Stable message-passing and convex surrogates for learning in Markov random fields

Martin Wainwright

[a]Department of EECS, and Department of Statistics, UC Berkeley, Berkeley, CA, USA

Key statistical problems that arise in applications of Markov random fields (MRFs) are computing means and marginal distributions (e.g., for performing prediction) and estimating model parameters from data. Although easily solved for trees, these problems are intractable for general MRFs, which motivates the use of approximate methods. For instance, the sum-product or belief propagation algorithm is widely used in many fields (e.g., statistics, communication theory, computer vision) to compute approximate means/marginals in MRFs with cycles.

We consider the problem of learning MRF parameters from data, with the ultimate goal of using the estimated MRF to predict or smooth a new set of noisy observations. We propose a computationally efficient approach, based on constructing convex surrogates to the likelihood, such that the initial parameter estimation and subsequent prediction are coupled. As a particular example, we construct a surrogate likelihood based on a convexified Bethe approximation. Interestingly, even though this surrogate is implicitly defined, the function value and its derivatives can be computed efficiently by a tree-reweighted version of the belief propagation algorithm. We then prove that the parameter estimates obtained by maximizing the surrogate likelihood are asymptotically Gaussian but inconsistent. Nonetheless, this inconsistency turns out to beneficial in the prediction setting. In particular, we provide theoretical bounds on the performance loss of our computationally tractable method relative to the unattainable Bayes optimum, which depend on the graph structure and potential strength. These bounds hinge on the Lipschitz stability of the tree-reweighted message-passing updates.