

Eigenvector localization, implicit regularization, and algorithmic anti-differentiation for large-scale graphs and networked data

Michael W. Mahoney, University of California, Berkeley

Although graphs and matrices are popular ways to model data drawn from many application domains, the size scale, sparsity properties, noise properties, and many other aspects of graphs and matrices arising in realistic machine learning and data analysis applications present fundamental challenges for traditional matrix and graph algorithms. Informally, this at root has to do with the observation that small-scale or local structure in many data sets is often better or more meaningful than large-scale or global structure, which is exactly the opposite of what many traditional asymptotic tools typically implicitly assume. For example, there might be meaningful small-scale clusters in social networks that are tree-like or expander-like at large size scales. Here, I will describe how eigenvector localization can be used to justify these sorts of claims, how localization can be engineered in a principled way into popular matrix and graph algorithms in order to characterize local properties in large data sets much more generally, and how seemingly-minor details in the approximate computation of localized (as well as non-localized) objectives can make a large difference in the usefulness of a given method in downstream applications. A common theme will be the critical role of what can be called implicit regularization---that is, regularization not in the usual sense of explicitly adding a regularization term and solving the modified objective exactly, but instead as an implicit side effect of heuristics that are often used to make approximation algorithms scale well---as well as the question of what is the objective function that an approximation algorithm implicitly solve exactly.