Learning Matrix Decomposition Structures

Bill Freeman
MIT

Many widely used models in unsupervised learning can be viewed as matrix decompositions, where the input matrix is expressed as sums and products of matrices drawn from a few simple priors. We present a unifying framework for matrix decompositions in terms of a context-free grammar which generates a wide variety of structures through the compositional application of a few simple rules. We use our grammar to generically and efciently infer latent components and estimate predictive likelihood for nearly 1000 structures using a small toolbox of reusable algorithms. Using best-rst search over our grammar, we can automatically choose the decomposition structure from raw data by evaluating only a tiny fraction of all models. This gives a recipe for selecting model structure in unsupervised learning situations. The proposed method almost always nds the right structure for synthetic data and backs off gracefully to simpler models under heavy noise. It learns plausible structures for datasets as diverse as image patches, motion capture, 20 Questions, and U.S. Senate votes, all using exactly the same code.

This work was first-authored by Roger Grosse, in collaboration with Ruslan Salakhutdinov, Josh Tenenbaum, and myself.