Fast Semidefinite Optimization with Latent Basis Learning

Georgina Hall
INSEAD
georgina.hall@insead.edu

When faced with a semidefinite program (SDP), it is often the case that we do not need to solve just one specific SDP, but rather a family of very similar problems with varying data, for example, when solving matrix completion problems to obtain movie recommendations for users with similar preferences. In this talk, I will present Fast Semidefinite Optimization (FSDO), a data-driven method to quickly solve SDPs coming from the same family. Our method learns a shared latent basis representation across the family, which is then used as input to a second-order cone program, whose solution constitutes an approximate solution to the original SDP. The learning is done using neural networks and leverages recent advances in differentiable convex optimization.