Scalable and Trustworthy Learning in Heterogeneous Networks

Tian Li

Computer Science Department and Data Science Institute, University of Chicago
litian@uchicago.edu

Motivated by the application of federated and collaborative learning, my work aims to develop principled methods for scalable and trustworthy data sharing in heterogeneous networks. In the talk, I discuss how heterogeneity affects federated optimization, and lies at the center of accuracy and trustworthiness constraints in federated learning. To address these concerns, I present scalable federated learning objectives and algorithms that rigorously account for and directly model the practical constraints. I will also explore trustworthy objectives and optimization methods for general ML problems beyond federated settings. I will conclude the talk by mapping out future research plans on optimization, data sharing, and privacy-preserving ML.

Large-scale optimization in heterogeneous environments. To fit data over practical (potentially) massive and unreliable networks, optimization methods must first and foremost deliver reasonably accurate solutions in an efficient manner. I developed FedProx, the first optimization framework with theoretical guarantees under realistic heterogeneity assumptions (Li et al. 2020a). The algorithmic modification we made is to add a proximal term to the local objective. This simple change makes the algorithm more amenable to theoretical analysis, and delivers more stable convergence. Federated optimization has received a significant amount of recent research attention, and FedProx (as well as another principled variant I develop named FedDANE (Li et al. 2019)) is a simple method that takes a first step towards analyzing and improving federated optimization under practical assumptions.

Scalable and principled trustworthy learning. Heterogeneity not only affects the convergence of optimization methods, but also poses challenges to the trustworthiness of the final model. For example, models trained jointly across massive clients can sometimes perform unfairly or even catastrophically on subsets of the network. I developed the q-Fair FL (q-FFL) objective, which minimizes an aggregate reweighted loss parameterized by q to effectively upweight clients with higher loss dynamically (Li et al. 2020b). This framework allows for a scalable, principled, and tunable fairness/utility tradeoff, and is easy to implement in a scalable fashion. Simultaneously satisfying the competing constraints of accuracy, fairness (performance uniformity), and robustness (against poisoning attacks) is exceptionally difficult under data heterogeneity. I proposed a simple multi-task learning method, Ditto, as a unified approach for satisfying the diverse constraints in federated learning (Li et al. 2021b). The Ditto work opened up many interesting questions regarding the connections between personalization and trustworthiness constraints.

The insights from my prior works can be extended to address deficiencies of empirical risk minimization (ERM) in broad ML problems, when there are shifts between training and test distributions. Motivated partially by exponential tilting and q-FFL, we proposed a *tilted* empirical risk minimization (TERM) objective, which introduces a flexible tilting hyperparameter to control the impact of individual losses in a principled manner (Li et al. 2021a, 2023). By tuning this parameter, we were able to *provably* decrease the influence of outliers, enable fairness or robustness, and provide better generalization. Lastly, privacy is a critical concern in both centralized and federated learning, but the study of differential privacy (DP) in machine learning has been largely limited to (variants) of the very basic SGD optimizer. I proposed new algorithms that leverage some structures of gradients to provably improve privacy-utility tradeoffs of adaptive optimization (such as AdaGrad and Adam).

References

Li, T.; Beirami, A.; Sanjabi, M.; and Smith, V. 2021a. Tilted Empirical Risk Minimization. In *ICLR*.

Li, T.; Beirami, A.; Sanjabi, M.; and Smith, V. 2023. On Tilted Losses in Machine Learning: Theory and Applications. *JMLR*.

Li, T.; Hu, S.; Beirami, A.; and Smith, V. 2021b. Ditto: Fair and Robust Federated Learning Through Personalization. In *ICML*.

Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2019. FedDANE: A Federated Newton-Type Method. In *Asilomar Conference*.

Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020a. Federated Optimization in heterogeneous networks. In *MLSys*.

Li, T.; Sanjabi, M.; Beirami, A.; and Smith, V. 2020b. Fair Resource Allocation in Federated Learning. In *ICLR*.