



Mathematics for Machine Learning Seminar

Manifold constrained steepest descent

Mr. Kaiwei Yang
Department of Mathematics, HKU

Abstract

Norm-constrained linear minimization oracle (LMO)-based optimizers such as spectral gradient descent and Muon are attractive in large scale learning, but extending them to manifold-constrained problems is nontrivial and often leads to nested-loop schemes that solve tangent-space subproblems iteratively. We propose Manifold Constrained Steepest Descent (MCSD), a single-loop framework for optimization over manifolds that selects a norm-induced steepest-descent direction via an LMO applied to the Riemannian gradient, and then returns to the manifold via projection. Under standard smoothness assumptions, we establish convergence guarantees for MCSD and a stochastic momentum variant. We further introduce SPEL, the spectral-norm specialization of MCSD on the Stiefel manifold, which admits scalable implementations via fast matrix sign computations. Experiments on PCA, orthogonality-constrained CNNs, and manifold-constrained LLM adapter tuning demonstrate improved stability and competitive performance relative to standard Riemannian baselines and existing manifold-aware LMO methods.

Members of the Department of Mathematics are welcome to attend. Any content designated as confidential by the presenter should be kept confidential.

Date:	May 7, 2026 (Thursday)
Time:	3:00 pm – 4:00 pm
Venue:	Room 210, Run Run Shaw Building HKU