Adaptive compressive noncoherent change detection: An AMP-based approach

Philip Schniter
Dept. of ECE, The Ohio State University, Columbus, OH 43210, Email: schniter@ece.osu.edu

***Abstract*—We propose a turbo approximate message passing (AMP) algorithm to detect spatially clustered changes in signal magnitude, relative to a reference signal, from compressive linear measurements. We then show how the Gaussian posterior approximations generated by this scheme can be used for mutual-information based measurement kernel adaptation. Numerical simulations show excellent performance.**

I. Summary

A. Compressive noncoherent change detection

In *change detection*, one observes noisy linear measurements $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{w} \in \mathbb{C}^M$ of a signal $\boldsymbol{x} \in \mathbb{C}^N$ and aims to detect changes in $\boldsymbol{x}$ relative to a known reference signal $\boldsymbol{r} \in \mathbb{C}^N$. Here, $\boldsymbol{A}$ represents a known measurement kernel and $\boldsymbol{w}$ represents white Gaussian noise.

Our focus is *noncoherent* change detection, where the phase difference between $\boldsymbol{r}$ and $\boldsymbol{x}$ may be significant even in the absence of a material change. In this case, the goal is to detect *changes in magnitude* between $\boldsymbol{x}$ and $\boldsymbol{r}$. An example application arises in radar, where small (e.g., wind-induced) movements in foliage can result in a large independent phase differences in each pair (x_n, r_n) even when the material present in pixel n has not changed.

We are particularly interested in the *compressive* case, where the number of measurements, M, is less than the signal length, N. Although we assume that the magnitude changes $|\boldsymbol{x}| - |\boldsymbol{r}|$ are sparse, and possibly even structured-sparse, we do not assume that the signals $\boldsymbol{x}$ and $\boldsymbol{r}$ themselves are sparse in a known basis, nor is their difference $\boldsymbol{x} - \boldsymbol{r}$. Note that, if (an estimate of) $\boldsymbol{x}$ was available, then standard techniques [1] could be applied to detect changes between $|\boldsymbol{x}|$ and $|\boldsymbol{r}|$. However, we do not observe $\boldsymbol{x}$, and the lack of sparsity in $\boldsymbol{x}$ (and $\boldsymbol{x} - \boldsymbol{r}$) prevents the use of standard compressed sensing techniques to recover $\boldsymbol{x}$ from $\boldsymbol{y}$. Thus, the problem is somewhat challenging.

Our approach exploits that fact that, under the sparse magnitude-change assumption, $|\boldsymbol{r}|$ does provide information about $|\boldsymbol{x}|$ that can aid in compressive recovery of $\boldsymbol{x}$ and—more importantly—*joint change detection and signal recovery*. For this, we model

$$x_n = s_n c_n + (1 - s_n)(r_n e^{j\theta_n} + d_n), \tag{1}$$

where $s_n \in \{0, 1\}$ indicates the presence of a change, $c_n \in \mathbb{C}$ represents the changed pixel value, $\theta_n \in [0, 2\pi)$ represents an unknown phase rotation, and $d_n \in \mathbb{C}$ represents a small deviation allowed in an "unchanged" pixel. We then assign the priors

$$\begin{aligned} c_n &\sim \mathcal{CN}(0, \nu^r) \text{ i.i.d with } \nu^r = \tfrac{1}{N} \textstyle\sum_{n=1}^N |r_n|^2 \\ \theta_n &\sim \mathcal{U}[0, 2\pi) \text{ i.i.d} \\ d_n &\sim \mathcal{CN}(0, \nu^d) \text{ i.i.d with } \nu^d \ll \nu^r \\ s_n &\sim \text{Markov}, \end{aligned} \tag{2}$$

where the Markov property on $\{s_n\}$ captures the fact that changes are often spatially clustered. Finally, we jointly infer the change pattern $\boldsymbol{s}$ and the signal $\boldsymbol{x}$ using the turbo extension [2] of the Bayesian approximate message passing (AMP) algorithm [3]. To our knowledge, the use of AMP with a signal prior of this form is novel.

With support from NSF grant CCF-1018368 and DARPA/ONR grant N66001-10-1-4090.

B. Measurement adaptation

We now allow the aforementioned approach multiple *adaptive* measurement steps, building on the work in [4]. In step $t = 1, \ldots, T$, the detector collects measurements $\boldsymbol{y}_t = \boldsymbol{A}_t \boldsymbol{x} + \boldsymbol{w}_t \in \mathbb{C}^{M_t}$ using a kernel $\boldsymbol{A}_t$ optimized around the uncertainty of $\boldsymbol{x}$ (or $\boldsymbol{s}$) that remains from inference based on the cumulative previous measurements $\underline{\boldsymbol{y}}_{t-1} \triangleq [\boldsymbol{y}_1^\mathsf{T}, \ldots, \boldsymbol{y}_{t-1}^\mathsf{T}]^\mathsf{T}$. When optimizing $\boldsymbol{A}_t$ for the recovery of $\boldsymbol{x}$, [4] suggested to maximize the mutual information (MI) between *Gaussian approximations* of the random vectors $\mathsf{x} \sim p(\boldsymbol{x}|\underline{\boldsymbol{y}}_{t-1})$ and $\mathsf{y}_t \sim p(\boldsymbol{y}_t|\underline{\boldsymbol{y}}_{t-1}; \boldsymbol{A}_t)$. Indeed, when x and y_t are jointly Gaussian, [4] established that the MI-maximizing $\boldsymbol{A}_t$ is computable using eigendecomposition and waterfilling. Conveniently, the necessary Gaussian approximation on x is an output of turbo AMP. For $\boldsymbol{s}$-adaptive kernel design, we now propose a similar approach based on a Gaussian approximation of $\mathsf{s} \sim p(\boldsymbol{s}|\underline{\boldsymbol{y}}_{t-1})$.

C. Numerical results

The left plot shows the normalized mean-squared error (NMSE) in recovering $\boldsymbol{x} \in \mathbb{C}^{200}$ versus cumulative number of measurements M, under 15 dB SNR and $\nu^d = 0.001$, averaged over 1000 realizations. All quantities were drawn according to (2), with the binary Markov chain for $\boldsymbol{s}$ activating 10% changes on average, clustered with an average run-length of 10. There, turbo-AMP with MI-x kernel adaptation performed best, approximately 2dB better than turbo-AMP with i.i.d-Gaussian $\boldsymbol{A}$, while LMMSE estimation of $\boldsymbol{x}$ with i.i.d-Gaussian $\boldsymbol{A}$ performed significantly worse. The right plot shows the corresponding normalized detection error rate (NDER), where turbo-AMP with MI-s kernel adaptation performed best, and significantly better than Bayes-optimal change detection using LMMSE-$\boldsymbol{x}$, even when change clustering was exploited. Although turbo-AMP with MI-s kernel adaptation did not work well for $\boldsymbol{x}$-recovery, we did not expect it to, since it was optimized for change detection.

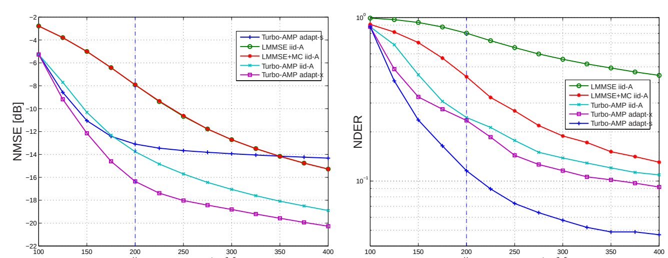

References

[1] H. V. Poor, *An Introduction to Signal Detection and Estimation*. New York: Springer, 2nd ed., 1994.

[2] P. Schniter, "Turbo reconstruction of structured sparse signals," in *Proc. Conf. Inform. Science & Syst.*, (Princeton, NJ), pp. 1–6, Mar. 2010.

[3] D. L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. Motivation and construction," in *Proc. Inform. Theory Workshop*, (Cairo, Egypt), pp. 1–5, Jan. 2010.

[4] P. Schniter, "Exploiting structured sparsity in Bayesian experimental design," in *Proc. IEEE Workshop Comp. Adv. Multi-Sensor Adaptive Process.*, (San Juan, Puerto Rico), pp. 357–360, Dec. 2011.