Fast Algorithms for Structured Sparsity*

Piotr Indyk

Abstract

Sparse representations of signals (i.e., representations that have only few non-zero or large coefficients) have emerged as powerful tools in signal processing theory, algorithms, machine learning and other applications. However, real-world signals often exhibit rich structure beyond mere sparsity. For example, a natural image, once represented in the wavelet domain, often has the property that its large coefficients occupy a subtree of the wavelet hierarchy, as opposed to arbitrary positions. A general approach to capturing this type of additional structure is to model the support of the signal of interest (i.e., the set of indices of large coefficients) as belonging to a particular family of sets. Computing a sparse representation of the signal then corresponds to the problem of finding the support from the family that maximizes the sum of the squares of the selected coefficients. Such a modeling approach has proved to be beneficial in a number of applications including compression, de-noising, compressive sensing and machine learning. However, the resulting optimization problem is often computationally difficult or intractable, which is undesirable in many applications where large signals and datasets are commonplace.

In this talk, I will outline some of the past and more recent algorithms for finding structured sparse representations of signals, including piecewise constant approximations, tree-sparse approximations and graph-sparse approximations. The algorithms borrow several techniques from combinatorial optimization (e.g., dynamic programming), graph theory, and approximation algorithms. For many problems the algorithms run in (nearly) linear time, which makes them applicable to very large datasets.

Joint work with Chinmay Hegde and Ludwig Schmidt

*Machine Learning External Seminar, Gatsby Unit, May 16, 2016.