Manifold identification for ultimately communicationefficient distributed optimization

LEE Ching-pei¹, LI Yu-Sheng², CHIANG Wei-Lin³

¹Academia Sinica, ²National Taiwan University, ³University of California Berlekey

Abstract

This work proposes a progressive manifold identification approach for distributed optimization with sound theoretical justifications to greatly reduce both the rounds of communication and the bytes communicated per round for partly-smooth regularized problems such as the ℓ_1 - and group-LASSO-regularized ones. Our two-stage method first uses an inexact proximal quasi-Newton method to iteratively identify a sequence of low-dimensional manifolds in which the final solution would lie, and restricts the model update within the current manifold to gradually lower the order of the perround communication cost from the problem dimension to the dimension of the manifold that contains a solution and makes the problem within it smooth. After identifying this manifold, we take superlinear-convergent truncated semismooth Newton steps computed by preconditioned conjugate gradient to largely reduce the communication rounds by improving the convergence rate from the existing linear or sublinear ones to a superlinear rate. Experiments show that our method can be orders of magnitudes lower in the communication cost and an order of magnitude faster in the running time than the state of the art. This is a joint work with LI Yu-Sheng and CHIANG Wei-Lin, and the paper is published in the proceedings of the international conference of machine learning, 2020.

Keyword: Communication-efficient distributed optimization, Partly smooth functions, Manifold identification, Nonsmooth optimization