DATA-DRIVEN GEOMETRY FOR CONVEX OPTIMIZATION

HONG YE TAN

Learning-to-optimize is an emerging framework that seeks to speed up the solution of certain optimization problems by leveraging training data. We propose a provably approximately convergent learning-to-optimize scheme for convex optimization based on a functional parameterization of the classical mirror descent algorithm. In particular, we model the underlying convex function with an input-convex neural network and derive corresponding convergence rate bounds. We demonstrate improved convergence rates on various convex image processing examples, as well as a dimensionality reduction method for neural network applications.