DYNAMIC SUBSPACE PURSUIT

Dave Zachariah, Saikat Chatterjee, Magnus Jansson

ACCESS Linnaeus Centre, School of Electrical Engineering
KTH - Royal Institute of Technology, Stockholm Sweden
dave.zachariah@ee.kth.se, sach@kth.se, magnus.jansson@ee.kth.se

ABSTRACT

For compressive sensing of dynamic sparse signals, we develop an iterative greedy search algorithm based on subspace pursuit (SP) that can incorporate sequential predictions, thereby taking advantage of its low complexity while improving recovery performance by exploiting correlations described by a state space model. The algorithm, which we call dynamic subspace pursuit (DSP), is presented and experimentally validated. It exhibits a graceful degradation at deteriorating signal conditions while capable of yielding substantial performance gains as conditions improve.

***Index Terms*—** Compressive sensing, recursive reconstruction, sparse reconstruction.

1. INTRODUCTION

Compressive Sensing (CS) [1] problems assume a sparse-signal model, undersampled by a linear measurement process. The algorithms for CS can be separated into three broad classes: convex relaxation, Bayesian inference, and iterative greedy search (IGS). For large-dimensional CS signal-reconstruction, IGS algorithms offer computationally efficient solutions by detecting and reconstructing the active signal coefficients in a least-squares framework. Examples of such IGS algorithms are orthogonal matching pursuit (OMP) [2] and subspace pursuit (SP) [3] and several variants of them. These algorithms may use some prior information, such as the maximum allowable cardinality of the 'support set' (defined as the set of non-zero signal coordinates of the underlying sparse signal).

A recent trend in CS is the recovery of dynamic sparse signals, exploiting temporal/spectral/spatial correlations as in the CS application scenarios of MRI [4] and spectrum sensing [5]. In [6], the overall methodology is sequential and can be seen as a two-step approach: (1) support-set detection of the sparse signal, and (2) reduced-order recovery using prior information on the detected support set. For a reasonable detection of support set, [6] uses convex relaxation algorithms. Then, a standard Kalman filter (KF) is employed to use prior information for sequential signal recovery. Without explicit support set detection, [7] uses KF to estimate the entire signal and enforces sparsity by imposing an approximate norm constraint. However, the work of [7] validates their algorithm for a signal with a static sparsity pattern (i.e. an unknown pattern that does not evolve over time). Similarly, [8] considers scenarios with static sparsity pattern and solves the reconstruction of a temporally evolving sparse signal with multiple measurement vectors in a batch Bayesian learning framework with unknown model parameters.

This work was partially supported by the Swedish Research Council under contract 621-2011-5847.

A notable omission in existing research work is the development of IGS algorithms that can use prior information to recover dynamic sparse signals. In this paper we are interested in generalizing the IGS algorithms for sequential estimation of such processes, thereby taking advantage of their low complexity. We develop such an algorithm based on the subspace pursuit and through experimental simulations we show that it provides a graceful degradation at higher measurement noise levels and/or lower measurement signal dimensions, while capable of yielding substantial gains at more favorable signal conditions.

Notation: $\|\mathbf{x}\|_0$ denotes l_0 'norm', i.e. the number of non-zero coefficients of the vector $\mathbf{x}$. $\mathbf{A} \oplus \mathbf{B}$ is the direct sum of matrices. $|S|$ and S^c are the cardinality and complement of set S, respectively. $(\cdot)^*$ is the Hermitian transpose operator. $\mathbf{A}^\dagger$ the Moore-Penrose pseudoinverse of matrix $\mathbf{A}$. $\mathbf{A}_{[\mathcal{I},\mathcal{J}]}$ denotes a submatrix of $\mathbf{A}$ with elements from row and column indices listed in ordered sets $\mathcal{I}$ and $\mathcal{J}$. Similarly, the column vector $\mathbf{x}_{[\mathcal{I}]}$ contains the elements of $\mathbf{x}$ with indices from set $\mathcal{I}$.

2. SIGNAL MODEL

Let the state vector $\mathbf{x}_t \in \mathbb{C}^N$ be sparse with a dynamically evolving sparsity pattern, represented by the 'support set' $I_{x,t} \subset \{1,\dots,N\}$. We will assume $\|\mathbf{x}_t\|_0 \equiv |I_{x,t}| \leq K$. Let λ_{ji} denote the state transition probability $j \to i$ of the non-zero signal coordinate j. The transition probabilities determine the transition $I_{x,t} \to I_{x,t+1}$.

For $\mathbf{x}_t$ we use an autoregressive (AR) process model with the transition of the non-zero signal coordinate $j \to i$,

$$x_{i,t+1} = \alpha_{ij} x_{j,t} + w_{i,t}, \tag{1}$$

where $x_{j,t}$ denotes the jth component of $\mathbf{x}_t$ and $w_{i,t}$ is the associated innovation. All active components are assumed to have variance σ_x^2. The process $\mathbf{x}_t$ can be written compactly as a linear state-space model with random transition matrices,

$$\mathbf{x}_{t+1} = \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \mathbf{w}_t, \tag{2}$$

where $\mathrm{E}[\mathbf{w}_t] = \mathbf{0}$, $\mathrm{E}[\mathbf{w}_t \mathbf{w}^*_{t-l}] = \mathbf{Q}_t \delta(l) \in \mathbb{C}^{N\times N}$. For all $j \in I_{x,t}$ and $i \in I_{x,t+1}$, the non-zero elements of $\mathbf{A}_t \in \mathbb{C}^{N\times N}$ and the diagonal matrix $\mathbf{B}_t \in \mathbb{C}^{N\times N}$ are $a_{ij,t} = \alpha_{ij,t}$ and $b_{ii,t} = 1$, respectively. The model parameters $\alpha_{ij,t}, \lambda_{ji}$ and $\mathbf{Q}_t$ are assumed to be known.

As an example consider a sparse process, $N = 200$, $K = 10$ and number of snapshots $T = 100$, with a slowly varying sparsity pattern $I_{x,t}$ following[1]

$$\lambda_{ji} = \begin{cases} 0.90 & i = j \\ 0.05 & i = j \pm 1, \text{ if } j \notin \{1, N\} \\ 0.10 & i = j + 1, \text{ if } j = 1 \\ 0.10 & i = j - 1, \text{ if } j = N \\ 0 & |i - j| > 1 \end{cases} . \quad (3)$$

A realization of such a dynamic sparse process is illustrated in figure 1. This choice of λ_{ji} is intended to model the strong temporal correlation of sparse signals exhibited in e.g. MRI [9].

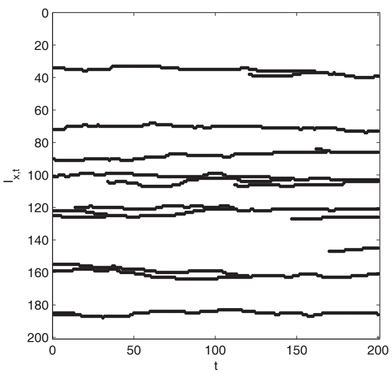

Fig. 1. Evolving sparsity pattern $I_{x,t}$ with $(N, K, T) = (200, 10, 200)$ and transition probabilities (3).

The standard CS measurement setup applies,

$$\mathbf{y}_t = \mathbf{H}\mathbf{x}_t + \mathbf{n}_t \in \mathbb{C}^M, \quad (4)$$

with $\mathrm{E}[\mathbf{n}_t] = \mathbf{0}$, $\mathrm{E}[\mathbf{n}_t\mathbf{n}_t^*] = \mathbf{R}_t$. The sensing matrix $\mathbf{H} = [\mathbf{h}_1 \quad \cdots \quad \mathbf{h}_N] \in \mathbb{C}^{M \times N}$, where $K < M < N$ and fulfills the restricted isometry property, which means that its columns are 'nearly' orthogonal [1]. Both $\mathbf{H}$ and $\mathbf{R}_t$ are given.

3. PREDICTIVE ITERATIVE GREEDY SEARCH

3.1. Incorporation of prior information

SP constructs a support set $I \subset \{1, \ldots, N\}$ by correlating measurement residuals $\mathbf{r} = \mathbf{y}_t - \mathbf{H}_{[\cdot,I]}\hat{\mathbf{x}}_{[I]}$ with the column vectors, or 'atoms' $\mathbf{h}_i$. The coefficients corresponding to a hypothesized support set $\tilde{I}$ are reconstructed based on a least-squares criterion, $\hat{\mathbf{x}}_{[\tilde{I}],t} = \mathbf{H}^\dagger_{[\cdot,\tilde{I}]}\mathbf{y}_t$. The indices $i \in \tilde{I}$ of the K_{max} largest magnitudes $|\hat{x}_{i,t}|$ are used to form I.

For sake of clarity SP is summarized in Algorithm 1. Here the function $\texttt{indexmax}(K, |z_i|, \mathcal{I})$ returns the indices of the K largest non-zero magnitudes $|z_i|$ for $i \in \mathcal{I}$.

Now assume prior statistical information of the state vector exists in the form of a prediction $\hat{\mathbf{x}}_t^-$ and its error covariance matrix $\mathbf{P}_t^-$. We use the error statistics of the prediction to improve the detection performance. The measurement residual based on the prediction is $\mathbf{r}^- = \mathbf{y}_t - \mathbf{H}\hat{\mathbf{x}}_t^- = \mathbf{H}\tilde{\mathbf{x}}_t^- + \mathbf{n}_t$, where $\tilde{\mathbf{x}}_t^-$ is the prediction error. The power of the correlation magnitude of atom i,

$$\mathrm{E}\left[\left|\mathbf{h}_i^*\mathbf{r}^-\right|^2\right] = \mathrm{E}\left[\left|\mathbf{h}_i^*\mathbf{h}_i\tilde{x}_{i,t}^- + \sum_{j \neq i}\mathbf{h}_i^*\mathbf{h}_j\tilde{x}_{j,t}^- + \mathbf{h}_i^*\mathbf{n}_t\right|^2\right],$$

[1]The two cases $j = 1$ or $j = N$ are necessary for the edge states.

Algorithm 1 Subspace Pursuit (SP)

Given: $\mathbf{y}_t$, $\mathbf{H}$ and K_{max}
Set $I = \texttt{indexmax}(K_{\text{max}}, |\mathbf{h}_i^*\mathbf{y}_t|, \cdot)$,
$k = 0$, $\mathbf{r}_0 = \mathbf{y}_t - \mathbf{H}_{[\cdot,I]}\mathbf{H}^\dagger_{[\cdot,I]}\mathbf{y}_t$
repeat
  $k := k + 1$
  $J = \texttt{indexmax}(K_{\text{max}}, |\mathbf{h}_i^*\mathbf{r}_{k-1}|, \cdot)$
  $\tilde{I} = I \cup J$
  $\hat{\mathbf{x}}_{[\tilde{I}],t} = \mathbf{H}^\dagger_{[\cdot,\tilde{I}]}\mathbf{y}_t$
  $I := \texttt{indexmax}(K_{\text{max}}, |\hat{x}_{i,t}|, \tilde{I})$
  $\hat{\mathbf{x}}_{[I],t} = \mathbf{H}^\dagger_{[\cdot,I]}\mathbf{y}_t$; $\hat{\mathbf{x}}_{[I^c],t} = \mathbf{0}$
  $\mathbf{r}_k = \mathbf{y}_t - \mathbf{H}_{[\cdot,I]}\hat{\mathbf{x}}_{[I],t}$
until $(\|\mathbf{r}_k\|_2 > \|\mathbf{r}_{k-1}\|_2)$
Output: $\hat{\mathbf{x}}_t$ and I

is approximated by

$$\gamma_i \triangleq p_{ii,t}^- + \mathbf{h}_i^*\mathbf{R}_t\mathbf{h}_i, \quad (5)$$

using $\mathbf{h}_i^*\mathbf{h}_j \approx \delta(i-j)$ and omitting the influence of the interference terms $\mathbf{h}_i^*\mathbf{h}_j\tilde{x}_{j,t}^-$. Hence γ_i embodies the uncertainty of predicting atom i, and applying the weighting $|\mathbf{h}_i^*\mathbf{r}|/\gamma_i$ will amplify correlation magnitudes with low uncertainty, and vice versa.

The prior can also be used for reconstruction, solving the weighed least-squares problem

$$\hat{\mathbf{x}}_{[I],t} = \underset{\mathbf{x}_{[I]} \in \mathbb{C}^{|I|}}{\arg\min} \left\| \begin{bmatrix} \mathbf{y}_t \\ \hat{\mathbf{x}}_{[I],t}^- \end{bmatrix} - \begin{bmatrix} \mathbf{H}_{[\cdot,I]} \\ \mathbf{I}_k \end{bmatrix} \mathbf{x}_{[I]} \right\|^2_{\mathbf{R}_t^{-1} \oplus \mathbf{S}_t^{-1}}, \quad (6)$$

where $\mathbf{S}_t = \mathbf{P}_{[I,I],t}^-$ and $k = |I|$, which is the linear MMSE estimator provided I is the correct support set [10]. Since I is not given but successively constructed, (6) solves an approximate problem by disregarding any correlations between the elements of sets I and I^c. For brevity we let $\texttt{MMSE-rec}$ denote a function that solves (6), and can be computed by a measurement update of form:

$$\hat{\mathbf{x}}_{[I],t} = \hat{\mathbf{x}}_{[I],t}^- + \mathbf{K}_t\left(\mathbf{y}_t - \mathbf{H}_{[\cdot,I]}\hat{\mathbf{x}}_{[I],t}^-\right),$$

where $\mathbf{K}_t = \mathbf{S}_t\mathbf{H}^*_{[\cdot,I]}(\mathbf{H}_{[\cdot,I]}\mathbf{S}_t\mathbf{H}^*_{[\cdot,I]} + \mathbf{R}_t)^{-1}$.

Combining the weighted correlations and the MMSE reconstruction, we develop a predictive SP (PrSP) in Algorithm 2. Our approach can be generalized to several other IGS algorithms.

3.2. Dynamic Subspace Pursuit

In order to predict the process (2) we model the random transition matrix $\mathbf{A}_t$ by $\mathbf{F}_t$, where the elements are set as $f_{ij,t} = \alpha_{ij,t}\lambda_{ji}$. This enables the formulation of a filtering problem.

Let I denote the final set of detected atoms after the application of a predictive greedy search algorithm then the updated matrix $\mathbf{P}_t$ is computed block-wise corresponding to the set I and its complement, I^c. First, $\mathbf{P}_{[I,I],t} = \left(\mathbf{I}_{|I|} - \mathbf{P}_{[I,I],t}^-\mathbf{H}^*_{[\cdot,I]}\mathbf{W}_t^{-1}\mathbf{H}_{[\cdot,I]}\right)\mathbf{P}_{[I,I],t}^-$ is the posterior error covariance [10], where

$$\mathbf{W}_t = (\mathbf{H}_{[\cdot,I]}\mathbf{P}_{[I,I],t}^-\mathbf{H}^*_{[\cdot,I]} + \mathbf{R}_t). \quad (7)$$

Second, the uncertainty of the zero coefficients is preserved by $\mathbf{P}_{[I^c,I^c],t} = \mathbf{P}_{[I^c,I^c],t}^-$. Finally, to ensure positive definiteness the cross-correlations between elements corresponding to I and I^c are

Algorithm 2 Predictive Subspace Pursuit (PrSP)

Given: $\mathbf{y}_t, \mathbf{H}, \mathbf{R}_t, \hat{\mathbf{x}}_t^-, \mathbf{P}_t^-$ and K_{max}
Compute γ_i using (5)
Set $I := \texttt{indexmax}(K_{\text{max}}, |\hat{x}_{i,t}^-|, \cdot)$
$J = \texttt{indexmax}(K_{\text{max}}, |\mathbf{h}_i^* \mathbf{y}_t|/\gamma_i, \cdot)$,
$\tilde{I} = I \cup J$
$\hat{\mathbf{x}}_{[\tilde{I}],t} = \texttt{MMSE-rec}(\mathbf{y}_t, \mathbf{H}, \mathbf{R}_t, \hat{\mathbf{x}}_t^-, \mathbf{P}_t^-, \tilde{I})$
$I := \texttt{indexmax}(K_{\text{max}}, |\hat{x}_{i,t}|, \tilde{I})$
$\hat{\mathbf{x}}_{[I],t} = \texttt{MMSE-rec}(\mathbf{y}_t, \mathbf{H}, \mathbf{R}_t, \hat{\mathbf{x}}_t^-, \mathbf{P}_t^-, I)$,
$k = 0, \mathbf{r}_0 = \mathbf{y}_t - \mathbf{H}_{[\cdot,I]} \hat{\mathbf{x}}_{[I],t}$
repeat
    $k := k + 1$
    $J = \texttt{indexmax}(K_{\text{max}}, |\mathbf{h}_i^* \mathbf{r}_{k-1}|/\gamma_i, \cdot)$
    $\tilde{I} = I \cup J$
    $\hat{\mathbf{x}}_{[\tilde{I}],t} = \texttt{MMSE-rec}(\mathbf{y}_t, \mathbf{H}, \mathbf{R}_t, \hat{\mathbf{x}}_t^-, \mathbf{P}_t^-, \tilde{I})$
    $I := \texttt{indexmax}(K_{\text{max}}, |\hat{x}_{i,t}|, \tilde{I})$
    $\hat{\mathbf{x}}_{[I],t} = \texttt{MMSE-rec}(\mathbf{y}_t, \mathbf{H}, \mathbf{R}_t, \hat{\mathbf{x}}_t^-, \mathbf{P}_t^-, I)$;
    $\hat{\mathbf{x}}_{[I^c],t} = \mathbf{0}$
    $\mathbf{r}_k = \mathbf{y}_t - \mathbf{H}_{[\cdot,I]} \hat{\mathbf{x}}_{[I],t}$
until $(\|\mathbf{r}_k\|_2 > \|\mathbf{r}_{k-1}\|_2)$
Output: $\hat{\mathbf{x}}_t$ and I

set $\mathbf{P}_{[I,I^c],t} = \mathbf{0}$ and $\mathbf{P}_{[I^c,I],t} = \mathbf{0}$, in line with the approximate MMSE reconstruction (6).

The reconstruction $\hat{\mathbf{x}}_t$ is subsequently used to linearly predict the state at $t+1$ by $\hat{\mathbf{x}}_{t+1}^- = \mathbf{F}_t \hat{\mathbf{x}}_t$, and the error covariance matrix is propagated by the equation, $\mathbf{P}_{t+1}^- = \mathbf{F}_t \mathbf{P}_t \mathbf{F}_t^* + \mathbf{Q}_t$. Putting these blocks together we develop a dynamic subspace pursuit (DSP), which can be thought of an instance of a broader class of dynamic IGS algorithms.

Algorithm 3 : Dynamic Subspace Pursuit (DSP)

Initialization $\hat{\mathbf{x}}_0^-$ and $\mathbf{P}_0^-$
for $t = 0, \ldots$ **do**
    %Measurement update
    $[\mathbf{x}_t, I] = \texttt{PrSP}(\mathbf{y}_t, \mathbf{H}, \mathbf{R}_t, \hat{\mathbf{x}}_t^-, \mathbf{P}_t^-, K_{\text{max}})$
    $\mathbf{W}_t = (\mathbf{H}_{[\cdot,I]} \mathbf{P}_{[I,I],t}^- \mathbf{H}_{[\cdot,I]}^* + \mathbf{R}_t)$
    $\mathbf{P}_{[I,I],t} = (\mathbf{I}_{|I|} - \mathbf{P}_{[I,I],t}^- \mathbf{H}_{[\cdot,I]}^* \mathbf{W}_t^{-1} \mathbf{H}_{[\cdot,I]}) \mathbf{P}_{[I,I],t}^-$
    $\mathbf{P}_{[I^c,I^c],t} = \mathbf{P}_{[I^c,I^c],t}^-$; $\mathbf{P}_{[I,I^c],t} = \mathbf{0}$; $\mathbf{P}_{[I^c,I],t} = \mathbf{0}$
    %Prediction
    $\hat{\mathbf{x}}_{t+1}^- = \mathbf{F}_t \hat{\mathbf{x}}_t$
    $\mathbf{P}_{t+1}^- = \mathbf{F}_t \mathbf{P}_t \mathbf{F}_t^* + \mathbf{Q}_t$
end for

4. EXPERIMENTS AND RESULTS

We evaluate DSP with respect to static SP and convex relaxation based basis pursuit denoising (BPDN) [1] algorithms. We also show the performance of a 'genie-aided' Kalman filter (KF) which provides a bound for MMSE-based reconstruction of linear processes. The genie-aided approach is given the sparsity pattern $I_{x,t}$ a priori, and hence the bound is not necessarily tight.

4.1. Signal generation

As an example we consider a sparse process, $N = 200$, $K = 10$ and number of snapshots $T = 100$, with oscillating coefficients according to an AR-model as in (1) with $\alpha_{ij} = \alpha \equiv -0.8$, and $\mathbf{Q}_t = \sigma_w^2 \mathbf{I}_N$. The sparsity pattern transitions, $I_{x,t} \to I_{x,t+1}$, are determined by transition probabilities λ_{ji} which are set in the experiments below. The transition of each active state $j \in I_{x,t}$ is generated by a first-order Markov chain with λ_{ji}. If two states in $I_{x,t}$ happen to transition into one, a new state is randomly assigned to $I_{x,t+1}$, to ensure that the sparsity level is constant in the experiment.

The entries of the sensing matrix, $\mathbf{H}$, are set by randomly drawing from a Gaussian distribution $\mathcal{N}(0,1)$ followed by unit-norm column scaling. The measurement noise covariance matrix has form $\mathbf{R}_t = \sigma_n^2 \mathbf{I}_M$. Process and measurement noise are generated as $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t)$ and $\mathbf{n}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$, respectively.

In the experiment, the signal-to-measurement noise level

$$\text{SMNR} \triangleq \frac{\mathrm{E}\left[\sum_t \|\mathbf{x}_t\|_2^2\right]}{\mathrm{E}\left[\sum_t \|\mathbf{n}_t\|_2^2\right]}, \tag{8}$$

is varied, while fixing $\mathrm{E}[\|\mathbf{x}_t\|_2^2] \equiv 1$ so that $\sigma_n^2 = \frac{1}{M \times \text{SMNR}}$. We also vary the fraction of measurements $\kappa = M/N$.

The signal-to-reconstruction error ratio, defined as

$$\text{SRER} \triangleq \frac{\mathrm{E}\left[\sum_t \|\mathbf{x}_t\|_2^2\right]}{\mathrm{E}\left[\sum_t \|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_2^2\right]}, \tag{9}$$

is used as a performance measure and is the inverse of the normalized MSE. Note that SRER = 0 dB, i.e. no reconstruction gain, is equivalent of using $\hat{\mathbf{x}}_t = \mathbf{0}$.

4.2. Algorithm initialization

For the predictive algorithms—DSP and genie-aided KF—tested below we use the mean and variance of an autoregressive process as initial values, $\hat{\mathbf{x}}_0^- = \mathbf{0}$ and $\mathbf{P}_0^- = \sigma_x^2 \mathbf{I}_N$ where $\sigma_x^2 = \frac{\sigma_w^2}{1-\alpha^2}$. In these algorithms we set $K_{\text{max}} = K$ for consistent comparisons, although strict equality is not a requirement.

Here we mention that BPDN [1] solves

$$\hat{\mathbf{x}}_t = \arg\min_{\mathbf{x}_t \in \mathbb{R}^N} \|\mathbf{x}_t\|_1 \text{ subject to } \|\mathbf{y}_t - \mathbf{H}\mathbf{x}_t\|_2 \leq \varepsilon, \tag{10}$$

where the slack parameter ε is determined by the measurement noise power, as $\varepsilon = \sqrt{\sigma_n^2 (M + (2\sqrt{2M}))}$ following [11, 12]. Note that BPDN does not provide a K-element solution. It is also unable to use prediction. The code for BPDN is taken from the l_1-magic toolbox.

4.3. Results

In the first experiment we set λ_{ji} according to (3). We ran 100 Monte Carlo simulations, where a new realization of $\{\mathbf{x}_t, \mathbf{y}_t\}_{t=1}^T$ and $\mathbf{H}$ was generated for each run. The performance of BPDN, static SP and DSP, when varying SMNR at a fixed fraction of measurements $\kappa = 0.25$, is compared in figure 2. DSP overtakes static BPDN and SP low SMNR levels, while exhibiting a similar graceful degradation as BPDN since both take into account the noise level. As SMNR increases the gains of DSP are substantial, from approximately +2 dB to +14 dB in the SMNR range 7 to 30 dB.

In figure 3 the SMNR-level is fixed to 15 dB and instead κ is varied. A threshold effect is visible as κ decreases below 0.2 and detection performance drops rapidly. DSP consistently better than static SP across all levels of κ, but diminishes as κ increases. BPDN, however, does not adapt well to improving signal conditions.

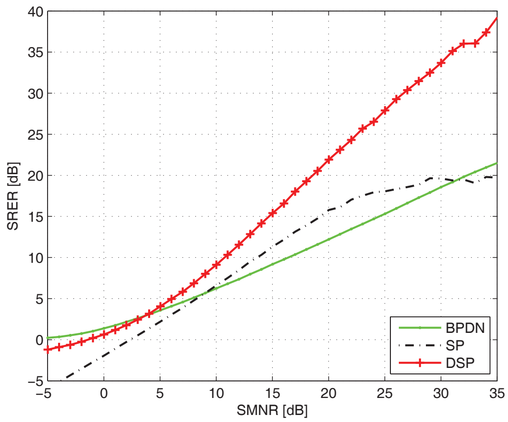

Fig. 2. SRER versus SMNR at $\kappa = 0.25$ and λ_{ji} according to (3).

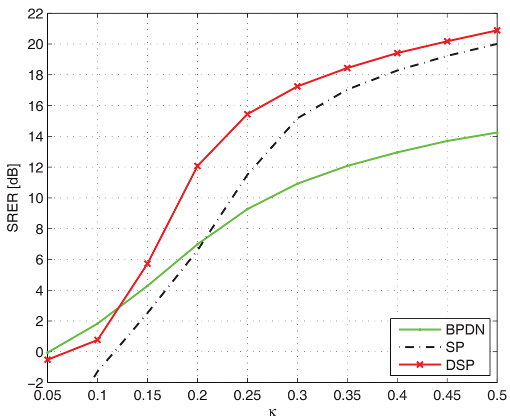

Fig. 3. SRER versus $\kappa = M/N$ at SMNR = 15 dB and λ_{ji} according to (3).

Finally we let $\lambda_{ji} = \delta(i - j)$, i.e. a fixed but unknown sparsity pattern, and compare DSP to the bound provided by the genie-aided KF, illustrated in figure 4. The performance of DSP is nearly identical as before and approaches that of the KF. At SMNR 20 dB, the gap is about 10.6 and 3.6 dB for SP and DSP respectively.

5. CONCLUSIONS

We have developed a dynamic subspace pursuit that can use sequential predictions for dynamic compressive sensing, as an instance of a broader class of dynamic IGS algorithms. It incorporates prior statistical information using linear MMSE reconstruction and weighted correlations. The algorithm was experimentally tested on a sparse signal with oscillating coefficients and evolving sparsity pattern. The results show that it exhibits graceful degradation at low SMNR regions while capable of yielding substantial performance gains as the SMNR level increases.

6. ACKNOWLEDGEMENT

The authors would like to thank E. Candes and J. Romberg making the l_1-magic toolbox available online.

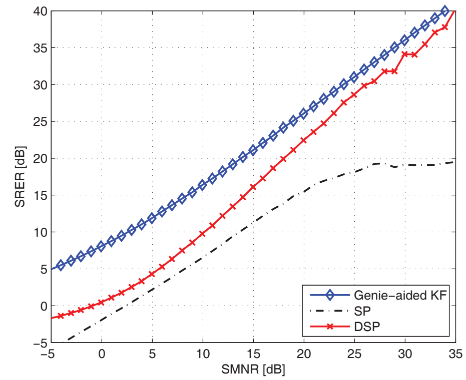

Fig. 4. SRER versus SMNR at $\kappa = 0.25$ and $\lambda_{ji} = \delta(i - j)$.

7. REFERENCES

[1] E.J. Candes and M.B. Wakin, “An introduction to compressive sampling,” *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, march 2008.

[2] J.A. Tropp and A.C. Gilbert, “Signal recovery from random measurements via orthogonal matching pursuit,” *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, dec. 2007.

[3] Wei Dai and O. Milenkovic, “Subspace pursuit for compressive sensing signal reconstruction,” *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2230 –2249, may 2009.

[4] Michael Lustig, David Donoho, and John M. Pauly, “Sparse mri: The application of compressed sensing for rapid mr imaging,” *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

[5] D. Sundman, S. Chatterjee, and M. Skoglund, “On the use of compressive sampling for wide-band spectrum sensing,” in *Signal Process. and Inf. Tech. (ISSPIT), IEEE Int. Symp. on*, dec. 2010, pp. 354 –359.

[6] N. Vaswani, “Kalman filtered compressed sensing,” in *15th IEEE Int. Conf. on Image Processing (ICIP).*, oct. 2008, pp. 893–896.

[7] A. Carmi, P. Gurfil, and D. Kanevsky, “Methods for sparse signal recovery using kalman filtering with embedded pseudo-measurement norms and quasi-norms,” *IEEE Trans. Signal Process.*, vol. 58, no. 4, pp. 2405 –2409, april 2010.

[8] Zhilin Zhang and B.D. Rao, “Sparse signal recovery in the presence of correlated multiple measurement vectors,” in *IEEE Int. Conf. Acoustics Speech and Signal Processing (ICASSP)*, march 2010, pp. 3986–3989.

[9] N. Vaswani, “Ls-cs-residual (ls-cs): Compressive sensing on least squares residual,” *IEEE Trans. Signal Process.*, vol. 58, no. 8, pp. 4108 –4120, aug. 2010.

[10] T. Kailath, A.H. Sayed, and B. Hassibi, *Linear Estimation*, Prentice Hall, 2000.

[11] E.J Candes, J. Romberg, and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements,” *Comm. Pure Appl. Math*, vol. 59, no. 8, pp. 1207–1223, 2006.

[12] S. Chatterjee, D. Sundman, and M. Skoglund, “Robust matching pursuit for recovery of gaussian sparse signal,” in *Digital Signal Process. Workshop and IEEE Signal Process. Education Workshop (DSP/SPE)*, jan. 2011, pp. 420–424.