

Martin Takac

Lehigh University, USA

takac.mt@gmail.com

Communication remains the most significant bottleneck in the performance of distributed optimization algorithms for large-scale machine learning. In this paper, we propose a communication-efficient framework, CoCoA, that uses local computation in a primal-dual setting to dramatically reduce the amount of necessary communication. We provide a strong convergence rate analysis for this class of algorithms, as well as experiments on real-world distributed datasets with implementations in Spark. In our experiments, we find that as compared to state-of-the-art mini-batch versions of SGD and SDCA algorithms, CoCoA converges to the same .001-accurate solution quality on average 25x as quickly.

Joint work with Martin Jaggi (ETH Zurich), Virginia Smith (UC Berkeley, USA), Jonathan Terhorst (UC Berkeley, USA), Sanjay Krishnan (UC Berkeley, USA), Thomas Hofmann (ETH Zurich) and Michael I. Jordan (UC Berkeley, USA).