(Meta) Bayesian Optimization with Privacy

Project Proposal for Master/Semester Theses

Overview

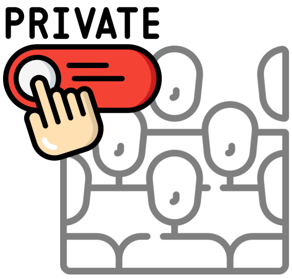

We consider Bayesian Optimization (BO) applications in which a server provides clients with a service by sequentially making recommendations, e.g., AutoML and recommender systems. In these applications, successfully recommending a choice/action, relies upon access to a context vector, which may include detailed data on user identity, preferences and interaction patterns. Access to such data further allows for harnessing the statistical patterns across the user pool, which in turn, improves the automated tailoring of the service to new users. However, exchanging such sensitive data with the BO server jeopardizes the clients' privacy. There is a clear trade-off between privacy and performance, which rises the questions

> Can we successfully perform BO without direct access to the user data?
> Can we privately transfer knowledge to enhance the experience of future users?

We are offering the following two projects, your contribution to which will primarily be of theoretical nature. This is accompanied by running experiments on academic data (synthetic, or easy benchmark datasets) as a proof of concept.

1 Private Meta-Learning for BO

Existing work on Representation Learning or Meta-Learning for BO [e.g., 6] are not private, since they either directly share sensitive data between users, or share the users' data with a server, e.g. a regression oracle. What is the cost of achieving privacy in this setting? How can the existing methods be privatized?

A starting point is F-LiBO; a federated algorithm for lifelong bandit optimization, but without privacy guarantees [9]. Your task would be to modify this algorithm so that it becomes differentially private (DP), and 1) provide the privacy guarantee 2) quantify the cost of privacy in the incurred regret. You would be incorporating DP tools from, e.g., [3] and [8]. This project contributes to the literature on federated meta-learning which focuses on private transfer of knowledge for supervised learning tasks [e.g., 1, 4].

2 Private Contextual BO

In contextual bandit optimization, at every step t, the server/agent observes a context x_t which contains sensitive information about the user, selects action a_t and receives a noisy reward of the form, $y_t = f(x_t, a_t) + \varepsilon_t$. To make this interaction private, instead we assume that at every step, the server only sees a distribution $p_t(x)$, which is a differentially-private copy of the user's true context x_t. What is the cost of this privacy? How can we reduce this cost via Meta-learning?

Your task would be to construct the DP copy of the context, and quantify the penalty of only having imperfect information, in terms of the incurred regret. Additionally, you may investigate if having access to a distribution/pool of users, can help reduce this penalty (i.e. the meta-learning setting). A starting point can be [7], which analyses BO with context distributions. This project also contributes to recent literature on federated learning which studies how agents can privately cooperate to solve a single task [2, 5, 10].

Contact

Before making an inquiry, please make sure that

– You are a Master's student,
– You have passed some machine learning, optimization and statistics courses,
– You are able to read (most of) the references and understand them down to the details.

If you are interested, please contact Parnian Kassraie (pkassraie@ethz.ch). If possible, **include a writing example in your email**. This can be the report of a course/semester project, your bachelor thesis, or a previous publication.

References

[1] Fei Chen, Mi Luo, Zhenhua Dong, Zhenguo Li, and Xiuqiang He. Federated meta-learning with fast convergence and efficient communication. *arXiv preprint*, 2018.

[2] Abhimanyu Dubey and AlexSandy Pentland. Differentially-private federated linear bandits. *NeurIPS*, 2020.

[3] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 2014.

[4] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *NeurIPS*, 2020.

[5] Ruiquan Huang, Weiqiang Wu, Jing Yang, and Cong Shen. Federated linear contextual bandits. *NeurIPS*, 2021.

[6] Parnian Kassraie, Jonas Rothfuss, and Andreas Krause. Meta-Learning Hypothesis Spaces for Sequential Decision-making. In *ICML*, 2022.

[7] Johannes Kirschner and Andreas Krause. Stochastic bandits with context distributions. *NeurIPS*, 2019.

[8] Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Ulfar Erlingsson. Scalable private learning with PATE. In *ICLR*, 2018.

[9] Felix Schur, Parnian Kassraie, Jonas Rothfuss, and Andreas Krause. Lifelong bandit optimization: No prior and no regret. *arXiv preprint*, 2022.

[10] Chengshuai Shi, Cong Shen, and Jing Yang. Federated multi-armed bandits with personalization. In *AISTATS*, 2021.