Distributed Algorithms for Graphical Models

Alan S. Willsky
Massachusetts Institute of Technology

In this talk we provide an overview of several topics related to message-passing algorithms for graphical models. After a brief review of the basic idea behind Belief Propagation (BP) algorithms, we describe a *walk-sum* interpretation of BP for Gaussian models, identifying explicitly how messages "collect" weighted walks through the graph and also making clear what walks BP collects, what walks it misses, and when it converges for graphs with cycles. We use these interpretations to describe a different algorithm based on the concept of a *Feedback Vertex Set* in which key nodes are identified that break all or the most important cycles in the graph and these are used for a multi-step messaging scheme that collects all walks if all cycles are broken or the most critical ones if only some are broken (in terms of collecting the most significant walks). This leads to some interesting current research directions that point out the underexplored idea of including protocol bits in messages, identifying more about the role and path those messages have taken. We then turn to problems of building graphical models with "nice" graphical structure (i.e., structure that leads to efficient inference). Topics to be discussed will be taken from the following: (a) a method based on hierarchical modeling and inspired by so-called *multipole* methods for the solution of partial differential equations that involves what we term *conjugate graphs*; (b) a constructive and provably consistent method for recovering latent tree models from sample data; and (c) a convex optimization methodology, extending ideas from the growing field of sparse modeling, to recover models that consist of a combination of a sparse graphical model together with a small number of nodes which are comparatively densely connected to the rest of the graph.