Greedy Brain Source Localization with Rank Constraints

Viviana Hernandez-Castañon, Steven Le Cam and Radu Ranta
CRAN, Univ. de Lorraine-CNRS, Nancy, F-54000, France

Keywords: Sparse Source Localization, Greedy Algorithms, Stepwise Regression, Rank Constraints, EEG.

Abstract: This paper introduces a new low rank matrix approximation model and a greedy algorithm from the iterative regression family to solve it. Unlike the classical Orthogonal Matching Pursuit (OMP) or Orthogonal Least Squares (OLS), the elements of the dictionary are not vectors but matrices. For reconstructing a measurement matrix from this dictionary, the regression coefficients are thus matrices, constrained to be low (unit) rank. The target application is the inverse problem in brain source estimation. On simulated data, the proposed algorithm shows better performances than classical solutions used for solving the mentioned inverse problem.

1 INTRODUCTION

This work finds its initial motivation in the inverse problem of brain source estimation from electrophysiological signals. It can be formulated as the more general problem of a constrained low rank matrix approximation: given a dictionary of matrices $\bar{\mathbf{K}}_i \in \mathbb{R}^{m\times q}, (q < m)$ and a data matrix $\mathbf{X} \in \mathbb{R}^{m\times n}, (m < n)$, find the rank 1 matrices $\bar{\mathbf{J}}_i \in \mathbb{R}^{q\times n}$ such as the $\tilde{\mathbf{X}} = \sum_i \bar{\mathbf{K}}_i \bar{\mathbf{J}}_i$ approximates closely $\mathbf{X}$. Before introducing the details of this model and the greedy algorithm we propose, we will briefly justify it by explaining its application in neuroscience.

Macroscopic brain electrophysiological signals $\mathbf{X}$ are recorded by sensors placed on the scalp (EEG), on the cortex (ECoG) or implanted in the brain (SEEG). The sources $\mathbf{J}_i$ that generate the measurements are neural populations modelled as electrical current dipoles, characterized by their positions, orientations and amplitudes (6 parameters per dipole). All possible positions are given as a grid of points inside the discretized brain volume. The sensors are distant to the active sources and thus they record a sum of propagated activities. The propagation is assumed to be instantaneous at the frequencies of interest (below 1kHz) and at the distances between sources and sensors (few cm) Logothethis et al. (2007); Ranta et al. (2017). The forward problem expressing the recorded potentials knowing the dipolar sources' characteristics and the propagation environment writes thus in the linear form introduced above, where the gain (lead-field) matrix $\mathbf{K}$ relating the dipolar sources $\mathbf{J}$ to the measurements $\mathbf{X}$ is known and it embeds the geometry and physical properties of the propagation environment (head-model).

The inverse problem aims to estimate the dipolar sources (positions, orientations and amplitudes) responsible for the measured signals. This is an ill-posed problem, mainly because of its dimensions: usually, only a few dozens of sensors are available (at most a few hundreds), while the possible source positions and orientations are far more numerous, depending on the numerical resolution of the head model. Noise and spatial redundancy add to the difficulties of the problem.

The literature on inverse problem solving in brain electrophysiology is extremely vast. Several techniques were proposed in the last 30 years (see *e.g.*, the reviews from Baillet et al. (2001); Michel et al. (2004); He et al. (2011, 2018)). We are particularly interested in this paper in sparse methods, which aim to estimate a low number of sources explaining most of the measured signals. Scanning methods such as MUSIC and beamforming, and in particular the iterative versions such as (T)RAP-MUSIC Mosher and Leahy (1999); Mäkelä et al. (2018), yield naturally sparse solutions. They can be seen as two-steps approaches: first, scan for the most plausible positions and orientations (*i.e.*, choose among the columns of the gain matrix the ones most likely to project the sources on the sensors) and second estimate the amplitudes minimizing the error between the actual measurements and their modeled counterpart. By construction (sparsity + minimization), they do not completely explain the data (unlike minimum norm type solutions Dale and Sereno (1993); Hämäläinen and Il-

Hernandez-Castañon, V., Le Cam, S. and Ranta, R.
Greedy Brain Source Localization with Rank Constraints.
DOI: 10.5220/0013386600003911

In *Proceedings of the 18th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2025) - Volume 1*, pages 819-826
ISBN: 978-989-758-731-3; ISSN: 2184-4305

moniemi (1994); Pascual-Marqui (2007)). They neither impose a fixed orientation for the estimated dipolar sources, *i.e.*, the three Cartesian projections of the dipole are allowed to vary independently in time.

This paper proposes a new greedy algorithm derived from the family of step-wise regressions on a given (forward) dictionary[1]. Unlike for its predecessors, the regressors are not column vectors, but matrices, representing 3D lead-field matrices associated to each possible source position. Moreover, the regression coefficients for each source (time varying dipolar amplitudes) are explicitly constrained to have estimated fixed orientations. Thus, compared to state-of-the-art methods based on distributed source models where the source orientations are fixed and known (e.g., Gramfort et al. (2013); Korats et al. (2016)), the proposed method explicitly estimates sources with fixed orientations starting from a freely oriented distributed model.

2 FORWARD MODEL

The data matrix $\mathbf{X}$ ($m \times n$, channels by time samples) is a linear mixture of the electrical fields of p dipolar sources. At any time instant the (no-noise[2]) forward model writes:

$$\mathbf{x} = \sum_{i=1}^{p} \begin{bmatrix} k_{x,i,1} & k_{y,i,1} & k_{z,i,1} \\ k_{x,i,2} & k_{y,i,2} & k_{z,i,2} \\ \vdots & \vdots & \vdots \\ k_{x,i,m} & k_{y,i,m} & k_{z,i,m} \end{bmatrix} \begin{bmatrix} j_{x,i} \\ j_{y,i} \\ j_{z,i} \end{bmatrix} = \sum_{i=1}^{p} \bar{\mathbf{K}}_i \mathbf{j}_i(t) = \mathbf{K}_I \mathbf{j}_I, \quad (1)$$

where $\mathbf{x}$ ($m \times 1$) is the vector of measured potentials (electrodes) at time t, $\bar{\mathbf{K}}_i$ ($m \times 3$) represent the propagation coefficients of the unit dipole situated at position i and $\mathbf{j}_i$ (3×1) are the 3D projections of the dipole at position i (3×1) at time t ($k_{a,b,c}$ is the propagation coefficient in the direction $a \in \{x,y,z\}$ of the dipole in position $b \in \{1 \dots p\}$ onto the sensor $c \in \{1 \dots m\}$). In compact form, $\mathbf{K}_I$ ($m \times 3p$) is the lead-field matrix of the active sources (the concatenation of p blocks $\bar{\mathbf{K}}_i$ $m \times 3$), $\mathbf{j}_I$ the dipolar projections ($3p \times 1$) and I the set of active dipoles. Note that the we place ourselves directly in a "sparse" configuration, *i.e.*, when the indices in I do not cover all the columns of the lead-field matrix or, in other words, when I encodes the positions of the active sources. The set of I indices and thus the set of $(m \times 3)$ $\bar{\mathbf{K}}_i$ matrices is supposed to belong to a much larger set (known dictionary) $\mathbf{K}$ encoding the propagation coefficients for all possible dipolar current sources (*i.e.*, , in the neuroscientific application we consider, distributed over all possible positions in the brain).

Considering time samples $t = 1 \dots n$, (1) becomes:

$$\mathbf{X} = \sum_{i=1}^{p} \left(\begin{bmatrix} k_{x,i,1} & k_{y,i,1} & k_{z,i,1} \\ k_{x,i,2} & k_{y,i,2} & k_{z,i,2} \\ \vdots & \vdots & \vdots \\ k_{x,i,m} & k_{y,i,m} & k_{z,i,m} \end{bmatrix} \begin{bmatrix} j_{x,i}(1) & \dots & j_{x,i}(n) \\ j_{y,i}(1) & \dots & j_{y,i}(n) \\ j_{z,i}(1) & \dots & j_{z,i}(n) \end{bmatrix} \right) = \sum_{i=1}^{p} \bar{\mathbf{K}}_i \bar{\mathbf{J}}_i = \mathbf{K}_I \mathbf{J}_I \quad (2)$$

with $\mathbf{X} : m \times n$ and $\mathbf{J}_I : 3p \times n$, *i.e.*, the model proposed in the introduction for $q = 3$. Without noise, the rank of $\mathbf{X}$ is upper bounded by $\min\{m, n, 3p\}$. In a spatially sparse setup, the number of active sources p is much smaller than the number of channels m. In most of the situations, one assumes that the number of available time samples is also sufficiently big ($n \gg m$). Adding noise to the data restores its rank to *a priori* $r = \min\{m,n\} \gg 3p$ and a low-rank ($3p$) approximation of $\mathbf{X}$ is supposed to recover the no-noise data and, in the same time, the set of position indices I and the time varying amplitudes $\mathbf{j}_i$, solving thus the brain sources inverse problem.

According to the model introduced in the previous equations, the matrices $\bar{\mathbf{J}}_i$ are not constrained: for a given dipole i, its projections $[j_{x,i}(t) \;\; j_{y,i}(t) \;\; j_{z,i}(t)]^T$ can vary independently with the time t. The dipole has thus a varying orientation in time and $\bar{\mathbf{J}}_i$ is of rank 3. Still, there are different justifications for a fixed orientation hypothesis: if a dipole theoretically represents a cortical column, it shouldn't change its orientation in time. Even if usually a dipole represents the activity of several columns (a "patch" of cortex), one can still consider that the dipoles are orthogonal to the cortical surface. Another popular fixed known orientation is orthogonal to the head surface: indeed, dipoles parallel to the surface are not seen by the closest surface electrodes (EEG), as they are in the "zero-field" of the dipoles. In both previous cases, one can consider the (normalized) vector of orientations $\mathbf{o}_i$ for each dipole as known and rewrite the model as:

[1] *e.g.*, the Matching Pursuit – MP and its derivatives OMP (Orthogonal Matching Pursuit), OLS (Orthogonal Least Squares) or the more recent Single Best Replacement – SBR Soussen et al. (2011, 2015).

[2] We do not explicitly consider the noise in this paper, as it will be implicitly taken into account by the sparse nature of the considered algorithms.

$$\mathbf{X}=\sum_{i=1}^{p}\left(\begin{bmatrix} k_{x,i,1} & k_{y,i,1} & k_{z,i,1} \\ k_{x,i,2} & k_{y,i,2} & k_{z,i,2} \\ \vdots & \vdots & \vdots \\ k_{x,i,m} & k_{y,i,m} & k_{z,i,m} \end{bmatrix}\begin{bmatrix} o_{x,i} \\ o_{y,i} \\ o_{z,i} \end{bmatrix}[s_i(1)\ \ldots\ s_i(n)]\right)$$

$$=\mathbf{K}_I\begin{bmatrix} \mathbf{o}_1 & 0 & \ldots & 0 \\ 0 & \mathbf{o}_2 & \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots \\ 0 & \ldots & 0 & \mathbf{o}_p \end{bmatrix}\begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \\ \vdots \\ \mathbf{s}_p \end{bmatrix}=\underbrace{\mathbf{K}_I\mathbf{O}_{B,I}}_{\mathbf{A}_I}\mathbf{S}_I \qquad (3)$$

with $\mathbf{O}_{B,I}$ a $3p \times p$ block diagonal matrix having on the diagonal the *known* orientations $\mathbf{o}_i$ of the p dipoles and $\mathbf{S}_I$ the $p \times n$ amplitudes matrix containing the stacked time courses $\mathbf{s}_i$ of these dipoles ($\mathbf{A}_I = \mathbf{K}_I\mathbf{O}_{B,I}$ is the collapsed lead-field matrix $m \times p$ that can be obtained for known orientations).

The known orientation hypothesis is not always justified: the "orthogonal to the head surface" orientation is very restrictive and ignores the folded shape of the cortical surface. Besides, it cannot be used for inverse problems based on intracerebral measurements as in SEEG. On the other hand, the "orthogonal to the cortical surface" orientation is subject to the imaging quality and to the resolution of the numerical models. In this paper, we make a relaxed assumption: the dipoles are supposed to have fixed but *unknown* orientations.

3 INVERSE PROBLEM

The inverse problem we aim to solve is the following: knowing the measurements $\mathbf{X}$ and the head model $\mathbf{K}$, estimate the positions $i \in I$, the *fixed unknown* orientations $\mathbf{o}_i$ and the time-varying amplitudes vectors $\mathbf{s}_i$, under the assumption of sparseness, *i.e.*, a low cardinal of I. By default, the *known* lead-field matrix $\mathbf{K}$ is "fat", *i.e.*, the total number of potential source positions is much higher that the number of electrodes.

3.1 MUSIC Type Algorithms

Although the problem has not been formulated as above, the MUSIC family of algorithms, and especially the recursive versions RAP-MUSIC Mosher and Leahy (1999) and TRAP-MUSIC Mäkelä et al. (2018), provide a possible solution. These methods start by finding a "spatial" basis $\mathbf{V}$ of the $\mathbf{X}$ measurements by SVD:

$$\mathbf{X} = \mathbf{V}\mathbf{D}\mathbf{W}^T$$

In other words, each column of $\mathbf{X}$ (the spatial distribution of the measured signals at a given time instant) is a weighted sum of the columns of the $\mathbf{V}$ basis. This basis is next truncated by grouping in a matrix $\mathbf{V}_r$ only the r singular vectors corresponding to the largest singular values (assuming that the others correspond to noise).

The next step consists in "scanning" the $\mathbf{K}$ matrix and evaluate for each sub-matrix $\bar{\mathbf{K}}_i$ ($m \times 3$, containing the projection coefficients of the dipole i on the m electrodes) its canonical correlation with $\mathbf{V}_r$ (the truncated basis)[3]. The chosen positions (the set of indices $i \in I$ of the dipoles included in the solution) are those corresponding to the maximums (in absolute value) of the canonical correlation coefficients (in the recursive versions mentioned above, these correlations are computed after a succession of re-projections between the sub-spaces of $\mathbf{V}_r$ and $\mathbf{K}$ and the set of indices I gets larger as the iterations proceed, see for details Mosher and Leahy (1999); Mäkelä et al. (2018)).

For every selected position i, the estimates of the dipolar orientations $\hat{\mathbf{o}}_i$ are in turn provided by the coefficients of the linear combination of the 3 columns of $\bar{\mathbf{K}}_i$ which generates a vector as close as possible (*i.e.*, with the smallest angle) to the base signal space $\mathbf{V}_r$ (or its residual after reprojection in the iterative versions). For every position, once the orientation is estimated, the 3-columns submatrix $\bar{\mathbf{K}}_i$ is collapsed into a single column vector by right multiplication with the estimated orientation vector $\hat{\mathbf{o}}_i$, yielding a column vector $\mathbf{a}_i$ as in model (3). For the p selected dipoles (positions), this procedure yields the $m \times p$ matrix $\hat{\mathbf{A}}$ (using the notations from (3), $\hat{\mathbf{A}} = \mathbf{K}_I\mathbf{O}_{B,I}$), which is finally pseudo-inverted in order to estimate (as in multiple regression) the dipole amplitudes $\mathbf{s}_i$ (note that in MUSIC-like algorithms, it is possible to select a specific position up to 3 times, in which case the dipoles at those specific positions have time varying orientations).

This solution is very elegant and effective. There are two issues that might diminish its accuracy: the estimation of the initial dimension of the source space r (done usually by truncating the eigenvalues of the covariance matrix of the data / the singular values of the data $\mathbf{X}$) and the fact that, if the initial orientation is not well estimated (because of noise for example), the algorithm is not correcting it during further iterations, but it selects again the same position and estimates a new orientation vector and its corresponding time varying amplitude, yielding thus a source with potentially time varying orientation (as it results, for that specific position, from a sum of 2 or 3 dipoles with potentially different amplitudes in time). In principle though, the (T)RAP-MUSIC algorithms can be seen as providing low-rank approximation of the data $\mathbf{X}$,

[3]We speak here of the correlation between the sub-spaces $\bar{\mathbf{K}}_i$ and $\mathbf{V}_r$ (in other words, the cosine of the principal angle between the two sub-spaces).

where the rank is somewhere between p (if only one orientation is estimated for each selected position i) and $3p$ (if the dipoles are freely rotating and thus each $\bar{\mathbf{J}}_i$ is of rank 3).

3.2 Unconstrained Block-OLS and -SBR Algorithms

One can formulate the problem described above as a regression: express the data $\mathbf{X}$ as a sparse linear combination of lead-field blocks $\bar{\mathbf{K}}_i$ (given regressors) multiplied by estimated regression coefficients, *i.e.*, the matrices $\bar{\mathbf{J}}_i$. We propose here a new family of algorithms belonging to the (step-wise) regression family. They re-estimate, at each iteration, the positions, the orientations and the amplitudes of all the dipoles from I (note that, in the iterative versions of MUSIC, only the orientation of the current dipole is estimated, while the ones already included in I maintain their previously estimated orientations). Using the notations from (2), we can write the optimization problem as:

$$\min\left(\|\hat{I}\|_0\right) \text{ u.c. } \|\mathbf{X}-\sum_{i\in\hat{I}}\bar{\mathbf{K}}_i\hat{\mathbf{J}}_i\|_2 = \|\mathbf{X}-\mathbf{K}_{\hat{I}}\hat{\mathbf{J}}_{\hat{I}}\|_2^2 < \varepsilon \tag{4}$$

Alternatively, we can state the optimization problem as in Soussen et al. (2011), using a penalization term. More precisely, the goal is to estimate a set $\hat{I}$ of source positions and their dipolar moments $\hat{\mathbf{J}}_{\hat{I}}$, while minimizing the reconstruction error and the cardinal of I:

$$\{\hat{I},\hat{\mathbf{J}}_{\hat{I}}\} = \underset{I,\mathbf{J}_I}{\operatorname{argmin}}(\|\mathbf{X}-\mathbf{K}_I\mathbf{J}_I\|_2^2+\lambda\|I\|_0) \tag{5}$$

where λ is a user parameter balancing between sparseness and reconstruction error, while $\|I\|_0$ is the cardinal of I.

We start by describing an OLS type procedure, adapted to a matrix dictionary $\mathbf{K}$ where, unlike in classical multiple regressions (as for example in (8)), every element of the dictionary is a sub-matrix of size $m\times 3$ (the 3-columns "lead-field" matrix for each position)[4]. Using this dictionary, a block version of OLS can be written. It consists in scanning the dictionary for choosing the element (the $m\times 3$ block) that provides, together with the elements already chosen (*i.e.*, the $m\times 3$ blocks selected during previous iterations), the best estimate of the data $\mathbf{X}$. At iteration i, by best estimate we understand the best linear combination of the $3i$ selected columns (i selected blocks). The weights of these columns are the estimated elementary dipolar moments $\hat{\mathbf{J}}$. This algorithm, corresponding to the optimization problem (4), has OLS optimality, the error is obviously monotonically decreasing and thus is convergent (as for any regression, adding more regressors diminishes the error).

In order to control sparsity, a subtraction step inspired from the SBR procedure proposed in Soussen et al. (2011) can be added. The SBR version also implies minimizing the objective function (5). In this second version, the cardinal of the solution $\hat{I}$ is controlled through the sparsity parameter λ, which has to be tuned or at least initialized and let to adapt automatically (CSBR algorithm Soussen et al. (2015)). Clearly, these solutions do not fulfill the fixed orientation condition we imposed: $\hat{\mathbf{J}}_{\hat{I}}$ is of dimension $3p\times n$, with *a priori* linearly independent rows. The low-rank approximation $\hat{\mathbf{X}}$ of $\mathbf{X}$ is potentially of rank $3p$.

3.3 Block-OLS-R1 and -SBR-R1 Algorithms

In order to introduce the mentioned constraints, one can rewrite the optimization problem (4) as:

$$\min\left(\|\hat{I}\|_0\right) \text{ u.c. } \begin{cases}\|\mathbf{X}-\sum_{i\in\hat{I}}\bar{\mathbf{K}}_i\hat{\mathbf{J}}_i\|_2<\varepsilon\\ \operatorname{rank}(\hat{\mathbf{J}}_i)=1\end{cases} \tag{6}$$

while (5) becomes:

$$\{\hat{I},\hat{\mathbf{O}}_{B,\hat{I}},\hat{\mathbf{S}}_{\hat{I}}\} = \underset{I,\mathbf{O}_{B,I},\mathbf{S}_I}{\operatorname{argmin}}(\|\mathbf{X}-\mathbf{K}_I\mathbf{O}_{B,I}\mathbf{S}_I\|_2^2+\lambda\|I\|_0) \tag{7}$$

with the notations used in (3).

Before proceeding with the description of these algorithms, we evacuate the problem of estimating the amplitudes $\hat{\mathbf{S}}_{\hat{I}}$. Indeed, for any given estimated positions $\hat{I}$ and orientations $\hat{\mathbf{O}}_{B,\hat{I}}$ (including those estimated at every iteration), the optimal amplitudes can be obtained trivially by regression/ least-squares (as also in (T)RAP-MUSIC):

$$\hat{\mathbf{S}}_{\hat{I}} = (\mathbf{K}_{\hat{I}}\hat{\mathbf{O}}_{B,\hat{I}})^{+}\mathbf{X} = (\hat{\mathbf{A}}_{\hat{I}})^{+}\mathbf{X} \tag{8}$$

where $^+$ is the pseudo-inverse and $\hat{\mathbf{A}}_{\hat{I}} = \mathbf{K}_{\hat{I}}\hat{\mathbf{O}}_{B,\hat{I}}$ is a collapsed (fixed orientation) dictionary matrix ($m\times\hat{p}$) ($\hat{p}$ the cardinal of the current $\hat{I}$). The difficulty is thus estimating the positions and, mainly, the fixed orientations.

The OLS-R1 (SBR-R1) algorithms will have the same structure as the unconstrained block versions described above, with a supplementary step. At every iteration, the block algorithms above choose the "best" positions and yield estimations of the current time varying orientations and amplitudes embedded in $\hat{\mathbf{J}}_{\hat{I}}$ ($3\hat{p}\times n$), *a priori* of rank $3\hat{p}$ (assuming that n is sufficiently high), with $\hat{p}$ the cardinal of $\hat{I}$. The best current approximation of $\hat{\mathbf{J}}_{\hat{I}}$ that fulfills the fixed orientation condition can be constructed by an SVD decomposition of each of the stacked $\hat{p}$ blocks of size

[4]This can be trivially generalized to dictionary matrices ($m\times q$).

$3 \times n$. By SVD, each block $\hat{\mathbf{J}}_{\hat{i}}$ can be written as a sum of 3 rank 1 matrices

$$\hat{\mathbf{J}}_{\hat{i}} = \sum_{k=1}^{3} \mathbf{u}_{k,\hat{i}} \sigma_{k,\hat{i}} \mathbf{w}_{k,\hat{i}}^T \tag{9}$$

where $\mathbf{u}, \mathbf{w}$ and σ are respectively the left singular vectors, the right singular vectors and the singular values. As the best approximation (in Frobenius norm) of $\hat{\mathbf{J}}_{\hat{i}}$ is given by the first term of this sum, the best approximation of $\hat{\mathbf{J}}_{\hat{I}}$ is given by:

$$\tilde{\mathbf{J}}_{\hat{I}} = \begin{bmatrix} \mathbf{u}_1 & 0 & \dots & 0 \\ 0 & \mathbf{u}_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & \mathbf{u}_{\hat{p}} \end{bmatrix} \begin{bmatrix} \sigma_1 \mathbf{w}_1^T \\ \sigma_2 \mathbf{w}_2^T \\ \vdots \\ \sigma_{\hat{p}} \mathbf{w}_{\hat{p}}^T \end{bmatrix} = \hat{\mathbf{O}}_{B,\hat{I}} \begin{bmatrix} \sigma_1 \mathbf{w}_1^T \\ \sigma_2 \mathbf{w}_2^T \\ \vdots \\ \sigma_{\hat{p}} \mathbf{w}_{\hat{p}}^T \end{bmatrix} \tag{10}$$

where, in order to simplify the notations, $\mathbf{u}_i$ is the first singular vector of the block $\hat{\mathbf{J}}_{\hat{i}}$ (same for $\mathbf{w}_i$ and σ_i). Through the last equality we propose here to estimate, at every iteration, the block diagonal orientation matrix $\hat{\mathbf{O}}_{B,\hat{I}}$. Note that the best estimate of the source amplitudes matrix $\hat{\mathbf{S}}_{\hat{I}}$ is not the last term in (10), but the one computed with (8). Indeed, for any given positions and orientations, the amplitudes obtained by (8) ensure the smallest reconstruction error. In other words, (10) is only used for estimating the orientations.

An important issue appears at this point: while the block OLS or SBR algorithms described in section 3.2 are convergent and monotonic by construction, reducing the dipolar moments matrices $\mathbf{J}$ to their rank 1 approximations by (9), (10) and (8) breaks the monotony. Indeed, once we choose the "best block" at iteration i and we recompute all regression coefficients/orientations $\mathbf{J}$, we modify this solution by using the rank 1 approximations, so we increase the error. In other words, at every iteration, the minimization of (4) or (5) without the constraint has better optimality (lower error) than when the constraint is imposed and we minimize (6) or (5).

Still, the convergence remains trivial: the algorithm stops in at most m iterations (when cardinal of I $p = m$, the error $\|\mathbf{X} - \hat{\mathbf{X}}\|_2^2 = 0$). Still, the solution is not monotonic, thus largely suboptimal.

Monotonicity Indeed, even if the convergence is ensured, there is no guaranteed monotonicity. This can be explained by noticing that the orientation matrix $\hat{\mathbf{O}}_{B,\hat{I}}$) change at every iteration. Consequently, the collapsed dictionary matrix $\hat{\mathbf{A}}_{\hat{I}} = \mathbf{K}_{\hat{I}} \hat{\mathbf{O}}_{B,\hat{I}}$ appearing in (8) changes at every iteration, which might result in an increase of the reconstruction error $\|\mathbf{X} - \hat{\mathbf{X}}\|_2$. As a non monotonic algorithm is certainly not optimal in terms of sparsity (it chooses in the solution regressors that increase the error), a way of improving the optimality is to guarantee that the reconstruction error $\|\mathbf{X} - \hat{\mathbf{X}}\|_2$ monotonically decreases at every iteration. The following algorithmic modification ensures monotonicity.

Consider the solution after iteration k. At iteration $k+1$, instead of recomputing all the orientations $\hat{\mathbf{O}}_{B,\hat{I}}$ by the procedure described above, the orientations computed at iteration k are preserved and the SVD procedure is applied only to compute the optimal orientation of the currently estimated source at position i_{k+1}. This is equivalent to add a new 3×1 block on the diagonal of the matrix $\hat{\mathbf{O}}_{B,\hat{I}}$ (10) obtained at the previous iteration k (increasing its size). Further, this is equivalent to adding a supplementary column to the matrix $\hat{\mathbf{A}}_{\hat{I}}$ while keeping the others unchanged. By construction (adding a new regressor to the dictionary), this will necessarily decrease the reconstruction error, so it ensures monotonicity.

This alternative solution is monotonic but, at a given iteration, it might decrease slower than the solution obtained directly by (10) and (8). A naturally arising solution is the following: at every iteration, compute both solutions (re-estimate all orientations vs. estimate only the last one) and choose the solution minimizing the reconstruction error. In this way, the reconstruction error is upper bounded by a monotonically decreasing sequence, thus it is monotonically decreasing itself. Note that a similar argument is proposed for convergence of the K-SVD algorithm Aharon et al. (2006) (this algorithm, which re-estimates iteratively the regression coefficients but also the regressor dictionary, cannot be applied in the brain inverse problem because the dictionary is constrained by the physical properties of the propagation medium).

We have implemented these different flavors of the step-wise regression including the fixed orientation constraint. The most significant results are presented in the next section.

4 RESULTS

We illustrate here the performances of the proposed methods on realistically simulated data. We used a forward FEM head model based on the MNI-ICBM152 template Fonov et al. (2011), with $P = 509$ source positions restricted to the gray matter and $m = 186$ SEEG sensors ($\mathbf{K} : (m \times 3P)$, see figure 1 and Le Cam et al. (2017) for a detailed description). Source positions and (fixed unknown) orientations were randomly chosen, with the number of sources

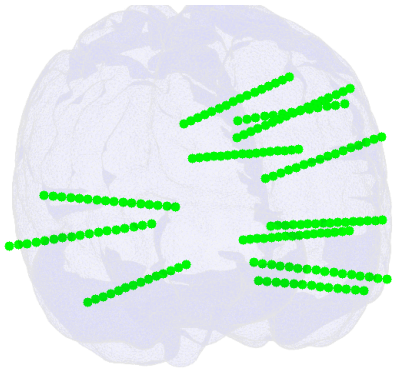

Figure 1: Brain mesh with the simulated SEEG electrodes.

$p = \{3,5,7\}$. The number of time instants was $n = \{10,50,100\}$. White Gaussian noise was added to the simulated clean data $\mathbf{X} : (m \times n)$, with signal to noise ratios $SNR = \{100,20,10,3,0\}$dB. One hundred simulations were executed for each combination $p \times n \times SNR$. The target clinical situation is the localization of sparse and short events (*i.e.*, few data points), as they can be encountered when localizing interictal epileptic spikes. A frequency domain application on highly structured data (*e.g.*, narrow band or periodic sources) is also possible Hernandez-Castañon (2023).

We compare two versions of our approach against four versions of MUSIC: Block-SBR (corresponding to algorithm described in section 3.2, with the subtraction step included but with no fixed orientation constraint) and Block-SBR-R1 (with the rank 1 constraint and the monotonic upper bound), versus RAP Mosher and Leahy (1999) and TRAP-MUSIC Mäkelä et al. (2018), each with the source space dimension estimated by MDL or overestimated as MDL+2 (as recommended *e.g.*, in Mosher and Leahy (1999); Mäkelä et al. (2018)). Three performance criteria were used: PPV (positive predictive value, *i.e.*, the proportion of true positives among the estimated source positions), REC (recall, *i.e.*, the number of true positives with respect to the number of simulated sources) and DLE Becker et al. (2014) (distance localization error, which increases both when spurious sources (false positives) are found and when true sources are missed (false negatives)). Good performances imply PPV and REC close to 1 and DLE close to 0 (or at least below 10mm).

We illustrate the algorithms' performances in figure 2, where we have selected the four extreme cases of the different simulations: from the sparsest setup, with low number of sources and few data points ($p = 3, n = 10$) to the most dense one, with a high numbers of active sources and data points ($p = 7, n = 100$). A first observation is that the constrained version of the SBR (Block-SBR-R1, red bar) is systematically more performant than the unconstrained version (Block-SBR, blue bar), for all criteria. Imposing the supplementary constraint of a fixed orientation (when the sources have indeed a fixed, yet unknown, orientation) improves the localization.

When analyzing MUSIC-type algorithms, one can notice than RAP-MUSIC seems to perform better in the tested situations than its truncated version TRAP-MUSIC. Indeed, comparing yellow vs. magenta bars (overestimated source space dimensions) or green vs. cyan bars (MDL estimated dimensions), indicate that RAP-MUSIC-o and RAP-MUSIC are better. Between these two versions, when evaluating the positive predictive value (PPV - the proportion of true sources among the found ones), the overestimated-MDL RAP-MUSIC (yellow bar) performs worse than strict-MDL based RAP-MUSIC (green bar), while when evaluating the recall (REC - the proportion of found true sources among the actual number of true sources), the two algorithms switch places. In other words, overestimating the dimension of the source space leads to the discovery of more true sources, with the cost of having more false detections.

A synthetic criterion can then be the DLE, which penalizes both non-detections and false-detections. Figure 3 focuses on this criterion and on the 3 best algorithms, namely Block-SBR-R1 (red curve), RAP-MUSIC-o (yellow) and RAP-MUSIC (green). As it can be seen, for small n (few data points), Block-SBR-R1 outperforms the two MUSIC versions for all SNR. The situation is less clear when the n is sufficiently big: for a low number of sources ($p = 3$), Block-SBR-R1 and MDL-based RAP-MUSIC have roughly similar performances, both better than the overestimated RAP-MUSIC-o, while when both n and p are big, the performances of all three algorithms are close and depend on the SNR. Note that a DLE below 10mm remains acceptable and, in our simulations, it indicates that the estimated sources are immediate neighbors of the true ones (the grid of source positions used in our simulations has approximately 10mm spacing between neighboring points).

The proposed constrained block-stepwise regression algorithm (Block-SBR-R1) exhibits significant improvements in localization performance compared to state-of-the-art methods, such as RAP-MUSIC and TRAP-MUSIC. By integrating fixed but unknown orientation constraints and ensuring monotonicity through iterative optimization, our approach achieves higher precision in sparse setups, particularly in scenarios with limited data points and high noise levels. These advancements highlight its potential superior applicability for clinical tasks (short

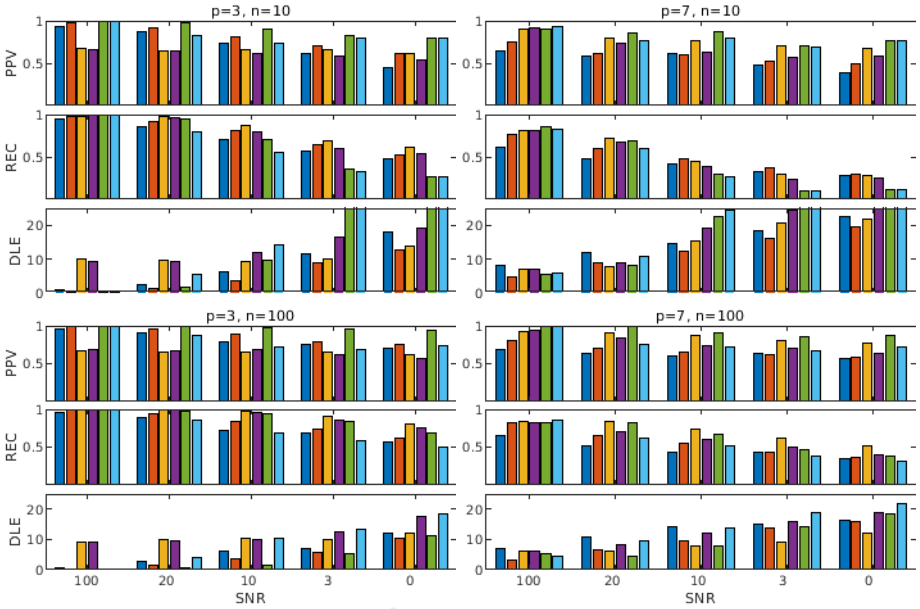

Figure 2: Comparison of localization performance for the six core algorithms: Block-SBR (blue bar), Block-SBR-R1 (red bar), RAP-MUSIC-o (yellow),TRAP-MUSIC-o (magenta), RAP-MUSIC (green), TRAP-MUSIC (cyan). Algorithms marked with '-o' indicate versions using an overestimated source space (see main text). The figure is structured in 4 panels, for different degrees and types of sparsity: up-left, few sources and few number of data points ($p = 3, n = 10$); up-right, more sources and few data points ($p = 7, n = 10$); bottom-left, few sources, more data points ($p = 3, n = 100$); bottom-right, more sources and more data points ($p = 7, n = 100$). For each situation/ panel, we present 3 performance metrics: Positive Predicted Value (PPV), Recall (REC), and Distance Localization Error (DLE). Block-SBR-R1 consistently demonstrates superior performance, particularly in sparse configurations (see text for a more detailed discussion).

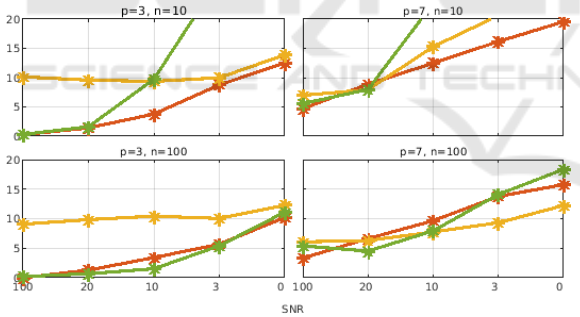

Figure 3: DLE criterion for the 3 best algorithms: Block-SBR-R1 (red curve), RAP-MUSIC-o (yellow) and RAP-MUSIC (green). The four panels illustrate different degrees of sparsity, in terms of number of active sources p and data points n. Values of DLE above 20 mm are discarded, the error is out of reasonably acceptable bounds.

events or narrow band sources localization), as well as for other sparse signal reconstruction problems.

Clearly, the proposed methods need to be validated on real data. Work is in progress on this topic, with the inherent limitations when doing source localization in brain problems: absence of ground truth (although, in some clinical setups, conjectures can be made using other imaging techniques such as fMRI), artifact and noise contamination, particularities of the lead-field matrices when considering surface or intracerebral EEG electrodes, etc. Some preliminary real data results can be found in Hernandez-Castañon (2023).

Despite its promising performance, the constrained block-stepwise regression algorithm is not without limitations. One consideration is the computational cost associated with iterative optimization and the computational demands of applying the rank constraint on regression matrices, which could impose challenges when scaling the method to very large datasets of high-resolution brain models. Future work should include developing strategies to enhance computational efficiency and assess the algorithm's scalability in such scenarios. The actual Matlab implementations are available upon request, but curated scripts and functions will be available to download together with detailed comparisons with other state of the art algorithms.

5 CONCLUSION

We have introduced in this paper a constrained block version of a stepwise regression algorithm aiming to tackle the brain sources localization problem in

sparse situations (few number of sources and few data points). We give convergence and sub-optimality results. The proposed algorithm is evaluated in comparison with classical state of the art methods in realistic simulated situations. More results, including in frequency domain on narrow band data, as well as on real (S)EEG signals will be presented elsewhere.

REFERENCES

Aharon, M., Elad, M., and Bruckstein, A. (2006). K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322.

Baillet, S., Mosher, J., and Leahy, R. (2001). Electromagnetic brain mapping. *IEEE Signal Processing Magazine*, 18(6):14–30.

Becker, H., Albera, L., Comon, P., Gribonval, R., Wendling, F., and Merlet, I. (2014). A performance study of various brain source imaging approaches. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5869–5873.

Dale, A. M. and Sereno, M. I. (1993). Improved Localizadon of Cortical Activity by Combining EEG and MEG with MRI Cortical Surface Reconstruction: A Linear Approach. *Journal of Cognitive Neuroscience*, 5(2):162–176.

Fonov, V., Evans, A. C., Botteron, K., Almli, C. R., McKinstry, R. C., and Collins, D. L. (2011). Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage*, 54(1):313 – 327.

Gramfort, A., Strohmeier, D., Haueisen, J., Hämäläinen, M., and Kowalski, M. (2013). Time-frequency mixed-norm estimates: Sparse M/EEG imaging with non-stationary source activations. *NeuroImage*, 70:410–422.

Hämäläinen, M. S. and Ilmoniemi, R. J. (1994). Interpreting magnetic fields of the brain: minimum norm estimates. *Medical & Biological Engineering & Computing*, 32(1):35–42.

He, B., Sohrabpour, A., Brown, E., and Liu, Z. (2018). Electrophysiological source imaging: A noninvasive window to brain dynamics. *Annual Review of Biomedical Engineering*, 20(1):171–196.

He, B., Yang, L., Wilke, C., and Yuan, H. (2011). Electrophysiological imaging of brain activity and connectivity - challenges and opportunities. *IEEE Transactions on Biomedical Engineering*, 58(7):1918–1931.

Hernandez-Castañon, V. (2023). *Localization of brain oscillatory sources from (S)EEG recordings*. PhD thesis, Université de Lorraine.

Korats, G., Cam, S. L., Ranta, R., and Louis-Dorr, V. (2016). A space-time-frequency dictionary for sparse cortical source localization. *IEEE Transactions on Biomedical Engineering*, 63(9):1966–1973.

Le Cam, S., Ranta, R., Caune, V., Korats, G., Koessler, L., Maillard, L., and Louis-Dorr, V. (2017). SEEG dipole source localization based on an empirical bayesian approach taking into account forward model uncertainties. *NeuroImage*, 153(Supplement C):1 – 15.

Logothethis, N., Kayser, C., and Oeltermann, A. (2007). In vivo measurement of cortical impedance spectrum in monkeys: Implications for signal propagation. *Neuron*, 55:809–823.

Mäkelä, N., Stenroos, M., Sarvas, J., and Ilmoniemi, R. J. (2018). Truncated RAP-MUSIC (TRAP-MUSIC) for MEG and EEG source localization. *NeuroImage*, 167:73–83.

Michel, C., Murray, M., Lantz, G., Gonzalez, S., Spinelli, L., and Grave de Peralta, R. (2004). EEG source imaging. *Clinical neurophysiology*, 115(10):2195–2222.

Mosher, J. C. and Leahy, R. M. (1999). Source localization using recursively applied and projected (RAP) MUSIC. *IEEE Transactions on Signal Processing*, 47(2):332–340.

Pascual-Marqui, R. D. (2007). Discrete, 3D distributed, linear imaging methods of electric neuronal activity. Part 1: exact, zero error localization. *arXiv:0710.3341 [math-ph], 2007-October-17, http://arxiv.org/pdf/0710.3341*, pages 1–16.

Ranta, R., Le Cam, S., Tyvaert, L., and Louis-Dorr, V. (2017). Assesing human brain impedance using simultaneous surface and intracerebral recordings. *Neuroscience*, 343:411–422.

Soussen, C., Idier, J., Brie, D., and Duan, J. (2011). From Bernoulli-Gaussian deconvolution to sparse signal restoration. *Signal Processing, IEEE Transactions on*, 59(10):4572–4584.

Soussen, C., Idier, J., Duan, J., and Brie, D. (2015). Homotopy based algorithms for ℓ_0-regularized least-squares. *IEEE Transactions on Signal Processing*, 63(13):3301–3316.