

Masters Internships, Inria Paris

Learning Theory and the Theory of Incentives

Michael Jordan

<https://people.eecs.berkeley.edu/~jordan/>

November 14, 2022

I wish to recruit students to work on topics at the intersection of machine learning and economic mechanism design. Specifically, I am interested in developing frameworks that combine statistical methods with ideas from contract theory, auction theory, and other areas in the theory of incentives. Consider a setting in which a principal wishes to accomplish some data-analysis task not by merely collecting data, but by interacting with one or more distributed agents who possess local data and who wish to obtain some benefit from supplying their data to the principal. The strategic agents may have partial knowledge of the quality of their local data, knowledge which they are not willing to simply reveal to the principal. Moreover, they may be tempted to "free-ride," letting other agents supply data and benefitting from the resulting data analysis without making any contribution. Such information asymmetries, as well as the free-riding problem, have been studied in contract theory in microeconomics. What is new here is that the principal is a statistical learning system, wishing to minimize generalization error with respect to a data-generating distribution. Thus we wish to design incentive structures that maximize data flow so as to minimize statistical risk. The project will focus on algorithms and theoretical analysis, but it will also aim to solve emerging real-world problems in domains such as federated learning. Applicants should have a strong background in areas such as statistics, game theory, optimization, and applied mathematics.

References

- Bates, S., Jordan, M. I., Sklar, M., and Soloff, J. A. (2022). Principal-agent hypothesis testing. [arXiv:2205.06812](https://arxiv.org/abs/2205.06812).
- Guo, W., Jordan, M. I., and Zampetakis, E. (2021). Robust learning of optimal auctions. In M. Ranzato, A. Beygelzimer, P. Liang, J. Wortman Vaughan, and Y. Dauphin (Eds), *Advances in Neural Information Processing Systems (NeurIPS)* 34.
- Haghtalab, N., Jordan, M. I., and Zhao, E. (2022). On-demand sampling: Learning optimally from multiple distributions. In A. Agarwal, A. Oh, D. Belgrave, and K. Cho (Eds), *Advances in Neural Information Processing Systems (NeurIPS)* 35.
- Jagadeesan, M., Wei, A., Wang, Y., Jordan, M. I., and Steinhardt, J. (to appear). Learning equilibria in matching markets with bandit feedback. *Journal of the ACM*.

Zhu, B., Bates, S., Yang, Z., Wang, Y., Jiao, J., and Jordan, M. I. (2022). The sample complexity of online contract design. arXiv:2211.05732.

Internship details

The interns will be located at Inria Paris, in the Sierra Project-team (<https://www.di.ens.fr/sierra/>), and co-advised by Francis Bach (<https://www.di.ens.fr/~fbach/>)

The proposed salary is 1329 Euros net per month.

The goal of the internships is to pursue as PhD students in Fall 2023 in the same team.

Application details

Please send a CV and the names of two or three references to jordan@cs.berkeley.edu and francis.bach@inria.fr.