Constrained adaptive sensing of sparse signals

Mark A. Davenport
Georgia Institute of Technology, Atlanta, U.S.A. – mdav@gatech.edu

Suppose that we wish to estimate a sparse vector from a small number of noisy linear measurements. In the setting where the measurements are selected in advance (independently of the signal) we now have a rich understanding of both practical algorithms and the theoretical limits on the performance of these algorithms. A typical result from this literature states that for a suitable measurement design, one can estimate a sparse vector with an accuracy that matches the minimax lower bound up to a constant factor. Such results have had a tremendous impact in a variety of practical settings. In particular, they provide the mathematical foundation for "compressive sensing" – a paradigm for efficient sampling that has inspired a range of new sensor designs over the last decade.

However, the sensing strategies typically promoted in this literature also suffer from an unfortunate sensitivity to noise when the number of measurements is severely limited. A possible way to overcome this problem is to consider the case where the measurements are acquired sequentially in an adaptive manner. Perhaps surprisingly, one can prove that in certain worst-case regimes, adaptivity does not allow for substantial improvement over standard nonadaptive techniques in terms of the minimax performance. Nonetheless, one can also show that there are important regimes where the benefits of adaptive sensing are clear and overwhelming, and can be achieved via relatively simple algorithms. Unfortunately, these algorithms cannot typically be applied in the common setting where our measurements are limited to satisfy certain natural constraints.

We will explore both the limitations and advantages of adaptive sensing in the setting where our measurements are constrained to be chosen from some pre-defined ensemble. We will show that in the constrained case, the room for improvement is dictated by the interaction between the sparsity basis for our signal and the measurement ensemble. For certain pairs of measurement ensembles and sparsity bases, adaptive sensing cannot offer significant improvements over nonadaptive techniques, regardless of the signal-to-noise ratio. On the other hand, we will also provide both theoretical and empirical evidence that there do exist practical scenarios where adaptivity can indeed result in the improved performance that we can obtain in the unconstrained case. We will then describe algorithms for constrained adaptive sensing that build on connections with the theory of optimal experimental design. Finally, we will show that these algorithms exhibit promising performance in some representative applications.

Keywords: sparsity; adaptivity; optimal experimental design; compressive sensing.