

Sparse Recovery Using Sparse Random Matrices

Piotr Indyk

Abstract

Over the recent years, a new linear method for compressing high-dimensional data (e.g., images) has been discovered. For any high-dimensional vector x , its sketch is equal to Ax , where A is an $m \times n$ matrix (possibly chosen at random). Although typically the sketch length m is much smaller than the number of dimensions n , the sketch contains enough information to recover an approximation to x . At the same time, the linearity of the sketching method is very convenient for many applications, such as data stream computing and compressed sensing.

The major sketching approaches can be classified as either combinatorial (using sparse sketching matrices) or geometric (using dense sketching matrices). They achieve different trade-offs, notably between the compression rate and the running time. Thus, it is desirable to understand the connections between them, with the goal of obtaining the "best of both worlds" solution. Several recent results established such connections, indicating that the two approaches are just different manifestations of the same underlying phenomenon. This enabled the development of novel algorithms, including the first algorithms that provably achieve the (asymptotically) optimal compression rate and near-linear recovery time simultaneously.

In this talk we give an overview of the results in the area, as well as look at some of them in more detail. In particular, we will describe a new algorithm, called "Sequential Sparse Matching Pursuit (SSMP)". In addition to having the aforementioned theoretical guarantees, the algorithm works well on real data, with the recovery quality often outperforming that of more complex algorithms, such as l_1 minimization.

Joint work with: Radu Berinde, Anna Gilbert, Howard Karloff, Milan Ruzic and Martin Strauss.