

Learning Discrete Markov Random Fields with Nearly Optimal Runtime and Sample Complexity
Adam Klivans, University of Texas at Austin

We give an algorithm for learning the structure of an undirected graphical model over any finite alphabet that has nearly optimal runtime and sample complexity. We make no assumptions on the structure of the graph. For Ising models, this subsumes and improves on all prior work. For higher-order MRFs, these are the first results of their kind.

Our approach is new and uses a multiplicative-weight update algorithm. Our algorithm-- Sparsitron-- is easy to implement (has only one parameter) and holds in the online setting. It also gives the first provably efficient solution to the problem of learning sparse Generalized Linear Models (GLMs).

Joint work with Raghu Meka