

A GENERALIZED FRAMEWORK FOR LEARNING AND RECOVERY OF STRUCTURED SPARSE SIGNALS

Justin Ziniel,⁽¹⁾ Sundeep Rangan,⁽²⁾ and Philip Schniter⁽¹⁾

⁽¹⁾ ECE Department, The Ohio State University, Columbus, Ohio

⁽²⁾ ECE Department, Polytechnic Institute of New York University, Brooklyn, New York

ABSTRACT

We report on a framework for recovering single- or multi-timestep sparse signals that can learn and exploit a variety of probabilistic forms of structure. Message passing-based inference and empirical Bayesian parameter learning form the backbone of the recovery procedure. We further describe an object-oriented software paradigm for implementing our framework, which consists of assembling modular software components that collectively define a desired statistical signal model. Lastly, numerical results for synthetic and real-world structured sparse signal recovery are provided.

Index Terms— compressed sensing, structured sparse signal recovery, multiple measurement vectors, structured sparsity, dynamic compressed sensing

1. INTRODUCTION

In recent years, a great deal of effort by the compressed sensing (CS) community has been directed at developing ways to incorporate additional signal structure beyond simple sparsity into recovery techniques [1]. Recent work into Bayesian approaches aimed at exploiting the low-dimensional structure inherent in many real-world signals has demonstrated that significant performance gains can be achieved, even when the structure must be learned (e.g., [2–6]).

In this work we present a flexible framework for performing empirical Bayesian estimation of structured sparse signals. Our approach follows the “turbo CS” [7] principle of breaking apart an intractable global inference problem into smaller sub-problems for which efficient and accurate inference is possible. By exchanging information between the sub-problems, we obtain a high-quality approximation of the solution to the global problem. Additionally, as a byproduct of solving these sub-problems, we are often able to learn model parameters iteratively from the data using an expectation-maximization (EM) algorithm.

2. A STRUCTURED CS SIGNAL MODEL

We consider the task of recovering a collection of sparse vectors $\{\mathbf{x}^{(t)}\}_{t=1}^T$ from a collection of measurement vectors $\{\mathbf{y}^{(t)}\}_{t=1}^T$, where $\mathbf{x}^{(t)} \in \mathbb{C}^N$, $\mathbf{y}^{(t)} \in \mathbb{C}^M$, and (typically) $M < N$. The relationship between $\mathbf{x}^{(t)}$ and $\mathbf{y}^{(t)}$ is determined as follows: Each $\mathbf{x}^{(t)}$ is transformed through the linear process

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

where $\mathbf{A}^{(t)} \in \mathbb{C}^{M \times N}$ is a known linear operator. Each “transform coefficient” $z_m^{(t)}$ (i.e., the m^{th} element of $\mathbf{z}^{(t)}$) is then observed

through an independent scalar “observation channel,” defined by the conditional distribution $p(y_m^{(t)} | z_m^{(t)})$, to yield a measurement $y_m^{(t)}$. Observe that the standard noisy CS model $\mathbf{y}^{(t)} = \mathbf{A}^{(t)} \mathbf{x}^{(t)} + \mathbf{w}^{(t)}$ is a special case of the aforementioned signal model when $p(\mathbf{w}^{(t)}) = \prod_{m=1}^M p_W(w_m^{(t)})$ and $p_W(y_m^{(t)} - z_m^{(t)}) = p(y_m^{(t)} | z_m^{(t)})$, as is the matrix CS model $\mathbf{Y} = \mathbf{A}\mathbf{X} + \mathbf{W}$. We assume w.l.o.g. that $\mathbf{x}^{(t)}$ is sparse in the canonical basis.

As mentioned in Section 1, $\{\mathbf{x}^{(t)}\}_{t=1}^T$ oftentimes exhibits substantial structure beyond simple sparsity. We will refer to the structure present within a single vector $\mathbf{x}^{(t)}$ as spatial structure, and structure across multiple such vectors as temporal structure. Also, we will use the overbar notation $\bar{\mathbf{x}} \triangleq \text{vec}([\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}])$ to denote the vectorization of all signal timesteps (with other overbar variables defined analogously). In order to model the spatial and temporal structure probabilistically, we introduce a set \mathcal{H} of hidden random variables such that $p(\bar{\mathbf{x}} | \mathcal{H})$ becomes separable, i.e., $p(\bar{\mathbf{x}} | \mathcal{H}) = \prod_{t=1}^T \prod_{n=1}^N p(x_n^{(t)} | \mathcal{H})$.

The utility of using a separable signal prior is twofold: it oftentimes simplifies the task of describing the structure probabilistically, and allows us to apply a powerful inference algorithm known as GAMP [8] within a turbo inference framework. For ease of explanation, we will henceforth focus on one particular choice of hidden variables, namely, where each signal coefficient $x_n^{(t)}$ can be expressed as the product of two hidden variables: $x_n^{(t)} = s_n^{(t)} \cdot \gamma_n^{(t)}$, where $s_n^{(t)} \in \{0, 1\}$ is an indicator of support set membership, and $\gamma_n^{(t)} \in \mathbb{C}$ expresses the amplitude of a non-zero $x_n^{(t)}$. This decomposition results in the following collection of hidden random variables:

$$\mathcal{H} \triangleq \{\mathbf{s}^{(t)}, \boldsymbol{\gamma}^{(t)}\}_{t=1}^T. \quad (2)$$

We stress that this is simply one of many choices of \mathcal{H} and $p(\bar{\mathbf{x}} | \mathcal{H})$; many others could be considered within the framework we propose. (See Section 5 for another example.)

Given the hidden variables in (2), we can model a wide variety of signal structures by choosing appropriate priors for $\bar{\mathbf{s}}$ and $\bar{\boldsymbol{\gamma}}$. Specifically, $p(\bar{\mathbf{s}})$ can model spatio-temporal structure in the support of $\bar{\mathbf{x}}$, such as block-, tree-, or clustered-sparsity, while $p(\bar{\boldsymbol{\gamma}})$ can be used to model spatio-temporal correlations in the amplitudes of $\bar{\mathbf{x}}$.

3. TURBO INFERENCE AND PARAMETER LEARNING

For the signal model of Section 2, the joint posterior distribution of all of the random variables, given the measurements, can be expressed using Bayes’ rule as:

$$p(\bar{\mathbf{x}}, \mathcal{H} | \bar{\mathbf{y}}; \mathcal{P}) \propto p(\bar{\mathbf{y}} | \bar{\mathbf{x}}; \mathcal{P}) p(\bar{\mathbf{x}} | \mathcal{H}; \mathcal{P}) p(\mathcal{H}; \mathcal{P}), \quad (3)$$

where \propto indicates equality up to a normalizing constant and \mathcal{P} denotes a set of model parameters that are used to parameterize the

Work supported in part by NSF grant CCF-1018368 and DARPA/ONR grant N66001-10-1-4090.

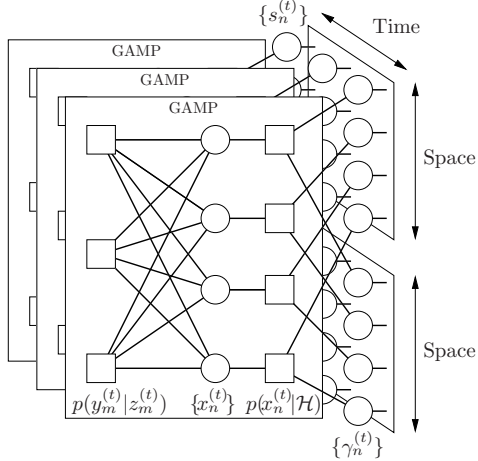


Fig. 1: A factor graph representation of the joint posterior distribution of (3). In this example, $M = T = 3$, $N = 4$, and $\mathcal{H} = \{s^{(t)}, \gamma^{(t)}\}_{t=1}^T$.

signal model, e.g., prior means, variances, etc., which we treat as deterministic unknowns. The objective of our inference procedure will be to obtain marginal posterior distributions of each random variable, e.g., $p(x_n^{(t)} | \bar{\mathbf{y}}; \mathcal{P})$, which can be used to produce an MMSE estimate of the signal. At the same time, we would like to learn the model parameters \mathcal{P} from the data in a principled fashion.

By taking advantage of the fact that $p(\bar{\mathbf{y}} | \bar{\mathbf{x}}; \mathcal{P}) (\equiv p(\bar{\mathbf{y}} | \bar{\mathbf{z}}; \mathcal{P}))$ is separable due to the independent observation assumption, and that $p(\bar{\mathbf{x}} | \mathcal{H}; \mathcal{P})$ is separable by definition of \mathcal{H} , we can conveniently describe (3) using a graphical model known as a *factor graph*. The sample factor graph shown in Fig. 1 expresses the probabilistic structure of the model (2), where $p(\mathcal{H}) = p(\bar{\mathbf{s}})p(\bar{\boldsymbol{\gamma}})$ and $p(x_n^{(t)} | \mathcal{H}) = p(x_n^{(t)} | s_n^{(t)}, \gamma_n^{(t)})$, using circles to denote random variables (here $\{x_n\}$, $\{s_n\}$, and $\{\gamma_n\}$) and squares to denote posterior factors. Although the subsequent discussion will focus on this example factor graph, our technique generalizes to any factor graph that describes separable signal and observation-channel priors.

A popular means of performing inference on probabilistic factor graphs is via belief propagation (BP) [9], whose objective is to compute the posterior marginals of all unobserved random variables. While implementing a conventional BP algorithm would be computationally intractable for the factor graph of Fig. 1, there exists an attractive approximate message passing algorithm known as GAMP [8]. GAMP’s appeal for our problem stems from several considerations: (i) GAMP supports arbitrary separable signal and observation-channel priors, (ii) inference is rapid and highly accurate, and (iii) theoretical analyses demonstrate that the behavior of GAMP can be accurately predicted by a set of state evolution equations [8].

To perform inference on the complete factor graph of Fig. 1, we employ the “turbo CS” framework of [7], which alternates¹ between exploiting the measurement structure (using GAMP) and exploiting the signal structure specified by $p(\mathcal{H}; \mathcal{P})$. This approach is reminiscent of modern turbo communications receivers, which alternate between channel equalization and decoding. The messages leaving the GAMP planes in Fig. 1 constitute beliefs about the hidden variables \mathcal{H} , given the measurements $\bar{\mathbf{y}}$. These messages act as inputs to the $\{s_n^{(t)}\}$ and $\{\gamma_n^{(t)}\}$ nodes in the rightmost portion of the factor

graph. There, inference can be performed using any technique (e.g., the forward-backward algorithm [9]) that provides extrinsic likelihoods of the hidden variables. These likelihoods act as updated beliefs about the hidden variables, given the underlying structure, and are fed back to GAMP, after which the entire process is repeated.

At the same time that turbo inference is being performed, the estimates of the unknown model parameters \mathcal{P} can be updated through an expectation-maximization (EM) [11] learning procedure. Oftentimes it makes sense to use the hidden variables \mathcal{H} as the “missing data,” in which case the EM update for \mathcal{P} , at iteration $k + 1$, can be expressed as the optimization problem

$$\mathcal{P}^{k+1} = \underset{\mathcal{P}}{\operatorname{argmax}} E_{\mathcal{H} | \bar{\mathbf{y}}} [\log p(\bar{\mathbf{y}}, \mathcal{H}; \mathcal{P}) | \bar{\mathbf{y}}; \mathcal{P}^k],$$

where \mathcal{P}^k is the estimated value of all model parameters as of iteration k . For many signal models, $\log p(\bar{\mathbf{y}}, \mathcal{H}; \mathcal{P})$ will decouple into a sum of many terms that depend only on small subsets of hidden variables and parameters. Therefore, it is often possible to obtain closed-form EM updates of the model parameters using only marginal or pairwise joint posteriors, e.g., $p(s_n^{(t)} | \bar{\mathbf{y}})$ and $p(\gamma_n^{(t)}, \gamma_n^{(t-1)} | \bar{\mathbf{y}})$ in [12]. Since belief propagation provides these posteriors, it is feasible to perform these EM updates as an auxiliary procedure to the main turbo inference process with very little additional cost.

4. OBJECT-ORIENTED SOFTWARE IMPLEMENTATION

One of the defining features of the framework we propose in Sections 2 and 3 is that it decouples the global inference problem of marginalizing (3) into smaller sub-problems that require only local, not global, information to complete their tasks. Furthermore, these sub-problems roughly correspond to inference on different regions of the factor graph. Consequently, the object-oriented programming (OOP) paradigm is useful for building a powerful and flexible software implementation of our approach, which we call EMturboGAMP.

We will now describe a software implementation² of EMturboGAMP that uses OOP principles to allow one to solve a variety of structured CS problems. At its core, our implementation relies on assembling objects of different classes in a modular fashion to specify a particular signal model. To date, we have defined four abstract classes: `Signal`, `Observation`, `SupportStruct`, and `AmplitudeStruct`, and a container class called `TurboOpt` that holds objects derived from these four classes.

At the highest level, EMturboGAMP consists of two steps performed repeatedly: in the first step, GAMP is run for a particular choice of “local” signal and observation-channel priors³. In the second step, the final state of GAMP messages is given to the `TurboOpt` object, which works together with the `Signal` and `Observation` objects to provide updated local signal and observation-channel priors to GAMP for the next turbo iteration.

Concrete implementations of these four abstract classes are responsible for overseeing specific tasks. A `Signal` class object defines the marginal prior distribution $p(x_n^{(t)}; \mathcal{P})$ and hidden variables \mathcal{H} . It delegates the task of exploiting the signal support structure to a `SupportStruct` object, which defines $p(\bar{\mathbf{s}}; \mathcal{P})$, and the task of exploiting the signal amplitude structure to an `AmplitudeStruct` object, which defines $p(\bar{\boldsymbol{\gamma}}; \mathcal{P})$. The

²Available at www.ece.osu.edu/~schniter/EMturboGAMP

³The GAMP algorithm is conventionally run with a fixed choice of signal and observation-channel priors. In the turbo framework, these priors (from GAMP’s perspective) are updated every iteration, thus we refer to them as “local priors”.

¹Another way to exploit GAMP for message passing on the complete factor graph of Fig. 1 is through “hybrid-GAMP” [10]. Unlike the “turbo” message-passing schedule adopted in this work, hybrid-GAMP employs a flooding schedule.




Fig. 2: Information flow between EMturboGAMP classes and GAMP.

Observation class object specifies the observation-channel prior $p(y_m^{(t)} | z_m^{(t)}; \mathcal{P})$. Each class is also responsible for performing EM updates of any model parameters $\in \mathcal{P}$ which they define. In Fig. 2 we summarize the relationship between the various classes and GAMP, illustrating how information flows between them.

One nice property of the OOP approach is that it makes it easy to incorporate new signal or observation-channel models without requiring one to code up an entirely new algorithm from scratch. A new marginal signal prior, for example, can be specified by creating a new sub-class of the **Signal** class. Each new sub-class must implement the handful of methods (functions) specified by its abstract super-class, but the manner in which they are implemented is entirely up to the programmer. This allows EMturboGAMP to be assured of a common interface without mandating the way in which the inference is performed. For further details, we invite the reader to explore the MATLAB code we have made publicly available² (and to contribute additional classes as well!).

5. NUMERICAL EXAMPLES

To demonstrate both the flexibility of our proposed framework, as well as the convenience of its OOP-based software implementation, we undertook a numerical study of EMturboGAMP on synthetic structured sparse signals, as well as a real-world compressively sensed audio example.

First, to understand the average performance of EMturboGAMP under a variety of test conditions, we empirically evaluated mean-squared error (MSE) performance on the sparsity-undersampling plane, calculating MSE at various combinations of the normalized sparsity ratio, β (i.e., the ratio of non-zero coefficients-to-measurements, K/M), and undersampling ratio, δ (i.e., the ratio of measurements-to-unknowns, M/N). In particular, for each (δ, β) pair, multiple independent signal realizations were recovered, and for each realization the time-averaged normalized MSE (TNMSE) was computed, where $\text{TNMSE}(\hat{\mathbf{x}}, \hat{\mathbf{x}}) \triangleq \frac{1}{T} \sum_{t=1}^T \|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t)}\|_2^2 / \|\mathbf{x}^{(t)}\|_2^2$, and $\hat{\mathbf{x}}^{(t)}$ is an estimate of $\mathbf{x}^{(t)}$.

As synthetic data, we considered a multiple measurement vector (MMV) problem in which each coefficient was marginally a Bernoulli-Gaussian-mixture of the form

$$p(x_n^{(t)}) = (1 - \lambda)\delta(x_n^{(t)}) + \sum_{i=1}^2 \frac{\lambda}{2} \mathcal{N}(x_n^{(t)}; \mu_i, \frac{1}{4}), \quad (4)$$

with $\mu_1 = 1$ and $\mu_2 = -1$. We wished to impose smooth variations in the amplitudes over time, in addition to the time-invariant support constraint inherent in the classical MMV model. To accomplish this, we defined our hidden variables as $\mathcal{H} \triangleq \{\mathbf{s}^{(t)}, \gamma_1^{(t)}, \gamma_2^{(t)}\}_{t=1}^T$, where $\gamma_{1,n}^{(t)}$ was marginally distributed $\mathcal{N}(+1, \frac{1}{4})$, $\gamma_{2,n}^{(t)}$ was marginally $\mathcal{N}(-1, \frac{1}{4})$, and $s_n^{(t)} \in \{0, 1, 2\}$ specified whether $x_n^{(t)}$ was non-zero, and if so, which component Gaussian it was drawn from. Independent Gauss-Markov processes with correlation ρ , (i.e., $\gamma_{d,n}^{(t)} = \rho \gamma_{d,n}^{(t-1)} + (1 - \rho)e_{d,n}^{(t)}$, $d = 1, 2$, where $e_{d,n}^{(t)}$ is a Gaussian perturbation process), were used to model amplitude correlations for both $\bar{\gamma}_1$ and $\bar{\gamma}_2$, while $p(\bar{\mathbf{s}})$ was chosen to enforce joint sparsity. As

an observation-channel model, we considered additive noise with a heavy-tailed distribution. Specifically, the observation-channel prior was $p(y_m^{(t)} | z_m^{(t)}) = (1 - \pi)\mathcal{N}(z_m^{(t)}; 0, \nu_0) + \pi\mathcal{N}(z_m^{(t)}; 0, \nu_1)$, where $\nu_1 = 1000 \cdot \nu_0$, and $\pi = 0.10$.

Iso-dB contours of median TNMSE performance for a signal model in which $N = 1024$, $T = 6$, $\rho = 0.95$, and $\text{SNR} = 25\text{dB}$ are plotted in Fig. 3 for four different recovery methods. In Fig. 3a, we show the performance of a support-aware genie smoother that had perfect knowledge of the support, and perfect knowledge of the signal model and its parameters, \mathcal{P} . In Fig. 3b, we show the performance of TurboGAMP with perfect signal model knowledge (i.e., no need for EM learning) but no support knowledge. Then, in Fig. 3c, we plot the performance of EMturboGAMP with EM learning of the model parameters. Finally, Fig. 3d shows the performance of a structure-agnostic GAMP recovery method, which had knowledge of $p(x_n^{(t)}; \mathcal{P})$ and $p(y_m^{(t)} | z_m^{(t)}; \mathcal{P})$, but no knowledge of $p(\bar{\gamma}_1)$, $p(\bar{\gamma}_2)$, or $p(\bar{\mathbf{s}})$. The structure-agnostic GAMP method was allowed to refine model parameters using EM learning to compensate for model mismatch. Despite this, the advantages of exploiting the additional signal structure are clearly evident.

We next used EMturboGAMP to recover an audio signal from sub-Nyquist samples. The audio clip is a 7 second recording of a trumpet solo, which presents a challenge for CS methods since the signal will not be perfectly 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. 4 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 five signal models, each constructed by combining various objects of the four classes described in Section 4: BG-GAMP (a structureless signal model with Bernoulli-Gaussian signal marginals), GM-GAMP (a structureless signal model with Bernoulli-Gaussian-mixture signal marginals with $D = 4$ Gaussian mixture components), DCS-BG-GAMP (a structured Bernoulli-Gaussian dynamic CS model described in [2]), MRF-BG-GAMP (Bernoulli-Gaussian GAMP with a spatio-temporal Markov random field structure on the signal support [13]), and DCS-GM-GAMP (a Bernoulli-Gaussian-mixture dynamic CS model with $D = 4$ mixture components). For each algorithm, unique model parameters were learned using EM at each timestep.

In Table 1 we present the results of applying each algorithm to the audio dataset for three different undersampling rates, δ . Overall, we see that performance improves at each undersampling rate as the signal model becomes more expressive. While GM-GAMP outperforms BG-GAMP at all undersampling rates, it is surpassed by DCS-BG-GAMP, MRF-BG-GAMP, and DCS-GM-GAMP, which offers the best TNMSE performance. Indeed, we observe that one can obtain comparable, or even better, performance with an undersampling rate of $\delta = \frac{1}{5}$ using any of the three structured sparsity techniques, with that obtained using BG-GAMP with an undersampling rate of $\delta = \frac{1}{3}$.

6. REFERENCES

- [1] M. F. Duarte and Y. C. Eldar, "Structured compressed sensing: From theory to applications," *IEEE Trans. Signal Process.*, vol.

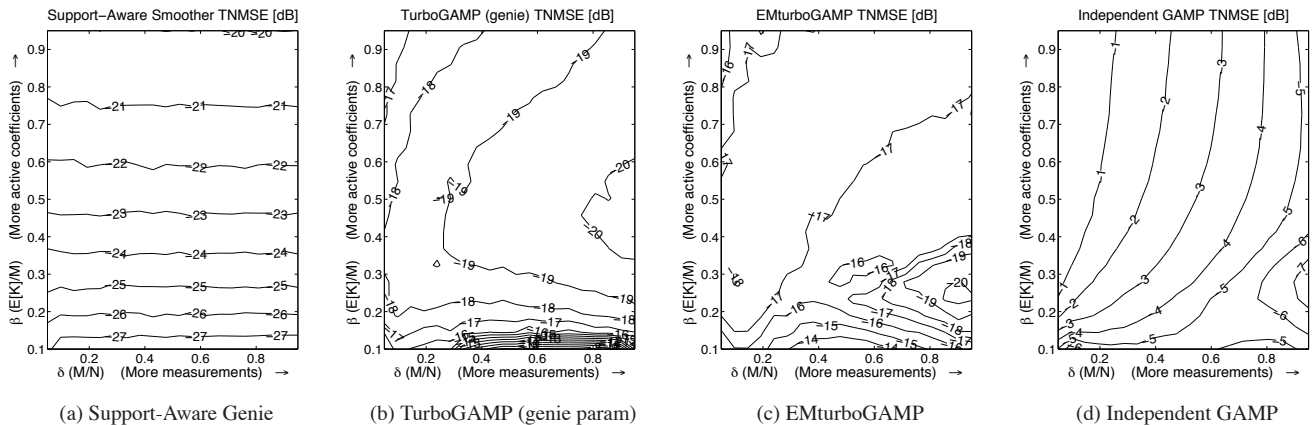


Fig. 3: Median TNMSE performance (in dB) across the sparsity-undersampling plane for four different recovery methods

		Undersampling Rate		
		$\delta = \frac{1}{2}$	$\delta = \frac{1}{3}$	$\delta = \frac{1}{5}$
Algorithm	BG-GAMP	-16.88 (dB) 9.11 (s)	-11.67 (dB) 8.27 (s)	-8.56 (dB) 6.63 (s)
	GM-GAMP ($D = 4$)	-17.49 (dB) 19.36 (s)	-13.74 (dB) 17.48 (s)	-10.23 (dB) 15.98 (s)
	DCS-BG-GAMP	-19.84 (dB) 10.20 (s)	-14.33 (dB) 8.39 (s)	-11.40 (dB) 6.71 (s)
	MRF-BG-GAMP	-21.15 (dB) 14.92 (s)	-16.25 (dB) 13.22 (s)	-12.35 (dB) 11.88 (s)
	DCS-GM-GAMP ($D = 4$)	-21.33 (dB) 20.34 (s)	-16.78 (dB) 18.63 (s)	-12.49 (dB) 10.13 (s)

Table 1: Performance on audio CS dataset (TNMSE (dB) | Runtime (s)) of two unstructured sparsity algorithms, BG-GAMP and GM-GAMP, and three structured sparsity algorithms, DCS-BG-GAMP, MRF-BG-GAMP, and DCS-GM-GAMP.




Fig. 4: DCT coefficient magnitudes (in dB) of an audio signal.

59, no. 9, pp. 4053–4085, 2011.

- [2] J. Ziniel, L. C. Potter, and P. Schniter, “Tracking and smoothing of time-varying sparse signals via approximate belief propagation,” in *Proc. 44th Asilomar Conf. Sig., Sys., & Comput. (SS&C)*, Pacific Grove, CA, Nov. 2010.
- [3] T. Faktor, Y. C. Eldar, and M. Elad, “Modeling statistical dependencies in sparse representations,” in *Wkshp. on Signal Process. w/ Adaptive Sparse Struct. Rep. (SPARS ’11)*, Edinburgh, UK, June 2011.
- [4] L. Yu, H. Sun, J. P. Barbot, and G. Zheng, “Compressive sensing for clustered sparse signals,” in *Intl. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Prague, Czech Republic, May 2011.
- [5] Z. Zhang and B. D. Rao, “Sparse signal recovery with temporally correlated source vectors using Sparse Bayesian Learning,” *IEEE J. Selected Topics Signal Process.*, vol. 5, no. 5, pp. 912–926, Sept. 2011.
- [6] D. Baron and M. F. Duarte, “Universal MAP estimation in compressed sensing,” in *Proc. 49th Allerton Conf. Comm., Control, & Comput.*, Monticello, IL, Sept. 2011.
- [7] P. Schniter, “Turbo reconstruction of structured sparse signals,” in *Conf. on Information Sciences and Systems (CISS)*, Princeton, NJ, Mar. 2010, pp. 1–6.
- [8] S. Rangan, “Generalized approximate message passing for estimation with random linear mixing,” in *Proc. IEEE Int. Symp. Inform. Theory*, St. Petersburg, Russia, Aug. 2011, pp. 2168–2172.
- [9] F. R. Kschischang, B. J. Frey, and H. A. Loeliger, “Factor graphs and the sum-product algorithm,” *IEEE Trans. Inform. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.
- [10] S. Rangan, A. K. Fletcher, V. K. Goyal, and P. Schniter, “Hybrid approximate message passing with applications to structured sparsity,” arXiv:1111.2581 [cs.IT], Nov. 2011.
- [11] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the EM algorithm,” *J. Roy. Statist. Soc., B*, vol. 39, pp. 1–38, 1977.
- [12] J. Ziniel and P. Schniter, “Efficient high-dimensional inference in the multiple measurement vector problem,” arXiv:1111.5272 [cs.IT], Nov. 2011.
- [13] S. Som and P. Schniter, “Approximate message passing for recovery of sparse signals with Markov-random-field support structure,” in *ICML Workshop on Structured Sparsity*, Bellevue, Wash., Jul. 2011.