



HAL
open science

From Bernoulli-Gaussian deconvolution to sparse signal restoration

Charles Soussen, Jérôme Idier, David Brie, Junbo Duan

► **To cite this version:**

Charles Soussen, Jérôme Idier, David Brie, Junbo Duan. From Bernoulli-Gaussian deconvolution to sparse signal restoration. IEEE Transactions on Signal Processing, 2011, Accepted Article. hal-00443842v3

HAL Id: hal-00443842

<https://hal.science/hal-00443842v3>

Submitted on 3 Feb 2010 (v3), last revised 17 Jun 2011 (v4)

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

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

From Bernoulli-Gaussian deconvolution to sparse signal restoration

Charles Soussen*, Jérôme Idier, *Member, IEEE*, David Brie, and Junbo Duan

Abstract

Formulated as a least-square problem under an ℓ_0 constraint, sparse signal restoration is a discrete optimization problem, known to be NP complete. Classical algorithms include, by increasing cost and efficiency, Matching Pursuit (MP), Orthogonal Matching Pursuit (OMP), Orthogonal Least Squares (OLS) and the exhaustive search algorithm. In inverse problems involving highly correlated dictionaries, OMP and OLS are not guaranteed to find the optimal solution. It is of interest to develop slightly slower sub-optimal search algorithms yielding better approximations. We revisit the Single Most Likely Replacement (SMLR) algorithm, developed in the mid-80's for Bernoulli-Gaussian signal restoration. We show that the formulation of sparse signal restoration as a limit case of Bernoulli-Gaussian signal restoration leads to an ℓ_0 -penalized least-square minimization problem, to which SMLR can be straightforwardly adapted. The resulting algorithm, called Single Best Replacement (SBR), can be interpreted as a forward-backward extension of OLS. A fast and stable implementation is proposed. The approach is illustrated on a deconvolution problem with a Gaussian impulse response and on the joint detection of discontinuities at different orders in a signal.

Index Terms

Sparse signal estimation; inverse problems; Bernoulli-Gaussian signal restoration; SMLR algorithm; mixed ℓ_2 - ℓ_0 criterion minimization; Orthogonal Least Squares; forward-backward greedy algorithms.

C. Soussen, D. Brie, and J. Duan are with the Centre de Recherche en Automatique de Nancy (CRAN, UMR 7039, Nancy-University, CNRS). Boulevard des Aiguillettes, B.P. 70239, F-54506 Vandœuvre-lès-Nancy, France. Tel: (+33)-3 83 68 44 71, Fax: (+33)-3 83 68 44 62. E-mail: {FirstName.SecondName}@cran.uhp-nancy.fr.

J. Idier is with the Institut de Recherche en Communications et Cybernétique de Nantes (IRCCyN, UMR CNRS 6597), BP 92101, 1 rue de la Noë, 44321 Nantes Cedex 3, France. Tel: (+33)-2 40 37 69 09, Fax: (+33)-2 40 37 69 30. E-mail: Jerome.Idier@irccyn.ec-nantes.fr.

I. INTRODUCTION

Sparse signal restoration arises in inverse problems such as Fourier synthesis, mono- and multidimensional deconvolution, and statistical regression. It consists in the decomposition of a given signal \mathbf{y} as a linear combination of a limited number of elements from a dictionary \mathbf{A} . While formally very similar, sparse signal restoration has to be distinguished from sparse signal approximation. The main difference is that in sparse signal restoration, the choice of the dictionary is imposed by the inverse problem at hand whereas in sparse signal approximation, the dictionary has to be chosen according to its ability to represent the data with a limited number of coefficients. A more subtle difference is that in sparse signal restoration, the focus is set on the estimation of the weights of the linear combination while in sparse signal approximation, the goal is to reproduce the data \mathbf{y} as well as possible at a given level of sparsity.

Sparse signal restoration can be formulated as the minimization of a least-square cost function of the form $\mathcal{E}(\mathbf{x}) = \|\mathbf{y} - \mathbf{Ax}\|^2$ under the constraint that the ℓ_0 pseudo-norm of \mathbf{x} , defined as the number of non-zero entries in \mathbf{x} , is lower than a given number k . This problem is often referred to as *subset selection*, because imposing the sparsity constraint consists in selecting a subset of columns of \mathbf{A} . This yields a discrete problem (since there are a finite number of possible subsets) which is known to be NP-complete [1]. In this paper, we focus on “difficult” problems in which some of the columns of \mathbf{A} are highly correlated, the unknown weight vector \mathbf{x}^* is only approximately sparse, and/or the data are noisy. Hereafter, we distinguish two approaches to address the subset selection problem in a fast and sub-optimal manner and we discuss their relevance for difficult problems.

The first approach, which has been the most popular in the last decade, approximates the subset selection problem by a continuous optimization problem, convex or not, that is easier to solve [2–5]. In particular, the approach that replaces the ℓ_0 -norm by the ℓ_1 -norm [4, 5] has been increasingly investigated, leading to the LASSO optimization problem. Its popularity relies on efficient algorithms, such as LARS which finds the set of solutions for all degrees of sparsity [6]. Several authors have provided sufficient conditions under which the ℓ_0 - and ℓ_1 -constrained least-square problems lead to solutions having the same support [5, 7]. These conditions typically state that the unknown weight signal has to be highly sparse, that the correlation between any pair of columns of \mathbf{A} must be sufficiently small, and that the noise level must be low. They are often not satisfied when dealing with real data.

The second approach addresses the *exact* subset selection problem using either thresholding algorithms [8, 9] or greedy search algorithms. The latter gradually increase or decrease by one the set of active columns. The simplest greedy algorithms are Matching Pursuit (MP) [10] and the improved version

Orthogonal Matching Pursuit (OMP) [11]. Both are referred to as forward greedy algorithms, since they start from an empty active set and then gradually increase it by one element. In contrast, the backward algorithm of Couvreur and Bresler [12] starts from a complete active set which is gradually decreased by one element. It is, however, only valid if the dictionary is not overcomplete. A few authors have introduced forward-backward algorithms in which insertions and removals of dictionary elements into the active set are both allowed [13, 14]. This strategy yields better recovery performance since an early wrong selection can be counteracted by its further removal from the active set. In contrast, the insertion of a wrong element is irreversible when using forward algorithms.

The choice of the algorithm depends on the amount of time available and on the structure of matrix \mathbf{A} . In favorable cases, the sub-optimal search algorithms described above (belonging to the first or the second approach) provide solutions having the same support as the exhaustive search solution. For example, if the unknown signal \mathbf{x}^* is highly sparse and if the correlation between any pair of columns of \mathbf{A} is low, the ℓ_1 -norm approximation provides optimal solutions [5, 7]. In other cases, however, the only guarantee to recover the optimal solution is to use the exhaustive search algorithm. When fast sub-optimal algorithms lead to unsatisfactory results, it is of great interest to develop slower sub-optimal algorithms providing more accurate solutions, but remaining very fast compared to the exhaustive search. The Orthogonal Least Squares algorithm (OLS) [15] which is sometimes confused with OMP [16], falls into this category. Both OLS and OMP share the same structure, the difference being that at each iteration, OLS solves a large number of least-square problems ($n - k$, where k is the cardinal of the current active set) while OMP only computes the $n - k$ inner products between the current residual $\mathbf{y} - \mathbf{A}\mathbf{x}$ and the candidate columns \mathbf{a}_i and chooses the column yielding the maximal inner product. OMP solves only one least-square problem per iteration, once the column to be inserted is selected (in order to update all the active weights). In the following, we propose a forward-backward extension of OLS allowing an insertion or a removal at each iteration, each iteration requiring to solve n least-square problems. It differs from the FoBa algorithm of Zhang [14] which is an OMP forward-backward extension. It is closer to the bidirectional OLS based algorithm of Haugland [13], the main differences relying on the search and implementation strategies.

The starting point of our forward-backward algorithm is the Single Most Likely Replacement (SMLR) algorithm, which proved to be a very efficient tool for the deconvolution of a sparse signal modeled as a Bernoulli-Gaussian process [17–20]. This approach relies on a Bayesian formulation of a deconvolution problem of the form $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$ (where \mathbf{A} denotes the convolution matrix) and on the maximum *a posteriori* (MAP) estimation of the sparse signal. The Bernoulli-Gaussian model is a probabilistic model for sparse signals, in which (binary) Bernoulli random variables are associated to the position of the

non zero entries in \mathbf{x} while the corresponding amplitudes are distributed according to an independent identically distributed (i.i.d.) centered Gaussian distribution of variance σ_x^2 . SMLR is a deterministic ascent algorithm which maximizes the posterior likelihood in a sub-optimal manner. It consists in updates (increase or decrease) of the support of \mathbf{x} by one element and the subsequent estimation of the non-zero amplitudes. Sparse signal restoration can be seen as a limit case of Bernoulli-Gaussian MAP restoration in which the variance σ_x^2 of the amplitudes is set to infinity. We will deduce an adaptation of SMLR to subset selection relying on a single insertion or a single removal of a column into/from the active set.

The paper is organized as follows. In Section II, we introduce the Bernoulli-Gaussian model and the Bayesian framework from which we formulate the sparse signal restoration problem. In Section III, we adapt the SMLR algorithm resulting in the so-called Single Best Replacement (SBR) algorithm. In Section IV, a fast and stable SBR implementation is proposed, based on the efficient update of the squared error when the active set is modified by one element. Finally, Sections V and VI illustrate the method on the sparse spike deconvolution with a Gaussian impulse response and on the joint detection of discontinuities at different orders in a signal.

II. FROM BERNOULLI-GAUSSIAN SIGNAL MODELING TO SPARSE SIGNAL REPRESENTATION

We consider the restoration of a sparse signal \mathbf{x} from a linear observation $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$, where \mathbf{n} stands for the observation noise. An acknowledged probabilistic model dedicated to sparse signals is the Bernoulli-Gaussian (BG) model [17, 18, 20]. For such model, deterministic optimization algorithms [20] and Markov chain Monte Carlo techniques [21] are used to compute the MAP and the posterior mean, respectively. We first recall the known BG models and the formulation of BG signal restoration in the Bayesian framework. Then, we extend this formulation to a more general representation of sparse signals.

A. Preliminary definitions and working assumptions

Given an observation vector $\mathbf{y} \in \mathbb{R}^m$ and a dictionary $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n] \in \mathbb{R}^{m \times n}$, a subset selection algorithm aims at computing a weight vector $\mathbf{x} \in \mathbb{R}^n$ yielding an accurate approximation $\mathbf{y} \approx \mathbf{A}\mathbf{x}$ of the observation. The columns \mathbf{a}_i of \mathbf{A} whose indices correspond to the non-zero components x_i of \mathbf{x} are referred to as the active (or selected) columns.

Throughout this paper, no assumption is made on the size of \mathbf{A} : m can be either smaller or larger than n . Here, \mathbf{A} is assumed to satisfy the unique representation property (URP). This assumption is usual in the case $m \leq n$ [22]. It is a stronger assumption than the full rank assumption. We now recall this definition and extend it to the case where $m > n$. Notation $\|\cdot\|$ refers to the Euclidean norm.

Definition 1 When $m \leq n$, \mathbf{A} satisfies the URP if and only if any selection of m columns of \mathbf{A} forms a family of linearly independent vectors. When $m > n$, \mathbf{A} satisfies the URP if and only if it is full rank.

Before going further, let us mention that this assumption can be relaxed providing that the search strategy can guarantee that the selected columns of \mathbf{A} result in a full rank matrix (see Section VI-C for details).

Under the URP assumption, when $m \leq n$, the system $\mathbf{y} = \mathbf{A}\mathbf{x}$ has a number of solutions whose ℓ_0 -norm are lower or equal to m since any selection of m columns yields a solution of the system. When $m > n$, there is generally no solution to $\mathbf{y} = \mathbf{A}\mathbf{x}$ but the least-square estimator $\mathbf{x} = (\mathbf{A}^t \mathbf{A})^{-1} \mathbf{A}^t \mathbf{y}$ is unique, although not necessarily sparse.

Definition 2 The support of a vector $\mathbf{x} \in \mathbb{R}^n$ is the set $\mathcal{S}(\mathbf{x}) \subseteq \{1, \dots, n\}$ defined by $i \in \mathcal{S}(\mathbf{x})$ if and only if $x_i \neq 0$.

Definition 3 We denote by $\mathcal{Q} \subseteq \{1, \dots, n\}$ the active set and we define the related vector $\mathbf{q} \in \{0, 1\}^n$ by $q_i = 1$ if and only if $i \in \mathcal{Q}$. Let $\mathbf{A}_{\mathcal{Q}}$ be the submatrix of size $m \times \text{Card}[\mathcal{Q}]$ formed of the active columns of \mathbf{A} ($\mathbf{a}_i, i \in \mathcal{Q}$). The observation model $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$ also reads $\mathbf{y} = \mathbf{A}_{\mathcal{Q}}\mathbf{t} + \mathbf{n}$, where the reduced vector \mathbf{t} of size $\text{Card}[\mathcal{Q}]$ gathers the values $\{x_i, i \in \mathcal{Q}\}$.

Definition 4 For all $\mathcal{Q} \subseteq \{1, \dots, n\}$ such that $\text{Card}[\mathcal{Q}] \leq \min(m, n)$, let $\mathbf{x}_{\mathcal{Q}}$ be the least-square solution and let $\mathcal{E}_{\mathcal{Q}}$ be the associated squared error:

$$\mathbf{x}_{\mathcal{Q}} \triangleq \arg \min_{\mathcal{S}(\mathbf{x}) \subseteq \mathcal{Q}} \{\mathcal{E}(\mathbf{x}) = \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2\} \quad (1)$$

$$\mathcal{E}_{\mathcal{Q}} \triangleq \mathcal{E}(\mathbf{x}_{\mathcal{Q}}) = \|\mathbf{y} - \mathbf{A}\mathbf{x}_{\mathcal{Q}}\|^2. \quad (2)$$

B. Bernoulli-Gaussian models

A BG process¹ \mathbf{x} can be defined by means of a Bernoulli random vector $\mathbf{q} \in \{0, 1\}^n$ corresponding to the active set, and a Gaussian random vector $\mathbf{r} \sim \mathcal{N}(\mathbf{0}, \sigma_x^2 \mathbf{I}_n)$, where \mathbf{I}_n stands for the identity matrix of size $n \times n$. Each sample x_i of \mathbf{x} is modeled as $x_i = q_i r_i$ [17, 18]. Thus, r_i code for the amplitudes of the nonzero entries in \mathbf{x} . In the vector form, \mathbf{x} reads $\mathbf{x} = \mathbf{\Delta}_{\mathbf{q}} \mathbf{r}$ where $\mathbf{\Delta}_{\mathbf{q}}$ is the diagonal matrix of size $n \times n$ whose diagonal elements are equal to q_i . The Bernoulli random variables q_i code for the presence ($q_i = 1$) or absence ($q_i = 0$) of signal at location i , the Bernoulli parameter $\rho = \Pr(q_i = 1)$ being the probability

¹For convenience, we use the same notations for random vectors and their realization.

of presence of signal. It is easy to check that the prior likelihood of \mathbf{q} reads $l(\mathbf{q}) = \rho^{\|\mathbf{q}\|_0} (1 - \rho)^{n - \|\mathbf{q}\|_0}$. Because \mathbf{q} and \mathbf{r} are independent random vectors, the prior likelihood of $\mathbf{x} = (\mathbf{q}, \mathbf{r})$ reads:

$$l(\mathbf{q}, \mathbf{r}) = l(\mathbf{r})l(\mathbf{q}) = g(\mathbf{r}; \sigma_x^2 \mathbf{I}_n) \rho^{\|\mathbf{q}\|_0} (1 - \rho)^{n - \|\mathbf{q}\|_0}, \quad (3)$$

where $g(\cdot; \mathbf{\Gamma})$ denotes the probability density of the centered Gaussian with covariance matrix $\mathbf{\Gamma}$.

C. Bayesian formulation of sparse signal restoration

The Bayesian formulation of the inverse problem $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$ consists in inferring the distribution of $\mathbf{x} = (\mathbf{q}, \mathbf{r})$ knowing \mathbf{y} . The MAP estimator of \mathbf{x} can be obtained by maximizing the marginal distribution $l(\mathbf{q} | \mathbf{y})$ [20] or the joint distribution $l(\mathbf{q}, \mathbf{r} | \mathbf{y})$ [18, 19]. Following [18], we focus on the joint likelihood $l(\mathbf{q}, \mathbf{r} | \mathbf{y})$ which leads to a cost function involving the squared error $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ and the ℓ_0 -norm of \mathbf{x} .

Assuming a Gaussian white noise $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I}_m)$, independent from \mathbf{x} , it is easy to obtain

$$\begin{aligned} \mathcal{L}(\mathbf{q}, \mathbf{r}) &\triangleq -2\sigma_n^2 \log[l(\mathbf{q}, \mathbf{r} | \mathbf{y})] \\ &= \|\mathbf{y} - \mathbf{A}\Delta_{\mathbf{q}}\mathbf{r}\|^2 + \frac{\sigma_n^2}{\sigma_x^2} \|\mathbf{r}\|^2 + \lambda \|\mathbf{q}\|_0 + C, \end{aligned} \quad (4)$$

where $\lambda = 2\sigma_n^2 \log(1/\rho - 1)$ and C is a constant. Given \mathbf{q} , let us split \mathbf{r} into two subvectors \mathbf{u} and \mathbf{t} indexed by the null and non-null entries of \mathbf{q} , respectively. Since $\mathbf{A}\Delta_{\mathbf{q}}\mathbf{r} = \mathbf{A}_{\mathcal{Q}}\mathbf{t}$ does not depend on \mathbf{u} , it is easy to check that $\min_{\mathbf{u}} \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{u}) = \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$. Finally, the joint MAP estimation problem reduces to the minimization of $\mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$ w.r.t. (\mathbf{q}, \mathbf{t}) .

D. Mixed ℓ_2 - ℓ_0 minimization as a limit case

A signal \mathbf{x} is sparse if some entries x_i are equal to 0. Since this definition does not impose constraints on the range of values of the non zero amplitudes, we choose to describe a sparse signal by a limit Bernoulli-Gaussian model in which the amplitude variance σ_x^2 is set to infinity. The minimization of $\mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$ thus rereads:

$$\min_{\mathbf{q}, \mathbf{t}} \{\mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0}) = \|\mathbf{y} - \mathbf{A}_{\mathcal{Q}}\mathbf{t}\|^2 + \lambda \|\mathbf{q}\|_0\}. \quad (5)$$

Theorem 1 *The above formulation (5) is equivalent to:*

$$\min_{\mathbf{x} \in \mathbb{R}^n} \{\mathcal{J}(\mathbf{x}; \lambda) = \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \lambda \|\mathbf{x}\|_0\} \quad (6)$$

which is referred to as the ℓ_0 -penalized least-square problem. The term “equivalent” means that given a minimizer (\mathbf{q}, \mathbf{t}) of (5), the related vector $\mathbf{x} = \{\mathbf{t}, \mathbf{0}\}$ is a minimizer of (6), and conversely, given

a minimizer \mathbf{x} of (6), the vectors \mathbf{q} and \mathbf{t} defined as the support of \mathbf{x} and its non-zero amplitudes, respectively, are such that (\mathbf{q}, \mathbf{t}) is a minimizer of (5).

Proof: To prove the equivalence, we first prove that $\min_{\mathbf{x}} \mathcal{J}(\mathbf{x}; \lambda) = \min_{\mathbf{q}, \mathbf{t}} \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$:

- Let \mathbf{x} be a minimizer of $\mathcal{J}(\cdot; \lambda)$. We set \mathbf{q} and \mathbf{t} to the support and the non zero amplitudes of \mathbf{x} , respectively. Obviously, $\mathcal{J}(\mathbf{x}; \lambda) = \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$. It follows that $\min_{\mathbf{x}} \mathcal{J}(\mathbf{x}; \lambda) \geq \min_{\mathbf{q}, \mathbf{t}} \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$.
- Let (\mathbf{q}, \mathbf{t}) be a minimizer of $\mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$. The vector $\mathbf{x} = \{\mathbf{t}, \mathbf{0}\}$ is such that $\mathbf{A}\mathbf{x} = \mathbf{A}_{\mathcal{Q}}\mathbf{t}$ and $\|\mathbf{x}\|_0 = \|\mathbf{t}\|_0 \leq \|\mathbf{q}\|_0$. Therefore, $\mathcal{J}(\mathbf{x}; \lambda) \leq \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$. It follows that $\min_{\mathbf{q}, \mathbf{t}} \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0}) \geq \min_{\mathbf{x}} \mathcal{J}(\mathbf{x}; \lambda)$.

We have shown that $\min_{\mathbf{x}} \mathcal{J}(\mathbf{x}; \lambda) = \min_{\mathbf{q}, \mathbf{t}} \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$, but also that the minimizers of both problems coincide, *i.e.*, are vectors describing identical signals. ■

In the following, we focus on the minimization problem (6) involving the penalization term $\|\mathbf{x}\|_0$. The hyperparameter λ is fixed. It controls the level of sparsity of the desired solution. The algorithm that will be developed is based on an efficient search of the support of \mathbf{x} . In that respect, the ℓ_0 -penalized least-square problem does not drastically differ from the ℓ_0 -constrained problem $\min \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ subject to $\|\mathbf{x}\|_0 \leq k$.

III. ADAPTATION OF SMLR TO ℓ_0 -PENALIZED LEAST-SQUARE OPTIMIZATION

We propose to adapt the SMLR algorithm to the minimization of the mixed ℓ_2 - ℓ_0 cost function $\mathcal{J}(\mathbf{x}; \lambda)$ defined in (6). To clearly distinguish SMLR which specifically aims at minimizing (4), the adapted algorithm will be termed as Single Best Replacement (SBR).

A. Principle of SMLR and main notations

The Single Most Likely Replacement algorithm [17] is a deterministic coordinatewise ascent algorithm to maximize log-likelihood functions of the form $l(\mathbf{q} | \mathbf{y})$ (marginal MAP estimation) or $l(\mathbf{q}, \mathbf{t}, \mathbf{0} | \mathbf{y})$ (joint MAP estimation). In the latter case, it is easy to check from (4) that given \mathbf{q} , the maximizer of $l(\mathbf{q}, \mathbf{t}, \mathbf{0} | \mathbf{y})$ w.r.t. \mathbf{t} has a closed form expression $\mathbf{t} = \mathbf{t}(\mathbf{q})$. Consequently, the joint MAP estimation reduces to the maximization of $l(\mathbf{q}, \mathbf{t}(\mathbf{q}), \mathbf{0} | \mathbf{y})$ w.r.t. \mathbf{q} . At each SMLR iteration, all the possible single replacements of the support \mathbf{q} (set $q_i = 1 - q_i$ while keeping the other $q_j, j \neq i$ unchanged) are tested, then the replacement yielding the maximal increase of $l(\mathbf{q}, \mathbf{t}(\mathbf{q}), \mathbf{0} | \mathbf{y})$ is chosen. This task is repeated until no single replacement can increase $l(\mathbf{q}, \mathbf{t}(\mathbf{q}), \mathbf{0} | \mathbf{y})$ anymore. The number of possible supports \mathbf{q} being finite (2^n) and SMLR being an ascent algorithm, it terminates after a finite number of iterations.

Before adapting SMLR, let us introduce some useful notations.

- We denote by $\mathcal{Q} \bullet i$ a single replacement, *i.e.*, the insertion or removal of an index i into/from the active set \mathcal{Q} :

$$\mathcal{Q} \bullet i \triangleq \begin{cases} \mathcal{Q} \cup \{i\} & \text{if } i \notin \mathcal{Q}, \\ \mathcal{Q} \setminus \{i\} & \text{otherwise.} \end{cases} \quad (7)$$

- If $\text{Card}[\mathcal{Q}] \leq \min(m, n)$, we define the cost functions:

$$\mathcal{J}_{\mathcal{Q}}(\lambda) \triangleq \mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda) = \mathcal{E}_{\mathcal{Q}} + \lambda \|\mathbf{x}_{\mathcal{Q}}\|_0 \quad (8)$$

$$\mathcal{K}_{\mathcal{Q}}(\lambda) \triangleq \mathcal{E}_{\mathcal{Q}} + \lambda \text{Card}[\mathcal{Q}] \quad (9)$$

where the least-square solution $\mathbf{x}_{\mathcal{Q}}$ and the corresponding error $\mathcal{E}_{\mathcal{Q}}$ have been defined in (1) and (2). Obviously, $\mathcal{J}_{\mathcal{Q}}(\lambda) = \mathcal{K}_{\mathcal{Q}}(\lambda)$ if and only if $\mathbf{x}_{\mathcal{Q}}$ has a support equal to \mathcal{Q} . In subsection III-B, we introduce a first version of SBR involving $\mathcal{J}_{\mathcal{Q}}(\lambda)$ only, and in subsection III-C, we present an alternative (simpler) version relying on the computation of $\mathcal{K}_{\mathcal{Q}}(\lambda)$ instead of $\mathcal{J}_{\mathcal{Q}}(\lambda)$ and we discuss in which extent both versions differ. Then, subsection III-D describes the behavior of SBR and states its main properties.

B. The Single Best Replacement algorithm (first version)

SMLR can be seen as an exploration strategy for discrete optimization rather than an algorithm specific to a posterior likelihood function. Here, we use the same strategy to minimize the cost function $\mathcal{J}(\mathbf{x}; \lambda)$. We rename the algorithm Single Best Replacement to remove any statistical connotation. The SBR algorithm works as follows. At each iteration, the n possible single replacements $\mathcal{Q} \bullet i$ ($i = 1, \dots, n$) are tested, then the best is selected, *i.e.*, the replacement yielding the maximal decrease of $\mathcal{J}(\mathbf{x}; \lambda)$. This task is repeated until $\mathcal{J}(\mathbf{x}; \lambda)$ cannot decrease anymore. Let us detail an SBR iteration.

Consider the current active set \mathcal{Q} . For each index i , we compute the minimizer $\mathbf{x}_{\mathcal{Q} \bullet i}$ of \mathcal{E} whose support is included in $\mathcal{Q} \bullet i$ and we keep in memory the value of $\mathcal{J}_{\mathcal{Q} \bullet i}(\lambda)$. If the minimum of $\{\mathcal{J}_{\mathcal{Q} \bullet i}(\lambda), i = 1, \dots, n\}$ is lower than $\mathcal{J}_{\mathcal{Q}}(\lambda)$, then we select the index yielding this minimal value:

$$\ell \in \arg \min_{i \in \{1, \dots, n\}} \mathcal{J}_{\mathcal{Q} \bullet i}(\lambda). \quad (10)$$

The next SBR iterate is thus defined as $\mathcal{Q}' = \mathcal{Q} \bullet \ell$, yielding the vector $\mathbf{x}_{\mathcal{Q}'}$.

Except when an initial support estimate (of cardinality lower than $\min(m, n)$) is available, we suggest to use an initial empty active set.

Remark 1 (Relationship between SBR and SMLR) *We introduced SBR as the application of the SMLR search strategy to the ℓ_0 -penalized least square cost function (6) which is obtained by taking the limit of*

the cost function (4) when σ_x tends towards infinity. Conversely, applying SMLR to (4) and then, taking the limit of the SMLR formula when σ_x tends to infinity also yields the SBR algorithm.

Actually, the main difference between SMLR and SBR is that SMLR (which can take several forms depending on the use of the joint distribution $l(\mathbf{q}, \mathbf{r} | \mathbf{y})$ or the marginal distribution $l(\mathbf{q} | \mathbf{y})$) involves the inversion of a matrix of the form $\mathbf{A}_{\mathcal{Q}}^t \mathbf{A}_{\mathcal{Q}} + \alpha \mathbf{I}_{\text{Card}[\mathcal{Q}]}$ whereas SBR involves the inverse of the Gram matrix $\mathbf{A}_{\mathcal{Q}}^t \mathbf{A}_{\mathcal{Q}}$. For this reason, instabilities may occur while using SBR when $\mathbf{A}_{\mathcal{Q}}$ is ill conditioned. The use of the term $\alpha \mathbf{I}_{\text{Card}[\mathcal{Q}]}$, which acts as a regularization on the amplitude values, avoids such instability while using SMLR at the price of handling the additional hyperparameter α .

C. Modified version of SBR (final version)

We introduce a slight modification of SBR by replacing (10) with:

$$\ell \in \arg \min_{i \in \{1, \dots, n\}} \mathcal{K}_{\mathcal{Q}, i}(\lambda). \quad (11)$$

We propose this modification because $\mathcal{K}_{\mathcal{Q}}(\lambda) = \mathcal{E}_{\mathcal{Q}} + \lambda \text{Card}[\mathcal{Q}]$ can be computed more efficiently than $\mathcal{J}_{\mathcal{Q}}(\lambda)$, the computation of $\mathbf{x}_{\mathcal{Q}}$ being no longer necessary. The use of $\mathcal{K}_{\mathcal{Q}}(\lambda)$ makes the penalization term very easy to update when \mathcal{Q} is modified by one element (add or remove λ), and the only necessary update is that of $\mathcal{E}_{\mathcal{Q}}$. We now show that there is almost surely no difference between both versions of SBR provided that the data \mathbf{y} are corrupted with “non degenerate” noise.

Theorem 2 *Let $\mathbf{y} = \mathbf{y}_0 + \mathbf{n}$, where \mathbf{y}_0 is a given vector of \mathbb{R}^m and \mathbf{n} is a random vector. We assume that \mathbf{n} is an absolute continuous random vector, i.e., admitting a probability density w.r.t. the Lebesgue measure. Then, when $\text{Card}[\mathcal{Q}] \leq \min(m, n)$, the probability that $\|\mathbf{x}_{\mathcal{Q}}\|_0 < \text{Card}[\mathcal{Q}]$ is equal to 0, i.e., $\|\mathbf{x}_{\mathcal{Q}}\|_0 = \text{Card}[\mathcal{Q}]$ almost surely.*

Proof: Let $k = \text{Card}[\mathcal{Q}]$ and let $\mathbf{t}_{\mathcal{Q}}$ be the minimizer of $\|\mathbf{y} - \mathbf{A}_{\mathcal{Q}} \mathbf{t}\|^2$ over \mathbb{R}^k . Obviously, $\|\mathbf{x}_{\mathcal{Q}}\|_0 = \|\mathbf{t}_{\mathcal{Q}}\|_0 \leq k$. Let $\mathbf{V}_{\mathcal{Q}} = (\mathbf{A}_{\mathcal{Q}}^t \mathbf{A}_{\mathcal{Q}})^{-1} \mathbf{A}_{\mathcal{Q}}^t$ be the matrix of size $k \times m$ such that $\mathbf{t}_{\mathcal{Q}} = \mathbf{V}_{\mathcal{Q}} \mathbf{y}$. Denoting by $\mathbf{v}^1, \dots, \mathbf{v}^k \in \mathbb{R}^m$ the row vectors of $\mathbf{V}_{\mathcal{Q}}$, $\|\mathbf{t}_{\mathcal{Q}}\|_0 < k$ if and only if there exists i such that $\langle \mathbf{y}, \mathbf{v}^i \rangle = 0$ (where $\langle \cdot, \cdot \rangle$ denotes the inner product). Because $\mathbf{A}_{\mathcal{Q}}$ is full rank, $\mathbf{V}_{\mathcal{Q}}$ is full rank and then $\forall i, \mathbf{v}^i \neq \mathbf{0}$. Denoting by $\mathcal{H}^{\perp}(\mathbf{v}^i)$ the hyperplane of \mathbb{R}^m which is orthogonal to \mathbf{v}^i , we have

$$\|\mathbf{x}_{\mathcal{Q}}\|_0 < k \iff \mathbf{y} \in \bigcup_{i=1}^k \mathcal{H}^{\perp}(\mathbf{v}^i). \quad (12)$$

Because the set $\bigcup_i \mathcal{H}^{\perp}(\mathbf{v}^i)$ has a Lebesgue measure equal to zero and the random vector \mathbf{y} admits a probability density, the probability of event (12) is zero, thus $\Pr(\|\mathbf{x}_{\mathcal{Q}}\|_0 < k) = 0$. ■

TABLE I

SBR ALGORITHM (FINAL VERSION). BY DEFAULT, THE INITIAL ACTIVE SET IS EMPTY: $\mathcal{Q}_1 = \emptyset$.

Input: \mathbf{A} , \mathbf{y} , λ and active set \mathcal{Q}_1 ($\text{Card}[\mathcal{Q}_1] \leq \min(m, n)$)
Step 1: Set $j = 1$.
Step 2: For $i \in \{1, \dots, n\}$, compute $\mathcal{K}_{\mathcal{Q}_j \bullet i}(\lambda)$.
Compute ℓ using (11).
If $\mathcal{K}_{\mathcal{Q}_j \bullet \ell}(\lambda) < \mathcal{K}_{\mathcal{Q}_j}(\lambda)$,
Set $\mathcal{Q}_{j+1} = \mathcal{Q}_j \bullet \ell$.
else,
Terminate SBR.
End if.
Set $j = j + 1$ and go to Step 2.
Output: active set $\mathcal{Q}_j = \text{SBR}(\mathcal{Q}_1; \lambda)$

Theorem 2 implies that when dealing with real noisy data, it is almost sure that no active component x_i is exactly equal to 0. Thus, the original and modified versions of SBR almost surely lead to exactly the same iterates. Even in the noiseless case, an active component is rarely numerically evaluated to 0 due to the round-off errors. In all cases, the modified version of SBR can be applied without restriction and the properties stated below (*e.g.*, termination after a finite number of iterations) remain valid even if an SBR iterate satisfies $\|\mathbf{x}_{\mathcal{Q}}\|_0 < \text{Card}[\mathcal{Q}]$.

We adopt the modified version of SBR in the rest of the paper. It is summarized in Table I.

D. Behavior and adaptations of SBR

Termination: SBR is a descent algorithm because the value of $\mathcal{K}_{\mathcal{Q}}(\lambda)$ is always decreasing. Consequently, a set \mathcal{Q} cannot be explored twice and similarly to SMLR, SBR terminates after a finite number of iterations. Notice that the size of \mathcal{Q} remains lower or equal to $\min(m, n)$. Indeed, if a set \mathcal{Q} of cardinality $\min(m, n)$ is reached, then $\mathcal{E}_{\mathcal{Q}}$ is equal to 0 due to the URP assumption. Hence, any replacement of the form $\mathcal{Q}' = \mathcal{Q} \cup \{i\}$ yields an increase of the cost function ($\mathcal{K}_{\mathcal{Q}'}(\lambda) = \mathcal{K}_{\mathcal{Q}}(\lambda) + \lambda$). We emphasize that no stopping condition is needed unlike many algorithms which require to set a maximum number of iterations (MP and variations, OLS) and/or a threshold on the squared error variation (CoSaMP, Iterative Hard Thresholding).

Proposition 1 *Under the assumptions of Theorem 2, each SBR iterate \mathbf{x}_Q is almost surely a local minimizer of $\mathcal{J}(\mathbf{x}; \lambda)$ and of the ℓ_0 -constrained problem $\min \mathcal{E}(\mathbf{x})$ subject to $\|\mathbf{x}\|_0 \leq k$, with $k = \text{Card}[Q]$. This property holds in particular for the SBR output.*

Proof: Let $\mathbf{x} = \mathbf{x}_Q$ be an SBR iterate. According to Theorem 2, the support $\mathcal{S}(\mathbf{x}) = Q$ almost surely. Setting $\varepsilon = \min_{i \in Q} |x_i| > 0$, it is easy to check that if $\mathbf{x}' \in \mathbb{R}^n$ satisfies $\|\mathbf{x}' - \mathbf{x}\| < \varepsilon$, then $\forall i \in Q, x'_i \neq 0$. Thus, $\|\mathbf{x}' - \mathbf{x}\| < \varepsilon$ implies that $\mathcal{S}(\mathbf{x}') \supseteq \mathcal{S}(\mathbf{x})$ and then $\|\mathbf{x}'\|_0 \geq k$.

Now, let \mathbf{x}' satisfy $\|\mathbf{x}' - \mathbf{x}\| < \varepsilon$.

- If $\|\mathbf{x}'\|_0 = k$, then necessarily, $\mathcal{S}(\mathbf{x}') = \mathcal{S}(\mathbf{x}) = Q$. By definition of $\mathbf{x} = \mathbf{x}_Q$, $\mathcal{E}(\mathbf{x}') \geq \mathcal{E}(\mathbf{x})$. It follows that \mathbf{x} is a local minimizer of the ℓ_0 -constrained problem. Also, $\mathcal{J}(\mathbf{x}'; \lambda) \geq \mathcal{J}(\mathbf{x}; \lambda)$ holds.
- If $\|\mathbf{x}'\|_0 > k$, then $\mathcal{J}(\mathbf{x}'; \lambda) = \mathcal{E}(\mathbf{x}') + \lambda \|\mathbf{x}'\|_0 \geq \mathcal{E}(\mathbf{x}') + \lambda(k+1)$. By continuity of \mathcal{E} , there exists a neighborhood $\mathcal{V}(\mathbf{x})$ of \mathbf{x} such that if $\mathbf{x}' \in \mathcal{V}(\mathbf{x})$, $|\mathcal{E}(\mathbf{x}') - \mathcal{E}(\mathbf{x})| < \lambda$. Thus, if $\mathbf{x}' \in \mathcal{V}(\mathbf{x})$ and $\|\mathbf{x}' - \mathbf{x}\| < \varepsilon$, $\mathcal{J}(\mathbf{x}'; \lambda) > \mathcal{E}(\mathbf{x}) + \lambda k = \mathcal{J}(\mathbf{x}; \lambda)$.

Finally, if $\mathbf{x}' \in \mathcal{V}(\mathbf{x})$ and $\|\mathbf{x}' - \mathbf{x}\| < \varepsilon$, then $\mathcal{J}(\mathbf{x}'; \lambda) > \mathcal{J}(\mathbf{x}; \lambda)$. ■

OLS as a special case: When $\lambda = 0$, SBR coincides with the well known Orthogonal Least Squares (OLS) algorithm [15, 23]. The removal operation never occurs, because it automatically leads to an increase of the squared error $\mathcal{K}_Q(0) = \mathcal{E}_Q$. Consequently, only insertions are worth being tested.

Empty solutions: We now characterize the λ -values for which SBR yields an empty solution.

Proposition 2 *The output of $\text{SBR}(\emptyset; \lambda)$ is the empty set if and only if $\lambda \geq \lambda_{\max}$, with $\lambda_{\max} \triangleq \max_i (\langle \mathbf{a}_i, \mathbf{y} \rangle^2 / \|\mathbf{a}_i\|^2)$.*

Proof: SBR stops during its first iteration if all the insertion trials fail: $\forall i, \mathcal{E}_{\{i\}} + \lambda \geq \mathcal{E}_\emptyset = \|\mathbf{y}\|^2$. Given i , $\|\mathbf{y} - x_i \mathbf{a}_i\|^2$ is minimal when $x_i = \langle \mathbf{a}_i, \mathbf{y} \rangle / \|\mathbf{a}_i\|^2$, leading to $\mathcal{E}_{\{i\}} = \|\mathbf{y}\|^2 - \langle \mathbf{a}_i, \mathbf{y} \rangle^2 / \|\mathbf{a}_i\|^2$. Thus, SBR stops during its first iteration *if and only if* $\forall i, \lambda \geq \langle \mathbf{a}_i, \mathbf{y} \rangle^2 / \|\mathbf{a}_i\|^2$, i.e., $\lambda \geq \lambda_{\max}$. ■

This result allows us to design an automatic procedure which sets a number of λ -values adaptively to the data in order to compute SBR solutions at different sparsity levels (see Section VI-D).

Reduced search: Instead of testing all the replacements $Q' = Q \bullet i$ at each SBR iteration, it is advantageous, if possible, to explore only a subset of these n replacements. We give two ideas to reduce the number of tests: the first is an acceleration of SBR, yielding the same iterates with a reduced search. The second idea is a modification of SBR.

Given an active set Q , a removal $Q' = Q \setminus \{i\}$ yields an increase of the squared error and a decrease of the penalty equal to λ . Hence, the maximum decrease of the ℓ_0 -penalized cost function which can be

expected with a removal is λ : $\mathcal{K}_{\mathcal{Q}}(\lambda) - \mathcal{K}_{\mathcal{Q}'}(\lambda) \leq \lambda$. Consequently, if a given insertion $\mathcal{Q}' = \mathcal{Q} \cup \{i\}$ is such as $\mathcal{K}_{\mathcal{Q}}(\lambda) - \mathcal{K}_{\mathcal{Q}'}(\lambda) > \lambda$, then no removal is worth being tested. The acceleration of SBR thus consists in testing all the insertions first, and if the best insertion yields a decrease larger than λ , selecting the best insertion. Otherwise, all the removals have to be tested as stated in Table I. This acceleration does not modify the SBR iterates. However, the gain is limited when the level of sparsity is high, *i.e.*, when the number of removals to be tested is reduced.

Haugland and Zhang pointed out that in a forward-backward strategy, it can be helpful to favor removals [13, 14]. Adapted to SBR, this idea leads to a modified algorithm in which removals are tested first, and the insertions are tested only if no removal yields a decrease of $\mathcal{K}_{\mathcal{Q}}(\lambda)$. If some removals decrease $\mathcal{K}_{\mathcal{Q}}(\lambda)$, then the removal yielding the maximal decrease is selected. In our experiments, the average qualitative performance of SBR and this modified version are quite comparable (there is no obvious gain or loss of quality nor a significant saving in computation time). Thus, in the following, we keep the version of SBR presented in Table I.

IV. IMPLEMENTATION ISSUES

Given the current active set \mathcal{Q} , an SBR iteration consists in computing the squared error $\mathcal{E}_{\mathcal{Q}'}$ for any replacement $\mathcal{Q}' = \mathcal{Q} \bullet i$, leading to the computation of $\mathcal{K}_{\mathcal{Q}'}(\lambda) = \mathcal{E}_{\mathcal{Q}'} + \lambda \text{Card}[\mathcal{Q}']$. We first describe a basic implementation in which $\mathcal{E}_{\mathcal{Q}'}$ is computed independently of the knowledge of $\mathcal{E}_{\mathcal{Q}}$. Then, we present an efficient implementation allowing a fast update when \mathcal{Q} is modified. We will denote by $k \triangleq \text{Card}[\mathcal{Q}]$ the cardinality of the active set ($k \leq \min(m, n)$).

A. Basic implementation

The minimization problem (1) reduces to the unconstrained minimization of $\|\mathbf{y} - \mathbf{A}_{\mathcal{Q}}\mathbf{t}\|^2$ w.r.t. $\mathbf{t} \in \mathbb{R}^k$. Because $\mathbf{A}_{\mathcal{Q}}$ is full rank, this problem has a unique minimizer that reads:

$$\mathbf{t}_{\mathcal{Q}} \triangleq \arg \min_{\mathbf{t}} \|\mathbf{y} - \mathbf{A}_{\mathcal{Q}}\mathbf{t}\|^2 = (\mathbf{A}_{\mathcal{Q}}^t \mathbf{A}_{\mathcal{Q}})^{-1} \mathbf{A}_{\mathcal{Q}}^t \mathbf{y} \quad (13)$$

and the minimal squared error reads:

$$\mathcal{E}_{\mathcal{Q}} = \|\mathbf{y} - \mathbf{A}_{\mathcal{Q}}\mathbf{t}_{\mathcal{Q}}\|^2 = \|\mathbf{y}\|^2 - \mathbf{y}^t \mathbf{A}_{\mathcal{Q}}\mathbf{t}_{\mathcal{Q}}. \quad (14)$$

Finally, given the active set \mathcal{Q} , an SBR iteration involves the computation of $\mathbf{t}_{\mathcal{Q}'}$ and $\mathcal{E}_{\mathcal{Q}'}$ for all possible replacements $\mathcal{Q}' = \mathcal{Q} \bullet i$, using (13) and (14).

B. Recursive implementation

At each SBR iteration, n least-square problems of the form (13) must be solved, each requiring the inversion of the Gram matrix $\mathbf{G}_Q \triangleq \mathbf{A}_Q^t \mathbf{A}_Q$ of size $k \times k$. The computation cost can be high since in the general case, a matrix inversion costs $\mathcal{O}(k^3)$ scalar operations. Following an idea widely spread in the subset selection literature, we propose to solve (13) in a recursive manner.

A first possibility is to use the Gram-Schmidt procedure [15, 23] which yields an orthogonal decomposition of $\mathbf{A}_Q = \mathbf{W}\mathbf{U}$, where \mathbf{W} is an $m \times k$ matrix with orthogonal columns and \mathbf{U} is a $k \times k$ upper triangular matrix. Although it yields an efficient updating strategy when including an index into the active set (leading to the update of $\mathbf{A}_{Q'} = [\mathbf{A}_Q, \mathbf{a}_i]$), the Gram-Schmidt procedure does not extend with the same level of efficiency when an index removal is considered [12].

An alternative is to use the block matrix inversion lemma [24] allowing an efficient update of \mathbf{G}_Q^{-1} for both index insertion and removal. An efficient SMLR implementation is proposed in [20], based on the recursive update of matrices of the form $(\mathbf{G}_Q + \alpha \mathbf{I}_k)^{-1}$. This approach can easily be adapted to SBR where the matrix to update is \mathbf{G}_Q^{-1} (see also [25] for the downdate step). However, we have observed numerical instabilities when the selected columns of \mathbf{A} are highly correlated.

A more stable solution is based on the Cholesky factorization $\mathbf{G}_Q = \mathbf{L}_Q \mathbf{L}_Q^t$, where \mathbf{L}_Q is a lower triangular matrix. Updating \mathbf{L}_Q is advantageous since it is better conditioned than \mathbf{G}_Q^{-1} . This update can be easily done in the insertion case [26]. It is less easy for removals, since they break the triangular structure of \mathbf{L}_Q . A recursive update of the Cholesky factor of \mathbf{G}_Q^{-1} was recently proposed [27]. Here, we introduce a simpler strategy relying on the factorization of \mathbf{G}_Q .

C. Efficient strategy based on the Cholesky factorization

First, we notice that a new column \mathbf{a}_i can be inserted at the last position in $\mathbf{A}_{Q \cup \{i\}}$ to compute the value of $\mathcal{E}_{Q \cup \{i\}}$. On the contrary, when removing a column, we do not know *a priori* the position of the column to be removed, thus it cannot be assumed to be the last column of \mathbf{A}_Q . We will hence describe the cases where:

- a non active element $i \notin Q$ is included after the other columns: $\mathbf{A}_{Q'} = [\mathbf{A}_Q, \mathbf{a}_i]$;
- an active element $i \in Q$ is to be removed, the column \mathbf{a}_i being in an arbitrary position.

\mathbf{G}_Q being a symmetric positive-definite matrix, it reads $\mathbf{G}_Q = \mathbf{L}_Q \mathbf{L}_Q^t$ where the Cholesky factor \mathbf{L}_Q is a lower triangular matrix of size $k \times k$. Applying (13), the least-square minimizer rereads $\mathbf{t}_Q = \mathbf{L}_Q^{-t} \mathbf{L}_Q^{-1} \mathbf{A}_Q^t \mathbf{y}$ where the superscript $-t$ refers to the inverse transposition operator, and (14) yields:

$$\mathcal{K}_Q(\lambda) = \mathcal{E}_Q(\lambda) + \lambda k = \|\mathbf{y}\|^2 - \|\mathbf{L}_Q^{-1} \mathbf{A}_Q^t \mathbf{y}\|^2 + \lambda k. \quad (15)$$

Given L_Q , $\mathcal{O}(k^2)$ scalar operations are required to solve the triangular system $L_Q^{-1}(A_Q^t \mathbf{y})$.

Insertion of a new column after the existing columns: Including a new column leads to $A_{Q'} = [A_Q, \mathbf{a}_i]$.

Thus, the new Gram matrix can be expressed as a 2×2 block matrix:

$$\mathbf{G}_{Q'} = \begin{bmatrix} \mathbf{G}_Q & A_Q^t \mathbf{a}_i \\ (A_Q^t \mathbf{a}_i)^t & \|\mathbf{a}_i\|^2 \end{bmatrix} \quad (16)$$

and the Cholesky factor of $\mathbf{G}_{Q'}$ can be straightforwardly updated:

$$\mathbf{L}_{Q'} = \begin{bmatrix} \mathbf{L}_Q & \mathbf{0} \\ \mathbf{l}_{Q,i}^t & \alpha_{Q,i} \end{bmatrix}, \quad (17)$$

with $\mathbf{l}_{Q,i} = L_Q^{-1} A_Q^t \mathbf{a}_i$ and $\alpha_{Q,i} = (\|\mathbf{a}_i\|^2 - \|\mathbf{l}_{Q,i}\|^2)^{1/2}$.

The computation of $\mathcal{K}_{Q'}(\lambda)$ using (15) requires two triangular system inversions (computation of $\mathbf{l}_{Q,i}$ and computation of $\mathcal{K}_{Q'}(\lambda)$). However, by computing

$$\mathcal{K}_{Q'}(\lambda) - \mathcal{K}_Q(\lambda) = \lambda - (\mathbf{l}_{Q,i}^t L_Q^{-1} A_Q^t \mathbf{y})^2 / \alpha_{Q,i}^2, \quad (18)$$

the cost can be reduced up to the pre-computation of $L_Q^{-1}(A_Q^t \mathbf{y})$ at the beginning of the SBR iteration.

The computation of $\mathcal{K}_{Q'}(\lambda)$ only requires one triangular system inversion (computation of $\mathbf{l}_{Q,i}$).

Removal of an arbitrary column: When removing a column \mathbf{a}_i , updating L_Q remains possible although slightly more expensive. This idea was first developed by Ge *et al.* [27] who update the Cholesky factorization of matrix \mathbf{G}_Q^{-1} . We adapt it to the direct (simpler) factorization of \mathbf{G}_Q . Let I be the index such that \mathbf{a}_i is the I -th column of A_Q (with $1 \leq I \leq k$). L_Q can be written in a block matrix form:

$$\mathbf{L}_Q = \begin{bmatrix} \mathbf{A} & \mathbf{0} & \mathbf{0} \\ \mathbf{b}^t & d & \mathbf{0} \\ \mathbf{C} & \mathbf{e} & \mathbf{F} \end{bmatrix}, \quad (19)$$

where the lowercase characters refer to the scalar (d) and vector quantities (\mathbf{b} , \mathbf{e}) which appear in the I -th row and in the I -th column. The computation of $\mathbf{G}_Q = L_Q L_Q^t$ and the removal of the I -th row and the I -th column in \mathbf{G}_Q leads to

$$\mathbf{G}_{Q'} = \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{C} & \mathbf{F} \end{bmatrix} \begin{bmatrix} A^t & C^t \\ \mathbf{0} & \mathbf{F}^t \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{e} \end{bmatrix} [\mathbf{0} \quad e^t]. \quad (20)$$

By identification of this expression with the Cholesky factorization $\mathbf{G}_{Q'} = L_{Q'} L_{Q'}^t$, and because the Cholesky factorization is unique, $L_{Q'}$ necessarily reads:

$$\mathbf{L}_{Q'} = \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{C} & \mathbf{X} \end{bmatrix}, \quad (21)$$

TABLE II
EFFICIENT IMPLEMENTATION OF AN SBR ITERATION.

Input: \mathcal{Q}, λ

Pre-computed quantities: $A^t \mathbf{y}$ and $\|\mathbf{a}_i\|^2$ for all i

Stored quantities: $\mathcal{K}_{\mathcal{Q}}(\lambda), \mathbf{L}_{\mathcal{Q}}$, and $\mathbf{L}_{\mathcal{Q}}^{-1}(A_{\mathcal{Q}}^t \mathbf{y})$

Set $\ell = 0$, $\text{least_cost} = \mathcal{K}_{\mathcal{Q}}(\lambda)$.

For $i = 1$ to n ,

If $i \notin \mathcal{Q}$, /* Insertion test */

Compute $\mathbf{l}_{\mathcal{Q},i} = \mathbf{L}_{\mathcal{Q}}^{-1} A_{\mathcal{Q}}^t \mathbf{a}_i$ and $\mathcal{K}_{\mathcal{Q}'}(\lambda)$ using (18).

else, /* Removal test */

Update the Cholesky decomposition: $\mathbf{X} = \text{cholupdate}(\mathbf{F}, \mathbf{e}, +)$.

Compute $\mathbf{L}_{\mathcal{Q}'}$ and $\mathcal{K}_{\mathcal{Q}'}(\lambda)$ using (21) and (15).

End if.

If $\mathcal{K}_{\mathcal{Q}'}(\lambda) < \text{least_cost}$,

Set $\ell = i$, $\text{least_cost} = \mathcal{K}_{\mathcal{Q}'}(\lambda)$.

End if.

End for.

If $\ell \neq 0$, /* Perform the single replacement */

Set $\mathcal{Q}' = \mathcal{Q} \bullet \ell$, $\mathcal{K}_{\mathcal{Q}'}(\lambda) = \text{least_cost}$.

Compute $\mathbf{L}_{\mathcal{Q}'}$ using (17) or (21), and then $\mathbf{L}_{\mathcal{Q}'}^{-1}(A_{\mathcal{Q}'}^t \mathbf{y})$.

else,

Terminate SBR.

End if.

Output: next iterate $\mathcal{Q}' = \mathcal{Q} \bullet \ell$, $\mathcal{K}_{\mathcal{Q}'}(\lambda)$, $\mathbf{L}_{\mathcal{Q}'}$, and $\mathbf{L}_{\mathcal{Q}'}^{-1}(A_{\mathcal{Q}'}^t \mathbf{y})$

where \mathbf{X} is a lower triangular matrix satisfying $\mathbf{X} \mathbf{X}^t = \mathbf{F} \mathbf{F}^t + \mathbf{e} \mathbf{e}^t$. The problem of computing \mathbf{X} from \mathbf{F} and \mathbf{e} is classical; it is known as a positive rank 1 Cholesky update and there exists a stable algorithm in $\mathcal{O}(f^2)$ operations, where $f = k - I$ is the size of \mathbf{F} [28].

Finally, the computation of $\mathcal{K}_{\mathcal{Q}'}(\lambda)$ involves a positive Cholesky update and a triangular system inversion in (15). Thus, its overall cost is in $\mathcal{O}(k^2)$. Notice that matrix \mathbf{F} is of size $k - I$. Therefore, the cost of a Cholesky update completely depends on the position I of the column \mathbf{a}_i to be removed. The larger I , the more expensive is the Cholesky update.

D. Memory requirements and computation burden

The efficient (fast and stable) procedure is summarized in Table II. Given the current active set \mathcal{Q} , the index ℓ defining the next SBR iterate $\mathcal{Q} \bullet \ell$ is chosen according to (11) and $\mathbf{L}_{\mathcal{Q} \bullet \ell}$ is finally updated. No

update of the amplitudes is necessary. The actual implementation may vary depending on the size and the structure of matrix \mathbf{A} . We briefly detail the main possible implementations and their requirements in terms of storage and computation. Regarding the computation burden, we count the number of elementary operations expressed in terms of scalar multiplications since the cost of a scalar addition is negligible with respect to that of a multiplication.

When \mathbf{A} is relatively small, the computation and storage of the full Gram matrix $\mathbf{A}^t\mathbf{A}$ prior to any SBR iteration (storage of n^2 scalar elements) avoids to recompute the vectors $\mathbf{A}_{\mathcal{Q}}^t\mathbf{a}_i$ which are needed when the insertion of \mathbf{a}_i into the active set is tested. Similarly, we store $\mathbf{A}^t\mathbf{y}$ and the values $\|\mathbf{a}_i\|^2$ in two 1D arrays of size n , prior to any SBR loop. The storage of the other quantities (mainly $\mathbf{L}_{\mathcal{Q}}$) that are being updated amounts to $\mathcal{O}(k^2)$ scalar elements and each test costs $\mathcal{O}(k^2)$ elementary operations, as it involves the inversion of a triangular system of size $k \times k$, plus a positive rank 1 Cholesky update in the removal case. This cost has to be compared with the $\mathcal{O}(k^3)$ scalar operations which are necessary when inverting the Gram matrix in the basic implementation of SBR.

When \mathbf{A} is larger, the storage of $\mathbf{A}^t\mathbf{A}$ is no longer possible and vectors $\mathbf{A}_{\mathcal{Q}}^t\mathbf{a}_i$ must be recomputed at any SBR iteration, for each insertion test $\mathcal{Q}' = \mathcal{Q} \cup \{i\}$. The computation of $\mathbf{A}_{\mathcal{Q}}^t\mathbf{a}_i$ costs km elementary operations and represents the most important cost of an insertion test. Indeed, the remaining part is in $\mathcal{O}(k^2)$ and for sparse representations, k is expected to be much lower than m . The cost of a single replacement finally amounts to $\mathcal{O}(k^2) + \mathcal{O}(km)$ elementary operations.

When the dictionary has some specific structure, the above storage limitation can be alleviated, enabling a fast implementation even when n is large. For instance, if a large number of pairs of columns of \mathbf{A} are orthogonal to each other, $\mathbf{A}^t\mathbf{A}$ can be stored as a sparse array. Also, finite impulse response deconvolution problems enable a fast implementation since $\mathbf{A}^t\mathbf{A}$ is then a Toeplitz matrix (save north-west and/or south-east submatrices, depending on the boundary conditions). The knowledge of the auto-correlation of the impulse response is sufficient to describe most of the Gram matrix.

All these variants have been implemented (Matlab codes are available to academic users from the authors upon request). In the following, we analyze the behavior of SBR on two difficult problems, in which the dictionaries are highly correlated: the deconvolution of a sparse signal with a Gaussian impulse response (Section V) and the joint detection of discontinuities at different orders in a signal (Section VI).

V. DECONVOLUTION OF A SPARSE SIGNAL WITH A GAUSSIAN IMPULSE RESPONSE

This is a typical problem for which SMLR was introduced [20]. It affords us to study the ability of SBR to perform an exact recovery in a simple noise-free case (separation of two Gaussian features from noise-free data) and to test the behavior of SBR in a noisy case (estimation of a larger number of Gaussian features). For simulated problems, we denote by \mathbf{x}^* the exact sparse signal and we generate noisy data according to $\mathbf{y} = \mathbf{y}^* + \mathbf{n} = \mathbf{A}\mathbf{x}^* + \mathbf{n}$, where $\mathbf{y}^* = \mathbf{A}\mathbf{x}^*$ denotes the noise-free data and \mathbf{n} stands for the observation noise. The dictionary columns \mathbf{a}_i are always normalized: $\|\mathbf{a}_i\|^2 = 1$. The signal to noise ratio (SNR) is defined by $\text{SNR} = 10 \log(P_Y/P_N)$, where $P_Y = \|\mathbf{y}^*\|^2/m$ is the average power of the noise-free data and P_N is the variance of the noise process \mathbf{n} .

A. Dictionary and simulated data

The impulse response \mathbf{h} is a Gaussian signal of standard deviation σ , sampled on a regular grid at integer locations. It is approximated by a finite impulse response of length 6σ by thresholding the smallest values, allowing a fast implementation even for large size problems (see subsection IV-D). The deconvolution problem leads to a Toeplitz matrix \mathbf{A} whose columns \mathbf{a}_i are obtained by shifting the signal \mathbf{h} . The dimension of \mathbf{A} is chosen to have any Gaussian feature resulting from the convolution $\mathbf{h} * \mathbf{x}^*$ belonging to the observation window $\{1, \dots, m\}$. This implies that \mathbf{A} is slightly undercomplete ($m > n$). Denoting by $n_h = 1 + 2\text{round}(3\sigma)$ the size of the support of \mathbf{h} , the data size reads $m = n + n_h - 1$.

B. Separation of two close Gaussian features

We first analyze the ability of SBR to separate two Gaussian features ($\|\mathbf{x}^*\|_0 = 2$) from noise-free data. The centers of both Gaussian features lay at a relative distance d (expressed as a number of samples) and their weights x_i^* are set to 1. We generate the corresponding noise-free data \mathbf{y}^* and we run $\text{SBR}(\emptyset; \lambda)$ with a number of predefined λ -values. We analyze the SBR outputs $\mathcal{Q}(\lambda; d)$ by computing their size $\text{Card}[\mathcal{Q}(\lambda; d)]$ and by testing if $\mathcal{Q}(\lambda; d)$ is equal to the true support $\mathcal{S}(\mathbf{x}^*)$. Table III shows the results obtained for a problem of size 300×270 ($m = 300, \sigma = 5$, and $n_h = 31$) with distances equal to $d = 20, 13$, and 6 samples. The grid of λ -values for which SBR is run is common to the three tests. The maximal value λ_0 is chosen in such a way that the output $\mathcal{Q}(\lambda_0; d)$ is empty (see Proposition 2) and the other values are set according to $\lambda_j = \lambda_0/10^j$. It is noticeable that the exact recovery is always reached provided that λ is sufficiently small. This result remains true even for smaller distances (from $d = 2$). When the Gaussian features strongly overlap, *i.e.*, for $d \leq 13$, the size of the support obtained as output first increases while λ decreases, and then for lower λ -values, removals start to occur, enabling

TABLE III

SEPARATION OF TWO GAUSSIAN FEATURES FROM NOISE-FREE DATA WITH SBR. d STANDS FOR THE DISTANCE BETWEEN THE GAUSSIAN FEATURES. WE DISPLAY THE SIZE OF THE SUPPORT $\mathcal{Q}(\lambda; d)$ OBTAINED WITH A SEQUENCE OF DECREASING λ -VALUES $\lambda_0 > \lambda_1 > \dots > \lambda_7$. THE LABEL * INDICATES AN EXACT RECOVERY FOR A SUPPORT OF CARDINALITY 2.

λ	λ_0	λ_1	λ_2	λ_3	λ_4	λ_5	λ_6	$\lambda \leq \lambda_7$
$d = 20$	0	0	2*	2*	2*	2*	2*	2*
$d = 13$	0	1	3	4	5	2*	2*	2*
$d = 6$	0	1	1	3	5	6	8	2*

the exact recovery. Similarly to SBR, forward algorithms such as OMP and OLS start by positioning a (wrong) Gaussian feature in between the two Gaussians in their first iteration but in the latter case, the early wrong detection disables an exact recovery.

C. Behavior of SBR for noisy data

We run SBR on more realistic noisy data and on a larger problem ($m = 3000$ samples). The unknown sparse signal \mathbf{x}^* is composed of 17 Gaussian features. The impulse response \mathbf{h} is of size $n_h = 301$ ($\sigma = 50$) yielding an observation matrix \mathbf{A} of size 3000×2700 , and the SNR is set to 20 dB.

Fig. 1 displays the simulated data and the SBR results obtained with a few λ -values. When λ decreases, the SBR approximations are of better quality but less sparse. For large λ -values, only the main Gaussian features are found, and then, when λ decreases, the smaller features are being recovered together with unwanted features. Removals rarely occur for coarse approximations. They occur more frequently when two estimated features are overlapping and for low λ -values. On the simulation of Fig. 1, removals occur for $\lambda \leq 0.15$, yielding approximations that are more accurate than those obtained with OLS and for the same cardinality (the residual $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ is lower), while when $\lambda > 0.15$, the SBR output coincides with the OLS iterate of same cardinality. Although the performance of SBR is at least equal to that of OLS, the exact support of \mathbf{x}^* is never found. However, it must be stressed that the problem is very difficult because the data are noisy and the neighboring columns of \mathbf{A} are highly correlated. In such difficult case, one needs to perform a wider exploration of the discrete set $\{0, 1\}^n$ by introducing moves that are more complex than single replacements. Such extensions were already proposed in the case of SMLR. One can for instance shift a detected spike x_i forwards or backwards [29] or update a block of neighboring components jointly (e.g., x_i and x_{i+1}) [30].

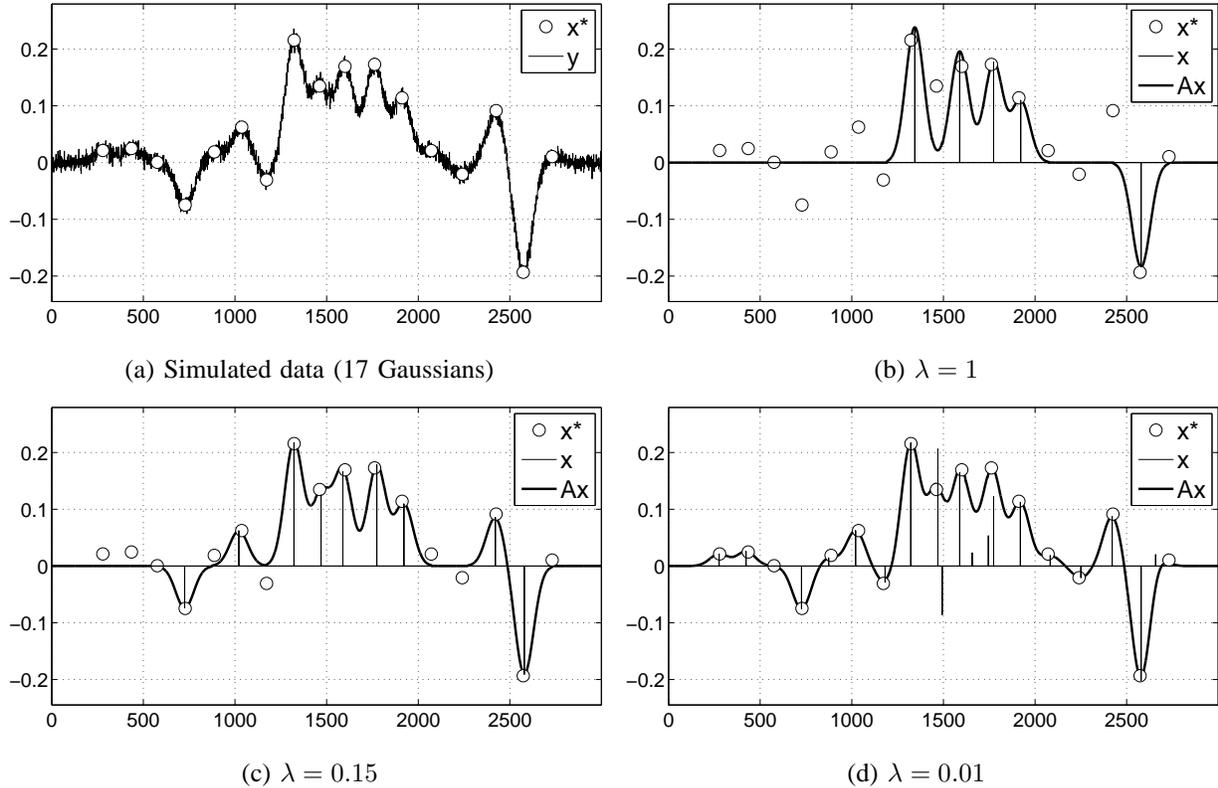


Fig. 1. Gaussian deconvolution results. Problem of size 3000×2700 ($\sigma = 50$). (a) Generated data, with 17 Gaussian features and with SNR = 20 dB. The exact locations x^* are labeled \circ . (b,c,d) SBR outputs and data approximations with empirical settings of λ . The estimated amplitudes x are shown with vertical spikes. The SBR outputs (supports) are of size 5, 9, and 19, respectively. The computation time always remains below 8 seconds (Matlab implementation).

VI. JOINT DETECTION OF DISCONTINUITIES AT DIFFERENT ORDERS IN A SIGNAL

We now consider another challenging problem, the joint detection of discontinuities at different orders in a signal. We process both simulated and real data and compare the performance of SBR with respect to other sparse approximation algorithms (OMP and OLS) in terms of discontinuity estimation, approximation accuracy, and computation time. Firstly, we formulate the detection of discontinuities at a single order p as a spline approximation problem. Then, we take advantage of this formulation to introduce the joint detection problem more easily.

A. Approximation of a spline of degree p

In the continuous case, a signal is a spline of degree p with k knots *if and only if* its $(p + 1)$ -th derivative is a stream of k weighted Diracs [31]. In the discrete case, we introduce the dictionary \mathbf{A}^p

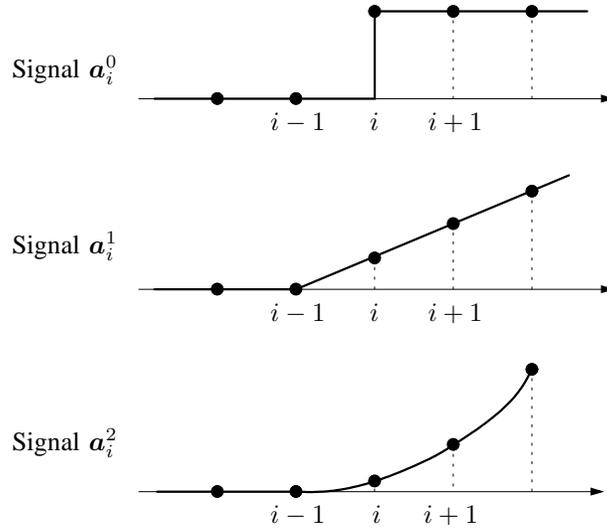


Fig. 2. Signals \mathbf{a}_i^p related to the p -th order discontinuities at location i . \mathbf{a}_i^0 is the Heaviside step function, \mathbf{a}_i^1 is the ramp function, and \mathbf{a}_i^2 is the one-sided quadratic function. Each signal is equal to 1 at location i and its support is equal to $\{i, \dots, m\}$.

formed of shifted versions of the one-sided power function $k \mapsto k_+^p \triangleq [\max(k, 0)]^p$ for all possible shifts (see Fig. 2). \mathbf{A}^p represents the integration operator of degree $p+1$. Denoting by $\{1, \dots, m\}$ the support of the data signal \mathbf{y} , the shifted signals \mathbf{a}_i^p (for $i \in \{1, \dots, m\}$) read

$$\forall k \in \{1, \dots, m\}, \mathbf{a}_i^p(k) = (k - i + 1)_+^p \quad (22)$$

and their support is equal to $\{i, \dots, m\}$. Finally, we form the dictionary $\mathbf{A}^p = [\mathbf{a}_1^p, \dots, \mathbf{a}_{m-p}^p]$ of size $m \times (m-p)$. It does not make sense to allow the occurrence of a p -th order discontinuity for the last samples (*i.e.*, to include \mathbf{a}_i^p for $i > m-p$) since the spline approximation would require to reconstruct a polynomial of degree p in the range $\{i, \dots, m\}$ from less than $p+1$ data samples.

We address the spline approximation problem as the sparse approximation of \mathbf{y} by the piecewise polynomial $\mathbf{g}^p = \mathbf{A}^p \mathbf{x}^p$ (actually, we impose as initial condition that the spline function is equal to 0 for $k \leq 0$). The sparse approximation consists in the detection of the discontinuity locations (also referred to as knots in the spline approximation literature) and the estimation of their amplitudes: x_i^p codes for the amplitude of a jump at location i ($p=0$), the change of slope at location i ($p=1$), *etc.* Here, the notion of sparsity is related to the number of discontinuity locations.

B. Approximation of a piecewise polynomial of maximum degree P

Following [31], we formulate this problem as the joint detection of discontinuities at orders $p = 0, \dots, P$. Let us append the elementary dictionaries \mathbf{A}^p in a global dictionary $\mathbf{A} = [\mathbf{A}^0, \dots, \mathbf{A}^P]$. The

approximation $\mathbf{g} = \mathbf{A}\mathbf{x}$ of a given signal rereads $\mathbf{g} = \sum_p \mathbf{A}^p \mathbf{x}^p$ where vector $\mathbf{x} = \{\mathbf{x}^0, \dots, \mathbf{x}^P\}$ gathers the p -th order amplitudes \mathbf{x}^p for all p . When \mathbf{x} is sparse, all vectors \mathbf{x}^p are sparse and the approximation signal \mathbf{g} is the sum of piecewise polynomials of degree lower than P with a limited number of pieces.

The dictionary \mathbf{A} is overcomplete since it is of size $m \times s$, with $s = (P + 1)(m - P/2) > m$ for all $P \geq 1$. Moreover, it is highly correlated: any column \mathbf{a}_i^p is strongly correlated with *all* other columns \mathbf{a}_j^q because their respective supports are the intervals $\{i, \dots, m\}$ and $\{j, \dots, m\}$, and hence overlap. The discontinuity detection problem is difficult, as most algorithms are very likely to position wrong discontinuities in their first iterations. For example, when approximating a signal with two discontinuities at distinct locations i and j , they start to position a first (wrong) discontinuity in between i and j , and forward algorithms cannot remove it (see Section VI-E and Fig. 5 for details).

C. Adaptation of SBR

It is important to notice that the dictionary defined above does not satisfy the unique representation property. For instance, the difference between two discrete ramps at locations i and $i + 1$ yields the discrete Heaviside function at location i : $\mathbf{a}_i^1 - \mathbf{a}_{i+1}^1 = \mathbf{a}_i^0$. More generally, for $p > 1$, $\mathbf{a}_i^p - \mathbf{a}_{i+1}^p$ reads as a linear combination of \mathbf{a}_i^0 and \mathbf{a}_{i+1}^q ($q = 1, \dots, p - 1$).

As mentioned in Section II, the SBR algorithm basically requires that the dictionary satisfies the URP to ensure that the Gram matrix $\mathbf{G}_{\mathcal{Q}} = \mathbf{A}_{\mathcal{Q}}^t \mathbf{A}_{\mathcal{Q}}$ is invertible, but this assumption can be relaxed provided that only full rank matrices $\mathbf{A}_{\mathcal{Q}}$ are explored. Here, SBR is slightly modified, based on the following proposition which gives a sufficient condition of invertibility of $\mathbf{G}_{\mathcal{Q}}$.

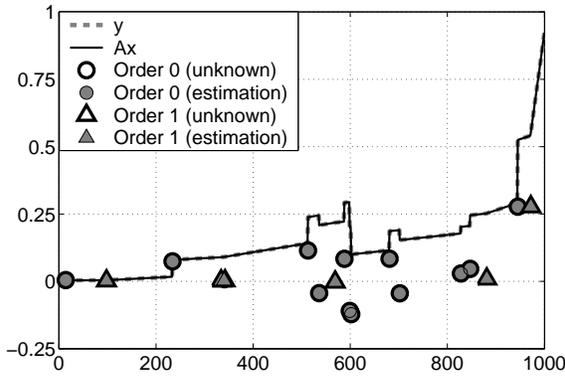
Proposition 3 *Let n_i denote the number of discontinuities \mathbf{a}_i^p , $p = 0, \dots, P$ which are being activated at sample i , i.e., for which $x_i^p \neq 0$. Let us define the binary condition $\mathcal{C}(i)$:*

- if $n_i = 0$, $\mathcal{C}(i) \triangleq 1$;
- if $n_i \geq 1$, $\mathcal{C}(i) \triangleq (\forall j \in \{1, \dots, n_i - 1\}, n_{i+j} = 0)$.

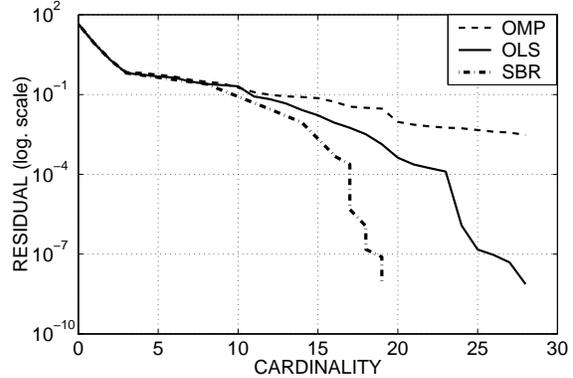
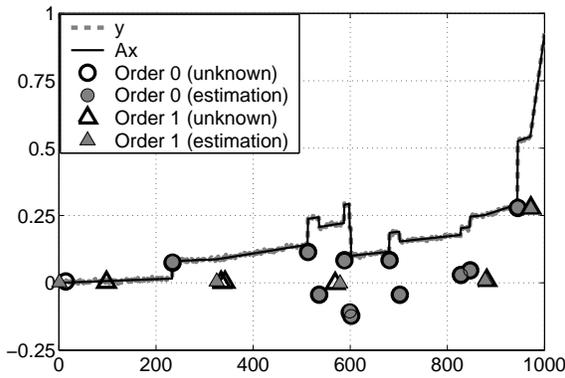
If \mathcal{Q} is such that for all i , $\mathcal{C}(i) = 1$, then $\mathbf{G}_{\mathcal{Q}}$ is invertible.

Proposition 3 is proved in Appendix A.

Basically, it states that we can allow several discontinuities to be active at the same location i , but then, the next samples $i + 1, \dots, i + n_i - 1$ must not host any discontinuity. This condition ensures that there are at most n_i discontinuities in the interval $\{i, \dots, i + n_i - 1\}$ of length n_i . The adaptation of SBR consists in testing insertions into the current active set only if the above condition remains true.



(a) Noise-free data and SBR approximation


 (b) “ ℓ_2 - ℓ_0 ” curves (noise-free data)


(c) Noisy data (SNR = 35 dB) and SBR approximation

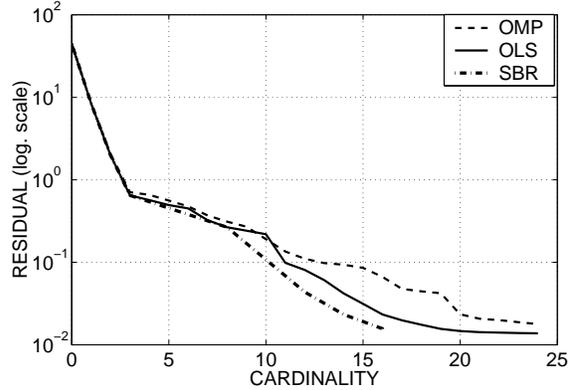

 (d) “ ℓ_2 - ℓ_0 ” curves (noisy data)

Fig. 3. Joint detection of discontinuities at orders 0 and 1. The dictionary is of size 1000×1999 and the data signal \mathbf{y} includes 18 discontinuities. The true and estimated discontinuity locations are represented with unfilled black and filled gray labels. The shape of the labels (circular or triangular) indicates the discontinuity order. The dashed gray and solid black curves represent the data signal \mathbf{y} and its approximation $\mathbf{A}\mathbf{x}$ for the least λ -value. (a) Signal approximation from noise-free data. The recovery is exact and both curves are superimposed. (b) “ ℓ_2 - ℓ_0 ” curves showing the squared residual versus the cardinality for SBR, OLS, and OMP. The SBR performance is expressed only for the λ -values that are larger than λ_{20} , because below this value, the recovery is exact and the log-residual is equal to $-\infty$. (c,d) Similar results for noisy data (SNR = 35 dB).

D. Numerical simulations

Let us first consider the case $P = 1$, leading to the joint detection of discontinuities of order zero and one, *i.e.*, the piecewise affine approximation problem. We simulate noise-free data $\mathbf{y}^* = \mathbf{A}\mathbf{x}^*$ of size $m = 1000$ and with $\|\mathbf{x}^*\|_0 = 18$ discontinuities (see Fig. 3(a)). We use the result of Proposition 2 to compute the value λ_{\max} below which the SBR output is not the empty set, and we run SBR with $\lambda_j = \lambda_{\max} 10^{-j/2}$ for $j = 0, \dots, J_{\max}$, with $J_{\max} = 20$. These λ -values provide a sequence of solutions at different sparsity levels. For comparison purpose, we also run 27 iterations of OMP and OLS.

The SBR approximation shown in Fig. 3(a) corresponds to the least λ -value $\lambda_{J_{\max}}$. The recovery is exact. The “ ℓ_2 - ℓ_0 ” curves represented on Fig. 3(b) express the squared residual $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ versus the cardinality $\|\mathbf{x}\|_0$ for the output of each algorithm. Whatever the level of sparsity, SBR yields the least residual. We did the same experiment with noisy data $\mathbf{y} = \mathbf{A}\mathbf{x}^* + \mathbf{n}$, setting the SNR to 35 dB (see Figs. 3(c,d)). Again, the “ ℓ_2 - ℓ_0 ” curve corresponding to SBR lays below the OMP and OLS curves. For most sparsity levels, SBR outperforms the other algorithms. Note that for more noisy data (e.g., SNR = 15 dB), the SBR and OLS curves coincide and still lay below the OMP curve.

E. AFM data processing

In Atomic Force Microscopy (AFM), a force curve measures the interatomic forces exerting between a probe associated to a cantilever and a nano-object. More precisely, the recorded signal $z \mapsto y(z)$ shows the force evolution versus the probe-sample distance z , expressed in nanometers. Researching discontinuities (location, order, and amplitude) in a force curve is a challenging task because they are used to provide a precise characterization of the physico-chemical properties of the nano-object (topography, energy of adhesion, etc.) [32].

The data displayed on Fig. 4(a) are related to a bacterial cell *Shewanella putrefaciens* laying in aqueous solution, interacting with the tip of the AFM probe [33]. A force curve is recorded in two steps. Firstly, the tip is positioned far away from the sample. It is moved towards the sample until the contact is reached and the surface of the bacterial cell is deformed (approach curve). Secondly, the tip is retracted from the sample until it loses contact. The experimental curve shown on Fig. 4(a) is a retraction curve composed of $m = 2167$ force measurements. From right to left, three regions of interest can be distinguished. The linear region on the right characterizes the rigid contact between the probe and the sample. It describes the mechanical interactions of the cantilever and the sample. The rigid contact is maintained until $z \approx -2840$ nm. The interactions occurring in the interval $z \in [-3050, -2840]$ nm are adhesion forces during the retraction of the tip. In the flat part on the left, no interaction occurs as the cantilever has lost contact with the sample.

We search for the discontinuities of orders 0, 1, and 2. Similarly to the processing of simulated data, we run SBR with $J_{\max} = 14$ λ -values and we run OLS and OMP until iteration 41. For each algorithm, we plot the “ ℓ_2 - ℓ_0 ” curve representing the squared residual $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ versus the cardinality $\|\mathbf{x}\|_0$, and a curve displaying the time of reconstruction versus the cardinality (see Figs. 4(b,c)). These figures show that the performance of SBR is at least equal and sometimes better than that of OLS. Both algorithms yield results that are far more accurate than OMP at the price of a larger computation time. However, notice

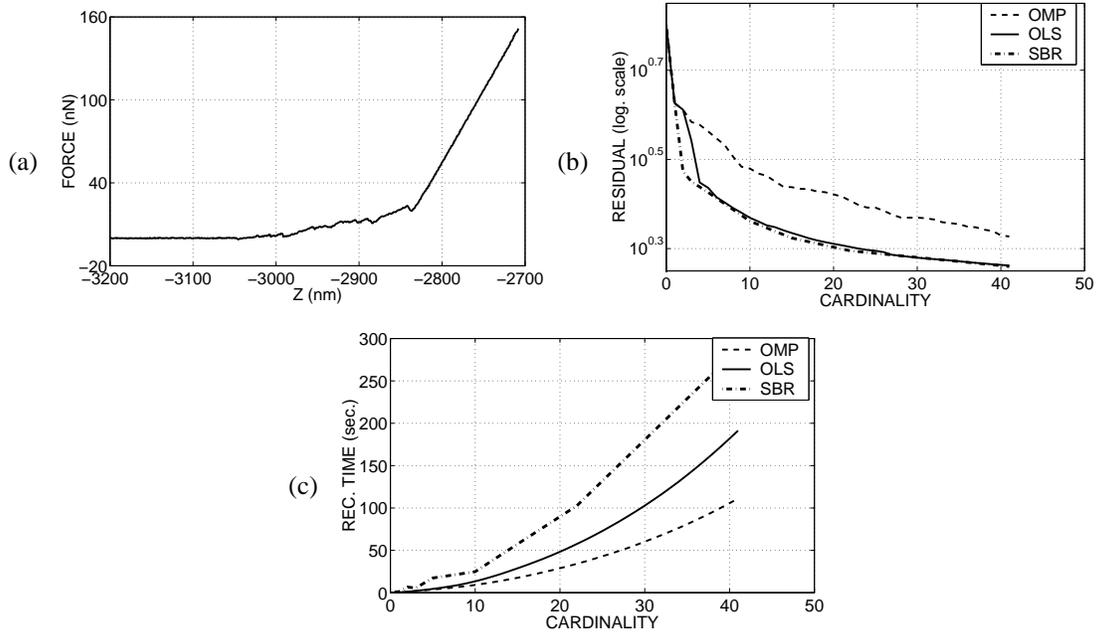


Fig. 4. AFM data processing: joint detection of discontinuities at orders 0, 1, and 2 (problem of size 2167×6498). (a) Experimental data showing the force evolution versus the probe-sample distance z . (b) Squared residual versus cardinality for the SBR, OLS, and OMP outputs. (c) Time of reconstruction versus cardinality for the three algorithms.

that the recorded computation time always remains below 350 seconds in the case of SBR (in a Matlab implementation taking advantage of the block Toeplitz structure of the dictionary: see Section IV-D).

Fig. 5 displays the approximations yielded by the three algorithms for supports of cardinality 2 and 5. SBR actually runs during 6 iterations (4 insertions and 2 removals are performed) to reach a support of cardinality 2. This approximation is very accurate compared to the OMP and OLS results obtained after 2 iterations (Figs. 5(a,b,c)). SBR provides a very precise localization of both first order discontinuities, which are crucial information for the physical interpretation of the data. On the contrary, OLS does not succeed after two iterations; it is able to locate accurately both discontinuities once 5 iterations have been performed (the desired discontinuities are the first and the last ones among the 5) while OMP fails even after 5 iterations (Figs. 5(d,e,f)). The residual yielded by the SBR approximation of cardinality 5 remains lower than the corresponding OLS and OMP residuals.

In order to better understand the forward and backward moves (respectively, insertions and removals) occurring during the SBR iterations, we display in Table IV the residual $\|\mathbf{y} - \mathbf{Ax}\|^2$ and the cardinality of each iterate for both SBR executions. Because SBR is a descent algorithm, the penalized cost $\mathcal{J}(\mathbf{x}; \lambda)$ keeps decreasing but when a removal occurs, $\|\mathbf{y} - \mathbf{Ax}\|^2$ increases. For the coarse approximation of

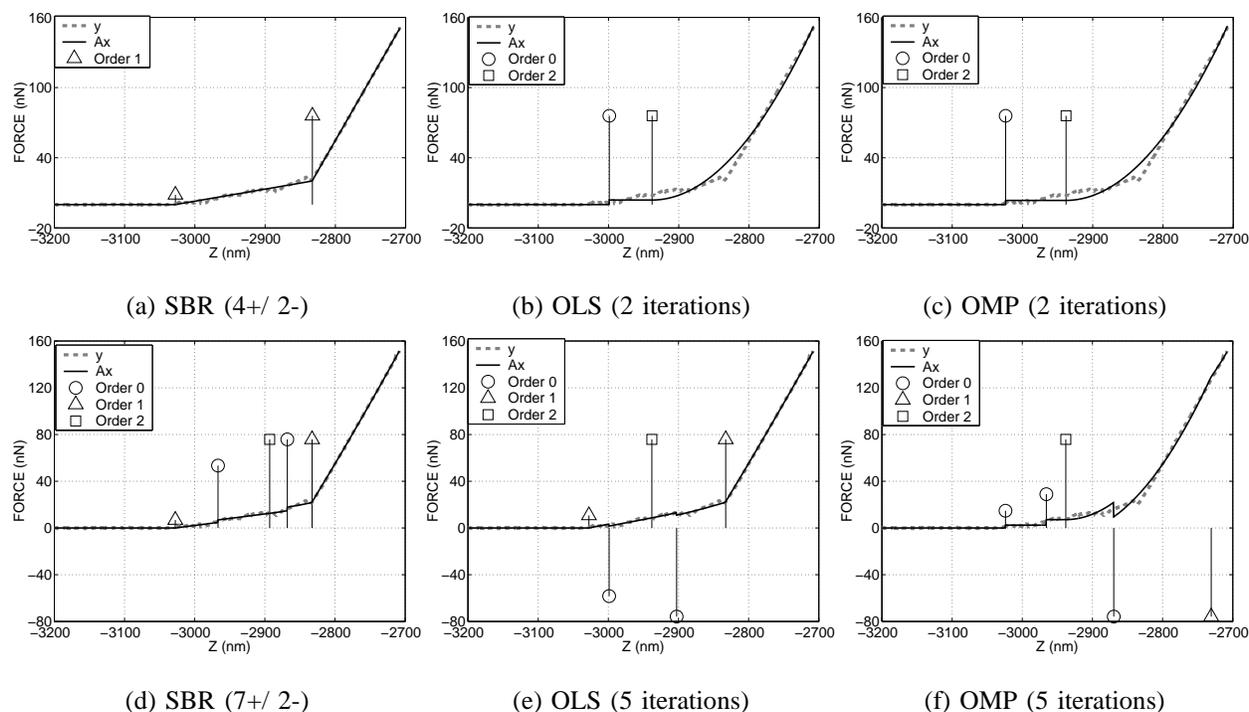


Fig. 5. AFM data processing: joint detection of discontinuities at orders 0, 1, and 2. The estimated discontinuities \mathbf{x} are represented with vertical spikes and with a label indicating the discontinuity order. The dashed gray and solid black curves represent the data signal \mathbf{y} and its approximation \mathbf{Ax} , respectively. (a) SBR output of cardinality 2: 4 insertions and 2 removals have been done ($\lambda = 2085$). (b,c) OLS and OMP outputs after 2 iterations. (d,e,f) Same simulation with a lower λ -value ($\lambda = 66$). The SBR output is of cardinality 5 (7 insertions and 2 removals) and we stop OLS and OMP after 5 iterations.

TABLE IV

BEHAVIOR OF THE SBR ITERATES FOR BOTH APPROXIMATIONS OF FIG. 5(A,D). THE TABLES DISPLAY THE SQUARED ERROR $\|\mathbf{y} - \mathbf{Ax}\|^2$ VERSUS THE CARDINALITY $\|\mathbf{x}\|_0$ FOR EACH ITERATE. i STANDS FOR THE ITERATION INDEX.

SBR
($\lambda = 2085, 4+/ 2-$)

i	$\ \mathbf{x}\ _0$	Error
0	0	2101.408
1	1	16.870
2	2	12.266
3	3	3.074
4	4	642
5	3	663
6	2	938

SBR
($\lambda = 66, 7+/ 2-$)

i	$\ \mathbf{x}\ _0$	Error
0	0	2101.408
1	1	16.870
2	2	12.266
3	3	3.074
4	4	642
5	3	663
6	4	555
7	5	480
8	4	532
9	5	464

Fig. 5(a), the residual is much smaller than the residual yielded by OLS since two configurations of cardinality 2 have been explored (see the left table and the rows in bold). We observed the same behavior for the finer approximation of cardinality 5 (right table).

F. Discussion

In the two previous subsections, we chose to compare SBR with OMP and OLS. We did not consider simpler algorithms like MP which are well suited to rapidly solve easier problems in which the dictionary columns are almost orthogonal. Because SBR involves more complex operations (matrix inversions), we chose to compare it with OMP and OLS which also require to solve at least one least-square problem per iteration. Their target is to provide results which are more accurate than the MP approximations in the case of difficult problems.

Up to our knowledge, the only minimization algorithm dedicated to the ℓ_0 -penalized cost function $\mathcal{J}(\mathbf{x}; \lambda) = \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \lambda\|\mathbf{x}\|_0$ is Blumensath and Davies' Iterative Hard Thresholding (IHT) [9]. It relies on gradient based iterations of the form $\mathbf{x}' = \mathbf{x} + \mathbf{A}^t(\mathbf{y} - \mathbf{A}\mathbf{x})$, followed by the thresholding to 0 of all the non-zero components x_i such that $|x_i| \leq \lambda^{0.5}$. We tested this version of IHT on both deconvolution and discontinuity detection problems and we observed that it is less efficient than the standard version of IHT, dedicated to the ℓ_0 -constrained problem. In the constrained version, the k components $|x_i|$ having the largest amplitudes are kept and the others are being thresholded. Generally speaking, we observed that IHT is competitive when the correlation between any pair of dictionary columns is limited, but for highly correlated dictionaries, IHT needs a very large number of iterations ($\mathcal{O}(m^2)$) to reach convergence. SBR seems to be better suited to such difficult problems. It is less sensitive to the initial solution and “skips” some local minimizers having a large cost $\mathcal{J}(\mathbf{x}; \lambda)$. We here recall that according to Proposition 1, each SBR iterate is almost surely a local minimizer of $\mathcal{J}(\mathbf{x}; \lambda)$.

In order to link up our approach to the forward-backward algorithm of [14], we also tested an OMP-like adaptation of SBR in which only one least-square problem is solved per iteration, instead of n . This adaptation consists in replacing the selection rule (11) in the following way. When an insertion $\mathcal{Q} \cup \{i\}$ is tested, all the active components x_j are kept constant and x_i is set to the minimizer of $\|\mathbf{y} - \mathbf{A}\mathbf{x}_{\mathcal{Q}} - x_i\mathbf{a}_i\|^2$. This leads to an approximation of $\mathcal{K}_{\mathcal{Q}, \bullet; i}(\lambda)$ without solving any least-square problem. Similarly, the removal test consists in setting x_i to 0 and leaving the other components x_j unchanged. In brief, this adapted version is an algorithm aimed at the minimization of $\mathcal{J}(\mathbf{x}; \lambda)$ at a cost which is comparable to that of OMP. In all our trials, SBR performs better than the OMP-like version except in very simple cases (limited correlation between the columns \mathbf{a}_i) where both versions yield the same

result. The performance of the OMP-like version fluctuates below or above that of OMP but is almost always far less accurate than the OLS and SBR approximations.

VII. CONCLUSION

We have evaluated the SBR algorithm on two problems in which the dictionary columns are highly correlated. SBR provides solutions which are at least as accurate as the OLS solutions and sometimes more accurate, with a cost of the same order of magnitude. For these difficult problems, MP and OMP provide poor approximations within a lower computation time. Compared to OLS, we believe that performing removals is the price to pay if one expects an enhanced quality approximation. Although removals rarely occur in comparison with the insertions, they play an important role in the enhancement of the approximation.

In the proposed approach, the main difficulty relies in the choice of the λ -value. If a specific sparsity level or approximation residual is desired, one can resort to a trial and error procedure in which a number of λ -values are tried until the desired approximation level is found. In [34], we sketched a continuation version in which a series of SBR solutions are successively estimated with a decreasing level of sparsity λ , and the λ -values are recursively computed. The first λ -value is set to $\lambda_0 = +\infty$, and at a given value λ_i , the initial solution (input of SBR) is set to the SBR output at $\lambda = \lambda_{i-1}$. This continuation version provides promising results and will be the subject of a future extended contribution. A similar perspective is actually proposed by Zhang to generalize his FoBa algorithm in a path-following algorithm (see the discussion section in [14]).

Another important perspective is to investigate whether SBR can guarantee exact recovery in the noise-free case under some conditions on matrix \mathbf{A} and on the unknown sparse signal \mathbf{x}^* . In the simulations done in Section V, we observed that SBR is able to exactly recover two close Gaussian features whatever their distance, provided that the hyperparameter λ is sufficiently small. This promising result is a first step towards a more general theoretical study. The FoBa algorithm [14] yields exact recovery results for problems satisfying the Restricted Isometry Property (RIP). Since the structure of SBR is somewhat close to that of FoBa, we expect that SBR shares similar theoretical properties. We will investigate whether the proofs provided in [14] are extendable to SBR.

APPENDIX A

PROOF OF PROPOSITION 3

The following lemma is a key element to prove Proposition 3.

Lemma 1 Consider an active set \mathcal{Q} satisfying the condition of Proposition 3, and let $i^- = \min\{i \mid n_i > 0\}$ denote the lowest location of an active entry. Up to a reordering of the columns, $\mathbf{A}_{\mathcal{Q}}$ rereads $\mathbf{A}_{\mathcal{Q}} = [\mathbf{A}_{i^-}, \mathbf{A}_{\mathcal{Q} \setminus \{i^-\}}]$ with obvious notations. If $\mathbf{A}_{\mathcal{Q} \setminus \{i^-\}}$ is full rank, then $\mathbf{A}_{\mathcal{Q}}$ is also full rank.

Proof: Let $I = n_{i^-}$ denote the number of discontinuities at location i^- and let $0 \leq p_1 < p_2 < \dots < p_I$ denote their order, sorted in the ascending order. Suppose that there exist two families of scalars $\{\mu_{i^-}^{p_1}, \dots, \mu_{i^-}^{p_I}\}$ and $\{\mu_i^p \mid i \neq i^- \text{ and } i \text{ is active at order } p\}$ such that

$$\sum_{j=1}^I \mu_{i^-}^{p_j} \mathbf{a}_{i^-}^{p_j} + \sum_{i \neq i^-} \sum_p \mu_i^p \mathbf{a}_i^p = \mathbf{0}. \quad (23)$$

Let us show that all μ -values are then equal to 0.

Rewriting the first I nonzero equations in this system and because \mathcal{Q} satisfies the condition of Proposition 3, we have, for all $k \in \{i^-, \dots, i^- + I - 1\}$, $\sum_{j=1}^I \mu_{i^-}^{p_j} (k + i^- - 1)^{p_j} = 0$. In other words, the polynomial $F(X) = \sum_{j=1}^I \mu_{i^-}^{p_j} X^{p_j}$ has I positive roots. It is shown in [35, p. 76] that a non-zero polynomial formed of I monomials of different degree has at most $I - 1$ positive roots. Therefore, F is the zero polynomial and all scalars $\mu_{i^-}^{p_j}$ are 0. We deduce from (23) and from the full rankness of $\mathbf{A}_{\mathcal{Q} \setminus \{i^-\}}$ that $\mu_i^p = 0$ for all (i, p) .

We have shown that the column vectors of $\mathbf{A}_{\mathcal{Q}}$ are linearly independent, *i.e.*, that $\mathbf{A}_{\mathcal{Q}}$ is full rank. ■ The proof of Proposition 3 directly results from a recursive application of Lemma 1. Starting from the empty set, all the indices, sorted by decreasing order, are included successively.

REFERENCES

- [1] B. K. Natarajan, “Sparse approximate solutions to linear systems”, *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, Apr. 1995.
- [2] B. D. Rao, K. Engan, S. F. Cotter, J. Palmer, and K. Kreutz-Delgado, “Subset selection in noise based on diversity measure minimization”, *IEEE Trans. Signal Processing*, vol. 51, no. 3, pp. 760–770, Mar. 2003.
- [3] G. H. Mohimani, M. Babaie-Zadeh, and C. Jutten, “A fast approach for overcomplete sparse decomposition based on smoothed ℓ^0 norm”, *IEEE Trans. Signal Processing*, vol. 57, no. 1, pp. 289–301, Jan. 2009.
- [4] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit”, *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.
- [5] D. L. Donoho and Y. Tsaig, “Fast solution of l_1 -norm minimization problems when the solution may be sparse”, *IEEE Trans. Inf. Theory*, vol. 54, no. 11, pp. 4789–4812, Nov. 2008.
- [6] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression”, *Annals Statist.*, vol. 32, no. 2, pp. 407–451, 2004.
- [7] J.-J. Fuchs, “Recovery of exact sparse representations in the presence of bounded noise”, *IEEE Trans. Inf. Theory*, vol. 51, no. 10, pp. 3601–3608, Oct. 2005.

- [8] D. Needell and J. A. Tropp, “CoSaMP: Iterative signal recovery from incomplete and inaccurate samples”, *Appl. Comp. Harmonic Anal.*, vol. 26, no. 3, pp. 301–321, May 2009.
- [9] T. Blumensath and M. E. Davies, “Iterative thresholding for sparse approximations”, *The Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 629–654, Dec. 2008.
- [10] S. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries”, *IEEE Trans. Signal Processing*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [11] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, “Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition”, in *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers*, Nov. 1993, vol. 1, pp. 40–44.
- [12] C. Couvreur and Y. Bresler, “On the optimality of the backward greedy algorithm for the subset selection problem”, *SIAM J. Matrix Anal. Appl.*, vol. 21, no. 3, pp. 797–808, Feb. 2000.
- [13] D. Haugland, *A Bidirectional Greedy Heuristic for the Subspace Selection Problem*, vol. 4638 of *Lecture Notes in Computer Science*, pp. 162–176, Springer Verlag, Berlin, Germany, Engineering stochastic local search algorithms. Designing, implementing and analyzing effective heuristics edition, 2007.
- [14] T. Zhang, “Adaptive forward-backward greedy algorithm for learning sparse representations”, Tech. Rep., Rutgers Statistics Department, Apr. 2008.
- [15] S. Chen, S. A. Billings, and W. Luo, “Orthogonal least squares methods and their application to non-linear system identification”, *Int. J. Control*, vol. 50, no. 5, pp. 1873–1896, Nov. 1989.
- [16] T. Blumensath and M. E. Davies, “On the difference between orthogonal matching pursuit and orthogonal least squares”, Tech. Rep., University of Edinburgh, Mar. 2007.
- [17] J. J. Kormylo and J. M. Mendel, “Maximum-likelihood detection and estimation of Bernoulli-Gaussian processes”, *IEEE Trans. Inf. Theory*, vol. 28, pp. 482–488, 1982.
- [18] J. M. Mendel, *Optimal Seismic Deconvolution*, Academic Press, New York, NY, 1983.
- [19] Y. Goussard, G. Demoment, and J. Idier, “A new algorithm for iterative deconvolution of sparse spike trains”, in *Proc. IEEE ICASSP*, Albuquerque, NM, Apr. 1990, pp. 1547–1550.
- [20] F. Champagnat, Y. Goussard, and J. Idier, “Unsupervised deconvolution of sparse spike trains using stochastic approximation”, *IEEE Trans. Signal Processing*, vol. 44, no. 12, pp. 2988–2998, Dec. 1996.
- [21] Q. Cheng, R. Chen, and T.-H. Li, “Simultaneous wavelet estimation and deconvolution of reflection seismic signals”, *IEEE Trans. Geosci. Remote Sensing*, vol. 34, pp. 377–384, Mar. 1996.
- [22] I. F. Gorodnitsky and B. D. Rao, “Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm”, *IEEE Trans. Signal Processing*, vol. 45, no. 3, pp. 600–616, Mar. 1997.
- [23] S. Chen and J. Wigger, “Fast orthogonal least squares algorithm for efficient subset model selection”, *IEEE Trans. Signal Processing*, vol. 43, no. 7, pp. 1713–1715, July 1995.
- [24] R. A. Horn and C. R. Johnson, *Matrix analysis*, Cambridge University Press, 1985.
- [25] S. J. Reeves, “An efficient implementation of the backward greedy algorithm for sparse signal reconstruction”, *IEEE Signal Processing Letters*, vol. 6, no. 10, pp. 266–268, Oct. 1999.
- [26] S. F. Cotter, J. Adler, B. D. Rao, and K. Kreutz-Delgado, “Forward sequential algorithms for best basis selection”, *IEE Proc. Vision, Image and Signal Processing*, vol. 146, no. 5, pp. 235–244, Oct. 1999.
- [27] D. Ge, J. Idier, and E. Le Carpentier, “Enhanced sampling schemes for MCMC based blind Bernoulli-Gaussian deconvolution”, Tech. Rep., Institut de Recherche en Communication et Cybernétique de Nantes, Sept. 2009.

- [28] P. E. Gill, G. H. Golub, W. Murray, and M. A. Saunders, “Methods for modifying matrix factorizations”, *Mathematics of Computation*, vol. 28, no. 126, pp. 505–535, Apr. 1974.
- [29] C. Y. Chi and J. M. Mendel, “Improved maximum-likelihood detection and estimation of Bernoulli-Gaussian processes”, *IEEE Trans. Inf. Theory*, vol. 30, pp. 429–435, Mar. 1984.
- [30] M. Allain and J. Idier, “Efficient binary reconstruction for non-destructive evaluation using gammagraphy”, *Inverse Problems*, vol. 23, no. 4, pp. 1371–1393, Aug. 2007.
- [31] M. Vetterli, P. Marziliano, and T. Blu, “Sampling signals with finite rate of innovation”, *IEEE Trans. Signal Processing*, vol. 50, no. 6, pp. 1417–1428, June 2002.
- [32] H.-J. Butt, B. Cappella, and M. Kappl, “Force measurements with the atomic force microscope: Technique, interpretation and applications”, *Surface Science Reports*, vol. 59, no. 1–6, pp. 1–152, Oct. 2005.
- [33] F. Gaboriaud, B. S. Parcha, M. L. Gee, J. A. Holden, and R. A. Strugnell, “Spatially resolved force spectroscopy of bacterial surfaces using force-volume imaging”, *Colloids and Surfaces B: Biointerfaces*, vol. 62, no. 2, pp. 206–213, Apr. 2008.
- [34] J. Duan, C. Soussen, D. Brie, and J. Idier, “A continuation approach to estimate a solution path of mixed L2-L0 minimization problems”, in *Signal Processing with Adaptive Sparse Structured Representations (SPARS workshop)*, Saint-Malo, France, Apr. 2009, pp. 1–6.
- [35] F. R. Gantmacher and M. G. Krein, *Oscillation matrices and kernels and small vibrations of mechanical systems*, AMS Chelsea Publishing, Providence, RI, revised edition, 2002.