

Dynamic Compressive Sensing of Time-Varying Signals Via Approximate Message Passing

Justin Ziniel, *Student Member, IEEE*, and Philip Schniter, *Senior Member, IEEE*

Abstract—In this work the dynamic compressive sensing (CS) problem of recovering sparse, correlated, time-varying signals from sub-Nyquist, non-adaptive, linear measurements is explored from a Bayesian perspective. While there has been a handful of previously proposed Bayesian dynamic CS algorithms in the literature, the ability to perform inference on high-dimensional problems in a computationally efficient manner remains elusive. In response, we propose a probabilistic dynamic CS signal model that captures both amplitude and support correlation structure, and describe an approximate message passing algorithm that performs soft signal estimation and support detection with a computational complexity that is linear in all problem dimensions. The algorithm, DCS-AMP, can perform either causal filtering or non-causal smoothing, and is capable of learning model parameters adaptively from the data through an expectation-maximization learning procedure. We provide numerical evidence that DCS-AMP performs within 3 dB of oracle bounds on synthetic data under a variety of operating conditions. We further describe the result of applying DCS-AMP to two real dynamic CS datasets, as well as a frequency estimation task, to bolster our claim that DCS-AMP is capable of offering state-of-the-art performance and speed on real-world high-dimensional problems.

Index Terms—Approximate message passing (AMP), belief propagation, compressed sensing, expectation-maximization algorithms, time-varying sparse signals, Kalman filters, dynamic compressive sensing, statistical signal processing.

I. INTRODUCTION

IN this work, we consider the *dynamic compressive sensing* (dynamic CS) problem, in which a sparse, vector-valued time series is recovered from a second time series of noisy, sub-Nyquist, linear measurements. Such a problem finds application in, e.g., dynamic MRI [2], high-speed video capture [3], and underwater channel estimation [4].

Framed mathematically, the objective of the dynamic CS problem is to recover the time series $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}\}$, where $\mathbf{x}^{(t)} \in \mathbb{C}^N$ is the signal at timestep t , from a time series of

measurements, $\{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(T)}\}$. Each $\mathbf{y}^{(t)} \in \mathbb{C}^M$ is obtained from the linear measurement process,

$$\mathbf{y}^{(t)} = \mathbf{A}^{(t)}\mathbf{x}^{(t)} + \mathbf{e}^{(t)}, \quad t = 1, \dots, T, \quad (1)$$

with $\mathbf{e}^{(t)}$ representing corrupting noise. The measurement matrix $\mathbf{A}^{(t)}$ (which may be time-varying or time-invariant, i.e., $\mathbf{A}^{(t)} = \mathbf{A} \forall t$) is known in advance, and is generally wide, leading to an underdetermined system of equations. The problem is regularized by assuming that $\mathbf{x}^{(t)}$ is sparse (or compressible),¹ having relatively few non-zero (or large) entries.

In many real-world scenarios, the underlying time-varying sparse signal exhibits substantial temporal correlation. This temporal correlation may manifest itself in two interrelated ways: (i) the support of the signal may change slowly over time [2], [3], [5], and (ii) the amplitudes of the large coefficients may vary smoothly in time.

In such scenarios, incorporating an appropriate model of temporal structure into a recovery technique makes it possible to drastically outperform structure-agnostic CS algorithms. From an analytical standpoint, Vaswani and Lu demonstrate that the restricted isometry property (RIP) sufficient conditions for perfect recovery in the dynamic CS problem are significantly weaker than those found in the traditional single measurement vector (SMV) CS problem when accounting for the additional structure [6]. In this work, we take a Bayesian approach to modeling this structure, which contrasts those dynamic CS algorithms inspired by convex relaxation, such as the Dynamic LASSO [5] and the Modified-CS algorithm [6]. Our Bayesian framework is also distinct from those hybrid techniques that blend elements of Bayesian dynamical models like the Kalman filter with more traditional CS approaches of exploiting sparsity through convex relaxation [2], [7] or greedy methods [8].

In particular, we propose a probabilistic model that treats the time-varying signal support as a set of independent binary Markov processes and the time-varying coefficient amplitudes as a set of independent Gauss-Markov processes. As detailed in Section II, this model leads to coefficient marginal distributions that are Bernoulli-Gaussian (i.e., “spike-and-slab”). Later, in Section V, we describe a generalization of the aforementioned model that yields Bernoulli-Gaussian-mixture coefficient marginals with an arbitrary number of mixture components. The models that we propose thus differ substantially from those used in other Bayesian approaches to dynamic CS, [9] and [10]. In particular, Sejdinović *et al.* [9] combine a linear Gaussian dynamical system model with a sparsity-promoting

¹Without loss of generality, we assume $\mathbf{x}^{(t)}$ is sparse/compressible in the canonical basis. Other sparsifying bases can be incorporated into the measurement matrix $\mathbf{A}^{(t)}$ without changing our model.

Manuscript received May 18, 2012; revised April 15, 2013; accepted June 29, 2013. Date of publication July 11, 2013; date of current version September 23, 2013. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Piotr Indyk. This work was supported in part by the NSF-I/UCRC Grant IIP-0968910, NSF Grant CCF-1018368, DARPA/ONR Grant N66001-10-1-4090, and an allocation of computing time from the Ohio Supercomputer Center. Portions of this paper were previously presented at the Asilomar Conference on Signals, Systems, and Computing, Pacific Grove, CA, USA, November 2010 [1].

The authors are with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210 USA (e-mail: zinielj@ece.osu.edu; schniter@ece.osu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2013.2273196

Gaussian-scale-mixture prior, while Shahrabi *et al.* [10] employ a particular spike-and-slab Markov model that couples amplitude evolution together with support evolution.

Our inference method also differs from those used in the alternative Bayesian dynamic CS algorithms [9] and [10]. In [9], Sejdinović *et al.* perform inference via a sequential Monte Carlo sampler [11]. Sequential Monte Carlo techniques are appealing for their applicability to complicated non-linear, non-Gaussian inference tasks like the Bayesian dynamic CS problem. Nevertheless, there are a number of important practical issues related to selection of the importance distribution, choice of the resampling method, and the number of sample points to track, since in principle one must increase the number of points exponentially over time to combat degeneracy [11]. Additionally, Monte Carlo techniques can be computationally expensive in high-dimensional inference problems. An alternative inference procedure that has recently proven successful in a number of applications is loopy belief propagation (LBP) [12]. In [10], Shahrabi *et al.* extend the conventional LBP method proposed in [13] for standard CS under a sparse measurement matrix \mathbf{A} to the case of dynamic CS under sparse $\mathbf{A}^{(t)}$. Nevertheless, the confinement to sparse measurement matrices is very restrictive, and, without this restriction, the methods of [10], [13] become computationally intractable.

Our inference procedure is based on the recently proposed framework of approximate message passing (AMP) [14], and in particular its “turbo” extension [15]. AMP, an unconventional form of LBP, was originally proposed for standard CS with a dense measurement matrix [14], and its noteworthy properties include: (i) a rigorous analysis (as $M, N \rightarrow \infty$ with M/N fixed, under i.i.d. sub-Gaussian \mathbf{A}) establishing that its solutions are governed by a state-evolution whose fixed points are optimal in several respects [16], and (ii) extremely fast runtimes (as a consequence of the fact that it needs relatively few iterations, each requiring only one multiplication by \mathbf{A} and its transpose). The turbo-AMP framework originally proposed in [15] offers a way to extend AMP to structured-sparsity problems such as compressive imaging [17], joint communication channel/symbol estimation [18], and—as we shall see in this work—the dynamic CS problem.

Our work makes several contributions to the existing literature on dynamic CS. First and foremost, the DCS-AMP algorithm that we develop offers an unrivaled combination of speed (e.g., its computational complexity grows only linearly in the problem dimensions M, N , and T) and reconstruction accuracy, as we demonstrate on both synthetic and real-world signals. Ours is the first work to exploit the speed and accuracy of loopy belief propagation (and, in particular, AMP) in the dynamic CS setting, accomplished by embedding AMP within a larger Bayesian inference algorithm. Second, we propose an expectation-maximization [19] procedure to automatically learn the parameters of our statistical model, as described in Section IV, avoiding a potentially complicated “tuning” problem. The ability to automatically calibrate algorithm parameters is especially important when working with real-world data, but is not provided by many of the existing dynamic CS algorithms (e.g., [2], [5]–[9]). In addition, our learned model parameters provide a convenient and interpretable characterization of time-varying signals in a way that, e.g., Lagrange

multipliers do not. Third, DCS-AMP provides a unified means of performing both filtering, where estimates are obtained sequentially using only past observations, and smoothing, where each estimate enjoys the knowledge of past, current, and future observations. In contrast, the existing dynamic CS schemes can support either filtering, or smoothing, but not both.

A. Notation

Boldfaced lower-case letters, e.g., \mathbf{a} , denote column vectors, while boldfaced upper-case letters, e.g., \mathbf{A} , denote matrices. The letter t is strictly used to index a timestep, $t = 1, 2, \dots, T$, the letter n is strictly used to index the coefficients of a signal, $n = 1, \dots, N$, and the letter m is strictly used to index the measurements, $m = 1, \dots, M$. The superscript (t) indicates a timestep-dependent quantity, while a superscript without parentheses, such as k , indicates a quantity whose value changes according to some algorithmic iteration index k . Subscript notations such as $x_n^{(t)}$ are used to denote the n th element of the vector $\mathbf{x}^{(t)}$, while set subscript notation, e.g., $\mathbf{x}_{\mathcal{S}}^{(t)}$, denotes the sub-vector of $\mathbf{x}^{(t)}$ consisting of indices contained in \mathcal{S} . The m th row of the matrix \mathbf{A} is denoted by \mathbf{a}_m^T , an M -by- M identity matrix is denoted by \mathbf{I}_M , and a length- N vector of ones is given by $\mathbf{1}_N$. Finally, $\mathcal{CN}(\mathbf{a}; \mathbf{b}, \mathbf{C})$ refers to the circularly symmetric complex normal distribution that is a function of the vector \mathbf{a} , with mean \mathbf{b} and covariance matrix \mathbf{C} .

II. SIGNAL MODEL

We assume that the measurement process can be accurately described by the linear model of (1). We further assume that $\mathbf{A}^{(t)} \in \mathbb{C}^{M \times N}$, $t = 1, \dots, T$, are measurement matrices known in advance, whose columns have been scaled to be of unit norm.² We model the noise as a stationary, circularly symmetric, additive white Gaussian noise (AWGN) process, with $\mathbf{e}^{(t)} \sim \mathcal{CN}(\mathbf{0}, \sigma_e^2 \mathbf{I}_M) \forall t$.

As noted in Section I, the sparse time series, $\{\mathbf{x}^{(t)}\}_{t=1}^T$, often exhibits a high degree of correlation from one timestep to the next. In this work, we model this correlation through a slow time-variation of the signal support, and a smooth evolution of the amplitudes of the non-zero coefficients. To do so, we introduce two hidden random processes, $\{\mathbf{s}^{(t)}\}_{t=1}^T$ and $\{\boldsymbol{\theta}^{(t)}\}_{t=1}^T$. The binary vector $\mathbf{s}^{(t)} \in \{0, 1\}^N$ describes the support of $\mathbf{x}^{(t)}$, denoted $\mathcal{S}^{(t)}$, while the vector $\boldsymbol{\theta}^{(t)} \in \mathbb{C}^N$ describes the amplitudes of the active elements of $\mathbf{x}^{(t)}$. Together, $\mathbf{s}^{(t)}$ and $\boldsymbol{\theta}^{(t)}$ completely characterize $\mathbf{x}^{(t)}$ as follows:

$$x_n^{(t)} = s_n^{(t)} \cdot \theta_n^{(t)} \quad \forall n, t. \quad (2)$$

Therefore, $s_n^{(t)} = 0$ sets $x_n^{(t)} = 0$ and $n \notin \mathcal{S}^{(t)}$, while $s_n^{(t)} = 1$ sets $x_n^{(t)} = \theta_n^{(t)}$ and $n \in \mathcal{S}^{(t)}$.

To model slow changes in the support $\mathcal{S}^{(t)}$ over time, we model the n th coefficient’s support across time, $\{s_n^{(t)}\}_{t=1}^T$, as a Markov chain defined by two transition probabilities: $p_{10} \triangleq \Pr\{s_n^{(t)} = 1 \mid s_n^{(t-1)} = 0\}$, and $p_{01} \triangleq \Pr\{s_n^{(t)} = 0 \mid s_n^{(t-1)} = 1\}$, and employ independent chains across $n = 1, \dots, N$. We further assume that each Markov chain operates in steady-state,

²Our algorithm can be generalized to support $\mathbf{A}^{(t)}$ without equal-norm columns, a time-varying number of measurements, $M^{(t)}$, and real-valued matrices/signals as well.

such that $\Pr\{s_n^{(t)} = 1\} = \lambda \forall n, t$. This steady-state assumption implies that these Markov chains are completely specified by the parameters λ and p_{01} , which together determine the remaining transition probability $p_{10} = \lambda p_{01}/(1 - \lambda)$. Depending on how p_{01} is chosen, the prior distribution can favor signals that exhibit a nearly static support across time, or it can allow for signal supports that change substantially from timestep to timestep. For example, it can be shown that $1/p_{01}$ specifies the average run length of a sequence of ones in the Markov chains.

The second form of temporal structure that we capture in our signal model is the correlation in active coefficient amplitudes across time. We model this correlation through independent stationary steady-state Gauss-Markov processes for each n , wherein $\{\theta_n^{(t)}\}_{t=1}^T$ evolves in time according to

$$\theta_n^{(t)} = (1 - \alpha) \left(\theta_n^{(t-1)} - \zeta \right) + \alpha w_n^{(t)} + \zeta, \quad (3)$$

where $\zeta \in \mathbb{C}$ is the mean of the process, $w_n^{(t)} \sim \mathcal{CN}(0, \rho)$ is an i.i.d. circular white Gaussian perturbation, and $\alpha \in [0, 1]$ controls the temporal correlation. At one extreme, $\alpha = 0$, the amplitudes are totally correlated, (i.e., $\theta_n^{(t)} = \theta_n^{(t-1)}$), while at the other extreme, $\alpha = 1$, the amplitudes evolve according to an uncorrelated Gaussian random process with mean ζ .

At this point, we would like to make a few remarks about our signal model. First, due to (2), it is clear that $p(x_n^{(t)} | s_n^{(t)}, \theta_n^{(t)}) = \delta(x_n^{(t)} - s_n^{(t)} \theta_n^{(t)})$, where $\delta(\cdot)$ is the Dirac delta function. By marginalizing out $s_n^{(t)}$ and $\theta_n^{(t)}$, one finds that

$$p(x_n^{(t)}) = (1 - \lambda) \delta(x_n^{(t)}) + \lambda \mathcal{CN}(x_n^{(t)}; \zeta, \sigma^2), \quad (4)$$

where $\sigma^2 \triangleq \frac{\alpha \rho}{2 - \alpha}$ is the steady-state variance of $\theta_n^{(t)}$. Equation (4) is a Bernoulli-Gaussian or “spike-and-slab” distribution, which is an effective sparsity-promoting prior due to the point-mass at $x_n^{(t)} = 0$. Second, we observe that the amplitude random process, $\{\theta_n^{(t)}\}_{t=1}^T$, evolves independently from the sparsity pattern random process, $\{s_n^{(t)}\}_{t=1}^T$. As a result of this modeling choice, there can be significant hidden amplitudes $\theta_n^{(t)}$ associated with inactive coefficients (those for which $s_n^{(t)} = 0$). Consequently, $\theta_n^{(t)}$ should be viewed as the amplitude of $x_n^{(t)}$ conditioned on $s_n^{(t)} = 1$. Lastly, we note that higher-order Markov processes and/or more complex coefficient marginals could be considered within the framework we propose, however, to keep development simple, we restrict our attention to first-order Markov processes and Bernoulli-Gaussian marginals until Section V, where we describe an extension of the above signal model that yields Bernoulli-Gaussian-mixture marginals.

III. THE DCS-AMP ALGORITHM

In this section we will describe the DCS-AMP algorithm, which efficiently and accurately estimates the marginal posterior distributions of $\{x_n^{(t)}\}$, $\{\theta_n^{(t)}\}$, and $\{s_n^{(t)}\}$ from the observed measurements $\{\mathbf{y}^{(t)}\}_{t=1}^T$, thus enabling both soft estimation and soft support detection. The use of soft support information is particularly advantageous, as it means that the algorithm need never make a firm (and possibly erroneous) decision about the support that can propagate errors across many timesteps. As mentioned in Section I, DCS-AMP can perform either filtering or smoothing.

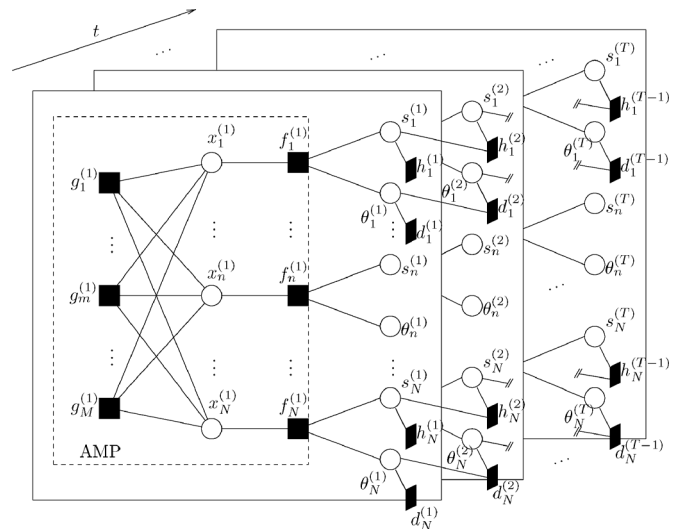


Fig. 1. Factor graph representation of the joint posterior distribution of (5).

The algorithm we develop is designed to exploit the statistical structure inherent in our signal model. By defining $\bar{\mathbf{y}}$ to be the collection of all measurements, $\{\mathbf{y}^{(t)}\}_{t=1}^T$ (and defining $\bar{\mathbf{x}}$, $\bar{\mathbf{s}}$, and $\bar{\boldsymbol{\theta}}$ similarly), the posterior joint distribution of the signal, support, and amplitude time series, given the measurement time series, can be expressed using Bayes' rule as

$$p(\bar{\mathbf{x}}, \bar{\mathbf{s}}, \bar{\boldsymbol{\theta}} | \bar{\mathbf{y}}) \propto \prod_{t=1}^T \left(\prod_{m=1}^M p(y_m^{(t)} | \mathbf{x}^{(t)}) \prod_{n=1}^N p(x_n^{(t)} | s_n^{(t)}, \theta_n^{(t)}) \times p(s_n^{(t)} | s_n^{(t-1)}) p(\theta_n^{(t)} | \theta_n^{(t-1)}) \right), \quad (5)$$

where \propto indicates proportionality up to a constant scale factor, $p(s_n^{(1)} | s_n^{(0)}) \triangleq p(s_n^{(1)})$, and $p(\theta_n^{(1)} | \theta_n^{(0)}) \triangleq p(\theta_n^{(1)})$. By inspecting (5), we see that the posterior joint distribution decomposes into the product of many distributions that only depend on small subsets of variables. A graphical representation of such decompositions is given by the *factor graph*, which is an undirected bipartite graph that connects the pdf “factors” of (5) with the random variables that constitute their arguments [20]. In Table I, we introduce the notation that we will use for the factors of our signal model, showing the correspondence between the factor labels and the underlying distributions they represent, as well as the specific functional form assumed by each factor. The associated factor graph for the posterior joint distribution of (5) is shown in Fig. 1, labeled according to Table I. Filled squares represent factors, while circles represent random variables.

As seen in Fig. 1, all of the variables needed at a given timestep can be visualized as lying in a plane, with successive planes stacked one after another in time. We will refer to these planes as “frames”. The temporal correlation of the signal supports is illustrated by the $h_n^{(t)}$ factor nodes that connect the $s_n^{(t)}$ variable nodes between neighboring frames. Likewise, the temporal correlation of the signal amplitudes is expressed by the interconnection of $d_n^{(t)}$ factor nodes and $\theta_n^{(t)}$ variable nodes. For visual clarity, these factor nodes have been omitted from the middle portion of the factor graph, appearing only at

TABLE I
THE FACTORS, UNDERLYING DISTRIBUTIONS, AND FUNCTIONAL FORMS ASSOCIATED WITH OUR SIGNAL MODEL

Factor	Distribution	Functional Form
$g_m^{(t)}(\mathbf{x}^{(t)})$	$p(y_m^{(t)} \mathbf{x}^{(t)})$	$\mathcal{CN}(y_m^{(t)}; \mathbf{a}_m^{(t)T} \mathbf{x}^{(t)}, \sigma_c^2)$
$f_n^{(t)}(x_n^{(t)}, s_n^{(t)}, \theta_n^{(t)})$	$p(x_n^{(t)} s_n^{(t)}, \theta_n^{(t)})$	$\delta(x_n^{(t)} - s_n^{(t)} \theta_n^{(t)})$
$h_n^{(1)}(s_n^{(1)})$	$p(s_n^{(1)})$	$(1 - \lambda)^{1-s_n^{(1)}} \lambda^{s_n^{(1)}}$
$h_n^{(t)}(s_n^{(t)}, s_n^{(t-1)})$	$p(s_n^{(t)} s_n^{(t-1)})$	$\begin{cases} (1 - p_{10})^{1-s_n^{(t)}} p_{10}^{s_n^{(t)}}, & s_n^{(t-1)} = 0 \\ p_{01}^{1-s_n^{(t)}} (1 - p_{01})^{s_n^{(t)}}, & s_n^{(t-1)} = 1 \end{cases}$
$d_n^{(1)}(\theta_n^{(1)})$	$p(\theta_n^{(1)})$	$\mathcal{CN}(\theta_n^{(1)}; \zeta, \sigma^2)$
$d_n^{(t)}(\theta_n^{(t)}, \theta_n^{(t-1)})$	$p(\theta_n^{(t)} \theta_n^{(t-1)})$	$\mathcal{CN}(\theta_n^{(t)}; (1 - \alpha)\theta_n^{(t-1)} + \alpha\zeta, \alpha^2\rho)$

indices $n = 1$ and $n = N$, but should in fact be present for all indices $n = 1, \dots, N$. Since the measurements $\{y_m^{(t)}\}$ are observed variables, they have been incorporated into the $g_m^{(t)}$ factor nodes.

The algorithm that we develop can be viewed as an approximate implementation of belief propagation (BP) [21], a message passing algorithm for performing inference on factor graphs that describe probabilistic models. When the factor graph is cycle-free, belief propagation is equivalent to the more general sum-product algorithm [20], which is a means of computing the marginal functions that result from summing (or integrating) a multivariate function over all possible input arguments, with one argument held fixed, (i.e., marginalizing out all but one variable). In the context of BP, these marginal functions are the marginal distributions of random variables. Thus, given measurements $\bar{\mathbf{y}}$ and the factorization of the posterior joint distribution $p(\bar{\mathbf{x}}, \bar{\mathbf{s}}, \bar{\boldsymbol{\theta}}|\bar{\mathbf{y}})$, DCS-AMP computes (approximate) posterior marginals of $x_n^{(t)}$, $s_n^{(t)}$, and $\theta_n^{(t)}$. In “filtering” mode, our algorithm would therefore return, e.g., $p(x_n^{(t)}|\{\mathbf{y}^{(t)}\}_{t=1}^t)$, while in “smoothing” mode it would return $p(x_n^{(t)}|\{\mathbf{y}^{(t)}\}_{t=1}^T)$. From these marginals, one can compute, e.g., minimum mean-squared error (MMSE) estimates. The factor graph of Fig. 1 contains many short cycles, however, and thus the convergence of loopy BP cannot be guaranteed [20].³ Despite this, loopy BP has been shown to perform extremely well in a number of different applications, including turbo decoding [26], computer vision [27], and compressive sensing [13]–[15], [17], [28]–[30].

A. Message Scheduling

In loopy factor graphs, there are a number of ways to schedule, or sequence, the messages that are exchanged between nodes. The choice of a schedule can impact not only the rate of convergence of the algorithm, but also the likelihood of convergence as well [31]. We propose a schedule (an evolution of the “turbo” schedule proposed in [15]) for DCS-AMP that is straightforward to implement, suitable for both filtering and smoothing applications, and empirically yields quickly converging estimates under a variety of diverse operating conditions.

Our proposed schedule can be broken down into four distinct steps, which we will refer to using the mnemonics **(into)**, **(within)**, **(out)**, and **(across)**. At a particular timestep t , the

(into) step involves passing messages that provide current beliefs about the state of the relevant support variables, $\{s_n^{(t)}\}_{n=1}^N$, and amplitude variables, $\{\theta_n^{(t)}\}_{n=1}^N$, laterally *into* the dashed AMP box within frame t . (Recall Fig. 1.) The **(within)** step makes use of these incoming messages, together with the observations available in that frame, $\{y_m^{(t)}\}_{m=1}^M$, to exchange messages *within* the dashed AMP box of frame t , thus generating estimates of the marginal posteriors of the signal variables $\{x_n^{(t)}\}_{n=1}^N$. Using these posterior estimates, the **(out)** step propagates messages *out* of the dashed AMP box, providing updated beliefs about the state of $\{s_n^{(t)}\}_{n=1}^N$ and $\{\theta_n^{(t)}\}_{n=1}^N$. Lastly, the **(across)** step involves transmitting messages *across* neighboring frames, using the updated beliefs about $\{s_n^{(t)}\}_{n=1}^N$ and $\{\theta_n^{(t)}\}_{n=1}^N$ to influence the beliefs about $\{s_n^{(t+1)}\}_{n=1}^N$ and $\{\theta_n^{(t+1)}\}_{n=1}^N$ (or $\{s_n^{(t-1)}\}_{n=1}^N$ and $\{\theta_n^{(t-1)}\}_{n=1}^N$).

The procedures for filtering and smoothing both start in the same way. At the initial $t = 1$ frame, steps **(into)**, **(within)** and **(out)** are performed in succession. Next, step **(across)** is performed to pass messages from $\{s_n^{(1)}\}_{n=1}^N$ and $\{\theta_n^{(1)}\}_{n=1}^N$ to $\{s_n^{(2)}\}_{n=1}^N$ and $\{\theta_n^{(2)}\}_{n=1}^N$. Then at frame $t = 2$ the same set of steps are executed, concluding with messages propagating to $\{s_n^{(3)}\}_{n=1}^N$ and $\{\theta_n^{(3)}\}_{n=1}^N$. This process continues until steps **(into)**, **(within)** and **(out)** have been completed at the terminal frame, T . At this point, DCS-AMP has completed what we call a single forward pass. If the objective was to perform filtering, DCS-AMP terminates at this point, since only causal measurements have been used to estimate the marginal posteriors. If instead the objective is to obtain smoothed, non-causal estimates, then information begins to propagate backwards in time, i.e., step **(across)** moves messages from $\{s_n^{(T)}\}_{n=1}^N$ and $\{\theta_n^{(T)}\}_{n=1}^N$ to $\{s_n^{(T-1)}\}_{n=1}^N$ and $\{\theta_n^{(T-1)}\}_{n=1}^N$. Steps **(into)**, **(within)**, **(out)**, and **(across)** are performed at frame $T - 1$, with messages bound for frame $T - 2$. This continues until the initial frame is reached. At this point DCS-AMP has completed what we term as a single forward/backward pass. Multiple such passes, indexed by the variable k , can be carried out until a convergence criterion is met or a maximum number of passes has been performed.

B. Implementing the Message Passes

We now provide some additional details as to how the above four steps are implemented. To aid our discussion, in Fig. 2 we summarize the form of the messages that pass between the various factor graph nodes, focusing primarily on a single coefficient index n at an intermediate frame t . Directed edges indicate the direction that messages are moving. In the **(across)**

³However, it is worth noting that in the past decade much work has been accomplished in identifying specific situations under which loopy BP is guaranteed to converge, e.g., [16], [22]–[25].

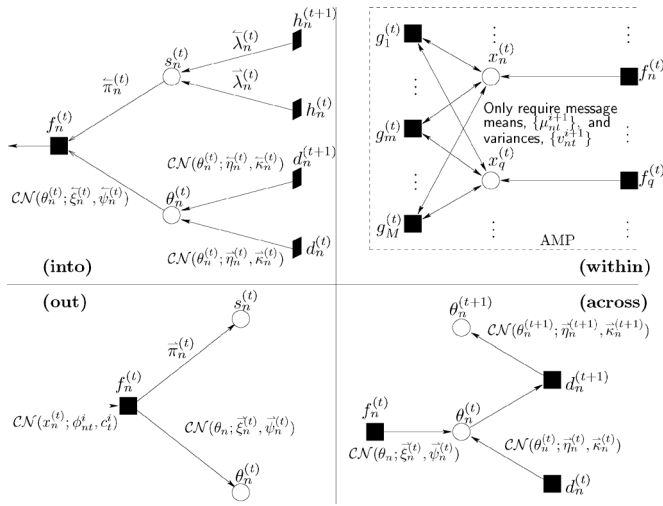


Fig. 2. A summary of the four message passing phases, including message notation and form. See the pseudocode of Table II for the precise message update computations.

phase, we only illustrate the messages involved in a forward pass for the amplitude variables, and leave out a graphic for the corresponding backward pass, as well as graphics for the support variable (**across**) phase. Note that, to be applicable at frame T , the factor node $d_n^{(t+1)}$ and its associated edge should be removed. The figure also introduces the notation that we adopt for the different variables that serve to parameterize the messages. We use the notation $\nu_{a \rightarrow b}(\cdot)$ to denote a message passing from node a to a connected node b . For Bernoulli message pdfs, we show only the non-zero probability, e.g., $\lambda_n^{-} = \nu_{h_n^{(t)} \rightarrow s_n^{(t)}}(s_n^{(t)} = 1)$.

To perform step (**into**), the messages from the factors $h_n^{(t)}$ and $h_n^{(t+1)}$ to $s_n^{(t)}$ are used to set $\bar{\pi}_n^{-}$, the message from $s_n^{(t)}$ to $f_n^{(t)}$. Likewise, the messages from the factors $d_n^{(t)}$ and $d_n^{(t+1)}$ to $\theta_n^{(t)}$ are used to determine the message from $\theta_n^{(t)}$ to $f_n^{(t)}$. When performing filtering, or the first forward pass of smoothing, no meaningful information should be conveyed from the $h_n^{(t+1)}$ and $d_n^{(t+1)}$ factors. This can be accomplished by initializing $\bar{\lambda}_n^{-}$, $\bar{\eta}_n^{-}$, $\bar{\kappa}_n^{-}$ with the values $(\frac{1}{2}, 0, \infty)$.

In step (**within**), messages must be exchanged between the $\{x_n^{(t)}\}_{n=1}^N$ and $\{g_m^{(t)}\}_{m=1}^M$ nodes. When $\mathbf{A}^{(t)}$ is not a sparse matrix, this will imply a dense network of connections between these nodes. Consequently, the standard sum-product algorithm would require us to evaluate multi-dimensional integrals of non-Gaussian messages that grow exponentially in number in both N and M . This approach is clearly infeasible for problems of any appreciable size, and thus we turn to a simplification known as *approximate message passing* (AMP) [14], [28].

At a high-level, AMP can be viewed as a simplification of loopy BP, employing central limit theorem arguments to approximate the sum of many non-Gaussian random variables as a Gaussian. Through a series of principled approximation steps (which become exact for sub-Gaussian \mathbf{A} matrices in the large-system limit [16]), AMP produces an iterative thresholding algorithm that requires only $\mathcal{O}(MN)$ operations, dominated by matrix-vector products, to obtain posteriors on the $\{x_n^{(t)}\}_{n=1}^N$

TABLE II
DCS-AMP STEPS FOR FILTERING MODE, OR THE FORWARD PORTION OF A SINGLE FORWARD/BACKWARD PASS IN SMOOTHING MODE. SEE FIG. 2 TO ASSOCIATE QUANTITIES WITH THE MESSAGES TRAVERSING THE FACTOR GRAPH

% Define soft-thresholding functions:	
$F_{nt}(\phi; c) \triangleq (1 + \gamma_{nt}(\phi; c))^{-1} \left(\frac{\bar{\psi}_n^{(t)} \phi + \bar{\xi}_n^{(t)} c}{\bar{\psi}_n^{(t)} + c} \right)$	(D1)
$G_{nt}(\phi; c) \triangleq (1 + \gamma_{nt}(\phi; c))^{-1} \left(\frac{\bar{\psi}_n^{(t)} c}{\bar{\psi}_n^{(t)} + c} \right) + \gamma_{nt}(\phi; c) F_{nt}(\phi; c) ^2$	(D2)
$F'_{nt}(\phi; c) \triangleq \frac{\partial}{\partial \phi} F_{nt}(\phi, c) = \frac{1}{c} G_{nt}(\phi; c)$	(D3)
$\gamma_{nt}(\phi; c) \triangleq \left(\frac{1 - \bar{\pi}_n^{(t)}}{\bar{\pi}_n^{(t)}} \right) \left(\frac{\bar{\psi}_n^{(t)} + c}{c} \right) \times \exp \left(- \left[\frac{\bar{\psi}_n^{(t)} \phi ^2 + \bar{\xi}_n^{(t)} * c \phi + \bar{\xi}_n^{(t)} c \phi^* - c \bar{\xi}_n^{(t)} ^2}{c(\bar{\psi}_n^{(t)} + c)} \right] \right)$	(D4)
% Begin passing messages . . .	
for $t = 1, \dots, T$:	
% Execute the (into) phase . . .	
$\bar{\pi}_n^{-} = \frac{\bar{\lambda}_n^{-} \cdot \bar{\lambda}_n^{-}}{(1 - \bar{\lambda}_n^{-}) \cdot (1 - \bar{\lambda}_n^{-}) + \bar{\lambda}_n^{-} \cdot \bar{\lambda}_n^{-}} \quad \forall n$	(A1)
$\bar{\psi}_n^{-} = \frac{\bar{\kappa}_n^{-} \cdot \bar{\kappa}_n^{-}}{\bar{\kappa}_n^{-} + \bar{\kappa}_n^{-}} \quad \forall n$	(A2)
$\bar{\xi}_n^{-} = \bar{\psi}_n^{-} \cdot \left(\frac{\bar{\eta}_n^{-}}{\bar{\kappa}_n^{-}} + \frac{\bar{\eta}_n^{-}}{\bar{\kappa}_n^{-}} \right) \quad \forall n$	(A3)
% Initialize AMP-related variables . . .	
$\forall m : z_{mt}^1 = y_m^{(t)}, \forall n : \mu_{nt}^1 = 0$, and $c_t^1 = 100 \cdot \sum_{n=1}^N \bar{\psi}_n^{(t)}$	
% Execute the (within) phase using AMP . . .	
for $i = 1, \dots, I$:	
$\phi_{nt}^i = \sum_{m=1}^M A_{mn}^{(t)} z_{mt}^i + \mu_{nt}^i \quad \forall n$	(A4)
$\mu_{nt}^{i+1} = F_{nt}(\phi_{nt}^i; c_t^i) \quad \forall n$	(A5)
$v_{nt}^{i+1} = G_{nt}(\phi_{nt}^i; c_t^i) \quad \forall n$	(A6)
$c_t^{i+1} = \sigma_c^2 + \frac{1}{M} \sum_{n=1}^N v_{nt}^{i+1}$	(A7)
$z_{mt}^{i+1} = y_m^{(t)} - \alpha_m^{(t)} \mathbf{T} \mu_t^{i+1} + \frac{z_{mt}^i}{M} \sum_{n=1}^N F'_{nt}(\phi_{nt}^i; c_t^i) \quad \forall m$	(A8)
end	
$\hat{x}_n^{(t)} = \mu_{nt}^{I+1} \quad \forall n$ % Store current estimate of $x_n^{(t)}$	(A9)
% Execute the (out) phase . . .	
$\bar{\pi}_n^{+} = \left(1 + \left(\frac{\bar{\pi}_n^{-}}{1 - \bar{\pi}_n^{-}} \right) \gamma_{nt}(\phi_{nt}^I; c_t^{I+1}) \right)^{-1} \quad \forall n$	(A10)
$(\bar{\xi}_n^{+}, \bar{\psi}_n^{+}) = \begin{cases} (\phi_n^I / \varepsilon, c_t^{I+1} / \varepsilon^2), & \bar{\pi}_n^{+} \leq \tau \\ (\phi_n, c_t^{I+1}), & \text{o.w.} \end{cases} \quad \forall n \quad (\varepsilon \ll 1)$	(A11)
% Execute the (across) phase forward in time . . .	
$\bar{\lambda}_n^{(t+1)} = \frac{p_{10}(1 - \bar{\lambda}_n^{(t)})(1 - \bar{\pi}_n^{(t)}) + (1 - p_{10}) \bar{\lambda}_n^{(t)} \bar{\pi}_n^{(t)}}{(1 - \bar{\lambda}_n^{(t)})(1 - \bar{\pi}_n^{(t)}) + \bar{\lambda}_n^{(t)} \bar{\pi}_n^{(t)}} \quad \forall n$	(A12)
$\bar{\eta}_n^{(t+1)} = (1 - \alpha) \left(\frac{\bar{\kappa}_n^{(t)} \bar{\psi}_n^{(t)}}{\bar{\kappa}_n^{(t)} + \bar{\psi}_n^{(t)}} \right) \left(\frac{\bar{\eta}_n^{(t)}}{\bar{\kappa}_n^{(t)}} + \frac{\bar{\xi}_n^{(t)}}{\bar{\psi}_n^{(t)}} \right) + \alpha \zeta \quad \forall n$	(A13)
$\bar{\kappa}_n^{(t+1)} = (1 - \alpha)^2 \left(\frac{\bar{\kappa}_n^{(t)} \bar{\psi}_n^{(t)}}{\bar{\kappa}_n^{(t)} + \bar{\psi}_n^{(t)}} \right) + \alpha^2 \rho \quad \forall n$	(A14)
end	

variable nodes. The specifics of the iterative thresholding algorithm will depend on the signal prior under which AMP is operating [28], but it is assumed that the joint prior decouples into independent (but not necessarily i.i.d.) priors on each coefficient $x_n^{(t)}$. See Appendix A for additional background on AMP.

By viewing $\nu_{f_n^{(t)} \rightarrow x_n^{(t)}}(\cdot)$ as a “local prior”⁴ for $x_n^{(t)}$, we can readily apply an off-the-shelf AMP algorithm (e.g., [28], [32], [33]) as a means of performing the message passes within the portions of the factor graph enclosed within the dashed boxes of Fig. 1 (only one such box is visible). The use of AMP with decoupled local priors within a larger message passing algorithm that accounts for statistical dependencies between signal coefficients was first proposed in [15], and further studied in [17],

⁴The AMP algorithm is conventionally run with static, i.i.d., priors for each signal coefficient. When utilized as a sub-component of a larger message passing algorithm on an expanded factor graph, the signal priors (from AMP’s perspective) will be changing in response to messages from the rest of the factor graph. We refer to these changing AMP priors as *local priors*.

[18], [29], [34], [35]. Here, we exploit this powerful “turbo” inference approach to account for the strong temporal dependencies inherent in the dynamic CS problem.

The local prior for our signal model is a Bernoulli-Gaussian, namely

$$\nu_{f_n^{(t)} \rightarrow x_n^{(t)}}(x_n^{(t)}) = \left(1 - \frac{\bar{\pi}_n^{(t)}}{\pi_n}\right) \delta(x_n^{(t)}) + \frac{\bar{\pi}_n^{(t)}}{\pi_n} \mathcal{CN}(x_n^{(t)}; \xi_n, \psi_n).$$

The appropriate AMP message update equations for this local prior follow a straightforward extension of the derivations outlined in [15], which considered the special case of a zero-mean Bernoulli-Gaussian prior. The specific AMP updates for our model are given by (A4)–(A8) in Table II.

After employing AMP to manage the message passing between the $\{x_n^{(t)}\}_{n=1}^N$ and $\{g_m^{(t)}\}_{m=1}^M$ nodes in step (**within**), messages must be propagated out of the dashed AMP box of frame t (step (**out**)) and either forward or backward in time (step (**across**)). While step (**across**) simply requires a straightforward application of the sum-product message computation rules, step (**out**) imposes several difficulties which we must address. For the remainder of this discussion, we focus on a novel approximation scheme for specifying the message $\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}(\cdot)$. Our objective is to arrive at a message approximation that introduces negligible error while still leading to a computationally efficient algorithm. A Gaussian message approximation is a natural choice, given the marginally Gaussian distribution of $\theta_n^{(t)}$. As we shall soon see, it is also a highly justifiable choice.

A routine application of the sum-product rules to the $f_n^{(t)}$ -to- $\theta_n^{(t)}$ message would produce the following expression:

$$\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}^{\text{exact}}(\theta_n^{(t)}) \triangleq \left(1 - \frac{\bar{\pi}_n^{(t)}}{\pi_n}\right) \mathcal{CN}(0; \phi_{nt}^i, c_t^i) + \frac{\bar{\pi}_n^{(t)}}{\pi_n} \mathcal{CN}(\theta_n^{(t)}; \phi_{nt}^i, c_t^i). \quad (6)$$

Unfortunately, the term $\mathcal{CN}(0; \phi_{nt}^i, c_t^i)$ prevents us from normalizing $\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}^{\text{exact}}(\theta_n^{(t)})$, because it is constant with respect to $\theta_n^{(t)}$. Therefore, the distribution on $\theta_n^{(t)}$ represented by (6) is improper. To provide intuition into why this is the case, it is helpful to think of $\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}(\theta_n^{(t)})$ as a message that conveys information about the value of $\theta_n^{(t)}$ based on the values of $x_n^{(t)}$ and $s_n^{(t)}$. If $s_n^{(t)} = 0$, then by (2), $x_n^{(t)} = 0$, thus making $\theta_n^{(t)}$ unobservable. The constant term in (6) reflects the uncertainty due to this unobservability through an infinitely broad, uninformative distribution for $\theta_n^{(t)}$.

To avoid an improper pdf, we modify how this message is derived by regarding our assumed signal model, in which $s_n^{(t)} \in \{0, 1\}$, as a limiting case of the model with $s_n^{(t)} \in \{\varepsilon, 1\}$ as $\varepsilon \rightarrow 0$. For any fixed positive ε , the resulting message $\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}(\cdot)$ is proper, given by

$$\begin{aligned} \nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}^{\text{mod}}(\theta_n^{(t)}) &= \left(1 - \Omega\left(\frac{\bar{\pi}_n^{(t)}}{\pi_n}\right)\right) \mathcal{CN}\left(\theta_n^{(t)}; \frac{1}{\varepsilon} \phi_{nt}^i, \frac{1}{\varepsilon^2} c_t^i\right) \\ &\quad + \Omega\left(\frac{\bar{\pi}_n^{(t)}}{\pi_n}\right) \mathcal{CN}\left(\theta_n^{(t)}; \phi_{nt}^i, c_t^i\right), \end{aligned} \quad (7)$$

where

$$\Omega(\pi) \triangleq \frac{\varepsilon^2 \pi}{(1 - \pi) + \varepsilon^2 \pi}. \quad (8)$$

The pdf in (7) is that of a binary Gaussian mixture. If we consider $\varepsilon \ll 1$, the first mixture component is extremely broad, while the second is more “informative,” with mean ϕ_{nt}^i and variance c_t^i . The relative weight assigned to each component Gaussian is determined by the term $\Omega(\frac{\bar{\pi}_n^{(t)}}{\pi_n})$. Notice that the limit of this weighting term is the simple indicator function

$$\lim_{\varepsilon \rightarrow 0} \Omega(\pi) = \begin{cases} 0 & \text{if } 0 \leq \pi < 1, \\ 1 & \text{if } \pi = 1. \end{cases} \quad (9)$$

Since we cannot set $\varepsilon = 0$, we instead fix a small positive value, e.g., $\varepsilon = 10^{-7}$. In this case, (7) could then be used as the outgoing message. However, this presents a further difficulty: propagating a binary Gaussian mixture forward in time would lead to an exponential growth in the number of mixture components at subsequent timesteps. This difficulty is a familiar one in the context of switched linear dynamical systems based on conditional Gaussian models, since such models are not closed under marginalization [36]. To avoid the exponential growth in the number of mixture components, we collapse our binary Gaussian mixture to a single Gaussian component, an approach sometimes referred to as a Gaussian sum approximation [37], [38]. This can be justified by the fact that, for $\varepsilon \ll 1$, $\Omega(\cdot)$ behaves nearly like the indicator function in (9), in which case one of the two Gaussian components will typically have negligible mass.

To carry out the Gaussian sum approximation, we propose the following two schemes. The first is to simply choose a threshold τ that is slightly smaller than 1 and, using (9) as a guide, threshold $\frac{\bar{\pi}_n^{(t)}}{\pi_n}$ to choose between the two Gaussian components of (7). The resultant message is thus

$$\nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}(\theta_n^{(t)}) = \mathcal{CN}\left(\theta_n^{(t)}; \xi_n, \psi_n\right), \quad (10)$$

with ξ_n and ψ_n chosen according to

$$\left(\xi_n, \psi_n\right) = \begin{cases} \left(\frac{1}{\varepsilon} \phi_{nt}^i, \frac{1}{\varepsilon^2} c_t^i\right), & \frac{\bar{\pi}_n^{(t)}}{\pi_n} \leq \tau \\ \left(\phi_{nt}^i, c_t^i\right), & \frac{\bar{\pi}_n^{(t)}}{\pi_n} > \tau \end{cases}. \quad (11)$$

The second approach is to perform a second-order Taylor series approximation of $-\log \nu_{f_n^{(t)} \rightarrow \theta_n^{(t)}}^{\text{mod}}(\theta_n^{(t)})$ with respect to $\theta_n^{(t)}$. The resultant quadratic form in $\theta_n^{(t)}$ can be viewed as the logarithm of a Gaussian kernel with a particular mean and variance, which can be used to parameterize a single Gaussian message, as described in [29]. The latter approach has the advantage of being parameter-free. Empirically, we find that this latter approach works well when changes in the support occur infrequently, e.g., $p_{01} < 0.025$, while the former approach is better suited to more dynamic environments.

In Table II we provide a pseudo-code implementation of our proposed DCS-AMP algorithm that gives the explicit message update equations appropriate for performing a single forward pass. The interested reader can find an expanded derivation of the messages in [39]. The primary computational burden

of DCS-AMP is computing the messages passing between the $\{x_n^{(t)}\}$ and $\{g_n^{(t)}\}$ nodes, a task which can be performed efficiently using matrix-vector products involving $\mathbf{A}^{(t)}$ and $\mathbf{A}^{(t)H}$. The resulting overall complexity of DCS-AMP is therefore $\mathcal{O}(TMN)$ flops (flops-per-pass) when filtering (smoothing).⁵ The storage requirements are $\mathcal{O}(N)$ and $\mathcal{O}(TN)$ complex numbers when filtering and smoothing, respectively.

IV. LEARNING THE SIGNAL MODEL PARAMETERS

The signal model of Section II is specified by the Markov chain parameters λ , p_{01} , the Gauss-Markov parameters ζ , α , ρ , and the AWGN variance σ_e^2 . It is likely that some or all of these parameters will require tuning in order to best match the unknown signal. To this end, we develop an expectation-maximization (EM) [19] algorithm that works together with the message passing procedure described in Section III-A to learn all of the model parameters in an iterative fashion from the data.

The EM algorithm is appealing for two principal reasons. First, the EM algorithm is a well-studied and principled means of parameter estimation. At every EM iteration, the likelihood is guaranteed to increase until convergence to a local maximum occurs [40]. For multimodal likelihoods, local maxima will, in general, not coincide with the global maximum, but a judicious initialization of parameters can help in ensuring the EM algorithm reaches the global maximum [40]. The second appealing feature of the EM algorithm lies in the fact that its expectation step leverages quantities that have already been computed in the process of executing DCS-AMP, making the EM procedure computationally efficient.

We let $\Gamma \triangleq \{\lambda, p_{01}, \zeta, \alpha, \rho, \sigma_e^2\}$ denote the set of all model parameters, and let Γ^k denote the set of parameter estimates at the k th EM iteration. The objective of the EM procedure is to find parameter estimates that maximize the data likelihood $p(\bar{\mathbf{y}} | \Gamma)$. Since it is often computationally intractable to perform this maximization, the EM algorithm incorporates additional “hidden” data and iterates between two steps: (i) evaluating the conditional expectation of the log likelihood of the hidden data given the observed data, $\bar{\mathbf{y}}$, and the current estimates of the parameters, Γ^k , and (ii) maximizing this expected log likelihood with respect to the model parameters. For all parameters except the noise variance, σ_e^2 , we use $\bar{\mathbf{s}}$ and $\bar{\boldsymbol{\theta}}$ as the hidden data, while for σ_e^2 we use $\bar{\mathbf{x}}$.

Before running DCS-AMP, the model parameters are initialized using any available prior knowledge. If operating in smoothing mode, DCS-AMP performs an initial forward/backward pass, as described in Section III-A. Upon completing this first pass, estimates of the marginal posterior distributions are available for each of the underlying random variables. Additionally, belief propagation can provide pairwise joint posterior distributions, e.g., $p(s_n^{(t)}, s_n^{(t-1)} | \bar{\mathbf{y}})$, for any variable nodes connected by a common factor node [41]. With these marginal, and pairwise joint, posterior distributions, it is possible to produce closed-form solutions for performing steps (i)

⁵When they exist, fast implicit $\mathbf{A}^{(t)}$ operators can provide significant computational savings in high-dimensional problems. Implementing a Fourier transform as a fast Fourier transform (FFT) subroutine, for example, would drop DCS-AMP's complexity from $\mathcal{O}(TMN)$ to $\mathcal{O}(TN \log_2 N)$.

TABLE III
EM UPDATE EQUATIONS FOR THE SIGNAL MODEL PARAMETERS OF SECTION II

% Define key quantities obtained from AMP-MMV at iteration k :	
$E[s_n^{(t)} \bar{\mathbf{y}}] = \frac{\left(\frac{\bar{\lambda}_n^{(t)}}{\lambda_n^{(t)}} \frac{\bar{\pi}_n^{(t)}}{\pi_n^{(t)}} \frac{\bar{\lambda}_n^{(t)}}{\lambda_n^{(t)}}\right)}{\left(\frac{\bar{\lambda}_n^{(t)}}{\lambda_n^{(t)}} \frac{\bar{\pi}_n^{(t)}}{\pi_n^{(t)}} \frac{\bar{\lambda}_n^{(t)}}{\lambda_n^{(t)}} + (1 - \frac{\bar{\lambda}_n^{(t)}}{\lambda_n^{(t)}})(1 - \frac{\bar{\pi}_n^{(t)}}{\pi_n^{(t)}})(1 - \frac{\bar{\lambda}_n^{(t)}}{\lambda_n^{(t)}})\right)}$	(Q1)
$E[s_n^{(t)} s_n^{(t-1)} \bar{\mathbf{y}}] = p(s_n^{(t)} = 1, s_n^{(t-1)} = 1 \bar{\mathbf{y}})$	(Q2)
$\tilde{v}_n^{(t)} \triangleq \text{var}\{\theta_n^{(t)} \bar{\mathbf{y}}\} = \left(\frac{1}{\kappa_n^{(t)}} + \frac{1}{\psi_n^{(t)}} + \frac{1}{\kappa_n^{(t)}}\right)^{-1}$	(Q3)
$\tilde{\mu}_n^{(t)} \triangleq E[\theta_n^{(t)} \bar{\mathbf{y}}] = \tilde{v}_n^{(t)} \cdot \left(\frac{\eta_n^{(t)}}{\kappa_n^{(t)}} + \frac{\xi_n^{(t)}}{\psi_n^{(t)}} + \frac{\zeta_n^{(t)}}{\kappa_n^{(t)}}\right)$	(Q4)
$v_n^{(t)} \triangleq \text{var}\{x_n^{(t)} \bar{\mathbf{y}}\}$	% See (A6) of Table II
$\mu_n^{(t)} \triangleq E[x_n^{(t)} \bar{\mathbf{y}}]$	% See (A5) of Table II
% EM update equations:	
$\lambda^{k+1} = \frac{1}{N} \sum_{n=1}^N E[s_n^{(1)} \bar{\mathbf{y}}]$	(E1)
$p_{01}^{k+1} = \frac{\sum_{t=2}^T \sum_{n=1}^N E[s_n^{(t-1)} \bar{\mathbf{y}}] - E[s_n^{(t)} s_n^{(t-1)} \bar{\mathbf{y}}]}{\sum_{t=2}^T \sum_{n=1}^N E[s_n^{(t-1)} \bar{\mathbf{y}}]}$	(E2)
$\zeta^{k+1} = \left(\frac{N(T-1)}{\rho^k} + \frac{N}{(\sigma_e^2)^k}\right) \left(\frac{1}{(\sigma_e^2)^k} \sum_{n=1}^N \tilde{\mu}_n^{(1)} + \sum_{t=2}^T \sum_{n=1}^N \frac{1}{\alpha^k \rho^k} (\tilde{\mu}_n^{(t)} - (1 - \alpha^k) \tilde{\mu}_n^{(t-1)})\right)$	(E3)
$\alpha^{k+1} = \frac{1}{4N(T-1)} \left(b - \sqrt{b^2 + 8N(T-1)c}\right)$	(E4)
where:	
$b \triangleq \frac{2}{\rho^k} \sum_{t=2}^T \sum_{n=1}^N \Re\{E[\theta_n^{(t)*} \theta_n^{(t-1)} \bar{\mathbf{y}}]\} - \Re\{(\tilde{\mu}_n^{(t)} - \tilde{\mu}_n^{(t-1)})^* \zeta^k\} - \tilde{v}_n^{(t-1)} - \tilde{\mu}_n^{(t-1)} ^2$	
$c \triangleq \frac{2}{\rho^k} \sum_{t=2}^T \sum_{n=1}^N \tilde{v}_n^{(t)} + \tilde{\mu}_n^{(t)} ^2 + \tilde{v}_n^{(t-1)} + \tilde{\mu}_n^{(t-1)} ^2 - 2\Re\{E[\theta_n^{(t)*} \theta_n^{(t-1)} \bar{\mathbf{y}}]\}$	
$\rho^{k+1} = \frac{1}{(\alpha^k)^2 N(T-1)} \sum_{t=2}^T \sum_{n=1}^N \tilde{v}_n^{(t)} + \tilde{\mu}_n^{(t)} ^2 + (\alpha^k)^2 \zeta^k ^2 - 2(1 - \alpha^k) \Re\{E[\theta_n^{(t)*} \theta_n^{(t-1)} \bar{\mathbf{y}}]\} - 2\alpha^k \Re\{\tilde{\mu}_n^{(t)*} \zeta^k\} + 2\alpha^k (1 - \alpha^k) \Re\{\tilde{\mu}_n^{(t-1)*} \zeta^k\} + (1 - \alpha^k)(\tilde{v}_n^{(t-1)} + \tilde{\mu}_n^{(t-1)} ^2)$	(E5)
$(\sigma_e^2)^{k+1} = \frac{1}{TM} \left(\sum_{t=1}^T \ \mathbf{y}^{(t)} - \mathbf{A}\boldsymbol{\mu}^{(t)}\ ^2 + \mathbf{1}_N^T \mathbf{v}^{(t)}\right)$	(E6)

and (ii) above. We adopt a Gauss-Seidel scheme, performing coordinate-wise maximization, e.g.,

$$\lambda^{k+1} = \arg \max_{\lambda} E_{\bar{\mathbf{s}}, \bar{\boldsymbol{\theta}} | \bar{\mathbf{y}}} [\log p(\bar{\mathbf{y}}, \bar{\mathbf{s}}, \bar{\boldsymbol{\theta}}; \lambda, \Gamma^k \setminus \{\lambda^k\}) | \bar{\mathbf{y}}; \Gamma^k].$$

The EM procedure is performed after each forward/backward pass, leading to a convergent sequence of parameter estimates. If operating in filtering mode, the procedure is similar, however the EM procedure is run after each recovered timestep using only causally available posterior estimates.

In Table III, we provide the EM update equations for each of the parameters of our signal model, assuming DCS-AMP is operating in smoothing mode. A complete derivation of each update can be found in [39].

V. INCORPORATING ADDITIONAL STRUCTURE

In Sections II–IV we described a signal model for the dynamic CS problem and summarized a message passing algorithm for making inferences under this model, while iteratively learning the model parameters via EM. We also hinted that the model could be generalized to incorporate additional, or more complex, forms of structure. In this section we will elaborate on this idea, and illustrate one such generalization.

Recall that, in Section II, we introduced hidden variables $\bar{\mathbf{s}}$ and $\bar{\boldsymbol{\theta}}$ in order to characterize the structure in the signal coefficients. An important consequence of introducing these hidden

variables was that they made each signal coefficient $x_n^{(t)}$ conditionally independent of the remaining coefficients in $\bar{\mathbf{x}}$, given $s_n^{(t)}$ and $\theta_n^{(t)}$. This conditional independence served an important algorithmic purpose since it allowed us to apply the AMP algorithm, which requires independent local priors, within our larger inference procedure.

One way to incorporate additional structure into the signal model of Section II is to generalize our choices of $p(\bar{\mathbf{s}})$ and $p(\bar{\boldsymbol{\theta}})$. As a concrete example, pairing the temporal support model proposed in this work with the Markovian model of wavelet tree inter-scale correlations described in [17] through a more complex support prior, $p(\bar{\mathbf{s}})$, could enable even greater undersampling in a dynamic MRI setting. Performing inference on such models could be accomplished through the general algorithmic framework proposed in [35]. As another example, suppose that we wish to expand our Bernoulli-Gaussian signal model to one in which signal coefficients are marginally distributed according to a Bernoulli-Gaussian-mixture, i.e.,

$$p(x_n^{(t)}) = \lambda_0^{(t)} \delta(x_n^{(t)}) + \sum_{d=1}^D \lambda_d^{(t)} \mathcal{CN}(x_n^{(t)}; \zeta_d, \sigma_d^2),$$

where $\sum_{d=0}^D \lambda_d^{(t)} = 1$. Since we still wish to preserve the slow time-variations in the support and smooth evolution of non-zero amplitudes, a natural choice of hidden variables is $\{\bar{\mathbf{s}}, \bar{\boldsymbol{\theta}}_1, \dots, \bar{\boldsymbol{\theta}}_D\}$, where $s_n^{(t)} \in \{0, 1, \dots, D\}$, and $\theta_{d,n}^{(t)} \in \mathbb{C}$, $d = 1, \dots, D$. The relationship between $x_n^{(t)}$ and the hidden variables then generalizes to:

$$p(x_n^{(t)} | s_n^{(t)}, \theta_{1,n}^{(t)}, \dots, \theta_{D,n}^{(t)}) = \begin{cases} \delta(x_n^{(t)}), & s_n^{(t)} = 0, \\ \delta(x_n^{(t)} - \theta_{d,n}^{(t)}), & s_n^{(t)} = d \neq 0. \end{cases}$$

To model the slowly changing support, we specify $p(\bar{\mathbf{s}})$ using a $(D+1)$ -state Markov chain defined by the transition probabilities $p_{0d} \triangleq \Pr\{s_n^{(t)} = 0 | s_n^{(t-1)} = d\}$ and $p_{d0} \triangleq \Pr\{s_n^{(t)} = d | s_n^{(t-1)} = 0\}$, $d = 1, \dots, D$. In this work, we assume that state transitions cannot occur between active mixture components, i.e., $\Pr(s_n^{(t)} = d | s_n^{(t-1)} = e) = 0$ when $d \neq e \neq 0$.⁶ For each amplitude time-series we again use independent Gauss-Markov processes to model smooth evolutions in the amplitudes of active signal coefficients, i.e.,

$$\theta_{d,n}^{(t)} = (1 - \alpha_d) (\theta_{d,n}^{(t-1)} - \zeta_d) + \alpha_d w_{d,n}^{(t)} + \zeta_d,$$

where $w_{d,n}^{(t)} \sim \mathcal{CN}(0, \rho_d)$.

As a consequence of this generalized signal model, a number of the message computations of Section III-B must be modified. For steps **(into)** and **(across)**, it is largely straightforward to extend the computations to account for the additional hidden variables. For step **(within)**, the modifications will affect the AMP thresholding equations defined in (D1)–(D4) of Table II.

⁶By relaxing this restriction on active-to-active state transitions, we can model signals whose coefficients tend to enter the support set at small amplitudes that grow larger over time through the use of a Gaussian mixture component with a small variance that has a high probability of transitioning to a higher variance mixture component.

Details on a Bernoulli-Gaussian-mixture AMP algorithm can be found in [30]. For the **(out)** step, we will encounter difficulties applying standard sum-product update rules to compute the messages $\{\nu_{f_n^{(t)} \rightarrow \theta_{d,n}^{(t)}}(\cdot)\}_{d=1}^D$. As in the Bernoulli-Gaussian case, we consider a modification of our assumed signal model that incorporates an $\varepsilon \ll 1$ term, and use Taylor series approximations of the resultant messages to collapse a $(D+1)$ -ary Gaussian mixture to a single Gaussian.

VI. EMPIRICAL STUDY

We now describe the results of an empirical study of DCS-AMP.⁷ The primary performance metric that we used in all of our experiments, which we refer to as the time-averaged normalized MSE (TNMSE), is defined as

$$\text{TNMSE}(\bar{\mathbf{x}}, \hat{\bar{\mathbf{x}}}) \triangleq \frac{1}{T} \sum_{t=1}^T \frac{\|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t)}\|_2^2}{\|\mathbf{x}^{(t)}\|_2^2},$$

where $\hat{\mathbf{x}}^{(t)}$ is an estimate of $\mathbf{x}^{(t)}$.

Unless otherwise noted, the following settings were used for DCS-AMP in our experiments. First, DCS-AMP was run as a smoother, with a total of 5 forward/backward passes. The number of inner AMP iterations I for each **(within)** step was $I = 25$, with a possibility for early termination if the change in the estimated signal, $\boldsymbol{\mu}_t^i$, fell below a predefined threshold from one iteration to the next, i.e., $\frac{1}{N} \|\boldsymbol{\mu}_t^i - \boldsymbol{\mu}_t^{i-1}\|_2 < 10^{-5}$. Equation (A9) of Table II was used to produce $\hat{\mathbf{x}}^{(t)}$, which corresponds to an MMSE estimate of $\mathbf{x}^{(t)}$ under DCS-AMP's estimated posteriors $p(x_n^{(t)} | \bar{\mathbf{y}})$. The amplitude approximation parameter ε from (7) was set to $\varepsilon = 10^{-7}$, while the threshold τ from (11) was set to $\tau = 0.99$. In our experiments, we found DCS-AMP's performance to be relatively insensitive to the value of ε provided that $\varepsilon \ll 1$. The choice of τ should be selected to provide a balance between allowing DCS-AMP to track amplitude evolutions on signals with rapidly varying supports and preventing DCS-AMP from prematurely gaining too much confidence in its estimate of the support. We found that the choice $\tau = 0.99$ works well over a broad range of problems. When the estimated transition probability $p_{01} < 0.025$, DCS-AMP automatically switched from the threshold method to the Taylor series method of computing (10), which is advantageous because it is parameter-free.

When learning model parameters adaptively from the data using the EM updates of Table III, it is necessary to first initialize the parameters at reasonable values. Unless domain-specific knowledge suggests a particular initialization strategy, we advocate using the following simple heuristics: The initial sparsity rate, λ^1 , active mean, ζ^1 , active variance, $(\sigma^2)^1$, and noise variance, $(\sigma_e^2)^1$, can be initialized according to the procedure described in ([30], Section V).⁸ The Gauss-Markov correlation parameter, α , can be initialized as

$$\alpha^1 = 1 - \frac{1}{T-1} \sum_{t=1}^{T-1} \frac{|\mathbf{y}^{(t)H} \mathbf{y}^{(t+1)}|}{\lambda^1 (\sigma^2)^1 |\text{tr}\{\mathbf{A}^{(t)} \mathbf{A}^{(t+1)H}\}|}. \quad (12)$$

⁷Code for reproducing our results is available at <http://www.ece.osu.edu/~schniter/turboAMPdcs>.

⁸For problems with a high degree of undersampling and relatively non-sparse signals, it may be necessary to threshold the value for λ^1 suggested in [30] so that it does not fall below, e.g., 0.10.




Fig. 3. A plot of the TNMSE (in dB) of (from left) the SKS, DCS-AMP, EM-DCS-AMP, and BG-AMP across the sparsity-undersampling plane, for temporal correlation parameters $p_{01} = 0.05$ and $\alpha = 0.01$.

The active-to-inactive transition probability, p_{01} , is difficult to gauge solely from sample statistics involving the available measurements, $\bar{\mathbf{y}}$. We used $p_{01}^1 = 0.10$ as a generic default choice, based on the premise that it is easier for DCS-AMP to adjust to more dynamic signals once it has a decent “lock” on the static elements of the support, than it is for it to estimate relatively static signals under an assumption of high dynamicity.

A. Performance Across the Sparsity-Undersampling Plane

Two factors that have a significant effect on the performance of any CS algorithm are the sparsity $|\mathcal{S}^{(t)}|$ of the underlying signal, and the number of measurements M . Consequently, much can be learned about an algorithm by manipulating these factors and observing the resulting change in performance. To this end, we studied DCS-AMP’s performance across the sparsity-undersampling plane [42], which is parameterized by two quantities, the *normalized sparsity ratio*, $\beta \triangleq E[|\mathcal{S}^{(t)}|]/M$, and the *undersampling ratio*, $\delta \triangleq M/N$. For a given (δ, β) pair (with N fixed at 1500), sample realizations of $\bar{\mathbf{s}}$, $\bar{\boldsymbol{\theta}}$, and $\bar{\mathbf{e}}$ were drawn from their respective priors, and elements of a time-varying $\mathbf{A}^{(t)}$ were drawn from i.i.d. zero-mean complex circular Gaussians, with all columns subsequently scaled to have unit ℓ_2 -norm, thus generating $\bar{\mathbf{x}}$ and $\bar{\mathbf{y}}$.

As a performance benchmark, we used the support-aware Kalman smoother. In the case of linear dynamical systems with jointly Gaussian signal and observations, the Kalman filter (smoother) is known to provide MSE-optimal causal (non-causal) signal estimates [43]. When the signal is Bernoulli-Gaussian, the Kalman filter/smoother is no longer optimal. However, a lower bound on the achievable MSE can be obtained using the support-aware Kalman filter (SKF) or smoother (SKS). Since the classical state-space formulation of the Kalman filter does not easily yield the support-aware bound, we turn to an alternative view of Kalman filtering as an instance of message passing on an appropriate factor graph [44]. For this, it suffices to use the factor graph of Fig. 1 with $\{s_n^{(t)}\}$ treated as fixed, known quantities. Following the standard sum-product algorithm rules results in a message passing algorithm in which all messages are Gaussian, and no message approximations are required. Then, by running loopy Gaussian belief propagation until convergence, we are guaranteed that the resultant posterior means constitute the MMSE estimate of $\bar{\mathbf{x}}$ ([22], Claim 5).

To quantify the improvement obtained by exploiting temporal correlation, signal recovery was also explored using the Bernoulli-Gaussian AMP algorithm (BG-AMP) independently at each timestep (i.e., ignoring temporal structure in the support and amplitudes), accomplished by passing messages only within the dashed boxes of Fig. 1 using $p(x_n^{(t)})$ from (4) as AMP’s prior.⁹

In Fig. 3, we present four plots from a representative experiment. {The TNMSE across the (logarithmically scaled) sparsity-undersampling plane is shown for (working from left to right) the SKS, DCS-AMP, EM-DCS-AMP (DCS-AMP with EM parameter tuning), and BG-AMP. In order to allow EM-DCS-AMP ample opportunity to converge to the correct parameter values, it was allowed up to 300 EM iterations/smoothing passes, although it would quite often terminate much sooner if the parameter initializations were reasonably close. The results shown were averaged over more than {300} independent trials at each (δ, β) pair. For this experiment, signal model parameters were set at $N = 1500$, $T = 25$, $p_{01} = 0.05$, $\zeta = 0$, $\alpha = 0.01$, $\sigma^2 = 1$, and a noise variance, σ_e^2 , chosen to yield a signal-to-noise ratio (SNR) of 25 dB. (M, λ) were set based on specific (δ, β) pairs, and p_{10} was set so as to keep the expected number of active coefficients constant across time. It is interesting to observe that the performance of the SKS and (EM-)DCS-AMP are only weakly dependent on the undersampling ratio δ . In contrast, the structure-agnostic BG-AMP algorithm is strongly affected. This is one of the principal benefits of incorporating temporal structure; it makes it possible to tolerate more substantial amounts of undersampling, particularly when the underlying signal is less sparse.

B. Performance vs p_{01} and α

The temporal correlation of our time-varying sparse signal model is largely dictated by two parameters, the support transition probability p_{01} and the amplitude forgetting factor α . Therefore, it is worth investigating how the performance of (EM-)DCS-AMP is affected by these two parameters. In an experiment similar to that of Fig. 3, we tracked the performance of (EM-)DCS-AMP, the SKS, and BG-AMP across a plane of (p_{01}, α) pairs. The active-to-inactive transition probability

⁹Experiments were also run that compared performance against Basis Pursuit Denoising (BPDN) [45] with genie-aided parameter tuning (solved using the SPGL1 solver [46]). However, this was found to yield higher TNMSE than BG-AMP, and at higher computational cost.

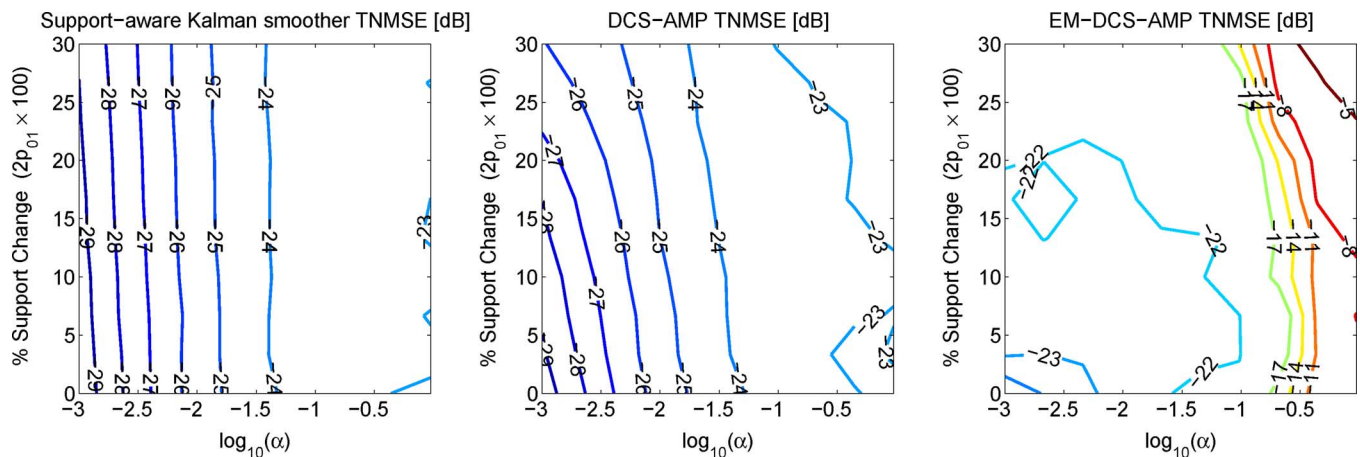


Fig. 4. TNMSE (in dB) of (from left) the SKS, DCS-AMP, and EM-DCS-AMP as a function of the model parameters p_{01} and α , for undersampling ratio $\delta = 1/3$ and sparsity ratio $\beta = 0.45$. BG-AMP achieved a TNMSE of -5.9 dB across the plane.

p_{01} was swept linearly over the range $[0, 0.15]$, while the Gauss-Markov amplitude forgetting factor α was swept logarithmically over the range $[0.001, 0.95]$. To help interpret the meaning of these parameters, we note that the fraction of the support that is expected to change from one timestep to the next is given by $2p_{01}$, and that the Pearson correlation coefficient between temporally adjacent amplitude variables is $1 - \alpha$.

In Fig. 4 we plot the TNMSE (in dB) of the SKS and (EM-)DCS-AMP as a function of the percentage of the support that changes from one timestep to the next (i.e., $2p_{01} \times 100$) and the logarithmic value of α for a signal model in which $\delta = 1/5$ and $\beta = 0.60$, with remaining parameters set as before. Since BG-AMP is agnostic to temporal correlation, its performance is insensitive to the values of p_{01} and α . Therefore, we do not include a plot of the performance of BG-AMP, but note that it achieved a TNMSE of -5.9 dB across the plane. For the SKS and (EM-)DCS-AMP, we see that performance improves with increasing amplitude correlation (moving leftward). This behavior, although intuitive, is in contrast to the relationship between performance and correlation found in the MMV problem [29], [47], primarily due to the fact that the measurement matrix is static for all timesteps in the classical MMV problem, whereas here it varies with time. Since the SKS has perfect knowledge of the support, its performance is only weakly dependent on the rate of support change. DCS-AMP performance shows a modest dependence on the rate of support change, but nevertheless is capable of managing rapid temporal changes in support, while EM-DCS-AMP performs very near the level of the noise when $\alpha < 0.10$.

C. Recovery of an MRI Image Sequence

While the above simulations demonstrate the effectiveness of DCS-AMP in recovering signals generated according to our signal model, it remains to be seen whether the signal model itself is suitable for describing practical dynamic CS signals. To address this question, we tested the performance of DCS-AMP on a dynamic MRI experiment first performed in [48]. The experiment consists of recovering a sequence of 10 MRI images of the larynx, each 256×256 pixels in dimension. (See Fig. 5.) The measurement matrices were never stored explicitly due to the prohibitive sizes involved, but were instead treated as the




Fig. 5. Frames 1, 2, 5, and 10 of the dynamic MRI image sequence of (from top to bottom): the fully sampled dataset, Basis Pursuit, Modified-CS, and DCS-AMP, with increased sampling rate at initial timestep.

composition of three linear operations, $\mathbf{A} = \mathbf{MFW}^T$. The first operation, \mathbf{W}^T , was the synthesis of the underlying image from a sparsifying 2-D, 2-level Daubechies-4 wavelet transform representation. The second operation, \mathbf{F} , was a 2-D Fourier transform that yielded the k-space coefficients of the image. The third operation, \mathbf{M} , was a sub-sampling mask that kept only a fraction of the available k-space data.

Since the image transform coefficients are compressible rather than sparse, the SKF/SKS no longer serves as an appropriate algorithmic benchmark. Instead, we compare performance against Modified-CS [6], as well as timestep-independent Basis Pursuit.¹⁰ As reported in [6], Modified-CS demonstrates that substantial improvements can be obtained over temporally agnostic methods.

Since the statistics of wavelet coefficients at different scales are often highly dissimilar (e.g., the coarsest-scale approximation coefficients are usually much less sparse than those at finer

¹⁰Modified-CS is available at <http://home.engineering.iastate.edu/~luwei/modcs/index.html>. Basis Pursuit was solved using the ℓ_1 -MAGIC equality-constrained primal-dual solver (chosen since it is used as a subroutine within Modified-CS), available at <http://users.ece.gatech.edu/~justin/l1magic/>.

TABLE IV
PERFORMANCE ON DYNAMIC MRI DATASET FROM [48] WITH INCREASED SAMPLING RATE AT INITIAL TIMESTEP

Algorithm	TNMSE (dB)	Runtime
Basis Pursuit	-17.22	47 min
Modified-CS	-34.30	7.39 hrs
DCS-AMP (Filter)	-34.62	8.08 sec

TABLE V
PERFORMANCE ON DYNAMIC MRI DATASET FROM [48] WITH IDENTICAL SAMPLING RATE AT EVERY TIMESTEP

Algorithm	TNMSE (dB)	Runtime
Basis Pursuit	-16.83	47.61 min
Modified-CS	-17.18	7.78 hrs
DCS-AMP (Filter)	-29.51	7.27 sec

scales, and are also substantially larger in magnitude), we allowed our EM procedure to learn different parameters for different wavelet scales. Using \mathcal{G}_1 to denote the indices of the coarsest-scale “approximation” coefficients, and \mathcal{G}_2 to denote the finer-scale “wavelet” coefficients, DCS-AMP was initialized with the following parameter choices: $\lambda_{\mathcal{G}_1} = 0.99$, $\lambda_{\mathcal{G}_2} = 0.01$, $p_{01} = 0.01$, $\zeta_{\mathcal{G}_1} = \zeta_{\mathcal{G}_2} = 0$, $\alpha_{\mathcal{G}_1} = \alpha_{\mathcal{G}_2} = 0.05$, $\rho_{\mathcal{G}_1} = 10^5$, $\rho_{\mathcal{G}_2} = 10^4$, and $\sigma_e^2 = 0.01$, and run in filtering mode with $I = 10$ inner AMP iterations.

We note that our initializations were deliberately chosen to be agnostic, but reasonable, values. In particular, observing that the coarsest-scale approximation coefficients of a wavelet decomposition are almost surely non-zero, we initialized the associated group’s sparsity rate at $\lambda_{\mathcal{G}_1} = 0.99$, while the finer scale detail coefficients were given an arbitrary sparsity-promoting rate of $\lambda_{\mathcal{G}_2} = 0.01$. The choices of α and ρ were driven by an observation that the variance of coefficients across wavelet scales often differs by an order-of-magnitude. The noise variance is arguably the most important parameter to initialize properly, since it balances the conflicting objectives of fitting the data and adhering to the assumed signal model. Our rule-of-thumb for initializing this parameter was that it is best to err on the side of fitting the data (since the SNR in this MRI data collection was high), and thus we initialized the noise variance with a small value.

In Table IV we summarize the performance of three different estimators: timestep-independent Basis Pursuit, which performs independent ℓ_1 minimizations at each timestep, Modified-CS, and DCS-AMP (operating in filtering mode). In this experiment, per the setup described in [48], the initial timestep was sampled at 50% of the Nyquist rate, i.e., $M = N/2$, while subsequent timesteps were sampled at 16% of the Nyquist rate. Both Modified-CS and DCS-AMP substantially outperform Basis Pursuit with respect to TNMSE, with DCS-AMP showing a slight advantage over Modified-CS. Despite the similar TNMSE performance, note that DCS-AMP runs in seconds, whereas Modified-CS takes multiple hours. In Fig. 5, we plot the true images along with the recoveries for this experiment, which show severe degradation for Basis Pursuit on all but the initial timestep.

In practice, it may not be possible to acquire an increased number of samples at the initial timestep. We therefore repeated the experiment while sampling at 16% of the Nyquist rate at every timestep. The results, shown in Table V, show that

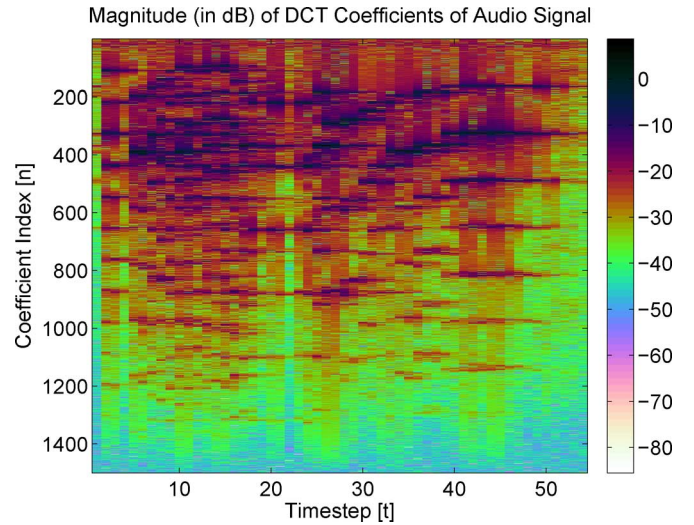


Fig. 6. DCT coefficient magnitudes (in dB) of an audio signal.

DCS-AMP’s performance degrades by about 5 dB, while Modified-CS suffers a 14 dB reduction, illustrating that, when the estimate of the initial support is poor, Modified-CS struggles to outperform Basis Pursuit.

D. Recovery of a CS Audio Sequence

In another experiment using real-world data, we used DCS-AMP to recover an audio signal from sub-Nyquist samples. In this case, we employ the Bernoulli-Gaussian-mixture signal model proposed for DCS-AMP in Section V. The audio clip is a 7 second recording of a trumpet solo, and contains a succession of rapid changes in the trumpet’s pitch. Such a recording presents a challenge for CS methods, since the signal will be only compressible, and not sparse. The clip, sampled at a rate of 11 kHz, was divided into $T = 54$ non-overlapping segments of length $N = 1500$. Using the discrete cosine transform (DCT) as a sparsifying basis, linear measurements were obtained using a time-invariant i.i.d. Gaussian sensing matrix.

In Fig. 6 we plot the magnitude of the DCT coefficients of the audio signal on a dB scale. Beyond the temporal correlation evident in the plot, it is also interesting to observe that there is a non-trivial amount of frequency correlation (correlation across the index $[n]$), as well as a large dynamic range. We performed recoveries using four techniques: BG-AMP, GM-AMP (a temporally agnostic Bernoulli-Gaussian-mixture AMP algorithm with $D = 4$ Gaussian mixture components), DCS-(BG)-AMP, and DCS-GM-AMP (the Bernoulli-Gaussian-mixture dynamic CS model described in Section V, with $D = 4$). For each algorithm, EM learning of the model parameters was performed using straightforward variations of the procedure described in Section IV, with model parameters initialized automatically using simple heuristics described in [30]. Moreover, unique model parameters were learned at each timestep (with the exception of support transition probabilities). Furthermore, since our model of hidden amplitude evolutions was poorly matched to this audio signal, we fixed $\alpha = 1$.

In Table VI we present the results of applying each algorithm to the audio dataset for three different undersampling rates, δ .

TABLE VI
PERFORMANCE ON AUDIO CS DATASET (TNMSE (dB) | RUNTIME (S)) OF TWO TEMPORALLY INDEPENDENT ALGORITHMS, BG-AMP AND GM-AMP, AND TWO TEMPORALLY STRUCTURED ALGORITHMS, DCS-BG-AMP AND DCS-GM-AMP

		Undersampling Rate		
		$\delta = \frac{1}{2}$	$\delta = \frac{1}{3}$	$\delta = \frac{1}{5}$
Algorithm	BG-AMP	-16.88 (dB) 09.11 (s)	-11.67 (dB) 08.27 (s)	-08.56 (dB) 06.63 (s)
	GM-AMP ($D = 4$)	-17.49 (dB) 19.36 (s)	-13.74 (dB) 17.48 (s)	-10.23 (dB) 15.98 (s)
	DCS-BG-AMP	-19.84 (dB) 10.20 (s)	-14.33 (dB) 08.39 (s)	-11.40 (dB) 06.71 (s)
	DCS-GM-AMP ($D = 4$)	-21.33 (dB) 20.34 (s)	-16.78 (dB) 18.63 (s)	-12.49 (dB) 10.13 (s)

For each algorithm, both the TNMSE in dB and the runtime in seconds are provided. Overall, we see that performance improves at each undersampling rate as the signal model becomes more expressive. While GM-AMP outperforms BG-AMP at all undersampling rates, it is surpassed by DCS-BG-AMP and DCS-GM-AMP, with DCS-GM-AMP offering the best TNMSE performance. Indeed, we observe that one can obtain comparable, or even better, performance with an undersampling rate $\delta = \frac{1}{5}$ using DCS-BG-AMP or DCS-GM-AMP, with that obtained using BG-AMP with an undersampling rate $\delta = \frac{1}{3}$.

E. Frequency Estimation

In a final experiment, we compared the performance of DCS-AMP against techniques designed to solve the problem of subspace identification and tracking from partial observations (SITPO) [49], [50], which bears similarities to the dynamic CS problem. In subspace identification, the goal is to learn the low-dimensional subspace occupied by multi-timestep data measured in a high ambient dimension, while in subspace tracking, the goal is to track that subspace as it evolves over time. In the partial observation setting, the high-dimensional observations are sub-sampled using a mask that varies with time. The dynamic CS problem can be viewed as a special case of SITPO, wherein the time- t subspace is spanned by a subset of the columns of an a priori known matrix $\mathbf{A}^{(t)}$. One problem that lies in the intersection of SITPO and dynamic CS is frequency tracking from partial time-domain observations.

For comparison purposes, we replicated the “direction of arrival analysis” experiment described in [50] where the observations at time t take the form

$$\mathbf{y}^{(t)} = \Phi^{(t)} \mathbf{V}^{(t)} \mathbf{a}^{(t)} + \mathbf{e}^{(t)}, \quad t = 1, 2, \dots, T \quad (13)$$

where $\Phi^{(t)} \in \{0, 1\}^{M \times N}$ is a selection matrix with non-zero column indices $\mathcal{Q}^{(t)} \subset \{1, \dots, N\}$, $\mathbf{V}^{(t)} \in \mathbb{C}^{N \times K}$ is a Vandermonde matrix of sampled complex sinusoids, i.e.,

$$\mathbf{V}^{(t)} \triangleq [\mathbf{v}(\omega_1^{(t)}), \dots, \mathbf{v}(\omega_K^{(t)})], \quad (14)$$

with $\mathbf{v}(\omega_k^{(t)}) \triangleq [1, e^{j2\pi\omega_k^{(t)}}, \dots, e^{j2\pi\omega_k^{(t)}(N-1)}]^\top$ and $\omega_k^{(t)} \in [0, 1)$. $\mathbf{a}^{(t)} \in \mathbb{R}^K$ is a vector of instantaneous amplitudes, and $\mathbf{e}^{(t)} \in \mathbb{R}^N$ is additive noise with i.i.d. $\mathcal{N}(0, \sigma_e^2)$ elements.¹¹ Here, $\{\Phi^{(t)}\}_{t=1}^T$ is known, while $\{\omega^{(t)}\}_{t=1}^T$ and $\{\mathbf{a}^{(t)}\}_{t=1}^T$ are

¹¹Code for replicating the experiment provided by the authors of [50]. Unless otherwise noted, specific choices regarding $\{\omega_k^{(t)}\}$ and $\{\mathbf{a}^{(t)}\}$ were made by the authors of [50] in a deterministic fashion, and can be found in the code.

unknown, and our goal is to estimate them. To assess performance, we report TNMSE in the estimation of the “complete” signal $\{\mathbf{V}^{(t)} \mathbf{a}^{(t)}\}_{t=1}^T$.

We compared DCS-AMP’s performance against two online algorithms designed to solve the SITPO problem: GROUSE [49] and PETRELS [50]. Both GROUSE and PETRELS return time-varying subspace estimates, which were passed to an ESPRIT algorithm to generate time-varying frequency estimates (as in [50]). Finally, time-varying amplitude estimates were computed using least-squares. For DCS-AMP, we constructed $\mathbf{A}^{(t)}$ using a $2 \times$ column-oversampled DFT matrix, keeping only those rows indexed by $\mathcal{Q}^{(t)}$. DCS-AMP was run in filtering mode for fair comparison with the “online” operation of GROUSE and PETRELS, with $I = 7$ inner AMP iterations.

The results of performing the experiment for three different problem configurations are presented in Table VII, with performance averaged over 100 independent realizations. All three algorithms were given the true value of K . In the first problem setup considered, we see that GROUSE operates the fastest, although its TNMSE performance is noticeably inferior to that of both PETRELS and DCS-AMP, which provide similar TNMSE performance and complexity. In the second problem setup, we reduce the number of measurements, M , from 30 to 10, leaving all other settings fixed. In this regime, both GROUSE and PETRELS are unable to accurately estimate $\{\omega_k^{(t)}\}$, and consequently fail to accurately recover $\mathbf{V}^{(t)} \mathbf{a}^{(t)}$, in contrast to DCS-AMP. In the third problem setup, we increased the problem dimensions from the first problem setup by a factor of 4 to understand how the complexity of each approach scales with problem size. In order to increase the number of “active” frequencies from $K = 5$ to $K = 20$, 15 additional frequencies and amplitudes were added uniformly at random to the 5 deterministic trajectories of the preceding experiments. Interestingly, DCS-AMP, which was the slowest at smaller problem dimensions, becomes the fastest (and most accurate) in the higher-dimensional setting, scaling much better than either GROUSE or PETRELS.

VII. CONCLUSION

In this work we proposed DCS-AMP, a novel approach to dynamic CS. Our technique merges ideas from the fields of belief propagation and switched linear dynamical systems, together with a computationally efficient inference method known as AMP. Moreover, we proposed an EM approach that learns all model parameters automatically from the data. In numerical experiments on synthetic data, DCS-AMP performed within 3 dB of the support-aware Kalman smoother bound across

TABLE VII
AVERAGE PERFORMANCE ON SYNTHETIC FREQUENCY ESTIMATION EXPERIMENT (TNMSE (dB) | RUNTIME (s)) OF
GROUSE, PETRELS, AND DCS-AMP. IN ALL CASES, $T = 4000$, $\sigma_e^2 = 10^{-6}$

		Problem Setup		
		$N = 256, M = 30, K = 5$	$N = 256, M = 10, K = 5$	$N = 1024, M = 120, K = 20$
Algorithm	GROUSE	-4.52 (dB) 6.78 (s)	2.02 (dB) 6.68 (s)	-4.51 (dB) 173.89 (s)
	PETRELS	-15.62 (dB) 29.51 (s)	0.50 (dB) 14.93 (s)	-7.98 (dB) 381.10 (s)
	DCS-AMP	-15.46 (dB) 34.49 (s)	-10.85 (dB) 28.42 (s)	-12.79 (dB) 138.07 (s)

the sparsity-undersampling plane. Repeating the dynamic MRI experiment from [48], DCS-AMP slightly outperformed Modified-CS in MSE, but required less than 10 seconds to run, in comparison to more than 7 hours for Modified-CS. For the compressive sensing of audio, we demonstrated significant gains from the exploitation of temporal structure and Gaussian-mixture learning of the signal prior. Lastly, we found that DCS-AMP can outperform recent approaches to Subspace Identification and Tracking from Partial Observations (SITPO) when the underlying problem can be well-represented through a dynamic CS model.

APPENDIX A

THE BASICS OF BELIEF PROPAGATION AND AMP

In this Appendix, we provide a brief primer on belief propagation and the Bayesian approximate message passing (AMP) algorithmic framework proposed by Donoho, Maleki, and Montanari [28]. In what follows, we consider the task of estimating a signal vector $\mathbf{x} \in \mathbb{C}^N$ from linearly compressed and AWGN-corrupted measurements:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e} \in \mathbb{C}^M. \quad (15)$$

AMP can be derived from the perspective of loopy belief propagation (LBP) [20], a Bayesian inference strategy that is based on a factorization of the signal posterior pdf, $p(\mathbf{x} | \mathbf{y})$, into a product of simpler pdfs that, together, reveal the probabilistic structure in the problem. Concretely, if the signal coefficients, \mathbf{x} , and noise samples, \mathbf{w} , in (15) are jointly independent such that $p(\mathbf{x}) = \prod_{n=1}^N p(x_n)$ and $p(\mathbf{y} | \mathbf{x}) = \prod_{m=1}^M \mathcal{CN}(y_m; \mathbf{a}_m^T \mathbf{x}, \sigma_e^2)$, then the posterior pdf factors as

$$p(\mathbf{x} | \mathbf{y}) \propto \prod_{m=1}^M \mathcal{CN}(y_m; \mathbf{a}_m^T \mathbf{x}, \sigma_e^2) \prod_{n=1}^N p(x_n), \quad (16)$$

yielding the factor graph in Fig. 7.

In belief propagation [21], messages representing beliefs about the unknown variables are exchanged amongst the nodes of the factor graph until convergence to a stable fixed point occurs. The set of beliefs passed into a given variable node are then used to infer statistical properties of the associated random variable, e.g., the posterior mode, or a complete posterior distribution. The sum-product algorithm [20] is perhaps the most well-known approach to belief propagation, wherein the messages take the form of probability distributions, and exact posteriors are guaranteed whenever the graph does not have cycles (“loops”). For graphs with cycles, exact inference is




Fig. 7. The factor graph representation of the decomposition of (16).

known to be NP-hard, and so LBP is not guaranteed to produce correct posteriors. Still, it has shown state-of-the-art performance on a wide array of challenging inference problems, as noted in Section III.B.

The conventional wisdom surrounding LBP says that accurate inference is possible only when the factor graph is locally tree-like, i.e., the girth of any cycle is relatively large. With (15), this would require that \mathbf{A} is an appropriately constructed sparse matrix, which precludes some of the most interesting CS problems. In a remarkable departure from convention, Donoho, Maleki, Montanari, and Bayati demonstrated that LBP-based compressive sensing is not only feasible [14], [28] for dense \mathbf{A} matrices, but provably accurate [16]. In particular, they established that, in the large-system limit (i.e., as $M, N \rightarrow \infty$ with M/N fixed) and under i.i.d. sub-Gaussian \mathbf{A} , the iterations of AMP are governed by a state-evolution whose fixed point—when unique—yields the true posterior means. Beyond its theoretical significance, AMP is important for its computational properties as well. As demonstrated in the original AMP work [14], not only can LBP solve the compressive sensing problem (15), but it can do so much faster, and more accurately, than other state-of-the-art methods, whether optimization-based, greedy, or Bayesian. To accomplish this feat, [14], [28] proposed a specific set of approximations that become accurate in the limit of large, dense \mathbf{A} matrices, yielding algorithms that give accurate results using only $\approx 2MN$ flops-per-iteration, and relatively few iterations (e.g., tens).

The specific implementation of any AMP algorithm will depend on the particular choices of likelihood and prior, but ultimately amounts to an iterative, scalar soft-thresholding procedure with a carefully chosen adaptive thresholding strategy. Deriving the appropriate thresholding functions for a particular signal model can be accomplished by computing scalar sum-product, or max-sum, updates of a simple form (see, e.g., [33], Table I).

REFERENCES

- [1] J. Ziniel, L. C. Potter, and P. Schniter, "Tracking and smoothing of time-varying sparse signals via approximate belief propagation," in *Asilomar Conf. on Signals, Syst. Comput.*, Pacific Grove, CA, USA, Nov. 2010, pp. 808–812.
- [2] N. Vaswani, "Kalman filtered compressed sensing," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, 2008, pp. 893–896, 12–15.
- [3] M. Salman Asif, D. Reddy, P. Boufounos, and A. Veeraraghavan, "Streaming compressive sensing for high-speed periodic videos," in *Proc. Int. Conf. Image Process. (ICIP)*, Hong Kong, Sep. 2010, pp. 3373–3376.
- [4] W. Li and J. C. Preisig, "Estimation of rapidly time-varying sparse channels," *IEEE J. Ocean. Eng.*, vol. 32, pp. 927–939, Oct. 2007.
- [5] D. Angelosante, G. B. Giannakis, and E. Grossi, "Compressed sensing of time-varying signals," in *Proc. Int. Conf. Digit. Signal Process.*, Jul. 2009, pp. 1–8.
- [6] N. Vaswani and W. Lu, "Modified-CS: Modifying compressive sensing for problems with partially known support," *IEEE Trans. Signal Process.*, vol. 58, pp. 4595–4607, Sep. 2010.
- [7] D. Angelosante, S. Roumeliotis, and G. Giannakis, "Lasso-Kalman smoother for tracking sparse signals," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, Pacific Grove, CA, USA, Nov. 2009, pp. 181–185.
- [8] W. Dai, D. Sejdinović, and O. Milenkovic, "Gaussian dynamic compressive sensing," presented at the Int. Conf. Sampling Theory Appl. (SampTA), Singapore, May 2011.
- [9] D. Sejdinović, C. Andrieu, and R. Piechocki, "Bayesian sequential compressed sensing in sparse dynamical systems," in *Proc. 48th Allerton Conf. Commun., Contr. Comput.*, Urbana, IL, USA, Nov. 2010, pp. 1730–1736.
- [10] B. Shahrasbi, A. Talari, and N. Rahnavard, "TC-CSBP: Compressive sensing for time-correlated data based on belief propagation," in *Proc. Conf. Inf. Sci. Syst. (CISS)*, Baltimore, MD, USA, Mar. 2011, pp. 1–6.
- [11] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed. New York, NY, USA: Springer, 2004.
- [12] B. J. Frey and D. J. C. MacKay, "A revolution: Belief propagation in graphs with cycles," *Neural Inf. Process. Syst. (NIPS)*, vol. 10, pp. 479–485, 1998.
- [13] D. Baron, S. Sarvotham, and R. G. Baraniuk, "Bayesian compressive sensing via belief propagation," *IEEE Trans. Signal Process.*, vol. 58, pp. 269–280, Jan. 2010.
- [14] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing," in *Proc. Nat. Acad. Sci.*, Nov. 2009, vol. 106, pp. 18914–18919.
- [15] P. Schniter, "Turbo reconstruction of structured sparse signals," in *Proc. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2010, pp. 1–6.
- [16] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, pp. 764–785, Feb. 2011.
- [17] S. Som and P. Schniter, "Compressive imaging using approximate message passing and a Markov-tree prior," *IEEE Trans. Signal Process.*, vol. 60, pp. 3439–3448, Jul. 2012.
- [18] P. Schniter, "A message-passing receiver for {BICM-OFDM} over unknown clustered-sparse channels," *IEEE J. Sel. Topics Signal Process.*, vol. 5, pp. 1462–1474, Dec. 2011.
- [19] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. R. Statist. Soc. B*, vol. 39, pp. 1–38, 1977.
- [20] F. R. Kschischang, B. J. Frey, and H. A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, pp. 498–519, Feb. 2001.
- [21] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA, USA: Morgan Kaufman, 1988.
- [22] Y. Weiss and W. T. Freeman, "Correctness of belief propagation in Gaussian graphical models of arbitrary topology," *Neural Comput.*, vol. 13, pp. 2173–2200, Oct. 2001.
- [23] S. C. Tatikonda and M. I. Jordan, "Loopy belief propagation and Gibbs measures," in *Proc. 18th Conf. Uncert. Artif. Intell. (UAI)*, 2002, pp. 493–500.
- [24] T. Heskes, "On the uniqueness of belief propagation fixed points," *Neural Comput.*, vol. 16, no. 11, pp. 2379–2413, 2004.
- [25] A. T. Ihler, J. W. Fisher, III, and A. S. Willsky, "Loopy belief propagation: Convergence and effects of message errors," *J. Mach. Learn. Res.*, vol. 6, pp. 905–936, 2005.
- [26] R. J. McEliece, D. J. C. MacKay, and J. Cheng, "Turbo decoding as an instance of Pearl's belief propagation algorithm," *IEEE J. Sel. Areas Commun.*, vol. 16, pp. 140–152, Feb. 1998.
- [27] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low-level vision," *Int. J. Comput. Vision*, vol. 40, pp. 25–47, Oct. 2000.
- [28] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. Motivation and construction," in *Proc. Inf. Theory Workshop*, Jan. 2010, pp. 1–5.
- [29] J. Ziniel and P. Schniter, "Efficient high-dimensional inference in the multiple measurement vector problem," *IEEE Trans. Signal Process.*, vol. 61, pp. 340–354, Jan. 2013.
- [30] J. P. Vila and P. Schniter, "Expectation-maximization Gaussian-mixture approximate message passing," *IEEE Trans. Signal Process.*, doi: 10.1109/TSP.2013.2272287, preprint.
- [31] G. Elidan, I. McGraw, and D. Koller, "Residual belief propagation: Informed scheduling for asynchronous message passing," in *Proc. 22nd Conf. Uncert. Artif. Intell. (UAI)*, 2006, pp. 165–173.
- [32] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inf. Theory*, St. Petersburg, Russia, Aug. 2011, pp. 2168–2172.
- [33] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," 2010 [Online]. Available: <http://arxiv.org/abs/1010.5141v1>
- [34] S. Som and P. Schniter, "Approximate message passing for recovery of sparse signals with Markov-random-field support structure," presented at the Proc. Int. Conf. Mach. Learn., Bellevue, WA, USA, Jul. 2011.
- [35] J. Ziniel, S. Rangan, and P. Schniter, "A generalized framework for learning and recovery of structured sparse signals," in *Proc. Stat. Signal Process. Wkshp*, Ann Arbor, MI, USA, Aug. 2012, pp. 325–328.
- [36] O. Zoeter and T. Heskes, "Change point problems in linear dynamical systems," *J. Mach. Learn. Res.*, vol. 6, pp. 1999–2026, Dec. 2005.
- [37] D. L. Alspach and H. W. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *IEEE Trans. Autom. Control*, vol. 17, pp. 439–448, Aug. 1972.
- [38] D. Barber and A. T. Cemgil, "Graphical models for time series," *IEEE Signal Process. Mag.*, vol. 27, pp. 18–28, Nov. 2010.
- [39] J. Ziniel, "Message passing approaches to compressive inference under structured signal priors," Ph.D. dissertation, The Ohio State Univ., Columbus, OH, USA, 2013, to be published.
- [40] T. K. Moon, "The expectation-maximization algorithm," *IEEE Signal Process. Mag.*, vol. 13, pp. 47–60, Nov. 1996.
- [41] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer-Verlag, 2006.
- [42] D. L. Donoho and J. Tanner, "Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing," *Philos. Trans. Roy. Soc. A*, vol. 367, pp. 4273–4293, 2009.
- [43] R. L. Eubank, *A Kalman Filter Primer*. Boca Raton, FL, USA: Chapman & Hall/CRC, 2006.
- [44] H. A. Loeliger, J. Dauwels, J. Hu, S. Korl, L. Ping, and F. R. Kschischang, "The factor graph approach to model-based signal processing," *Proc. IEEE*, vol. 95, no. 6, pp. 1295–1322, 2007.
- [45] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scientif. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.
- [46] E. van den Berg and M. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM J. Scientif. Comput.*, vol. 31, pp. 890–912, Nov. 2008.
- [47] Z. Zhang and B. D. Rao, "Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning," *IEEE J. Sel. Topics Signal Process.*, vol. 5, pp. 912–926, Sep. 2011.
- [48] W. Lu and N. Vaswani, "Modified compressive sensing for real-time dynamic MR imaging," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Nov. 2009, pp. 3045–3048.
- [49] L. Balzano, R. Nowak, and B. Recht, "Online identification and tracking of subspaces from highly incomplete information," in *Proc. Allerton Conf. Commun., Contr., Comput.*, Oct. 2010, pp. 704–711.
- [50] Y. Chi, Y. Eldar, and R. Calderbank, "PETRELS: Subspace estimation and tracking from partial observations," in *Proc. Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Kyoto, Japan, Mar. 2012, pp. 3301–3304.



Justin Ziniel (S'11) received the B.S. and M.S. degrees in electrical and computer engineering from The Ohio State University (OSU), Columbus, in 2007 and 2012, respectively.

He is a Ph.D. student in the Information Processing Systems Lab, Department of Electrical and Computer Engineering, OSU. His research interests include statistical signal processing, machine learning, and interdisciplinary “big data” inference problems.



Philip Schniter (SM'05) received the B.S. and M.S. degrees in electrical and computer engineering from the University of Illinois at Urbana-Champaign in 1992 and 1993, respectively, and the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, in 2000.

From 1993 to 1996, he was a Systems Engineer with Tektronix Inc., Beaverton, OR. After receiving the Ph.D. degree, he joined the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, where he is currently a Professor and a Member of the Information Processing Systems (IPS) Lab. During 2008–2009, he was a Visiting Professor at EURECOM, Sophia Antipolis, France, and Supélec, Gif-sur-Yvette, France. His areas of interest include statistical signal processing, wireless communications and networks, and machine learning.

Dr. Schniter received the National Science Foundation CAREER Award, in 2003.