

Low-rank tensor completion by Riemannian optimization

Daniel Kressner

In tensor completion, the goal is to fill in missing entries of a partially known tensor under low-rank constraints. We consider low-rank formats that form Riemannian manifolds, such as the Tucker or the tensor train (TT) format. This allows for the application of Riemannian optimization techniques for solving the tensor completion problem. In particular, the nonlinear CG algorithm can be implemented such that it scales linearly in the size of the tensor. We illustrate the use of this algorithm for approximating multivariate functions as well as for parameter-dependent and stochastic partial differential equations.