Sensing matrix design*

Mark A. Davenport

This work is produced by OpenStax-CNX and licensed under the
Creative Commons Attribution License 3.0†

Abstract

This module provides an overview of the sensing matrix design problem in compressive sensing.

In order to make the discussion more concrete, we will restrict our attention to the standard finite-dimensional compressive sensing (CS) model. Specifically, given a signal $x \in \mathbb{R}^N$, we consider measurement systems that acquire M linear measurements. We can represent this process mathematically as

$$y = \Phi x, \tag{1}$$

where Φ is an $M \times N$ matrix and $y \in \mathbb{R}^M$. The matrix Φ represents a *dimensionality reduction*, i.e., it maps $\mathbb{R}^N$, where N is generally large, into $\mathbb{R}^M$, where M is typically much smaller than N. Note that in the standard CS framework we assume that the measurements are *non-adaptive*, meaning that the rows of Φ are fixed in advance and do not depend on the previously acquired measurements. In certain settings adaptive measurement schemes can lead to significant performance gains.

Note that although the standard CS framework assumes that x is a finite-length vector with a discrete-valued index (such as time or space), in practice we will often be interested in designing measurement systems for acquiring continuously-indexed signals such as continuous-time signals or images. For now we will simply think of x as a finite-length window of Nyquist-rate samples, and we temporarily ignore the issue of how to directly acquire compressive measurements without first sampling at the Nyquist rate.

There are two main theoretical questions in CS. First, how should we design the sensing matrix Φ to ensure that it preserves the information in the signal x? Second, how can we recover the original signal x from measurements y? In the case where our data is sparse or compressible, we will see that we can design matrices Φ with $M \ll N$ that ensure that we will be able to recover the original signal accurately and efficiently using a variety of practical algorithms.

We begin in this part of the course[1] by first addressing the question of how to design the sensing matrix Φ. Rather than directly proposing a design procedure, we instead consider a number of desirable properties that we might wish Φ to have (including the null space property, the restricted isometry property, and bounded coherence). We then provide some important examples of matrix constructions that satisfy these properties.

*Version 1.6: Apr 14, 2011 6:33 pm +0000

†http://creativecommons.org/licenses/by/3.0/

[1] <http://cnx.org/content/col11133/latest/>