Private and Decentralized Machine Learning

Aurélien Bellet, Jan Ramon

November 3, 2016

Team and contact

- Équipe Magnet, INRIA/CRIStAL: http://team.inria.fr/magnet
- Aurélien Bellet (aurelien.bellet@inria.fr), Jan Ramon (jan.ramon@inria.fr)

Keywords

Machine Learning, Decentralized Algorithms, Privacy, Convex Optimization.

Context

Increasing amounts of data are being produced by interconnected devices such as mobile phones, connected objects, sensors, *etc*. For instance, history logs are generated when a smartphone user browses the web, gives product ratings and executes various applications. The currently dominant approach to extract useful information from such data is to collect all users' personal data on a server (or a tightly coupled system hosted in a data center) and apply centralized machine learning and data mining techniques. However, this centralization poses a number of issues, such as the need for users to "surrender" their personal data to the service provider without much control on how the data will be used, while incurring potentially high bandwidth and device battery costs.

In this internship, we are interested in the alternative setting of *decentralized machine learning*: a set of learning agents organized in a peer-to-peer network collaborate to learn a model based on the union of their personal datasets, without any central entity required for coordination or aggregation. Most existing approaches for decentralized learning rely on decentralized variants of optimization algorithms such as gradient descent [5, 6], the Alternating Direction Method of Multipliers (ADMM) [8, 7] or dual averaging [8, 2]. In these algorithms, agents exchange model parameters and/or gradient updates. While this is arguably better than sharing personal data directly, it is well-known that exchanging such information can still leak some sensitive information about the data used to compute these parameters/gradients (see e.g. [4, 1]).

Objectives

The goal of this internship is to propose and analyze privacy-preserving mechanisms for decentralized machine learning. In machine learning, the most popular notion of privacy is differential privacy [4, 1], which gives strong probabilistic guarantees. Differential privacy can be achieved by adding noise (drawn from a Gaussian or Laplace distribution) to various quantities: either the data itself, the model updates, the objective function, or the output [4, 1, 3]. This internship will study how to apply such strategies to decentralized learning, and try to assess their relative merits and drawbacks in terms of convergence rate and communication cost.

The tentative work-plan is as follows:

1. Review the relevant literature on decentralized learning and privacy-preserving mechanisms.
2. Propose, analyze and evaluate some protocols for differentially private decentralized learning.
3. If time permits, investigate some further questions, such as (i) how the network topology may influence privacy guarantees, (ii) how to design privacy-preserving mechanisms which reduce the amount of communication, (iii) how to make use of other private techniques (secure multiparty computation, homomorphic encryption).

Skills

Basics in machine learning, algorithms and complexity, linear algebra and probability.

References

[1] K. Chaudhuri, C. Monteleoni, and C. Monteleoni. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.

[2] I. Colin, A. Bellet, J. Salmon, and S. Clémençon. Gossip Dual Averaging for Decentralized Optimization of Pairwise Functions. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.

[3] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy Aware Learning. In *Advances in Neural Information Processing Systems 25*, pages 1439–1447, 2012.

[4] C. Dwork. Differential privacy: A survey of results. In *Proceedings of the 5th International Conference on Theory and Applications of Models of Computation*, pages 1–19, 2008.

[5] A. Nedic and A. E. Ozdaglar. Distributed Subgradient Methods for Multi-Agent Optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.

[6] S. S. Ram, A. Nedic, and V. V. Veeravalli. Distributed Stochastic Subgradient Projection Algorithms for Convex Optimization. *Journal of Optimization Theory and Applications*, 147(3):516–545, 2010.

[7] P. Vanhaesebrouck, A. Bellet, and M. Tommasi. Decentralized Collaborative Learning of Personalized Models over Networks. Technical report, arXiv:1610.05202, October 2016.

[8] E. Wei and A. E. Ozdaglar. Distributed Alternating Direction Method of Multipliers. In *Proceedings of the 51th IEEE Conference on Decision and Control*, pages 5445–5450, 2012.