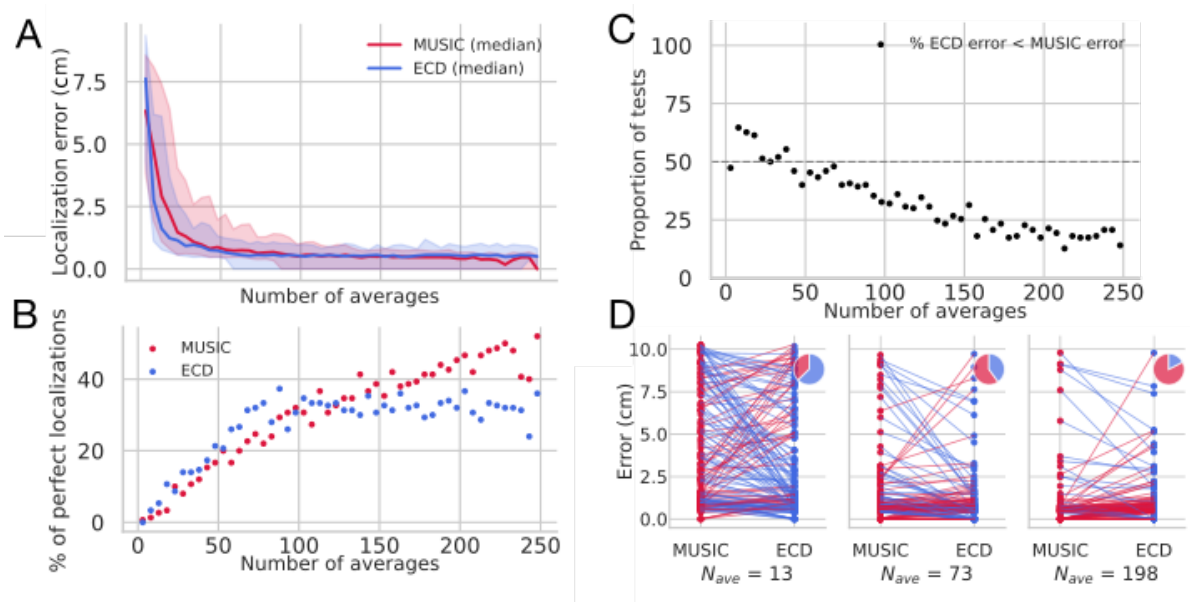

Figure 2. Comparison of the robustness of MUSIC and dipole fitting (ECD) to low signal-to-noise ratios. **A:** Median localization error (in cm) as a function of the number of averaged trials. Shaded areas represent inter-subject variability. MUSIC performs worse than ECD at low N_{ave}, but consistently outperforms it as the number of averages increases. The horizontal axis is shared with **B**. **B:** Percentage of perfect localizations as a function of N_{ave}. Both methods perform similarly at low averaging levels, but ECD plateaus early while MUSIC continues to improve with additional averaging. **C:** Percentage of cases in which ECD yields smaller localization error than MUSIC across averaging levels. The frequency of ECD outperforming MUSIC decreases with increasing N_{ave}. **D:** Paired comparisons of localization errors for MUSIC and ECD at three representative values of N_{ave} (13: low-SNR regime; 73: intermediate; 198: high-SNR). Each line connects the results of a single test for MUSIC (left) and ECD (right). Pie charts indicate the proportion of cases where each method was superior, in line with trends from **C**.

regimes. Notably, MUSIC outperforms dipole fitting for higher SNR levels, as evidenced by a greater number of perfect localizations. A Wilcoxon signed-rank test revealed that MUSIC significantly outperformed ECD (at $p < 0.01$) when $N_{\text{ave}} > 83$, while ECD was significantly better than MUSIC only in the range $N_{\text{ave}} \in [13, 33]$ (see supplementary figures).

In the theoretical supplement, I demonstrate that the performance similarity arises from deep mathematical connections between the two methods—particularly in the single-source case—even though they can yield different predictions in practice. The better performance of MUSIC in low-noise regimes suggests that isolating the signal subspace prior to evaluating source alignment increases robustness relative to direct fitting.

Interestingly, the minimum value of the MUSIC loss function becomes substantially shallower in noisy settings compared to noise-free conditions (e.g., roughly 1×10^{-2} in the noise-free case vs. 0.15 with 100 averaged trials), even when the localization is exact. A more systematic analysis showed that shallow minima can still correspond to accurate localizations, but deep minima were never associated with incorrect results (see supplementary figures). This indicates that the depth of the MUSIC loss minimum is a conservative but reliable proxy for localization confidence, assuming a perfect forward model.

Improvements using other linear analysis methods We show that, beyond simple PCA, other linear analysis methods can improve signal subspace detection by incorporating additional assumptions about the signal. While PCA is Bayes-optimal under the assumption of additive spherical noise and no prior knowledge about the signal, it does not leverage domain-specific hypotheses—such as temporal structure—that could guide signal extraction.

As an illustrative example, we constructed a scenario in which the signal-to-noise ratio was too low for standard methods to succeed, but where prior knowledge of the signal's frequency content could be exploited. A single source was placed arbitrarily in the auditory cortex and assigned a 75 Hz sine-wave time course, mimicking Gamma-band neural activity. The signal amplitude was reduced such that, even after averaging over 100 trials, both MUSIC and ECD failed to localize the source accurately.

To overcome this, we performed signal subspace extraction using Joint Decorrelation (see Methods), applying a 50–100 Hz band-pass filter as a temporal transformation. The extracted components correspond to orthogonal spatial directions whose variance is maximally increased by this filtering. Localizing this signal subspace using the standard MUSIC loss function resulted in a perfectly accurate dipole localization.

Sensitivity to forward modelling accuracy We complemented the previous results by comparing the robustness of MUSIC and Dipole Fitting to inaccuracies in the forward model. Two types of model mismatch were considered: one with the same anatomical geometry used for data generation but a reduced-resolution source model (0.01 downsampling ratio), and another with the same source model but a simplified 3-compartment spherical head model.

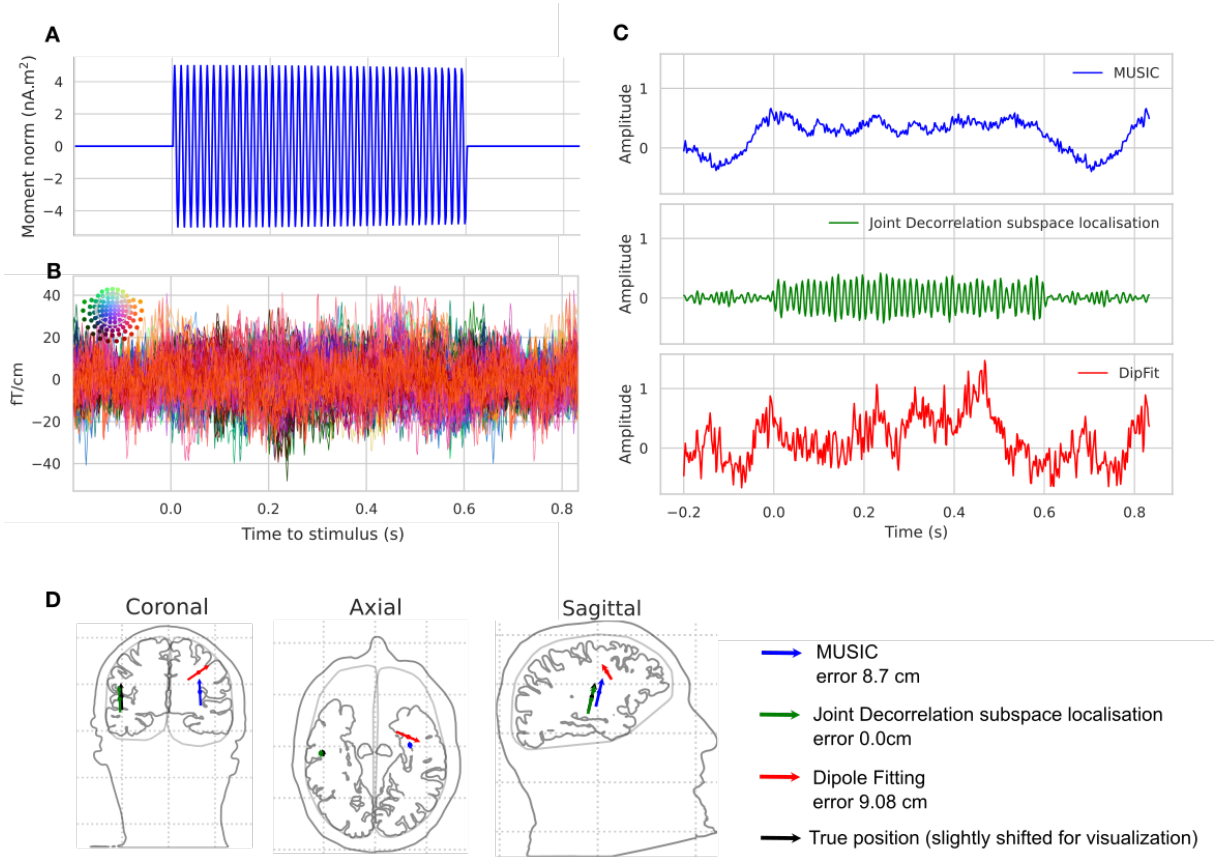

Figure 3. Band-pass Joint Decorrelation analysis improves localization of a blurred signal with known frequency. **A:** Simulated dipole moment time course with oscillatory activity at 75 Hz. **B:** MEG sensor-level data showing magnetic field fluctuations across time and channels. The signal is not visible to the naked eye, indicating a very low signal-to-noise ratio. **C:** Reconstructed source amplitude time courses using three localization methods: MUSIC (blue), Joint Decorrelation subspace localization with a 50–100 Hz band-pass filter (green), and Dipole Fitting (red). While standard methods fail due to the low SNR, joint decorrelation successfully reconstructs the source activity. Bottom panels: estimated dipole positions and orientations for each method, shown in coronal, axial, and sagittal brain slices. Colored arrows indicate localized dipole directions: MUSIC (blue), Joint Decorrelation (green), and Dipole Fitting (red), highlighting the improved accuracy of the frequency-informed method.

In a single-source, noise-free simulation, both inaccurate forward models led MUSIC to localization errors on the order of 1 cm. Contrary to expectations from previous findings [18], switching from the normal-orientation hypothesis to the free-orientation model further degraded localization accuracy (see supplementary figures).

The MUSIC loss at the estimated location remained slightly higher than in the ideal forward model case but was still on the order of 10^{-2}, indicating that moderate loss values do not necessarily imply poor head modelling. Instead, shallow minima in the MUSIC loss landscape are more likely to reflect poor signal subspace estimation, typically due to noise or artifacts.

Interestingly, Dipole Fitting yielded exactly the same predictions as MUSIC when using

identical forward models. This result is consistent with theoretical expectations: in the single-dipole, noise-free case, both methods are strictly equivalent (see the theoretical complement for a detailed proof).

2.6. Multiple dipole case

MUSIC vs. Dipole Fitting in the multiple dipole case I compared the ability of both MUSIC and Dipole Fitting methods to handle multiple dipoles using recursive application algorithms. After each dipole localization, the data explained by the previously identified source was projected out, and the localization algorithm was reapplied to the residual data (RAP MUSIC and RAP Dipole Fitting).

Two configurations involving two dipoles were considered. In the first scenario, one dipole was placed in the left auditory cortex and the other in the right visual cortex, resulting in near-orthogonal leadfields (cosine similarity < 0.05). In this case, source time courses are expected to be easily separable, and recursive localization algorithms are expected to perform well. In the second scenario, both dipoles were located in the left auditory cortex, leading to substantial spatial correlation between their leadfields (cosine similarity > 0.2).

To account for variability due to specific source positions, 100 random configurations were generated in each scenario. No noise was added to the simulations, allowing us to isolate algorithmic performance. As shown in Figure 4, RAP MUSIC significantly outperformed RAP Dipole Fitting when the dipoles were spatially correlated, particularly for the second localization: RAP ECD produced errors greater than 1 cm in 75% of trials, whereas RAP MUSIC achieved errors below 1 cm in 75% of cases ($p < 0.001$). For the first localized dipole in this configuration, RAP MUSIC was also significantly more accurate ($p < 0.01$).

When dipoles were spatially uncorrelated, both methods yielded comparable results with generally low localization errors. Nonetheless, RAP MUSIC still significantly outperformed RAP ECD for the first localization ($p < 0.001$).

Overall, these findings suggest that RAP MUSIC is more robust than RAP Dipole Fitting in the presence of spatially correlated sources. In the noise-free two-dipole case, MUSIC loss minima were higher than in the single-dipole case but remained below 0.05 for both sources, indicating that multiple sources do not inherently cause shallow minima—provided the subspace dimension is correctly set and out-projections are properly applied.

ICA as an alternative to RAP—or the dangers of source separation When applying PCA in the multiple-source case, and if the sources' leadfields are not perfectly orthogonal, source mixing is inevitable. This occurs because PCA produces an orthogonal basis for the signal subspace, which cannot perfectly separate correlated sources. RAP and TRAP MUSIC address this issue by spatially filtering out the activity associated with previously identified sources, effectively isolating new sources in the residual data. However, as discussed in the paragraph on beamforming, spatial filtering cannot perfectly isolate punctual sources,

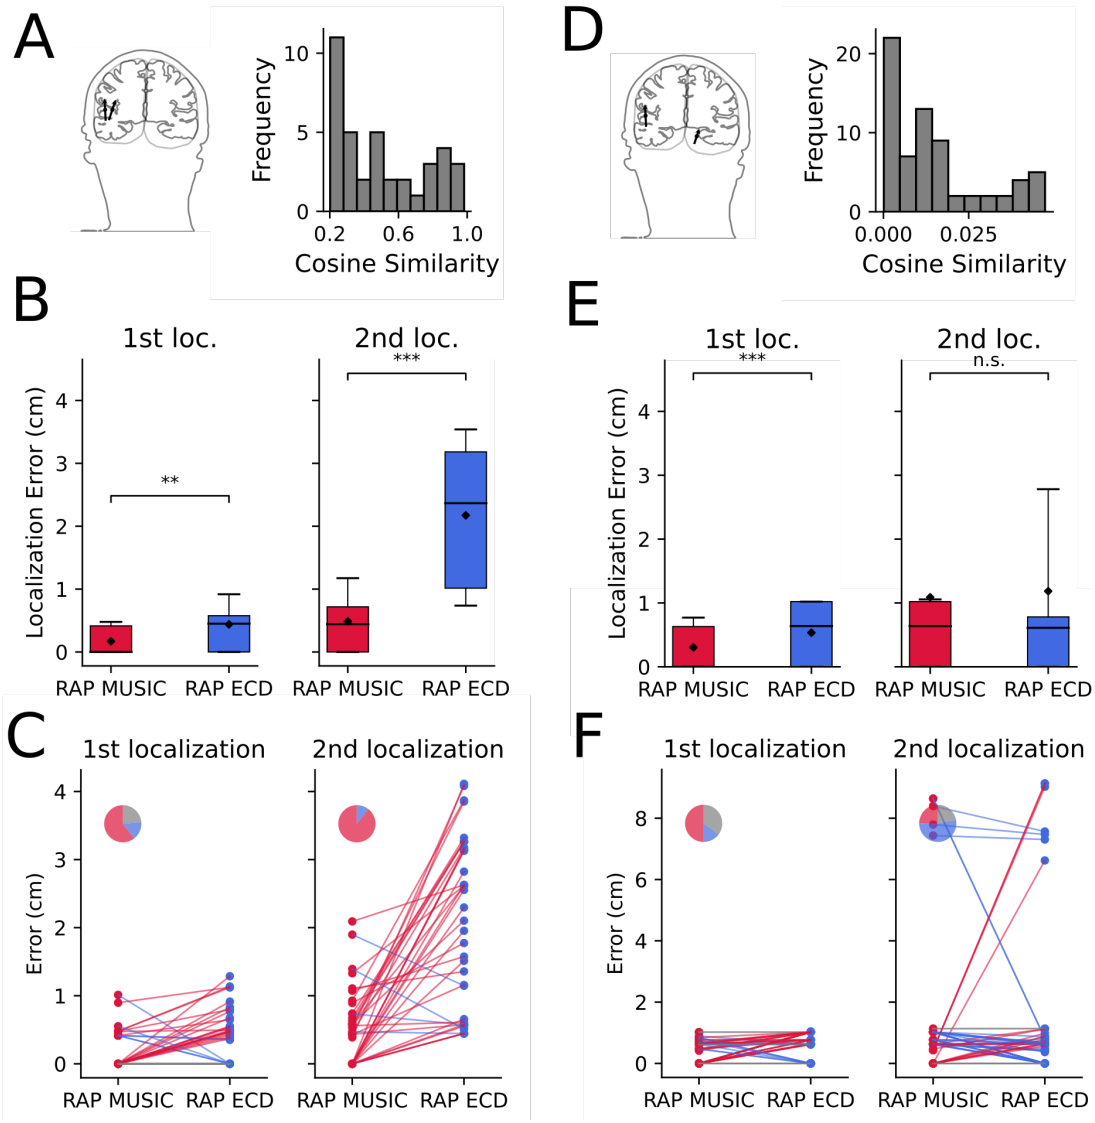

Figure 4. Comparison of RAP MUSIC and RAP Dipole Fitting in two-dipole configurations. **A:** Left: example source locations in distant areas (left auditory cortex and right visual cortex). Right: distribution of leadfield correlations between the two sources. **B:** Localization errors for both methods in the distant-source scenario. Diamond: mean; central line: median; whiskers: 10th and 90th percentiles. Significance: *** $p < 0.001$, ** $p < 0.01$, n.s. not significant (Wilcoxon signed-rank test). **C:** Pairwise comparison for the same configurations. Pie charts indicate the proportion of simulations where each method performed better. **D–F:** Same analyses for the correlated-source scenario (both sources in the left auditory cortex).

especially in the presence of spatial correlations. Consequently, residual source leakage can introduce localization errors when sources are too spatially correlated.

One might conjecture that if a linear decomposition method perfectly recovers the time courses of each source, then localizing the source associated with each component independently—using a subspace-based loss function—should yield accurate spatial localization. To challenge this reasoning, we compared ICA-based component-wise subspace localization (see Methods) with the standard RAP MUSIC and RAP Dipole Fitting approaches in a

scenario involving correlated dipole leadfields.

As shown in Figure 5, ICA effectively separates the individual source time courses (panel F), whereas PCA yields mixed components (panel C). RAP MUSIC estimates time courses by projecting the data onto the leadfield of each localized source: the first component is derived from a projection onto the first estimated leadfield, while the second is obtained by projecting the residual data—orthogonalized with respect to the first leadfield—onto the second leadfield (panel D). As a result, only the second time course is correctly unmixed due to the explicit outprojection. Similar behavior is observed for RAP Dipole Fitting (panel E).

Counterintuitively, panel G reveals that ICA component-wise localization fails to correctly localize either dipole. The theoretical explanation is that projecting the data onto the correct leadfield does not guarantee unmixed time courses, particularly when leadfields are correlated. In this example, even the time course estimated using RAP MUSIC with the correct leadfield (panel D) shows evidence of source mixing. This behavior is expected due to the correlation between dipoles. While ICA successfully recovers time courses by excluding projections aligned with the leadfields of other sources, this means that the gain vectors associated with each ICA component are no longer aligned with the true leadfields. As a consequence, localization via ICA-derived subspaces points to erroneous locations—neither corresponding to the true sources.

Therefore, ICA should be understood as a tool for unmixing source time courses rather than estimating spatial signal subspaces. A model-based and computationally efficient strategy to improve time course unmixing in RAP MUSIC would be to recursively project out the leadfields of all other sources from each subspace used for time course estimation. In panel D, this was done only for the second source. To improve the estimate of the first source's time course, the leadfield of the second source should also be projected out from the subspace used to reconstruct the first source's activity.

2.7. Real data example

To demonstrate the applicability of the MUSIC algorithm to real data, we used the MNE sample dataset, which consists of averaged auditory evoked fields (AEFs) recorded following a right-ear auditory stimulus. The data represent an average over 73 trials. We first performed principal component analysis (PCA) on the evoked data and compared the resulting eigenvalues with the theoretical upper bound of the Marčenko–Pastur distribution, which describes the limiting spectrum of a white noise random matrix with the same aspect ratio as our data matrix. Three components exceeded this threshold, suggesting a three-dimensional signal subspace.

We then applied RAP MUSIC using these three top principal components to estimate successive source locations. After each source was localized, its associated leadfield was projected out of the subspace before proceeding to the next step. The values of the subspace-based loss function at the optimal positions were 0.33 for the first dipole, 0.35 for the second,

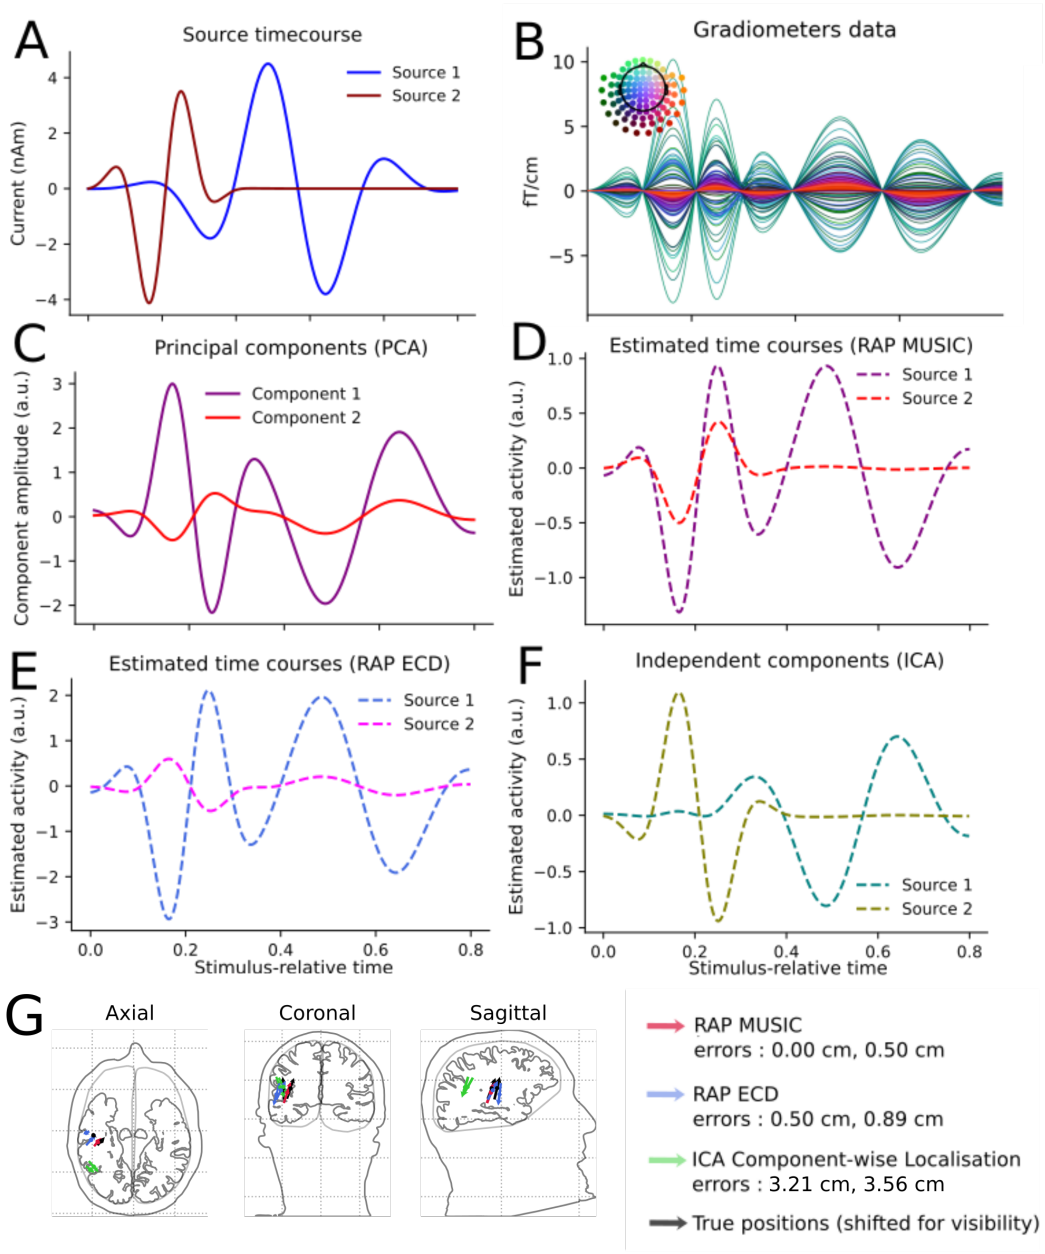

Figure 5. Source mixing in localization methods. **A:** Simulated time courses. The leadfields of the two dipoles had a correlation of 0.2 (cosine similarity). **B:** Resulting sensor data. Similar sensors are activated throughout the epoch. **C:** Principal component analysis of the data. Time courses are mixed across components. **D:** Time courses estimated using RAP MUSIC. The first component is obtained by projection onto the first leadfield; the second is extracted from residual data after outprojecting the first leadfield. Only the second time course is correctly unmixed. **E:** Time courses estimated using RAP ECD. As with MUSIC, the second time course is derived from residual data. Both time courses exhibit mixing. **F:** Time courses obtained from ICA closely match the ground truth with minimal mixing. **G:** Source localization results. RAP MUSIC yields accurate positions with minor error on the second dipole; RAP ECD performs poorly, especially on the second source; ICA fails to correctly localize either source.

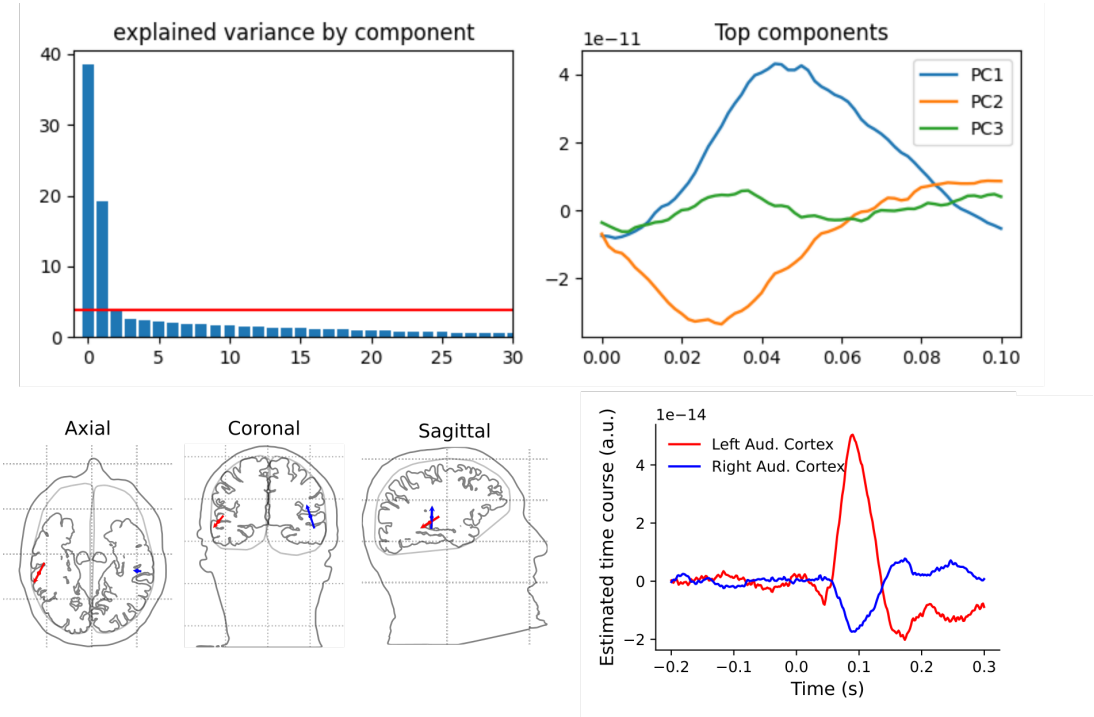

Figure 6. MUSIC applied to real data. Top left: eigenvalues associated with the top components of the data. Top right: first few principal components of the data (cropped around the region of highest variance). Bottom left: dipole locations for the first two steps of RAP MUSIC. Bottom right: third estimated dipole location, with shallower loss indicating lower reliability.

and 0.58 for the third. The substantial increase in the loss minimum after the second dipole suggests that the third localization is less reliable, likely corresponding to noise or model misfit rather than a true dipole. We therefore retained only the first two localized dipoles.

The two identified sources were located in the left and right auditory cortices, respectively. This is consistent with the expected bilateral auditory processing pathway. To estimate the time courses of the localized sources, we followed the approach described in Section 3.2: for each dipole, the evoked data were projected onto its leadfield after removing the contribution from the other source via orthogonal projection. The resulting time courses revealed a strong positive deflection in the right auditory cortex and a smaller negative deflection in the left auditory cortex, which aligns with the known predominance of contralateral auditory projections.

3. Theoretical Complement

3.1. Subspace-based methods as dipole fitting on pre-processed data

Reformulation of the subspace-based loss In this section, I demonstrate that subspace-based approaches such as MUSIC can be interpreted within the framework of dipole fitting. Intuitively, minimizing the subspace loss function in Equation (11) amounts to identifying a forward model vector g_k that is maximally orthogonal to the noise subspace. Equivalently,

this corresponds to finding the direction with the largest projection onto the signal subspace. I provide an algebraic justification for this interpretation, concluding that selecting the best dipole location using the signal subspace is effectively equivalent to fitting a dipole to denoised data. For clarity, the proof is restricted to the fixed-orientation dipole case, although the generalization to free orientations is straightforward.

Assume a k-dimensional signal subspace. Let $\mathcal{P} = \{p_1, \dots, p_k\}$ denote an orthonormal basis for the signal subspace and $\mathcal{N} = \{n_1, \dots, n_{J-k}\}$ for the noise subspace. Together, $\mathcal{S} = \mathcal{N} \cup \mathcal{P}$ forms an orthonormal basis of the full sensor space. Any vector $u \in \mathbb{R}^J$ can then be decomposed as:

$$u = \sum_{s \in \mathcal{S}} \langle s, u \rangle s = \sum_{n \in \mathcal{N}} \langle n, u \rangle n + \sum_{p \in \mathcal{P}} \langle p, u \rangle p = Nu + Pu, \tag{16}$$

where N and P are the orthogonal projection matrices onto the noise and signal subspaces, respectively (i.e., whose rows are the vectors of $\mathcal{N}$ and $\mathcal{P}$).

Using the Pythagorean theorem, the squared norm of u can be expressed in terms of its projections onto the signal and noise subspaces:

$$\|u\|_2^2 = \sum_{s \in \mathcal{S}} \langle s, u \rangle^2 = \|Nu\|_2^2 + \|Pu\|_2^2. \tag{17}$$

This identity allows us to rewrite the subspace loss defined in Equation (11) as:

$$\ell(k) = \frac{\|Ng_k\|_2^2}{\|g_k\|_2^2} = 1 - \frac{\|Pg_k\|_2^2}{\|g_k\|_2^2}. \tag{18}$$

Hence, minimizing the loss amounts to maximizing the energy of the forward model projected onto the signal subspace. In other words, the optimal dipole is the one whose leadfield vector lies most strongly within the signal subspace.

Dipole fitting loss maximizes average cosine similarity between data and leadfield vector
By combining Equations (5) and (4), the dipole fitting loss at time t can be expressed as:

$$\ell(\theta)_t = \left\| m(t) - G(\theta) \left[G(\theta)^\top G(\theta) \right]^{-1} G(\theta)^\top m(t) \right\|_2^2, \tag{19}$$

where $G(\theta)$ is the leadfield matrix associated with spatial parameters θ, and $m(t)$ is the sensor data at time t.

In the constrained orientation case, $G(\theta)$ reduces to a single column vector g_k, and θ corresponds solely to the position index k. The loss becomes:

$$\ell(k)_t = \left\| m(t) - \frac{1}{\|g_k\|_2^2} \langle g_k, m(t) \rangle g_k \right\|_2^2. \tag{20}$$

Expanding and simplifying the expression yields:

$$\ell(k)_t = \|m(t)\|_2^2 - \frac{1}{\|g_k\|_2^2} \langle g_k, m(t) \rangle^2. \tag{21}$$

Extending this to the time-fixed dipole case, the Frobenius norm over the entire data

matrix $M = [m(1), \dots, m(T)] \in \mathbb{R}^{J \times T}$ allows summing across all time points:

$$\ell(k) = \|M\|_F^2 - \frac{1}{\|g_k\|_2^2} \|g_k^\top M\|_2^2. \tag{22}$$

The loss thus ranges from 0 (when g_k is collinear to all $m(t)$) to $\|M\|_F^2$ (when g_k is orthogonal to all $m(t)$).

We now define the normalized loss:

$$\tilde{\ell}(k) = \frac{\ell(k)}{\|M\|_F^2} = 1 - \frac{\|g_k^\top M\|_2^2}{\|g_k\|_2^2 \|M\|_F^2}. \tag{23}$$

This normalized loss resembles the reformulated subspace-based loss in Equation (18), with two key differences: dipole fitting seeks to maximize alignment between the forward model and the observed data, whereas subspace-based methods maximize alignment between the forward model and the estimated signal subspace basis vectors.

Equivalence up to a pre-processing step I now show that applying a specific pre-processing transformation A, derived from the signal subspace projection matrix P, followed by dipole fitting, yields results equivalent to the MUSIC method. Let $M' = AM$ denote the pre-processed data. The normalized dipole fitting loss on M' reads:

$$\tilde{\ell}'(k) = 1 - \frac{\|g_k^\top PAM\|_2^2 + \|g_k^\top NAM\|_2^2}{\|g_k\|_2^2 \left(\|NAM\|_2^2 + \|PAM\|_2^2\right)}, \tag{24}$$

where P and N are the projection matrices onto the signal and noise subspaces, respectively.

Let W be the data whitening matrix such that the whitened data has identity covariance:

$$WM(WM)^\top = I. \tag{25}$$

Now define the pre-processing matrix as:

$$A = P^\top PW, \tag{26}$$

i.e., whitening the data, projecting it onto the signal subspace, and projecting back to sensor space. Applying dipole fitting to this low-rank pre-processed data yields:

$$\tilde{\ell}'(k) = 1 - \frac{\|g_k^\top PP^\top PWM\|_2^2 + \|g_k^\top NP^\top PWM\|_2^2}{\|g_k\|_2^2 \left(\|NP^\top PWM\|_2^2 + \|PP^\top PWM\|_2^2\right)}. \tag{27}$$

Since $PP^\top = I$ (orthogonal projection) and $NP^\top = 0$ (signal and noise subspaces are orthogonal), the expression simplifies to:

$$\tilde{\ell}'(k) = 1 - \frac{\|g_k^\top PWM\|_2^2}{\|g_k\|_2^2 \|PWM\|_2^2}. \tag{28}$$

Finally, since WM has identity covariance and thus does not alter norms up to a constant factor, the loss becomes:

$$\tilde{\ell}'(k) = 1 - \frac{\|g_k^\top P\|_2^2}{\|g_k\|_2^2 \|P\|_F^2}, \tag{29}$$

where $\|P\|_F^2$ is constant across all k. Therefore, minimizing this loss is equivalent (up to a constant scaling) to minimizing the reformulated subspace-based loss in Equation (18).

This result confirms that MUSIC and similar subspace-based methods can be interpreted as dipole fitting performed on whitened, low-rank data constrained to the signal subspace.

Subspace-based dipole fitting as a weighted subspace-based method In the previous paragraph, I showed that whitening the data before projecting it onto a low-dimensional subspace is necessary to retrieve the subspace-based loss formulation. This whitening step removes the influence of the data's variance structure, which is a key aspect of the dipole fitting method. While this allowed us to bridge the two approaches, it also highlighted a fundamental difference: subspace-based methods like MUSIC search for directions lying in the signal manifold, whereas dipole fitting applied to low-rank approximations also accounts for the data's variance structure along each direction of that manifold.

Let us now consider the low-rank approximation of the data without whitening, namely:

$$M' = P^\top P M.$$

The arguments leading to Equation (**??**) still hold, yielding:

$$\tilde{\ell}'(k) = 1 - \frac{\|g_k^\top P M\|_2^2}{\|g_k\|_2^2 \|P M\|_2^2}. \tag{30}$$

However, in this case, M cannot be factored out of the denominator. Expanding the numerator reveals the role of each direction $p \in \mathcal{P}$ of the signal subspace:

$$\|g_k^\top P M\|_2^2 = \sum_{p \in \mathcal{P}} \langle g_k, p \rangle^2 \cdot \|p^\top M\|_2^2. \tag{31}$$

Thus, minimizing the loss in Equation (30) is equivalent to maximizing the following expression:

$$\frac{1}{\|g_k\|_2^2} \sum_{p \in \mathcal{P}} \langle g_k, p \rangle^2 \cdot \|p^\top M\|_2^2. \tag{32}$$

This formulation clarifies the critical distinction between performing dipole fitting on subspace-reduced data and standard subspace-based localization. The term in Equation (32) corresponds to a weighted version of the subspace-based loss, where alignment with each component p is scaled by the corresponding projected energy $\|p^\top M\|_2^2$. In a classical MUSIC framework where the signal subspace is derived via PCA, this energy corresponds to the eigenvalue associated with eigenvector p.

Empirically, our results showed that the variance-based weighting induce a bias in the multiple source case leading to potential mislocations of the dipoles.

3.2. Similarities and differences between MUSIC and dipole fitting

In the previous section, I showed that subspace-based methods share structural similarities with dipole fitting. In the specific case of the standard PCA-based MUSIC algorithm, and under the assumption of a single dipole source, the link becomes particularly close. MUSIC first identifies the top principal component of the data and selects the leadfield vector that aligns best with this component. In contrast, dipole fitting directly finds the leadfield vector that aligns best with the full data. If the set of candidate leadfield vectors densely spanned the entire sensor space, both methods would yield identical results. However, this is not the case in MEG source localization, where leadfield vectors lie on a discrete and constrained manifold within the sensor space. This can lead to differences: dipole fitting may select the vector best aligned with the full data matrix, while MUSIC may prefer the one best aligned with the top component of variance.

To illustrate this potential divergence, consider a toy example in a 2-dimensional sensor space with two time steps. Let the data matrix $M \in \mathbb{R}^{2\times 2}$ have covariance eigenvalues $\lambda_1 = 2$, $\lambda_2 = 1$, with corresponding eigenvectors $v_1 = (1, 0)$, $v_2 = (0, 1)$. Assume the data has full rank due to additive noise.

Suppose we evaluate two unit-norm candidate leadfields:

$$g_1 = v_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad g_2 = \begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix}, \quad \theta = 60°.$$

Dipole fitting compares the squared alignment of each leadfield with the full data:

$$\|g_1^\top M\|_2^2 = 1, \qquad \|g_2^\top M\|_2^2 = 2\cos^2\theta + \sin^2\theta = 1.25.$$

It thus selects g_2 as the better-aligned leadfield.

By contrast, MUSIC evaluates the alignment with the top principal component v_1:

$$\langle g_1, v_1\rangle^2 = 1, \qquad \langle g_2, v_1\rangle^2 = \cos^2\theta = 0.25,$$

and selects g_1 instead.

This example highlights the key reason for divergence between the methods: 1) the leadfield space does not span all directions of the sensor space, and 2) the variance in some directions may originate from noise, not signal. Whether this divergence benefits one method over the other is context-dependent. If the estimated principal component is misaligned with the true signal direction (due to low SNR), MUSIC may localize the source less accurately. In contrast, dipole fitting can be misled by overfitting to noise directions.

Results from random matrix theory support this trade-off: in high-dimensional settings, PCA is known to exhibit a sharp phase transition in its ability to recover a low-rank signal buried in Gaussian noise. When the SNR drops below a critical threshold, the top eigenvector of the data covariance matrix becomes uninformative, leading to failure of both MUSIC

the underlying time courses, it yielded poor localization when applied component-wise, showing that signal subspace identification does not equate to time course unmixing. This emphasized the importance of operating on the full signal subspace when performing source localization.

A real data application using the MNE sample dataset illustrated the practicality of MUSIC in applied contexts. Subspace dimension estimation via eigenvalue thresholding allowed reliable localization of bilateral auditory sources. The loss minima were shallow but interpretable, and source locations matched well-established neurophysiological responses to auditory stimuli.

The theoretical complement formalized the connection between subspace methods and dipole fitting, emphasizing that their differences arise mainly from the variance-weighting applied in dipole fitting versus the directionality constraints used in MUSIC. A toy example illustrated how these approaches can yield different predictions when leadfields do not densely span sensor space or when data variance structure misaligns with actual signal.

Limitations Although the results consistently show MUSIC's strengths, several limitations must be acknowledged. First, MUSIC's performance is contingent on accurate estimation of the signal subspace, which is sensitive to noise and model misspecification. In practice, artifacts, head movement, or unexpected background activity can bias subspace estimation, reducing localization accuracy. While pre-whitening aims to mitigate these effects, its robustness under realistic non-stationary noise remains unclear.

Moreover, the advantages of MUSIC in the multiple-source case were established under idealized, noiseless conditions. These benefits may diminish when using imperfect forward models or in the presence of spatially correlated noise. Since MUSIC assumes a strict separation between signal and noise subspaces, it may be more vulnerable to model mismatch than dipole fitting, which flexibly adapts to the empirical variance structure. A comprehensive analysis that combines structured noise, model mismatch and artifacts should be considered.

Another limitation lies in the interpretation of loss depth. Although synthetic data showed that loss minima correlate with localization success, this relationship is less clear in real-world settings. Inaccurate head models, poor subspace estimation, or temporally variable sources can also lead to shallow minima, reducing the diagnostic utility of loss values alone.

Future directions The results of this study point toward several promising extensions. One immediate direction is improving signal subspace estimation under realistic conditions. Our analysis implicitly assumes access to well-separated signal and noise directions, but this may not hold in practice. Although many methods for robust or structured PCA exist, our findings suggest that tailoring subspace estimation to known signal features (e.g., frequency, latency, topography) could be more effective than general-purpose methods.

Another avenue concerns the statistical characterization of localization estimates. Cur-

and PCA-based approaches. In practice, MEG datasets are mesoscopic in size—with tens to hundreds of sensors and a comparable number of time points—resulting in smoother transitions and intermediate misalignment of the first component. In such cases, MUSIC can sometimes amplify errors compared to dipole fitting, as illustrated in Figure 2. However, in the high-SNR regime, MUSIC may outperform dipole fitting due to its robustness against noise-induced misalignments.

In the multi-dipole case, further differences arise. RAP dipole fitting and RAP MUSIC both rely on recursively projecting out the contribution of previously localized sources. However, TRAP MUSIC includes additional regularization and orthogonalization steps that increase divergence between the methods. These approximations, while designed to reduce overfitting, can further distinguish the behaviors and performance of MUSIC-based and fitting-based algorithms.

4. Discussion

In this work, I presented a comprehensive investigation of subspace-based methods for MEG source localization, with a particular focus on the MUSIC algorithm and its relation to traditional dipole fitting approaches. I first detailed the mathematical formulation of these methods, highlighting the conceptual parallels between subspace scanning and dipole fitting, which were further developed in the theoretical complement.

Empirical results in the single dipole case confirmed that both MUSIC and dipole fitting are capable of accurate localization when the signal-to-noise ratio (SNR) is sufficiently high. However, MUSIC proved more robust in high-SNR regimes, achieving a higher proportion of perfect localizations and maintaining lower variance in error distributions. This was attributed to its ability to isolate signal directions prior to fitting, thereby avoiding overfitting to noise. Dipole fitting, in contrast, was only found to outperform MUSIC in scenarios where noise was too strong for either method to succeed precisely, suggesting that MUSIC is generally a better greedy heuristic when any structure can be recovered.

Further experiments demonstrated the benefit of incorporating prior knowledge about the temporal properties of the signal. In particular, we showed that applying Joint Decorrelation with a band-pass filter could recover dipole locations in low-SNR conditions where standard PCA-based approaches failed. This highlights the flexibility of the subspace framework, which is not restricted to principal component analysis. Methods like Common Spatial Patterns (CSP) or other task-informed projections could similarly be adapted to optimize source extraction under particular hypotheses.

In the multiple dipole case, I evaluated recursive approaches such as RAP MUSIC and RAP dipole fitting. While both leverage a similar projection principle to reduce signal mixing, RAP MUSIC significantly outperformed dipole fitting when dipoles had spatially correlated leadfields. ICA was also tested as a decomposition method. Although it correctly unmixed

rent methods return deterministic solutions, but do not quantify uncertainty. Our findings imply that MUSIC, despite its robustness, can produce unreliable estimates when subspace estimation fails. This motivates the use of confidence metrics or Bayesian extensions of subspace localization to flag or correct potentially spurious results.

Additionally, our comparison with ICA-based localization raises important implications: unmixing time courses is not equivalent to identifying the signal subspace. Future work should explore how subspace methods and time-course separation can be combined without conflating their goals. This may involve iterative procedures or joint models that separate sources while maintaining spatial interpretability.

In summary, this work contributes both theoretical clarity and empirical insight to the understanding of subspace-based methods. It suggests that while MUSIC and related algorithms are powerful tools, their practical reliability hinges on accurate subspace estimation, careful forward modeling, and context-aware interpretation of results.

References

[1] P. Nunez, *Electrical fields oj the brain*, 1981.

[2] P. Gloor, "Neuronal generators and the problem of localization in electroencephalography: Application of volume conductor theory to electroencephalography," *Journal of clinical neurophysiology*, vol. 2, no. 4, pp. 327–354, 1985.

[3] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE transactions on antennas and propagation*, vol. 34, no. 3, pp. 276–280, 1986.

[4] J. Sarvas, "Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem," *Physics in Medicine & Biology*, vol. 32, no. 1, p. 11, 1987.

[5] J. C. Mosher, P. S. Lewis, and R. M. Leahy, "Multiple dipole modeling and localization from spatio-temporal meg data," *IEEE transactions on biomedical engineering*, vol. 39, no. 6, pp. 541–557, 1992.

[6] M. Hämäläinen, R. Hari, R. J. Ilmoniemi, J. Knuutila, and O. V. Lounasmaa, "Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain," *Reviews of modern Physics*, vol. 65, no. 2, p. 413, 1993.

[7] K. Sekihara, S. Miyauchi, and H. Koizumi, "Covariance incorporated meg-music algorithm and its application to detect si and sii when large background brain activity exists," *NeuroImage*, vol. 3, no. 3, S92, 1996.

[8] S.-G. Kim, W. Richter, and K. Uğurbil, "Limitations of temporal resolution in functional mri," *Magnetic resonance in medicine*, vol. 37, no. 4, pp. 631–636, 1997.

[9] M. A. Uusitalo and R. J. Ilmoniemi, "Signal-space projection method for separating meg or eeg into components," *Medical and biological engineering and computing*, vol. 35, pp. 135–140, 1997.

[10] J. Mosher and R. Leahy, "Source localization using recursively applied and projected (rap) music," *IEEE Transactions on Signal Processing*, vol. 47, no. 2, pp. 332–340, 1999. DOI: 10.1109/78.740118.

[11] J. Mosher, R. Leahy, and P. Lewis, "Eeg and meg: Forward solutions for inverse methods," *IEEE Transactions on Biomedical Engineering*, vol. 46, no. 3, pp. 245–259, 1999. DOI: 10.1109/10.748978.

[12] P. L. Nunez and R. B. Silberstein, "On the relationship of synaptic activity to macroscopic measurements: Does co-registration of eeg with fmri make sense?" *Brain topography*, vol. 13, pp. 79–96, 2000.

[13] K. Sekihara, S. S. Nagarajan, D. Poeppel, *et al.*, "Estimating neural sources from each time-frequency component of magnetoencephalographic data," *IEEE Transactions on Biomedical Engineering*, vol. 47, no. 5, pp. 642–653, 2000.

[14] S. Baillet, J. C. Mosher, and R. M. Leahy, "Electromagnetic brain mapping," *IEEE Signal processing magazine*, vol. 18, no. 6, pp. 14–30, 2001.

[15] K. Sekihara, S. Nagarajan, D. Poeppel, and Y. Miyashita, "Reconstructing spatio-temporal activities of neural sources from magnetoencephalographic data using a vector beamformer," in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*, IEEE, vol. 3, 2001, pp. 2021–2024.

[16] B. Fischl, D. H. Salat, A. J. Van Der Kouwe, *et al.*, "Sequence-independent segmentation of magnetic resonance images," *Neuroimage*, vol. 23, S69–S84, 2004.

[17] A. Hillebrand and G. R. Barnes, "Beamformer analysis of meg data," *International review of neurobiology*, vol. 68, pp. 149–171, 2005.

[18] F.-H. Lin, J. W. Belliveau, A. M. Dale, and M. S. Hämäläinen, "Distributed current estimates using cortical orientation constraints," *Human brain mapping*, vol. 27, no. 1, pp. 1–13, 2006.

[19] J. Dammers, M. Schiek, F. Boers, *et al.*, "Integration of amplitude and phase statistics for complete artifact removal in independent components of neuromagnetic recordings," *IEEE transactions on biomedical engineering*, vol. 55, no. 10, pp. 2353–2362, 2008.

[20] S. Baillet, "The dowser in the fields: Searching for meg sources," *MEG: An introduction to methods*, pp. 83–123, 2010.

[21] R. D. Pascual-Marqui, D. Lehmann, M. Koukkou, *et al.*, "Assessing interactions in the brain with exact low-resolution electromagnetic tomography," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 369, no. 1952, pp. 3768–3784, 2011.

[22] A. Gramfort, M. Luessi, E. Larson, *et al.*, "MEG and EEG data analysis with MNE-Python," *Frontiers in Neuroscience*, vol. 7, no. 267, pp. 1–13, 2013. DOI: 10.3389/fnins.2013.00267.

[23] A. Gramfort, M. Luessi, E. Larson, *et al.*, "MNE software for processing MEG and EEG data," *NeuroImage*, vol. 86, pp. 446–460, 2014. DOI: 10.1016/j.neuroimage.2013.10.027.

[24] D. A. Engemann and A. Gramfort, "Automated model selection in covariance estimation and spatial whitening of meg and eeg signals," *NeuroImage*, vol. 108, pp. 328–342, 2015.

[25] N. Mäkelä, M. Stenroos, J. Sarvas, and R. J. Ilmoniemi, "Truncated rap-music (trap-music) for meg and eeg source localization," *NeuroImage*, vol. 167, pp. 73–83, 2018.

Supplementary figures

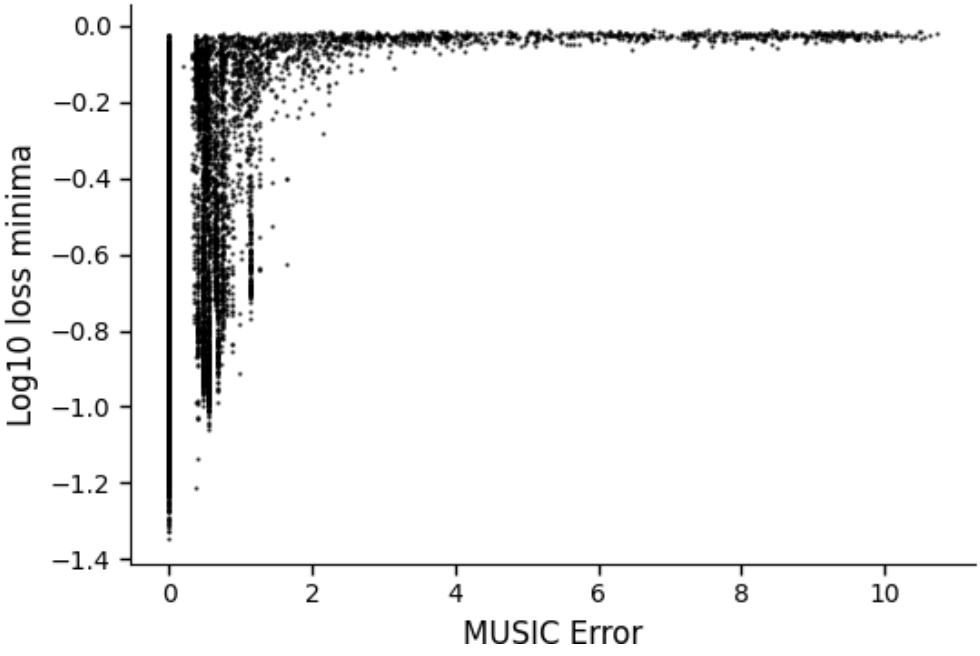

Figure 7. Loss minima deepness as a function of localization error (cm) for the MUSIC method (single dipole case)

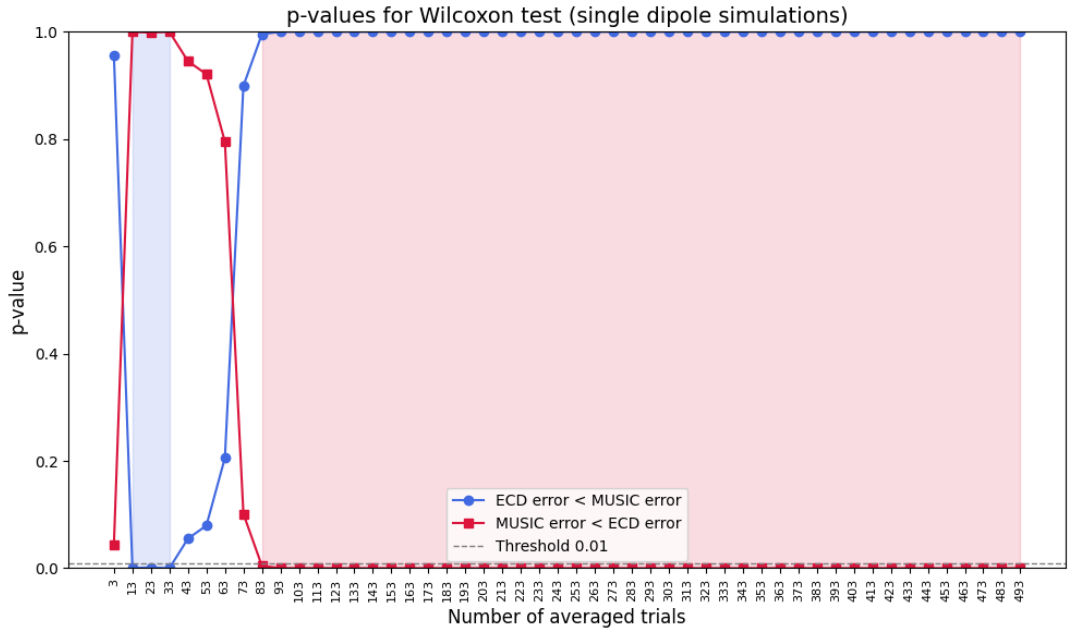

Figure 8. p-values for the Wilcoxon signed-rank test as a function of the number of averages Shaded areas correspond to areas were the p-value was lower than the threshold.

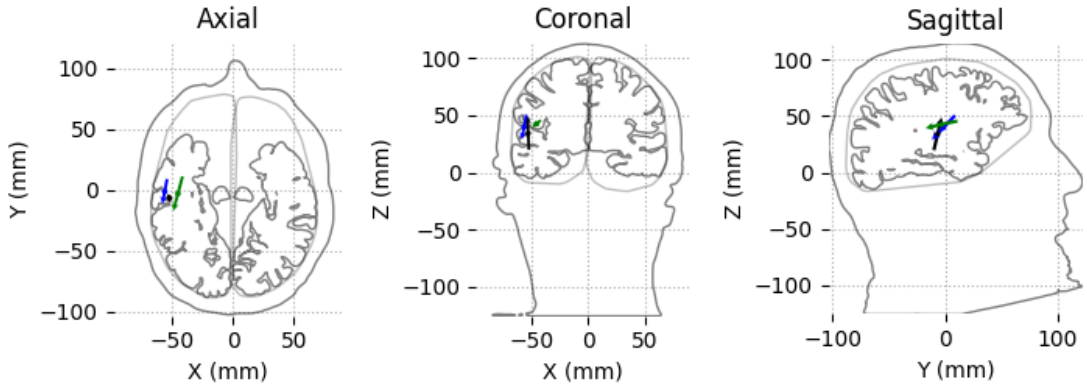

Figure 9. Localizations errors due to forward under the normal orientation hypothesis Black: true dipole location. Green: high-resolution source model paired with a bad sphere conductor model. Blue: low-resolution source model with exact conductivity modelling.

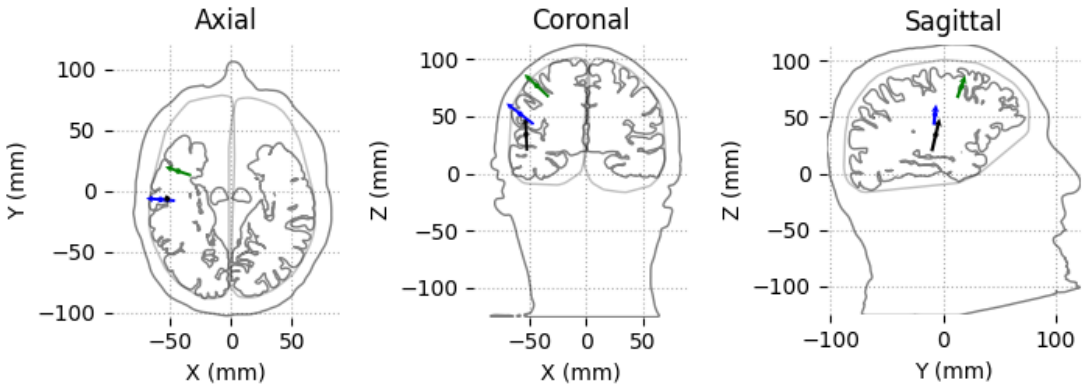

Figure 10. Localizations errors due to forward modelling without the normal orientation hypothesis Black: true dipole location. Green: high-resolution source model paired with a bad sphere conductor model. Blue: low-resolution source model with exact conductivity modelling.