Andrey Lokhov:
Learning Discrete Graphical Models with Neural Networks

Graphical models are widely used to represent joint probability distributions with an underlying conditional dependence structure. Learning of graphical models that describe the statistics of discrete variables is a particularly challenging problem, for which the maximum likelihood approach is intractable. We introduce the first sample-efficient method based on the Interaction Screening framework, GRISE, that allows one to provably learn fully general discrete factor models with node-specific discrete alphabets and multi-body interactions, specified in an arbitrary basis. A rigorous sample complexity analysis of GRISE shows systematic improvement compared to all previously known bounds. However, the computational cost of learning a graphical model can become prohibitive when a parsimonious basis representation of the energy function is unknown. We show how to overcome this limitation by using a neural-net representation for the conditional energy of the graphical model, acting as a function approximator. The resulting estimator NN-GRISE finds a sparse basis representation without being fed any prior information about the true model, and can also be used to learn the underlying structure of the graphical model, as well as learn a neural-net representation of the full energy function.