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), stepwise regression algorithms and the exhaustive search. 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 sharing similarities with stepwise regression algorithms. Some structural properties of SBR are put forward. A fast and stable implementation is proposed. The approach is illustrated on two inverse problems involving highly correlated dictionaries. We show that SBR is very competitive with popular sparse algorithms in terms of trade-off between accuracy and computation time.

Copyright (c) 2011 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

This work was carried out in part while C. Soussen was visiting IRCCyN during the academic year 2010-2011 with the financial support of CNRS.

C. Soussen and D. Brie are with the Centre de Recherche en Automatique de Nancy (CRAN, UMR 7039, Nancy-University, CNRS). Campus Sciences, 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, F-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.

J. Duan was with CRAN. He now is with the Department of Biomedical Engineering and Biostatistics, Tulane University, 1440 Canal Street, Suite 2051, New Orleans, LA 70112, USA. Tel: (+1)-504 988 1341, Fax: (+1)-504 988 1706. E-mail: jduan@tulane.edu.

Index Terms

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

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 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. In sparse signal restoration, the choice of the dictionary is imposed by the inverse problem at hand whereas in sparse approximation, the dictionary has to be chosen according to its ability to represent the data with a limited number of coefficients.

Sparse signal restoration can be formulated as the minimization of the squared error $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ (where $\|\cdot\|$ refers to the Euclidean norm) under the constraint that the ℓ_0 pseudo-norm of $\mathbf{x}$, defined as the number of non-zero entries in $\mathbf{x}$, is small. This problem is often referred to as subset selection because it 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” situations 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. To address subset selection in a fast and sub-optimal manner, two approaches can be distinguished.

The first one, 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 [2–7]. In particular, the ℓ_1 relaxation of the ℓ_0 -norm has been increasingly investigated [2, 3], leading to the LASSO optimization problem.

The second approach addresses the *exact* subset selection problem using either iterative thresholding [8–11] 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) [12] and the improved version Orthogonal Matching Pursuit (OMP) [13]. Both are referred to as forward greedy algorithms since they start from the empty active set and then gradually increase it by one element. In contrast, the backward algorithm of Cuvreur and Bresler [14] starts from a complete active set which is gradually decreased by one element. It is, however, only valid for undercomplete dictionaries. Forward-backward algorithms (also

known as stepwise regression algorithms) in which insertions and removals of dictionary elements are both allowed, are known to yield better recovery performance since an early wrong selection can be counteracted by its further removal from the active set [15–18]. In contrast, the insertion of a wrong element is irreversible when using forward algorithms. We refer the reader to [18, Chapter 3] for an overview of the forward-backward algorithms in subset selection.

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 belonging to the first or the second approach provide solutions having the same support as the exhaustive search solution. Specifically, if the unknown signal 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 [3]. But when fast algorithms are unsatisfactory, it is relevant to consider slower algorithms being more accurate and remaining very fast compared to the exhaustive search. The Orthogonal Least Squares algorithm (OLS) [19] which is sometimes confused with OMP [20], falls into this category. Both OLS and OMP share the same structure, the difference being that at each iteration, OLS solves as many least square problems as there are non-active columns while OMP only performs one linear inversion. In this paper, we derive a forward-backward extension of OLS allowing an insertion or a removal per iteration, each iteration requiring to solve n least square problems, where n is the size of $\mathbf{x}$.

The proposed forward-backward extension of OLS can be viewed as a new member of the family of stepwise regression algorithms. The latter family traces back to 1960 [15], and other popular algorithms were proposed in the 1980's [18] and more recently [21]. Note that forward-backward extensions of OMP have also been proposed [22, 23]. In contrast with the other stepwise regression algorithms, our approach relies on a bi-objective formulation in order to handle the trade-off between low residual and low cardinality. This formulation reads as the minimization of the ℓ_0 -penalized least square cost function $\|\mathbf{y} - \mathbf{Ax}\|^2 + \lambda\|\mathbf{x}\|_0$. Then, we design a heuristic algorithm to minimize this cost function in a suboptimal way. While the other forward-backward strategies [15–17, 21, 22] aim at handling the same trade-off, most of them are not expressed as optimization algorithms, but rather as empirical schemes without any connexion with an objective function. Moreover, some of them involve discrete search parameters that control variable selection or de-selection [15, 16, 22] while others do not involve any parameter [17, 21]. An exception can be made for Broersen's algorithm [17] since it aims at minimizing $\|\mathbf{y} - \mathbf{Ax}\|^2 + \lambda\|\mathbf{x}\|_0$ for a specific λ value corresponding to Mallows' C_p statistic. However, it is only valid for undercomplete problems. On the contrary, our proposed algorithm is general and valid for any λ value. It does not necessitate to tune any other parameters (e.g., stopping parameters).

Our starting point is the Single Most Likely Replacement (SMLR) algorithm which proved to be a very efficient tool for the deconvolution of a Bernoulli-Gaussian signal [24–27]. We show that sparse signal restoration can be seen as a limit case of maximum *a posteriori* (MAP) Bernoulli-Gaussian restoration which results in an adaptation of SMLR to subset selection. 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 SMLR resulting in the so-called Single Best Replacement (SBR) algorithm. In Section IV, we propose a fast and stable SBR implementation. 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. SPARSE SIGNAL ESTIMATION USING A LIMIT BERNOULLI-GAUSSIAN MODEL

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 $\mathbf{y} \approx \mathbf{A}\mathbf{x}$. The columns $\mathbf{a}_i$ corresponding to the non-zero weights x_i 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 . $\mathbf{A}$ is assumed to satisfy the unique representation property (URP): any $\min(m, n)$ columns of $\mathbf{A}$ are linearly independent. This assumption is usual when $m \leq n$; it is stronger than the full rank assumption [28]. When $m \geq n$, it amounts to the full rank assumption. Although URP was originally introduced to guarantee uniqueness of sparse solutions [28], we use this assumption to propose a valid algorithm. It can actually be relaxed provided that the search strategy guarantees that the selected columns are linearly independent (see Section VI-C for details).

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 $x_i \neq 0$. We denote by $\mathcal{Q} \subseteq \{1, \dots, n\}$ the active set and by $\mathbf{q} \in \{0, 1\}^n$ the related vector defined by $q_i = 1$ if and only if $i \in \mathcal{Q}$. When $\text{Card}[\mathcal{Q}] \leq \min(m, n)$, let $\mathbf{A}_{\mathcal{Q}}$ be the submatrix of size $m \times \text{Card}[\mathcal{Q}]$ formed of the active columns of $\mathbf{A}$. We define the least square solution and the related 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. Bayesian formulation of sparse signal restoration

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 [24, 25, 27]. For such model, deterministic optimization algorithms [27] and Markov chain Monte Carlo techniques [29] are used to compute the MAP and the posterior mean, respectively. Hereafter, we define the BG model and then consider its estimation in the joint MAP sense.

A BG process can be defined using a Bernoulli random vector $\mathbf{q} \in \{0, 1\}^n$ coding for the support and a Gaussian random vector $\mathbf{r} \sim \mathcal{N}(\mathbf{0}, \sigma_x^2 \mathbf{I}_n)$, with $\mathbf{I}_n$ the identity matrix of size n . Each sample x_i of $\mathbf{x}$ is modeled as $x_i = q_i r_i$ [24, 25]. The Bernoulli parameter $\rho = \Pr(q_i = 1)$ is the probability of presence of signal and σ_x^2 controls the variance of the nonzero amplitudes $x_i = r_i$. The Bayesian formulation consists in inferring $\mathbf{x} = (\mathbf{q}, \mathbf{r})$ knowing $\mathbf{y}$. The MAP estimator can be obtained by maximizing the marginal likelihood $l(\mathbf{q} | \mathbf{y})$ [27] or the joint likelihood $l(\mathbf{q}, \mathbf{r} | \mathbf{y})$ [25, 26]. Following [25] and assuming a Gaussian white noise $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I}_m)$, independent from $\mathbf{x}$, Bayes' rule leads to:

$$\begin{aligned} \mathcal{L}(\mathbf{q}, \mathbf{r}) &\triangleq -2\sigma_n^2 \log[l(\mathbf{q}, \mathbf{r} | \mathbf{y})] \\ &= \|\mathbf{y} - \mathbf{A}\Delta_{\mathbf{q}}\mathbf{r}\|^2 + \frac{\sigma_n^2}{\sigma_x^2} \|\mathbf{r}\|^2 + \lambda \|\mathbf{q}\|_0 + c \end{aligned} \quad (3)$$

where $\lambda = 2\sigma_n^2 \log(1/\rho - 1)$, $\Delta_{\mathbf{q}}$ is the diagonal matrix of size n whose diagonal elements are q_i ($\mathbf{x} = \Delta_{\mathbf{q}}\mathbf{r}$), and c is a constant.

Now, 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 the non zero amplitudes, we choose to use a limit Bernoulli-Gaussian model in which the amplitude variance σ_x^2 is set to infinity. Note that a parallel limit development was done, independently from our work, in the conference paper [23]. In Appendix A, we show that the minimization of $\mathcal{L}$ w.r.t. $\mathbf{x} = (\mathbf{q}, \mathbf{r})$ rereads:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \{\mathcal{J}(\mathbf{x}; \lambda) = \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \lambda \|\mathbf{x}\|_0\}. \quad (4)$$

This formulation is close to that obtained in the Bayesian subset selection literature [18, Chapter 7] using an alternative Bernoulli-Gaussian model. In the latter model, the Gaussian prior relies on $\mathbf{R}_{\mathbf{Q}}\mathbf{r}$ instead of $\mathbf{r}$, with $\mathbf{R}_{\mathbf{Q}}$ the Cholesky factor of the Gram matrix $\mathbf{A}_{\mathbf{Q}}^t \mathbf{A}_{\mathbf{Q}}$. This leads to a cost function of the form (4), the difference being that λ depends on the amplitude variance σ_x^2 and tends to infinity as σ_x^2 tends to infinity [30, 31].

Remark 1 (Noise-free case) *The Bayesian development above is valid for noisy data. In the noise-free case, we define the sparse solution as the limit of $\arg \min_{\mathbf{x}} \mathcal{J}(\mathbf{x}; \lambda)$ when λ tends towards 0. According*

to classical results in optimization [32, Chapter 17], if $\{\lambda_k\}$ is a sequence decreasing towards 0 and $\mathbf{x}_k$ is an exact global minimizer of $\mathcal{J}(\mathbf{x}; \lambda_k)$, then every limit point of the sequence $\{\mathbf{x}_k\}$ is a solution of $\arg \min_{\mathbf{x}} \|\mathbf{x}\|_0$ s.t. $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ is minimal. In Appendix B, we derive a more precise result: “the set of minimizers of $\mathcal{J}(\mathbf{x}; \lambda)$ is constant when λ is close enough to 0 ($\lambda \neq 0$). It is equal to the set of sparsest solutions to $\mathbf{y} = \mathbf{A}\mathbf{x}$ in the overcomplete case, and to the unconstrained least-squares solution in the undercomplete case.”

In the following, we focus on the minimization problem (4). The hyperparameter λ is fixed. It controls the level of sparsity of the desired solution. The algorithm that will be developed relies on an efficient search of the support of $\mathbf{x}$. The search strategy is based on the definition of a neighborhood relationship between two supports: two supports are neighbors if one is nested inside the other and the largest support has one more element.

III. SINGLE BEST REPLACEMENT ALGORITHM

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 (4). To clearly distinguish SMLR which specifically aims at minimizing (3), the adapted algorithm will be termed as Single Best Replacement (SBR).

A. Principle of SMLR and main notations

SMLR [24] is a deterministic coordinatewise ascent algorithm to maximize likelihood functions of the form $l(\mathbf{q} | \mathbf{y})$ (marginal MAP estimation) or $l(\mathbf{q}, \mathbf{r} | \mathbf{y})$ (joint MAP estimation). In the latter case, it is easy to check from (3) that given $\mathbf{q}$, the minimizer of $\mathcal{L}(\mathbf{q}, \mathbf{r})$ w.r.t. $\mathbf{r}$ has a closed form expression $\mathbf{r} = \mathbf{r}(\mathbf{q})$. Consequently, the joint MAP estimation reduces to the minimization of $\mathcal{L}(\mathbf{q}, \mathbf{r}(\mathbf{q}))$ 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 decrease of $\mathcal{L}(\mathbf{q}, \mathbf{r}(\mathbf{q}))$ is chosen. This task is repeated until no single replacement can decrease $\mathcal{L}(\mathbf{q}, \mathbf{r}(\mathbf{q}))$ anymore. The number of possible supports $\mathbf{q}$ being finite and SMLR being a descent 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., an insertion or removal into/from the active set $\mathcal{Q}$:

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

When $\text{Card}[\mathcal{Q}] \leq \min(m, n)$, we define the cost function

$$\mathcal{J}_{\mathcal{Q}}(\lambda) \triangleq \mathcal{E}_{\mathcal{Q}} + \lambda \text{Card}[\mathcal{Q}] \quad (5)$$

involving the squared error $\mathcal{E}_{\mathcal{Q}}$ defined in (2). By definition of $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda) = \mathcal{E}_{\mathcal{Q}} + \lambda \|\mathbf{x}_{\mathcal{Q}}\|_0$, $\mathcal{J}_{\mathcal{Q}}(\lambda)$ coincides with $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda)$ when the support of $\mathbf{x}_{\mathcal{Q}}$ is equal to $\mathcal{Q}$.

Although it aims at minimizing $\mathcal{J}(\mathbf{x}; \lambda)$, the proposed SBR algorithm involves the computation of $\mathcal{J}_{\mathcal{Q}}(\lambda)$ rather than $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda)$. We make this choice because $\mathcal{J}_{\mathcal{Q}}(\lambda)$ can be computed and updated more efficiently, the computation of $\mathbf{x}_{\mathcal{Q}}$ being no longer necessary. In subsection III-C, we show that for noisy data, the replacement of $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda)$ by $\mathcal{J}_{\mathcal{Q}}(\lambda)$ has a negligible effect.

B. The Single Best Replacement algorithm

SMLR can be seen as an exploration strategy for discrete optimization rather than an algorithm specific to a posterior likelihood function. Here, we use this strategy to minimize $\mathcal{J}(\mathbf{x}; \lambda)$. We rename the algorithm Single Best Replacement to remove any statistical connotation.

SBR works as follows. Consider the current support $\mathcal{Q}$. The n single replacements $\mathcal{Q} \bullet i$ are tested, i.e., we compute the squared errors $\mathcal{E}_{\mathcal{Q} \bullet i}$ and we memorize the values of $\mathcal{J}_{\mathcal{Q} \bullet i}(\lambda)$. If the minimum of $\mathcal{J}_{\mathcal{Q} \bullet i}(\lambda)$ is lower than $\mathcal{J}_{\mathcal{Q}}(\lambda)$, then we select the index yielding this minimum value:

$$\ell \in \arg \min_{i \in \{1, \dots, n\}} \mathcal{J}_{\mathcal{Q} \bullet i}(\lambda). \quad (6)$$

The next SBR iterate is thus defined as $\mathcal{Q}' = \mathcal{Q} \bullet \ell$. This task is repeated until $\mathcal{J}_{\mathcal{Q}}(\lambda)$ cannot decrease anymore. By default, we use the initial empty support. The algorithm is summarized in Table I.

C. Case where some active amplitudes are zero

We show that this case almost surely never arises when the data $\mathbf{y}$ are corrupted with “non degenerate” noise.

Theorem 1 *Let $\mathbf{y} = \mathbf{y}_0 + \mathbf{n}$ where $\mathbf{y}_0 \in \mathbb{R}^m$ is fixed and $\mathbf{n}$ 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.*

Proof: Let $k = \text{Card}[\mathcal{Q}]$ and $\mathbf{t}_{\mathcal{Q}}$ be the minimizer of $\|\mathbf{y} - \mathbf{A}_{\mathcal{Q}}\mathbf{t}\|^2$ over $\mathbb{R}^k$. $\mathbf{t}_{\mathcal{Q}}$ reads $\mathbf{t}_{\mathcal{Q}} = \mathbf{V}_{\mathcal{Q}}\mathbf{y}$ where matrix $\mathbf{V}_{\mathcal{Q}} = (\mathbf{A}_{\mathcal{Q}}^t \mathbf{A}_{\mathcal{Q}})^{-1} \mathbf{A}_{\mathcal{Q}}^t$ is of size $k \times m$, and $\|\mathbf{x}_{\mathcal{Q}}\|_0 = \|\mathbf{t}_{\mathcal{Q}}\|_0 \leq k$. 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

TABLE I
SBR ALGORITHM. BY DEFAULT, $\mathcal{Q}_1 = \emptyset$.

Input: $\mathbf{A}$, $\mathbf{y}$, λ and support $\mathcal{Q}_1$ ($\text{Card}[\mathcal{Q}_1] \leq \min(m, n)$)
Step 1: Set $j = 1$.
Step 2: For $i \in \{1, \dots, n\}$, compute $\mathcal{J}_{\mathcal{Q}_j \bullet i}(\lambda)$.
Compute ℓ using (6).
If $\mathcal{J}_{\mathcal{Q}_j \bullet \ell}(\lambda) < \mathcal{J}_{\mathcal{Q}_j}(\lambda)$,
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: support $\mathcal{Q}_j = \text{SBR}(\mathcal{Q}_1; \lambda)$

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

$$\|\mathbf{x}_{\mathcal{Q}}\|_0 < k \iff \mathbf{y} \in \bigcup_{i=1}^k \mathcal{H}^\perp(\mathbf{v}^i). \quad (7)$$

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 (7) is zero. ■

Theorem 1 implies that when dealing with real noisy data, it is almost sure that all active coefficients x_i are non-zero. Hence, each SBR iterate $\mathcal{Q}$ almost surely satisfies $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda) = \mathcal{J}_{\mathcal{Q}}(\lambda)$. In any case, SBR can be applied without restriction and the properties stated below (e.g., termination after a finite number of iterations) remain valid when an SBR iterate satisfies $\|\mathbf{x}_{\mathcal{Q}}\|_0 < \text{Card}[\mathcal{Q}]$.

D. Properties of SBR

Proposition 1 *Under the assumptions of Theorem 1, each SBR iterate $\mathbf{x}_{\mathcal{Q}}$ is almost surely a local minimizer of $\mathcal{J}(\mathbf{x}; \lambda)$. In particular, the SBR output satisfies this property.*

Proof: Let $\mathbf{x} = \mathbf{x}_{\mathcal{Q}}$ be an SBR iterate. According to Theorem 1, the support $\mathcal{S}(\mathbf{x}) = \mathcal{Q}$ almost surely. Setting $\varepsilon = \min_{i \in \mathcal{Q}} |x_i| > 0$, it is easy to check that if $\mathbf{x}' \in \mathbb{R}^n$ satisfies $\|\mathbf{x}' - \mathbf{x}\| < \varepsilon$, then $\mathcal{S}(\mathbf{x}') \supseteq \mathcal{S}(\mathbf{x}) = \mathcal{Q}$, thus $\|\mathbf{x}'\|_0 \geq \|\mathbf{x}\|_0$. Assume that $\mathbf{x}'$ satisfies $\|\mathbf{x}' - \mathbf{x}\| < \varepsilon$.

- If $\mathcal{S}(\mathbf{x}') = \mathcal{Q}$, then, by definition of $\mathbf{x} = \mathbf{x}_{\mathcal{Q}}$, we have $\mathcal{E}(\mathbf{x}') \geq \mathcal{E}(\mathbf{x})$. Thus, $\mathcal{J}(\mathbf{x}'; \lambda) \geq \mathcal{J}(\mathbf{x}; \lambda)$.

- Otherwise, $\mathcal{J}(\mathbf{x}'; \lambda) = \mathcal{E}(\mathbf{x}') + \lambda \|\mathbf{x}'\|_0 \geq \mathcal{E}(\mathbf{x}') + \lambda(\|\mathbf{x}\|_0 + 1)$. 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 $\mathbf{x}' \in \mathcal{V}(\mathbf{x})$, $\|\mathbf{x}' - \mathbf{x}\| < \varepsilon$ and $\mathcal{S}(\mathbf{x}') \supset \mathcal{Q}$, then $\mathcal{J}(\mathbf{x}'; \lambda) > \mathcal{E}(\mathbf{x}) + \lambda \|\mathbf{x}\|_0 = \mathcal{J}(\mathbf{x}; \lambda)$.

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

Termination: Because SBR is a descent algorithm, a support $\mathcal{Q}$ cannot be explored twice and SBR terminates after a finite number of iterations. We emphasize that no stopping condition is needed unlike many algorithms which require to set a maximum number of iterations and/or a threshold on the squared error variation (CoSaMP, Subspace Pursuit, Iterative Hard Thresholding, Iterative Reweighted ℓ_1).

OLS as a special case: When $\lambda = 0$, SBR coincides with the well known OLS algorithm [19, 33]. The removal operation never occurs because it yields an increase of the squared error $\mathcal{J}_{\mathcal{Q}}(0) = \mathcal{E}_{\mathcal{Q}}$.

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

Remark 2 *SBR($\emptyset; \lambda$) yields the empty set if and only if $\lambda \geq \lambda_{\max} \triangleq \max_i (\langle \mathbf{a}_i, \mathbf{y} \rangle^2 / \|\mathbf{a}_i\|^2)$.*

This result directly follows from checking that any insertion trial fails, *i.e.*, $\forall i, \mathcal{E}_{\{i\}} + \lambda \geq \mathcal{E}_{\emptyset}$. It 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).

Relation between SBR and SMLR: The main difference between both algorithms is that SMLR 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 computes the inverse of $\mathbf{A}_{\mathcal{Q}}^t \mathbf{A}_{\mathcal{Q}}$. In the case of SMLR, the term $\alpha \mathbf{I}_{\text{Card}[\mathcal{Q}]}$ acts as a regularization on the amplitude values. It avoids instabilities when $\mathbf{A}_{\mathcal{Q}}$ is ill conditioned at the price of handling the additional hyperparameter α . On the contrary, instabilities may occur while using SBR. In the next section, we focus on this issue and propose a stable implementation.

IV. IMPLEMENTATION ISSUES

Given the current support $\mathcal{Q}$, 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 $\mathcal{J}_{\mathcal{Q}'}(\lambda) = \mathcal{E}_{\mathcal{Q}'} + \lambda \text{Card}[\mathcal{Q}']$. Our implementation is inspired by the fast implementation of the homotopy algorithm for ℓ_1 regression [3, 34]. It consists in maintaining the Cholesky factorization of the Gram matrix $\mathbf{G}_{\mathcal{Q}} \triangleq \mathbf{A}_{\mathcal{Q}}^t \mathbf{A}_{\mathcal{Q}}$ when $\mathcal{Q}$ is modified by one element. The Cholesky factorization takes the form $\mathbf{G}_{\mathcal{Q}} = \mathbf{L}_{\mathcal{Q}} \mathbf{L}_{\mathcal{Q}}^t$ where $\mathbf{L}_{\mathcal{Q}}$ is a lower triangular matrix of size $k = \text{Card}[\mathcal{Q}]$. Also, $\mathbf{L}_{\mathcal{Q}}$ is better conditioned than $\mathbf{G}_{\mathcal{Q}}$, improving the stability of matrix inversion. We now give the main updating equations. Full detailed derivation can be found in Appendix C.

A. Efficient strategy based on the Cholesky factorization

The replacement tests only rely on the current matrix L_Q and do not require its update.

1) *Single replacement tests*: An insertion test $Q' = Q \cup \{i\}$ takes the form:

$$\mathcal{J}_{Q'}(\lambda) - \mathcal{J}_Q(\lambda) = \lambda - \frac{(l_{Q,i}^t L_Q^{-1} A_Q^t \mathbf{y} - \mathbf{a}_i^t \mathbf{y})^2}{\|\mathbf{a}_i\|^2 - \|l_{Q,i}\|^2} \quad (8)$$

with $l_{Q,i} = L_Q^{-1} A_Q^t \mathbf{a}_i$. This computation mainly requires a triangular system inversion (computation of $l_{Q,i}$ in $\mathcal{O}(k^2)$ elementary operations) up to the pre-computation of $L_Q^{-1}(A_Q^t \mathbf{y})$ at the beginning of the current SBR iteration.

According to [18, 35], a removal test $Q' = Q \setminus \{i\}$ reads $\mathcal{J}_{Q'}(\lambda) - \mathcal{J}_Q(\lambda) = x_Q(i)^2 / \gamma_i - \lambda$ where $x_Q(i)$ is the i th element in vector $\mathbf{x}_Q$ and γ_i is the diagonal element of $\mathbf{G}_Q^{-1}$ corresponding to the position of $\mathbf{a}_i$ in $\mathbf{A}_Q$. The overall removal tests mainly amount to the inversion of the triangular matrix L_Q (in $\mathcal{O}(k^3)$ operations) as the computation of γ_i for all i and of $\mathbf{G}_Q^{-1} A_Q^t \mathbf{y}$ (i.e., the values of $x_Q(i)$) from L_Q^{-1} are both in $\mathcal{O}(k^2)$.

Note that insertion and removal tests can be easily done in parallel. In Matlab, this parallel implementation leads to a significant save of computation time due to the SIMD capabilities of Matlab.

2) *Updating the Cholesky factorization*: The update of L_Q can be easily done in the insertion case by adding the new column $\mathbf{a}_i$ at the last position in $\mathbf{A}_{Q \cup \{i\}}$. The new matrix $L_{Q'}$ is a 2×2 block matrix whose upper left block is L_Q (see Appendix C). The removal case requires more care since a removal breaks the triangular structure of L_Q . The update can be done by performing either a series of Givens planar rotations [21] or a positive rank 1 Cholesky update [36]. We describe the latter strategy in Appendix C. The Cholesky factorization update is in $\mathcal{O}(k^2)$ in the insertion case and in $\mathcal{O}((k - I)^2)$ in the removal case where I denotes the position of the column to be removed in $\mathbf{A}_Q$.

B. Reduced search

Additionally, we propose an acceleration of SBR yielding the same iterates with a reduced search. We notice that a column 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 $\mathcal{J}_Q(\lambda)$ which can be expected is λ . The acceleration of SBR consists in testing insertions first. If any insertion leads to $\mathcal{J}_Q(\lambda) - \mathcal{J}_{Q'}(\lambda) > \lambda$, then removals are not worth being tested. Otherwise, the removals have to be tested as stated in Table I. We have implemented this acceleration systematically.

C. Memory requirements and computation burden

The actual implementation may vary depending on the size and the structure of matrix $\mathbf{A}$. We briefly describe the main possible implementations.

When the size of $\mathbf{A}$ is relatively small, the computation and storage of the Gram matrix $\mathbf{A}^t \mathbf{A}$ prior to any SBR iteration (storage of n^2 scalar elements) avoids to recompute the vectors $\mathbf{A}_Q^t \mathbf{a}_i$ which are needed when the insertion of $\mathbf{a}_i$ into the active set is tested. The storage of the other quantities (mainly $\mathbf{L}_Q$) that are being updated amounts to $\mathcal{O}(k^2)$ scalar elements and a replacement test costs $\mathcal{O}(k^2)$ elementary operations in average.

When $\mathbf{A}$ is larger, the storage of $\mathbf{A}^t \mathbf{A}$ is no longer possible, thus $\mathbf{A}_Q^t \mathbf{a}_i$ must be recomputed for any SBR iteration. This computation costs km elementary operations and now represents the most important part of an insertion test. When the dictionary has some specific structure, this limitation can be alleviated, enabling a fast implementation even for large n . 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 implemented¹. In the following, we analyze the behavior of SBR for two difficult problems involving highly correlated dictionaries: 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 [27]. It affords us to study the ability of SBR to perform an exact recovery in a simple noise-free case (separation of two Gaussian signals) and to test SBR in a noisy case (estimation of a larger number of Gaussians) and compare it with other algorithms. For simulated problems, we denote by $\mathbf{x}^*$ the exact sparse signal, the data reading $\mathbf{y} = \mathbf{A}\mathbf{x}^* + \mathbf{n}$. The dictionary columns 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 $P_y = \|\mathbf{A}\mathbf{x}^*\|^2/m$ is the average power of the noise-free data and P_n is the variance of the noise process $\mathbf{n}$.

¹Matlab codes provided by the authors can be downloaded at <http://ieeexplore.org>. In our Matlab implementation, the insertion and removal tests are done in parallel.

TABLE II

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 OBTAINED FOR 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*

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 for fast implementation even for large size problems (see subsection IV-C). The deconvolution problem leads to a Toeplitz matrix $\mathbf{A}$ whose columns 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 ($m > n$).

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 analyze the SBR outputs for decreasing λ -values by computing their cardinality and testing whether they coincide with the true support $\mathcal{S}(\mathbf{x}^*)$. Table II shows the results obtained for a problem of size 300×270 ($\sigma = 5$) with distances equal to $d = 20, 13$, and 6 samples. It is noticeable that the exact recovery always occurs 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 output support first increases while λ decreases, and then removals start to occur, enabling the exact recovery for lower λ 's.

C. Behavior of SBR for noisy data

We consider a more realistic simulation in which the data are of larger size ($m = 3000$ samples) and noisy. The impulse response $\mathbf{h}$ is of size 301 ($\sigma = 50$) yielding a matrix $\mathbf{A}$ of size 3000×2700 , and the SNR is set to 20 dB. Fig. 1(a) displays the generated data. The unknown sparse signal $\mathbf{x}^*$ is

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, 12, and 18, respectively. The computation time always remains below 3 seconds (Matlab implementation).

composed of 17 spikes that are uniformly located in $\{1, \dots, n\}$. The non-zero amplitudes x_i^* are drawn according to an i.i.d. Laplacian distribution. Let us remark that the limit Bernoulli-Gaussian model is not a proper probabilistic model so that one cannot use it to design simulated data. We choose a Laplacian distribution since the non-zero amplitudes are more heterogeneous than with a Gaussian distribution with finite variance.

On Fig. 1(b-d), we display the SBR results for three λ -values. For large λ 's, only the main Gaussian features are found. When λ decreases, the smaller features are being recovered together with spurious features. Removals occur for $\lambda \leq 0.8$ yielding approximations that are more accurate than those obtained with OLS and for the same cardinality (the residual $\|y - Ax\|^2$ is lower) while when $\lambda > 0.8$, the SBR output coincides with the OLS solution of same cardinality. Note that the theoretical value of λ obtained from (3) is equal to 0.3 yielding a support of cardinality 18. The residual is slightly lower than that obtained with $\lambda = 0.5$. The exact support of x^* is never found because the data are noisy

Fig. 2. Comparison of sparse algorithms in terms of trade-off between accuracy ($\mathcal{J}(\mathbf{x}; \lambda)$) and CPU time for the deconvolution problem of Fig. 1. SBR($\lambda = 0.5$) is run first yielding a support of cardinality $k_{\text{sbr}}=18$. Then, we run OLS(k_{sbr}), OMP(k_{sbr}), homotopy for ℓ_1 regression [39], and $\text{IR}\ell_1(\lambda)$ [40]. The ℓ_1 result is the homotopy iterate of cardinality k_{sbr} yielding the least value of $\mathcal{J}(\mathbf{x}; \lambda)$.

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 an existing spike x_i forwards or backwards [37] or update a block of neighboring amplitudes jointly (e.g., x_i and x_{i+1}) [38]. Various search strategies are also reported in [18, Chapter 3].

D. Comparison of SBR with other sparse algorithms

We compared SBR with classical and recent sparse algorithms: OMP, OLS, CoSaMP [8], Subspace Pursuit [9], Iterative Hard Thresholding (IHT) [10, 11], ℓ_1 regression [3] and Iterative Reweighted ℓ_1 ($\text{IR}\ell_1$) [5, 40]. A general trend is that thresholding algorithms perform poorly when the dictionary columns are strongly correlated. CoSaMP and Subspace Pursuit yield the worst results: they stop after a very few iterations as the squared error increases from one iteration to the next. On the contrary, IHT guarantees that the squared error decreases but the convergence is very slow and the results remain poor in comparison with SBR. In the simulation of Fig. 1(c), SBR performs 12 iterations (only insertions) leading to a support of cardinality 12. Meanwhile, the number of iterations of IHT before convergence is huge: both versions of IHT presented in [10] require at least 10,000 iterations to converge, leading to an overall computation time (22 and 384 seconds) that is much larger than the SBR computation time (3 seconds).

Fig. 2 is a synthetic view of the performance of SBR, OLS, OMP, ℓ_1 regression, and $\text{IR}\ell_1$ for a given sparsity level λ . The computation time and the value of $\mathcal{J}(\mathbf{x}; \lambda)$ are shown on the horizontal and

vertical axes, respectively. This enables us to define several categories of algorithms depending on their locations on the 2D plane: the outputs of fast algorithms (OMP and ℓ_1) lay in the upper left region whereas slower but more efficient algorithms (OLS, SBR, and $\text{IR}\ell_1$) yield points laying in the lower right region. We chose not to represent the outputs of thresholding algorithms since they yield poorer performance, *i.e.*, points located either in the upper right (IHT) or upper left (CoSaMP, Subspace Pursuit) regions. In details, we observed that ℓ_1 regression tends to overestimate the support cardinality and to place several spikes at very close locations. We used Donoho's homotopy implementation [3, 39] and found that it requires many iterations: homotopy runs during 200 iterations before reaching a support of cardinality 18 when processing the data of Fig. 1 (we recall that homotopy starts from the empty set and performs a single support replacement per iteration). The performance of ℓ_1 regression fluctuates around that of OMP depending on the trials and the sparsity level. Regarding $\text{IR}\ell_1$, we used the Adaptive LASSO implementation from Zou [40] since it is dedicated to the minimization of $\mathcal{J}(\mathbf{x}; \lambda)$. We stopped the algorithm when two successive ℓ_1 iterates share the same support. For the simulation of Fig. 1, $\text{IR}\ell_1$ and SBR yield comparable results in that one algorithm does not outperform the other for all λ values, but $\text{IR}\ell_1$ generally performs slightly better (Fig. 2). We designed other simulations in which the nonzero weights x_i^* are spread over a wider interval. In this case, SBR most often yields the best approximations.

Fig. 2 is representative of the empirical results obtained while performing many trials. Obviously, the figure may significantly change depending on several factors among which the λ -value and the tuning parameters of $\text{IR}\ell_1$. The goal is definitely not to conclude that an algorithm *always* outperforms the others but rather to sketch a classification of groups of algorithms according to the trade-off between accuracy and computation time.

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 [41, 42]. We process both simulated and real data and compare the performance of SBR with respect to OMP, Bayesian OMP (BOMP) which is an OMP based forward-backward algorithm [23], OLS, ℓ_1 regression [3], and $\text{IR}\ell_1$ [5, 7, 40]. Firstly, we formulate the detection of discontinuities at a single order as a spline approximation problem. Then, we take advantage of this formulation to introduce the joint detection problem.

Fig. 3. 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 $\{i, \dots, m\}$.

A. Approximation of a spline of degree p

Following [41], we introduce the dictionary $\mathbf{A}^p$ of size $m \times (m - p)$ formed of shifted versions of the one-sided power function $k \mapsto [\max(k, 0)]^p$ for all possible shifts (see Fig. 3) and we address the sparse approximation of $\mathbf{y}$ by the piecewise polynomial $\mathbf{A}^p \mathbf{x}^p$ (actually, we impose as initial condition that the spline function is equal to 0 for $k \leq 0$). It 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. Piecewise polynomial approximation

We formulate the joint detection of discontinuities of orders $p = 0, \dots, P$ by appending the elementary dictionaries $\mathbf{A}^p$ in a global dictionary $\mathbf{A} = [\mathbf{A}^0, \dots, \mathbf{A}^P]$. The product $\mathbf{A} \mathbf{x}$ yields a 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 $P \geq 1$. Moreover, any column $\mathbf{a}_i^p$ of $\mathbf{A}^p$ overlaps *all* other columns $\mathbf{a}_j^q$ because their respective supports are the intervals $\{i, \dots, m\}$ and $\{j, \dots, m\}$. 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 , greedy algorithms start to position a first (wrong) discontinuity in between i and j , and forward greedy algorithms cannot remove it.

C. Adaptation of SBR

The above defined dictionary does not satisfy the unique representation property. Indeed, it is easy to check that the difference between two discrete ramps at locations i and $i+1$ yields the discrete Heaviside function at location i : $\mathbf{a}_i^1 - \mathbf{a}_{i+1}^1 = \mathbf{a}_i^0$. We thus need to slightly modify SBR in order to ensure that only full rank matrices $\mathbf{A}_Q$ are explored. The modification is based on the following proposition which gives a sufficient condition for full rankness of $\mathbf{A}_Q$.

Proposition 2 *Let n_i denote the number of columns $\mathbf{a}_i^p$, $p \in \{0, \dots, P\}$ which are active for sample i . Let us define the binary condition $\mathcal{C}(i)$:*

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

If Q satisfies $\forall i, \mathcal{C}(i) = 1$, then $\mathbf{A}_Q$ is full rank.

Proposition 2 is proved in Appendix D. 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 SBR adaptation consists in testing an insertion only when the new support $Q' = Q \cup \{(i, p)\}$ satisfies the above condition.

D. Numerical simulations

We first set $P = 1$ leading to the piecewise affine approximation problem. The noise-free data $\mathbf{y} = \mathbf{A}\mathbf{x}^*$ of Fig. 4(a) are of size $m = 1000$ with $\|\mathbf{x}^*\|_0 = 18$ discontinuities. According to Remark 2, we compute the value $\lambda_{\max}$ above which the SBR output is the empty set, and we run SBR with $\lambda_j = \lambda_{\max} 10^{-j/2}$ for $j = 0, \dots, 20$. For the least λ -value, SBR yields an exact recovery (see Fig. 4(a)). For comparison purpose, we also run 27 iterations of OMP and OLS. The “ ℓ_2 - ℓ_0 ” curves represented on Fig. 4(b) express the squared residual $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ versus the cardinality $\|\mathbf{x}\|_0$ for each algorithm (we plot the first 27 iterates of OMP and OLS and for all j , we plot the output of $\text{SBR}(\lambda_j)$ after full convergence of SBR). Whatever the cardinality, SBR yields the least residual. For noisy data, the “ ℓ_2 - ℓ_0 ” curve corresponding to SBR still lays below the OMP and OLS curves for most cardinalities. In the next paragraph, we also consider the Bayesian OMP, ℓ_1 regression, and $\text{IR}\ell_1$ algorithms for further comparisons.

(a) Noise-free data and SBR approximation

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

(c) Noisy data (SNR = 20 dB) and SBR approximation

(d) “ ℓ_2 - ℓ_0 ” curves (noisy data)

Fig. 4. Joint detection of discontinuities of orders 0 and 1. The dictionary is of size 1000×1999 and the data signal 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 y and its approximation Ax for the least λ -value. (a) Approximation from noise-free data. The recovery is exact. (b) “ ℓ_2 - ℓ_0 ” curves showing the squared residual versus the cardinality for the SBR, OLS, and OMP solutions. (c,d) Similar results for noisy data (SNR = 20 dB).

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. Specifically, 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.*) [43].

The data displayed on Fig. 5(a) are related to a bacterial cell *Shewanella putrefaciens* laying in aqueous solution, interacting with the tip of the AFM probe [44]. A retraction force curve is recorded by positioning the tip in contact with the bacterial cell, and then gradually retracting the tip from the sample until it

Fig. 5. Joint detection of discontinuities of orders 0, 1, and 2 (problem of size 2167×6498). (a) Experimental AFM data showing the force evolution versus the probe-sample distance z . (b) Squared residual versus cardinality for the SBR, OLS, and OMP solutions. (c) Time of reconstruction versus cardinality.

loses contact. In the retraction curve shown on Fig. 5(a), three regions of interest can be distinguished from right to left. 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 tip retraction. 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. Similar to the processing of simulated data, we run SBR with 14 λ -values and we run OLS and OMP until iteration 41. For each algorithm, we plot the “ ℓ_2 - ℓ_0 ” curve and the curve displaying the time of reconstruction versus the cardinality (Figs. 5(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.

Fig. 6 displays the approximations yielded by the three algorithms together with the BOMP, ℓ_1 , and $\text{IR}\ell_1$ approximations. For the largest value λ_1 , SBR runs during 6 iterations (4 insertions and 2 removals) yielding a support of cardinality 2. SBR performs better than other algorithms (Figs. 6(a-f)). Although

Fig. 6. 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. (a) SBR output of cardinality 2: 4 insertions and 2 removals have been done ($\lambda_1 = 120$). (b-f) OLS and OMP outputs after 2 iterations, BOMP and $\text{IR}\ell_1$ [40] outputs for $\lambda = \lambda_1$, homotopy iterate (LASSO) leading to the minimal value of $\mathcal{J}(\mathbf{x}; \lambda_1)$. (g-l) Same simulation with a lower λ -value ($\lambda_2 = 8.5$). The SBR output is of cardinality 5 (7 insertions and 2 removals).

$\text{IR}\ell_1$ yields the most accurate approximation, it relies on 4 dictionary columns leading to a larger value of $\mathcal{J}(\mathbf{x}; \lambda_1)$. We observed the same behavior for the lowest value λ_2 (subfigures (g-l)). Again, SBR yields the least value of $\mathcal{J}(\mathbf{x}; \lambda_2)$ among all algorithms. Moreover, SBR provides a very precise localization of both first order discontinuities (subfigure (a)) which are crucial information for the physical interpretation of the data. On the contrary, all other algorithms fail for the highest sparsity level, and some do not even succeed for the lowest. Specifically, OLS accurately locates both first order discontinuities when 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. LASSO and BOMP yield very poor approximations for the

highest sparsity level and approximations with many dictionary columns for the lowest sparsity level. In terms of value of the cost function $\mathcal{J}(\mathbf{x}; \lambda)$, BOMP and LASSO fluctuate around OMP but they are far outperformed by OLS, SBR, and $\text{IR}\ell_1$.

VII. CONCLUSION

A. Discussion

We performed comparisons for two problems involving highly correlated dictionary columns. SBR is at least as accurate as OLS and sometimes more accurate, with a slightly larger cost of computation. We also considered sparse algorithms that are slower than OLS. SBR was found to be very competitive in terms of trade-off between accuracy and computation time. Although OLS based forward-backward algorithms yield a relatively large computational cost per iteration, we have noticed that for correlated dictionaries, the number of SBR iterations (*i.e.*, of elementary modifications of the support) is much lower than the number of support modifications performed by several other algorithms. Typically, IHT and $\text{IR}\ell_1$ can often be more expensive than SBR. Additionally, SBR terminates within a finite number of iterations, thus it does not require to tune any empirical stopping parameter. The limitation of SBR in terms of speed arises when the dictionary $\mathbf{A}$ is unstructured and the size of $\mathbf{A}$ is too large to store $\mathbf{A}^t \mathbf{A}$. The inner products $\mathbf{a}_i^t \mathbf{a}_j$ must then be recomputed for each iteration, which is relatively burdensome.

In the recent literature, it is often acknowledged that the cost function $\mathcal{J}(\mathbf{x}; \lambda)$ has a large number of local minimizers therefore discouraging its direct optimization [5, 7]. Many authors thus choose to minimize an approximate cost function in which the ℓ_0 norm $|\mathbf{x}_i|_0$ is replaced with a nonconvex continuous function $\varphi(x_i)$. However, when the range of values of the (expected) nonzero amplitudes $x_i \neq 0$ is wide, it is difficult to find a good approximation $\varphi(x_i)$ of $|\mathbf{x}_i|_0$ for all x_i . Selecting an appropriate φ function generally relies on the introduction of a degree of freedom whose tuning is not obvious [5, 6]. For instance, the $\text{IR}\ell_1$ algorithm can be interpreted as an approximate ℓ_2 - ℓ_0 minimization method where the ℓ_0 norm is replaced with $\varphi(x_i; \varepsilon) = \log(|x_i| + \varepsilon)$ [5, 7]. The parameter ε controls the “degree of nonconvexity” of the surrogate function φ ⁽²⁾.

Although $\mathcal{J}(\mathbf{x}; \lambda)$ has a large number of local minima, we have found that SBR is often as accurate as algorithms based on the nonconvex approximation of $\mathcal{J}$. Moreover, SBR is simple to use. The good behavior of SBR is somehow related to the result of Proposition 1 which states that any SBR iterate is

²In the comparisons with SBR, we set $\varepsilon = 0$ following [40].

almost surely a local minimizer of $\mathcal{J}$. We conclude that SBR is actually capable to “skip” local minima with a large cost $\mathcal{J}(\mathbf{x}; \lambda)$.

B. Perspectives

In the proposed approach, the main difficulty relies in the choice of the λ -value. If a specific cardinality 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 [45], we sketched a continuation version in which a series of SBR solutions are computed for decreasing levels of sparsity λ , and the λ -values are recursively computed. This continuation version is showing promising results and will be the subject of a future extended contribution. A similar perspective was actually proposed by Zhang to generalize his FoBa algorithm in a path-following algorithm (see the discussion section in [22]).

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 $\mathbf{x}^*$. According to Remark 1, we will study the behavior of SBR when $\lambda \rightarrow 0$. In the simulations done in Sections V and VI, we observed that SBR is able to perform exact recoveries provided that λ is sufficiently small. This promising result is a first step towards a more general theoretical study.

APPENDIX A

DETAILED DEVELOPMENT OF LIMIT BG SIGNAL RESTORATION

Consider the Bernoulli-Gaussian model $\mathbf{x} = (\mathbf{q}, \mathbf{r})$ introduced in Section II-B and the joint MAP formulation (3) involving the cost function $\mathcal{L}(\mathbf{q}, \mathbf{r})$. 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 $\|\mathbf{r}\|^2 = \|\mathbf{t}\|^2$ and $\mathbf{A}\Delta_{\mathbf{q}}\mathbf{r} = \mathbf{A}_{\mathcal{Q}}\mathbf{t}$ do not depend on $\mathbf{u}$, we have $\min_{\mathbf{u}} \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{u}) = \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$. Thus, the joint MAP estimation problem reduces to the minimization of $\mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0})$ w.r.t. $(\mathbf{q}, \mathbf{t})$. In the limit case $\sigma_x^2 \rightarrow \infty$, this problem rereads:

$$\min_{\mathbf{q}, \mathbf{t}} \{ \mathcal{L}(\mathbf{q}, \mathbf{t}, \mathbf{0}) = \|\mathbf{y} - \mathbf{A}_{\mathcal{Q}}\mathbf{t}\|^2 + \lambda \|\mathbf{q}\|_0 \}. \quad (9)$$

The equivalence between (9) and (4) directly follows from the change of variable $\mathbf{x} = \{\mathbf{q}, \mathbf{t}\}$ where $\mathbf{q}$ and $\mathbf{t}$ are the support and non-zero amplitudes of $\mathbf{x}$.

APPENDIX B

PROOF OF REMARK 1

The proof of the result stated in Remark 1 is based on the two following lemmas.

Lemma 1 For $\lambda > 0$, any minimizer of $\mathcal{J}(\mathbf{x}; \lambda)$ takes the form $\mathbf{x}_{\mathcal{Q}}$ with $\text{Card}[\mathcal{Q}] \leq \min(m, n)$.

Proof of lemma: According to the URP assumption, any $\min(m, n)$ columns of $\mathbf{A}$ yield an unconstrained minimizer of $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$. Let $\mathbf{x}_{\text{LS}}$ be such minimizer, with $\|\mathbf{x}_{\text{LS}}\|_0 \leq \min(m, n)$, and let $\mathbf{u}$ be a minimizer of $\mathcal{J}(\mathbf{x}; \lambda)$. $\mathcal{J}(\mathbf{u}; \lambda) \leq \mathcal{J}(\mathbf{x}_{\text{LS}}; \lambda)$ implies that $\|\mathbf{u}\|_0 \leq \|\mathbf{x}_{\text{LS}}\|_0 + (\mathcal{E}(\mathbf{x}_{\text{LS}}) - \mathcal{E}(\mathbf{u}))/\lambda \leq \|\mathbf{x}_{\text{LS}}\|_0 \leq \min(m, n)$.

We denote by $\mathcal{Q}$ the support of $\mathbf{u}$. The related least-square solution $\mathbf{x}_{\mathcal{Q}}$ obviously satisfies $\mathcal{E}(\mathbf{x}_{\mathcal{Q}}) \leq \mathcal{E}(\mathbf{u})$ and $\|\mathbf{x}_{\mathcal{Q}}\|_0 \leq \text{Card}[\mathcal{Q}] = \|\mathbf{u}\|_0$, thus $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda) \leq \mathcal{J}(\mathbf{u}; \lambda)$. Since $\mathbf{u}$ is a minimizer of $\mathcal{J}(\mathbf{x}; \lambda)$, we have $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda) = \mathcal{J}(\mathbf{u}; \lambda)$ hence $\mathcal{E}(\mathbf{x}_{\mathcal{Q}}) = \mathcal{E}(\mathbf{u})$. Because of the URP assumption, the least-squares minimizer over $\mathcal{Q}$ is unique, thus $\mathbf{u} = \mathbf{x}_{\mathcal{Q}}$. ■

Lemma 2 There exists $\lambda_{\min} > 0$ such that for $0 < \lambda \leq \lambda_{\min}$, the minimizers of $\mathcal{J}(\mathbf{x}; \lambda)$ are unconstrained minimizers of $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$.

Proof of lemma: When λ tends towards 0, we have for all $\mathcal{Q}$, $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda) = \mathcal{E}_{\mathcal{Q}} + \lambda\|\mathbf{x}_{\mathcal{Q}}\|_0 \rightarrow \mathcal{E}_{\mathcal{Q}}$. In particular, $\mathcal{J}(\mathbf{x}_{\mathcal{Q}_{\text{LS}}}; \lambda) \rightarrow \mathcal{E}_{\mathcal{Q}_{\text{LS}}}$ with $\mathbf{x}_{\mathcal{Q}_{\text{LS}}}$ an unconstrained minimizer of $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$ yielded by a subset $\mathcal{Q}_{\text{LS}}$ of cardinality $\min(m, n)$. Because the number of possible subsets $\mathcal{Q}$ is finite and for all $\mathcal{Q}$, $\mathcal{E}_{\mathcal{Q}} \geq \mathcal{E}_{\mathcal{Q}_{\text{LS}}}$, there exists $\lambda_{\min} > 0$ such that for $0 < \lambda \leq \lambda_{\min}$, the subsets $\mathcal{Q}^*$ minimizing $\mathcal{J}(\mathbf{x}_{\mathcal{Q}}; \lambda)$ satisfy $\mathcal{E}_{\mathcal{Q}^*} = \mathcal{E}_{\mathcal{Q}_{\text{LS}}}$. Consequently, the minimizers of $\mathcal{J}(\mathbf{x}; \lambda)$ are unconstrained least-squares solutions according to Lemma 1. ■

Proof of Remark 1: The proof directly follows from the application of Lemma 2. We denote by $\mathcal{X}_{\lambda}$ the set of minimizers of $\mathcal{J}(\mathbf{x}; \lambda)$.

In the undercomplete case, there is a unique unconstrained least-square minimizer $\mathbf{x}_{\text{LS}}$. Thus, $\mathcal{X}_{\lambda} = \{\mathbf{x}_{\text{LS}}\}$ for $0 \leq \lambda \leq \lambda_{\min}$.

In the overcomplete case, we denote by $\mathcal{X}^*$ the set of sparsest solutions to $\mathbf{y} = \mathbf{A}\mathbf{x}$. To show that $\mathcal{X}_{\lambda} = \mathcal{X}^*$ for $0 < \lambda \leq \lambda_{\min}$, we consider $\mathbf{x} \in \mathcal{X}^*$ and $\mathbf{x}' \in \mathcal{X}_{\lambda}$. According to Lemma 2, $\mathbf{x}'$ satisfies $\mathbf{y} = \mathbf{A}\mathbf{x}'$, then $\mathcal{J}(\mathbf{x}'; \lambda) = \lambda\|\mathbf{x}'\|_0$. By definition of $\mathcal{X}^*$, we have $\mathbf{y} = \mathbf{A}\mathbf{x}$ and $\mathcal{J}(\mathbf{x}; \lambda) = \lambda\|\mathbf{x}\|_0 \leq \mathcal{J}(\mathbf{x}'; \lambda)$. Because $\mathbf{x}' \in \mathcal{X}_{\lambda}$ is a minimizer of $\mathcal{J}$, we deduce that $\|\mathbf{x}'\|_0 = \|\mathbf{x}\|_0$, then $\mathbf{x}' \in \mathcal{X}^*$ and $\mathbf{x} \in \mathcal{X}_{\lambda}$. We have proved that $\mathcal{X}_{\lambda} = \mathcal{X}^*$ for $0 < \lambda \leq \lambda_{\min}$. ■

APPENDIX C

UPDATE OF THE CHOLESKY FACTORIZATION

At each SBR iteration, n linear systems of the form $\mathbf{t}_{\mathcal{Q}} \triangleq \mathbf{G}_{\mathcal{Q}}^{-1} \mathbf{A}_{\mathcal{Q}}^t \mathbf{y}$ must be solved, the corresponding squared errors reading $\mathcal{E}_{\mathcal{Q}} = \|\mathbf{y} - \mathbf{A}_{\mathcal{Q}} \mathbf{t}_{\mathcal{Q}}\|^2 = \|\mathbf{y}\|^2 - \mathbf{y}^t \mathbf{A}_{\mathcal{Q}} \mathbf{t}_{\mathcal{Q}}$. Using the Cholesky factorization

$G_Q = L_Q L_Q^t$, t_Q rereads $t_Q = L_Q^{-t} L_Q^{-1} A_Q^t y$, thus

$$\mathcal{E}_Q = \|y\|^2 - \|L_Q^{-1} A_Q^t y\|^2. \quad (10)$$

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

Thus, the new Gram matrix reads as a 2×2 block matrix:

$$G_{Q'} = \begin{bmatrix} G_Q & A_Q^t a_i \\ (A_Q^t a_i)^t & \|a_i\|^2 \end{bmatrix}$$

and the Cholesky factor of $G_{Q'}$ can be straightforwardly updated:

$$L_{Q'} = \begin{bmatrix} L_Q & \mathbf{0} \\ l_{Q,i}^t & \sqrt{\|a_i\|^2 - \|l_{Q,i}\|^2} \end{bmatrix} \quad (11)$$

with $l_{Q,i} = L_Q^{-1} A_Q^t a_i$. The update (8) of $\mathcal{J}_Q(\lambda) = \mathcal{E}_Q + \lambda \text{Card}[Q]$ directly follows from (10) and (11).

Removal of an arbitrary column: When removing a column a_i , updating L_Q remains possible although more complex. This idea was developed by Ge *et al.* [46] who update the Cholesky factorization of matrix G_Q^{-1} . We adapt it to the direct (simpler) factorization of G_Q . Let I be the position of a_i in A_Q (with $1 \leq I \leq \text{Card}[Q]$). L_Q can be written in a block matrix form:

$$L_Q = \begin{bmatrix} \Lambda & \mathbf{0} & \mathbf{0} \\ b^t & d & \mathbf{0} \\ C & e & F \end{bmatrix} \quad (12)$$

where the lowercase characters refer to the scalar (d) and vector quantities (b, e) appearing in the I th row and in the I th column. The computation of $G_Q = L_Q L_Q^t$ and the removal of the I th row and the I th column in G_Q lead to

$$G_{Q'} = \begin{bmatrix} \Lambda & \mathbf{0} \\ C & F \end{bmatrix} \begin{bmatrix} \Lambda^t & C^t \\ \mathbf{0} & F^t \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ e \end{bmatrix} \begin{bmatrix} \mathbf{0} & e^t \end{bmatrix}.$$

By identification with $G_{Q'} = L_{Q'} L_{Q'}^t$, and because the Cholesky factorization is unique, $L_{Q'}$ necessarily reads:

$$L_{Q'} = \begin{bmatrix} \Lambda & \mathbf{0} \\ C & X \end{bmatrix}, \quad (13)$$

where X is a lower triangular matrix satisfying $XX^t = FF^t + ee^t$. The problem of computing X from F and 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 $f = \text{Card}[Q] - I$ is the size of F [36].

APPENDIX D

PROOF OF PROPOSITION 2

Let us first introduce some notations specific to the piecewise polynomial dictionary problem. Consider a subset $\mathcal{Q}$ of columns $\mathbf{a}_i^p$ and let $i^- = \min\{i \mid n_i > 0\}$ denote the lowest location of an active entry (we recall that n_i denotes the number of active columns for sample i). Up to a reordering of the columns of $\mathbf{A}_{\mathcal{Q}}$, $\mathbf{A}_{\mathcal{Q}}$ rereads $\mathbf{A}_{\mathcal{Q}} = [\mathbf{A}_{i^-}, \tilde{\mathbf{A}}_{i^-}]$ where $\mathbf{A}_{i^-}$ gathers the n_{i^-} active columns $\mathbf{a}_i^p$ such that $i = i^-$ and $\tilde{\mathbf{A}}_{i^-}$ gathers the remaining active columns (with $i > i^-$). The following lemma is a key element to prove Proposition 2.

Lemma 3 *Assume that $\mathcal{Q}$ satisfies the condition of Proposition 2. If $\tilde{\mathbf{A}}_{i^-}$ is full rank, then $\mathbf{A}_{\mathcal{Q}}$ is 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 orders, 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

$$\sum_{j=1}^I \mu_{i^-}^{p_j} \mathbf{a}_{i^-}^{p_j} + \sum_{i \neq i^-} \sum_p \mu_i^p \mathbf{a}_i^p = \mathbf{0}. \quad (14)$$

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 2, we have, for all $k \in \{i^-, \dots, i^- + I - 1\}$, $\sum_{j=1}^I \mu_{i^-}^{p_j} (k + i^- - 1)^{p_j} = 0$. Hence, the polynomial $F(X) = \sum_{j=1}^I \mu_{i^-}^{p_j} X^{p_j}$ has I positive roots. Because any non-zero polynomial formed of I monomials of different degree has at most $I - 1$ positive roots [47, p. 76], F is the zero polynomial, thus all scalars $\mu_{i^-}^{p_j}$ are 0. We deduce from (14) and from the full rankness of $\tilde{\mathbf{A}}_{i^-}$ 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 2 directly results from the recursive application of Lemma 3. Starting from the empty set, all the indices, sorted by decreasing order, are successively included.

ACKNOWLEDGMENT

The authors would like to thank Dr. Grégory Francius from LCPME (UMR CNRS 7564, Nancy, France) for providing them with real AFM data.

REFERENCES

- [1] B. K. Natarajan, “Sparse approximate solutions to linear systems”, *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, Apr. 1995.
- [2] 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.
- [3] 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.
- [4] 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 Process.*, vol. 51, no. 3, pp. 760–770, Mar. 2003.
- [5] E. J. Candès, M. B. Wakin, and S. P. Boyd, “Enhancing sparsity by reweighted ℓ_1 minimization”, *J. Fourier Anal. Appl.*, vol. 14, no. 5-6, pp. 877–905, Dec. 2008.
- [6] 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 Process.*, vol. 57, no. 1, pp. 289–301, Jan. 2009.
- [7] D. P. Wipf and S. Nagarajan, “Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions”, *IEEE J. Sel. Top. Signal Process. (Special Issue on Compressive Sensing)*, vol. 4, no. 2, pp. 317–329, Apr. 2010.
- [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] W. Dai and O. Milenkovic, “Subspace pursuit for compressive sensing signal reconstruction”, *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2230–2249, May 2009.
- [10] T. Blumensath and M. E. Davies, “Iterative thresholding for sparse approximations”, *J. Fourier Anal. Appl.*, vol. 14, no. 5, pp. 629–654, Dec. 2008.
- [11] T. Blumensath and M. E. Davies, “Normalized iterative hard thresholding: Guaranteed stability and performance”, *IEEE J. Sel. Top. Signal Process.*, vol. 4, no. 2, pp. 298–309, Apr. 2010.
- [12] S. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries”, *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [13] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, “Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition”, in *Proc. 27th Asilomar Conf. on Signals, Systems and Computers*, Nov. 1993, vol. 1, pp. 40–44.
- [14] 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.
- [15] M. A. Efron, “Multiple regression analysis”, in *Mathematical Methods for Digital Computers*, A. Ralston and H. S. Wilf, Eds., vol. 1, pp. 191–203. Wiley, New York, NY, 1960.
- [16] K. N. Berk, “Forward and backward stepping in variable selection”, *J. Statist. Comput. Simul.*, vol. 10, no. 3-4, pp. 177–185, Apr. 1980.
- [17] P. M. T. Broersen, “Subset regression with stepwise directed search”, *J. R. Statist. Soc. C*, vol. 35, no. 2, pp. 168–177, 1986.
- [18] A. J. Miller, *Subset Selection in Regression*, Chapman and Hall, London, UK, 2nd edition, Apr. 2002.
- [19] 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.

- [20] T. Blumensath and M. E. Davies, “On the difference between Orthogonal Matching Pursuit and Orthogonal Least Squares”, Tech. Rep., University of Edinburgh, Mar. 2007.
- [21] D. Haugland, *A Bidirectional Greedy Heuristic for the Subspace Selection Problem*, vol. 4638 of *Lect. Notes Comput. Sci.*, pp. 162–176, Springer Verlag, Berlin, Germany, Engineering stochastic local search algorithms. Designing, implementing and analyzing effective heuristics edition, 2007.
- [22] T. Zhang, “Adaptive forward-backward greedy algorithm for learning sparse representations”, Tech. Rep., Rutgers Statistics Department, Apr. 2008.
- [23] C. Herzet and A. Drémeau, “Bayesian pursuit algorithms”, in *Proc. Eur. Sig. Proc. Conf.*, Aalborg, Denmark, Aug. 2010, pp. 1474–1478.
- [24] 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, May 1982.
- [25] J. M. Mendel, *Optimal Seismic Deconvolution*, Academic Press, New York, 1983.
- [26] 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.
- [27] F. Champagnat, Y. Goussard, and J. Idier, “Unsupervised deconvolution of sparse spike trains using stochastic approximation”, *IEEE Trans. Signal Process.*, vol. 44, no. 12, pp. 2988–2998, Dec. 1996.
- [28] I. F. Gorodnitsky and B. D. Rao, “Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm”, *IEEE Trans. Signal Process.*, vol. 45, no. 3, pp. 600–616, Mar. 1997.
- [29] 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.
- [30] E. I. George and D. P. Foster, “Calibration and empirical Bayes variable selection”, *Biometrika*, vol. 87, no. 4, pp. 731–747, 2000.
- [31] H. Chipman, E. I. George, and R. E. McCulloch, “The practical implementation of Bayesian model selection”, *IMS Lecture Notes – Monograph Series*, vol. 38, pp. 65–134, 2001.
- [32] J. Nocedal and S. J. Wright, *Numerical optimization*, Springer Series in Operations Research and Financial Engineering. Springer Verlag, New York, 1999.
- [33] S. Chen and J. Wigger, “Fast orthogonal least squares algorithm for efficient subset model selection”, *IEEE Trans. Signal Process.*, vol. 43, no. 7, pp. 1713–1715, July 1995.
- [34] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression”, *Ann. Statist.*, vol. 32, no. 2, pp. 407–451, 2004.
- [35] S. J. Reeves, “An efficient implementation of the backward greedy algorithm for sparse signal reconstruction”, *IEEE Signal Process. Lett.*, vol. 6, no. 10, pp. 266–268, Oct. 1999.
- [36] P. E. Gill, G. H. Golub, W. Murray, and M. A. Saunders, “Methods for modifying matrix factorizations”, *Math. Comp.*, vol. 28, no. 126, pp. 505–535, Apr. 1974.
- [37] 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.
- [38] 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.
- [39] D. L. Donoho, V. Stodden, and Y. Tsaig, “About SparseLab”, Tech. Rep., Stanford University, Mar. 2007.

- [40] H. Zou, “The adaptive Lasso and its oracle properties”, *J. Acoust. Soc. Amer.*, vol. 101, no. 476, pp. 1418–1429, Dec. 2006.
- [41] M. S. Smith and R. Kohn, “Nonparametric regression using Bayesian variable selection”, *J. Econometrics*, vol. 75, no. 2, pp. 317–343, Dec. 1996.
- [42] M. Vetterli, P. Marziliano, and T. Blu, “Sampling signals with finite rate of innovation”, *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1417–1428, June 2002.
- [43] H.-J. Butt, B. Cappella, and M. Kappl, “Force measurements with the atomic force microscope: Technique, interpretation and applications”, *Surf. Sci. Rep.*, vol. 59, no. 1–6, pp. 1–152, Oct. 2005.
- [44] 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 Surf. B.*, vol. 62, no. 2, pp. 206–213, Apr. 2008.
- [45] 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.
- [46] D. Ge, J. Idier, and E. Le Carpentier, “Enhanced sampling schemes for MCMC based blind Bernoulli-Gaussian deconvolution”, *Signal Process.*, vol. 91, no. 4, pp. 759–772, Apr. 2011.
- [47] 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.

Charles Soussen was born in France in 1972. He received the degree from the École Nationale Supérieure en Informatique et Mathématiques Appliquées, Grenoble, France, and the PhD degree in Physics from the Laboratoire des Signaux et Systèmes, Université de Paris-Sud, Orsay, France, in 1996 and 2000, respectively. He is currently an Assistant Professor at Nancy-University, France. He has been with the Centre de Recherche en Automatique de Nancy since 2005. His research interests are in inverse problems and sparse approximation.

Jérôme Idier was born in France in 1966. He received the diploma degree in electrical engineering from École Supérieure d’Électricité, Gif-sur-Yvette, France, in 1988 and the Ph.D. degree in physics from University of Paris-Sud, Orsay, France, in 1991.

Since 1991, he joined the Centre National de la Recherche Scientifique. He is currently a Senior Researcher at the Institut de Recherche en Communications et Cybernétique in Nantes. His major scientific interests are in probabilistic approaches to inverse problems for signal and image processing. He is serving as an Associate Editor for the IEEE Transactions on Signal Processing.

David Brie received the Ph.D. degree in 1992 and the *Habilitation à diriger des Recherches* degree in 2000, both from the Henri Poincaré University, Nancy, France. He is currently Professor at the telecommunication and network department from the *Institut Universitaire de Technologie*, Nancy-University. Since 1990, he has been with the *Centre de Recherche en Automatique de Nancy*. His research interests mainly concern inverse problems and multidimensional signal processing.

Junbo Duan was born in China in 1981. He received the B.S. degree in information engineering and M.S. degree in communication and information system from Xi'an Jiaotong University, China, in 2004 and 2007 respectively, and the Ph.D. degree in signal processing from Université Henry Poincaré, Nancy, France, in 2010. He is currently a postdoc in the Department of Biomedical Engineering and Biostatistics, Tulane University. His major research interests are in probabilistic approaches to inverse problems in bioinformatics.