Constrained adaptive sensing

Mark A. Davenport and Andrew K. Massimino
School of Electrical and Computer Engineering
Georgia Institute of Technology
Email: {mdav, massimino}@gatech.edu

Deanna Needell
Department of Mathematical Sciences
Claremont McKenna College
Email: dneedell@cmc.edu

Tina Woolf
Institute of Mathematical Sciences
Claremont Graduate University
Email: tina.woolf@cgu.edu

I. Adaptive sensing

We focus on the problem of estimating a vector $\mathbf{x} \in \mathbb{C}^n$ from a small number of noisy linear measurements of the form

$$y_i = \langle \mathbf{a}_i, \mathbf{x} \rangle + z_i, \quad i = 1, \dots, m \tag{1}$$

where $\|\mathbf{a}_i\|_2 = 1$ and $z_i \sim \mathcal{N}(0, \sigma^2)$. We will typically be interested in the case where the vector $\mathbf{x}$ is *sparse*, meaning that it has only s nonzeros with $s \ll n$. In such a case, one can obtain a significantly more accurate estimate of $\mathbf{x}$ by adaptively selecting the $\mathbf{a}_i$ based on the previous measurements (compared to the standard nonadaptive approach), provided that the signal-to-noise ratio (SNR) is sufficiently large (e.g., see [1–3]). In particular, a typical nonadaptive algorithm can produce an estimate $\widehat{\mathbf{x}}$ satisfying

$$\mathbb{E}\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \leq C \frac{n \log n}{m} s\sigma^2, \tag{2}$$

where $C > 1$ is a constant. One can show that this is essentially optimal (see [4] for further discussion and references). In contrast, provided that σ^2 is not too large relative to the nonzero entries of $\mathbf{x}$, a well-designed adaptive scheme, where the $\mathbf{a}_i$ are chosen sequentially as in [1], [2], will determine the support with high probability and hence has the potential to achieve

$$\mathbb{E}\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \leq C' \frac{s}{m} s\sigma^2, \tag{3}$$

which represents a substantial improvement when $s \ll n$.

Given the large potential for improvement over nonadaptive approaches, one would expect adaptive sensing schemes to be quite attractive in practical applications. However, in most applications, there are a number of *constraints* on what kinds of measurements can actually be acquired. For example, in a number of settings, physical hardware constraints dictate that Fourier coefficients are the only observable measurements. Thus, a natural question is whether adaptive sensing techniques can still be efficiently deployed in scenarios where the practitioner must respect these kinds of constraints. We prove that for certain measurement ensembles, constrained adaptive sensing offers little improvement over standard nonadaptive approaches, *no matter how large the SNR*. On the other hand, we provide both theoretical and empirical evidence that in other cases, adaptivity can still result in substantial improvement even in the constrained setting.

II. Limitations of constrained sensing

As a prototypical example of constrained sensing, we first consider the setting where the $\mathbf{a}_i$ must be selected from the Fourier measurement ensemble, i.e., each $\mathbf{a}_i$ is a row of the Discrete Fourier Transform (DFT) matrix $\mathbf{F} \in \mathbb{C}^{n \times n}$ with entries given by $F_{jk} = \frac{1}{\sqrt{n}} \exp(-2\pi\sqrt{-1}jk/n)$. In this constrained setting we show the following lower bound.

Theorem 1: Suppose that $\mathbf{x}$ is s-sparse and that samples are acquired via (1) using m vectors chosen from the rows of the DFT matrix $\mathbf{F}$. Then for *any* estimate $\widehat{\mathbf{x}}$ obtained by *any* adaptive scheme, we have

$$\mathbb{E}\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \geq \frac{n}{m} s\sigma^2. \tag{4}$$

This shows that even using an optimal choice of sensing vectors, the recovery error is still proportional to $\frac{n}{m} s\sigma^2$, which falls far short of the potential gains possible in the unconstrained setting shown in (3). This result is intuitive given the incoherence of the DFT and canonical bases. It is also possible to generalize this result to a somewhat broader class of measurement ensembles. This is somewhat reminiscent of existing results which show (in an unconstrained setting) that for a certain range of worst-case SNRs, adaptive schemes do not result in a substantial improvement in terms of estimation accuracy [3], [5]. However, these arguments only apply for a narrow range of SNRs, and if the SNR improves (by only a small constant factor), there is a dramatic transition and adaptivity yields significant improvements [1], [2]. In contrast, Theorem 1 applies *no matter how large the SNR*, and so in a sense is far more pessimistic.

III. Practical potential for constrained sensing

While Theorem 1 suggests that adaptivity can be of only limited benefit in certain constrained settings, it is important to note that this result applies only to certain specific classes of measurements/sparsifying bases (e.g., the case of DFT measurements of canonically sparse bases). In other constrained settings it may still be possible to realize significant performance gains via a carefully implemented adaptive scheme. To illustrate this potential, we propose practical algorithms for constrained adaptive sensing and show that these methods exhibit promising performance in a different setting where the signals are sparse in an alternative basis.

Our approach appeals to connections between our problem and the theory of optimal experimental design [6]. In particular, we note that if we knew the support of $\mathbf{x}$ in advance, then it is possible to obtain an optimal, continuous "weighting" of the elements in our measurement ensemble by solving a simple convex optimization problem (which corresponds to the A-optimality criterion in experimental design [6]). The challenge in applying this in practice is that (*i*) the support of $\mathbf{x}$ is not known, we must estimate this from the measurements, and (*ii*) an optimal "weighting" of the elements may not be practically useful—we may ultimately be required to select an integer number of (unweighted) measurements from our ensemble. We propose techniques for addressing these obstacles and evaluate these approaches in settings inspired by tomography, demonstrating significant performance gains over nonadaptive approaches.

As an example, Fig. 1 shows the results of a representative simulation in which DFT measurements are taken of signals which are sparse in a wavelet basis, in particular, whose Haar wavelet decomposition is sparsely supported on a tree. In the nonadaptive case, all measurements are taken without knowledge of the support according to variable density sampling (VDS) [7]. In the adaptive case, half of the measurements are taken using VDS and are used to estimate the support with ℓ_1 minimization. The other half are then chosen to optimize recovery error on this support. For comparison, the oracle adaptive error shows the performance when the support is known and all measurements are optimized. We observe that the adaptive algorithm significantly outperforms the nonadaptive algorithm and is competitive with an oracle estimate.

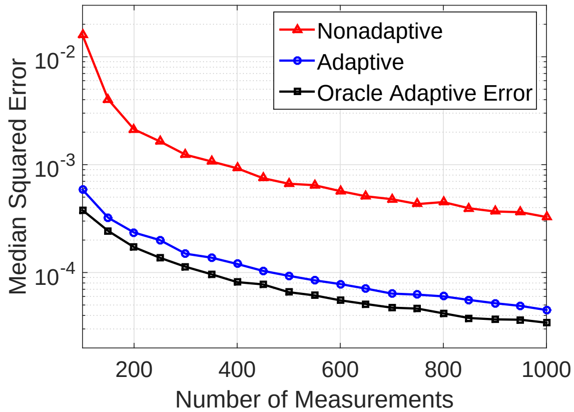

Fig. 1. Simulation results with Fourier measurements, Haar basis tree-sparse signals, $s = 10$, $n = 1024$, $\sigma = 0.01$.

Acknowledgment

M.A.D. and A.K.M. are supported by grants NRL N00173-14-2-C001, AFOSR FA9550-14-1-0342, NSF CCF-1409261 and CCF-1350616, and a gift from LexisNexis Risk Solutions. D.N. was partially supported by NSF Career DMS-1348721 and the Alfred P. Sloan Foundation. T.W. was partially supported by NSF Career DMS-1348721.

References

[1] M. Davenport and E. Arias-Castro, “Compressive binary search,” in *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, Cambridge, MA, Jul. 2012.

[2] M. Malloy and R. Nowak, “Near-optimal adaptive compressed sensing,” *IEEE Trans. Inform. Theory*, vol. 60, no. 7, pp. 4001–4012, 2014.

[3] R. Castro, “Adaptive sensing performance lower bounds for sparse signal detection and support estimation,” *Bernoulli*, vol. 20, no. 4, pp. 2217–2246, 2014.

[4] E. Candès and M. Davenport, “How well can we estimate a sparse vector?” *Appl. Comput. Harmon. Anal.*, vol. 34, no. 2, pp. 317–323, 2013.

[5] E. Arias-Castro, E. Candès, and M. Davenport, “On the fundamental limits of adaptive sensing,” *IEEE Trans. Inform. Theory*, vol. 59, no. 1, pp. 472–481, 2013.

[6] F. Pukelsheim, *Optimal Design of Experiments*. Society for Industrial and Applied Mathematics, 2006. [Online]. Available: http://epubs.siam.org/doi/abs/10.1137/1.9780898719109

[7] F. Krahmer and R. Ward, “Stable and robust sampling strategies for compressive imaging,” *IEEE Trans. Image Processing*, vol. 23, no. 2, pp. 612–622, 2014.