How to Collaboratively Learn to Taste Beer — Online Algorithms for Decentralized and Personalized Machine Learning on Graphs

Aurélien Bellet, Marc Tommasi, Claudio Gentile

October 6, 2017

Team and contact

- Équipe Magnet, INRIA/CRIStAL: http://team.inria.fr/magnet
- Aurélien Bellet (aurelien.bellet@inria.fr), Claudio Gentile (cla.gentile@gmail.com), Marc Tommasi (marc.tommasi@inria.fr)

Keywords

Machine Learning, Distributed Algorithms, Graph-based Learning, Model Propagation, Online Learning.

Context

Imagine a social network dedicated to beer lovers. Users have their own personal taste and they want to evaluate the numerous new beers that are created every day. Of course, they cannot taste all of them: they have joined the social network in order to collaborate with similar users and to learn from their experience. They discuss with their friends in the network, and sometimes by reading the forums they discover new people who share the same tastes. Hence, in order to gather information about new beers, there is a clear trade-off between drinking new beers, talking to friends, and reading the forums. Moreover, for privacy reasons, one may want that any user profile and data (consisting in annotated lists of tasted beers) remains hidden from other users.

This internship is concerned with Machine Learning approaches to personalized, decentralized and online learning problems, like the one illustrated by the above example. More formally, each node in the social network represents a learning agent with its own data distribution. The goal of each agent i is to learn a target function $f_i : \mathbb{R}^d \to \mathbb{R}$ which assigns a label to a data point $\mathbf{x} \in \mathbb{R}^d$. At each step, the agent receives a datum $\mathbf{x}$ drawn from its personalized distribution, and the value of $f_i(\mathbf{x})$ is revealed after the agent's prediction. Agents may also query other users in order to build a collaboration network with similar users. The additional information supported by this network graph can be leveraged so as to improve prediction accuracy and/or the speed of the learning process.

Objectives

The goal of this internship is to study the class of prediction problems we can derive from the setting described above. The tentative workplan is as follows:

1. Review of the relevant literature (see in particular the references listed below).

2. Combine approaches for personalized and decentralized learning with online learning on graphs to devise new algorithms.
3. Investigate some specific questions, such as: (i) What kind of loss better reflects the properties of the considered setting? (ii) What if we assume the distributions of inputs $\mathbf{x}$ are the same for all agents? (iii) What if users are constrained to maintain a list of at most k neighbors? (iv) What if we set bounds on the amount and kind of information exchanged among agents?

Ideally, the expected results should be some learning algorithms along with a theoretical and/or experimental analysis in terms of convergence and statistical accuracy under some appropriate assumptions.

Skills

Basics in Machine Learning, algorithms and complexity, linear algebra and probability, programming skills (e.g., Python, Matlab).

References

[1] A. Agarwal, S. Gerber, and H. Daume. Learning multiple tasks using manifold regularization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 46–54. 2010.

[2] A. Argyriou, S. Clémençon, and R. Zhang. Learning the Graph of Relations Among Multiple Tasks. Research report, Oct. 2013.

[3] A. Bellet, R. Guerraoui, M. Taziki, and M. Tommasi. Fast and Differentially Private Algorithms for Decentralized Collaborative Machine Learning. Technical report, arXiv:1705.08435, 2017.

[4] N. Cesa-Bianchi, C. Gentile, Y. Mansour, and A. Minora. Delay and cooperation in non-stochastic bandits. In V. Feldman, A. Rakhlin, and O. Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 605–622, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR.

[5] N. Cesa-Bianchi, C. Gentile, and G. Zappella. A gang of bandits. In *Advances in Neural Information Processing Systems (NIPS)*, pages 737–745. 2013.

[6] A. Saha, P. Rai, H. D. III, and S. Venkatasubramanian. Online learning of multiple tasks and their relationships. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 643–651, 2011.

[7] P. Vanhaesebrouck, A. Bellet, and M. Tommasi. Decentralized Collaborative Learning of Personalized Models over Networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.