

Smoothed Normalization for Efficient Distributed Private Optimization

Federated learning enables training machine learning models while preserving the privacy of participants. Surprisingly, there is no differentially private distributed method for smooth, non-convex optimization problems. The reason is that standard privacy techniques require bounding the participants' contributions, usually enforced via clipping of the updates. Existing literature typically ignores the effect of clipping by assuming the boundedness of gradient norms or analyzes distributed algorithms with clipping but ignores DP constraints. In this work, we study an alternative approach via smoothed normalization of the updates motivated by its favorable performance in the single-node setting. By integrating smoothed normalization with an error-feedback mechanism, we design a new distributed algorithm α -NormEC. We prove that our method achieves a superior convergence rate over prior works. By extending α -NormEC to the DP setting, we obtain the first differentially private distributed optimization algorithm with provable convergence guarantees. Finally, our empirical results from neural network training indicate robust convergence of α -NormEC across different parameter settings.

Author: SHULGIN, Egor (King Abdullah University of Science and Technology (KAUST))

Co-auteurs: Dr RICHTARIK, Peter (KAUST); Dr KHIRIRAT, Sarit (KAUST)

Orateur: SHULGIN, Egor (King Abdullah University of Science and Technology (KAUST))

Classification de Session: ML

Classification de thématique: Machine learning