"Dimensionality reduction for optimization problems"

Abstract:

Modern applications such as machine learning involve the solution of huge scale nonconvex optimization problems, sometimes with special structure. Motivated by these challenges, we investigate more generally, dimensionality reduction techniques in the variable/parameter domain for local and global optimization that rely crucially on random projections. We describe and use sketching results that allow efficient projections to low dimensions while preserving useful properties, as well as other tools from random matrix theory and conic integral geometry. We focus on functions with low effective dimensionality, a common occurrence in applications involving overparameterized models and that can serve as an insightful proxy for the training landscape in neural networks. We obtain algorithms that scale well with problem dimension, allow inaccurate information and biased noise, have adaptive parameters and benefit from high-probability complexity guarantees and almost sure convergence.