

Nearly Linear-Time Algorithms for Structured Sparsity

Piotr Indyk, Massachusetts Institute of Technology

Sparse representations have emerged as powerful tools in signal processing theory, algorithms, and applications. However, real-world signal and image ensembles often exhibit rich structure beyond mere sparsity; for example, a natural image can be modeled as (approximately) tree-sparse in the wavelet domain. A general approach for capturing this type of additional structure is to model the signal of interest as belonging to a particular union-of-subspaces in the signal space. Such a modeling approach has proved to be beneficial in a number of applications including compression, denoising, and compressive sensing. However, enforcing prior structure often leads to more cumbersome and resource-intensive algorithms. This is undesirable in many scenarios where high-dimensional signal and image datasets are commonplace.

In this talk, I will outline some of our recent efforts towards building highly efficient algorithms for structured sparse modeling. Interestingly, our algorithms borrow several elements from combinatorial optimization, graph theory, and approximation algorithms, and are in some sense fundamentally non-convex. Despite this feature, our algorithms enjoy a (nearly) linear-time computational complexity, thereby enabling their applicability to very large datasets. In particular, I will show that our methods lead to the first nearly-linear time algorithms for stable compressive sensing recovery of signals and images which are tree-sparse in the wavelet-domain, as well as their generalizations.

Joint work with Chinmay Hegde and Ludwig Schmidt.