Learning, Privacy and the Limits of Computation.

Christos Dimitrakakis
christos.dimitrakakis@gmail.com

April 12, 2017

1 Project summary.

The project deals with reinforcement learning[12] agents that act under privacy or computational constraints. The objective is to develop theory and algorithms to investigate the interaction between approximate inference and planning in constrained agents, and consequently bounds on reinforcement learning under these conditions. The goal is the development of provably efficient, and privacy-preserving algorithms for reinforcement learning.

On the theoretical side, this will be done by formalising computational limitations as approximate statistics[1,7,14] or differential privacy[6,8,9,13] constraints; two new areas in learning theory that are deeply connected to computational problems. Then we can obtain general bounds on problems with such constraints. We can also leverage approximate statistics to optimise the amount of computational effort used while planning, which will allow us to design efficient algorithms.

2 Candidate background

Essential attributes. The candidate is expected to hold a Masters degree (or equivalent) in Computer Science, Mathematics, or Statistics. Background knowledge in *probability and calculus*, mathematical maturity and demonstrable prior experience in research (through a master thesis or other project) is essential.

Desirable attributes. Courses or research experience in one or more of: machine learning, statistics, Bayesian inference, game theory, reinforcement learning and optimisation. International experience.

3 Knowledge acquired during the project.

The candidate's research is expected to draw upon the following mature and newer fields of research, for which textbooks are currently available: Markov decision processes[11], regret analysis[4] and concentration inequalities[3], reinforcement learning[2], statistical decision theory[5] and differential privacy[10].

4 Location, duration, duties and supervision.

The position is a 5-year fully-funded PhD position at Chalmers university of technology, Gothenbug, Sweden. 80% of the time is devoted to research and 20% to teaching. The main supervisor will be Dr. Christos Dimitrakakis, and the student will be embedded within his reinforcement learning group, currently including Aristide Tossou and Hannes Eriksson. During the course of the PhD the student may have the opportunity to visit the INRIA-Lille team SequeL, and the computer science department at Harvard.

5 How to apply.

For inquiries about this position, please send an email to `christos.dimitrakakis@gmail.com` with the subject "PhD thesis: Learning, Privacy and the Limits of Computation". For preliminary consideration, include:

- A CV detailing your experience.
- An example of research output (Master thesis, draft paper, ...).
- A motivation letter explaining why you are interested in doing a PhD in this specific area.

Details for formally submitting an application will be announced soon.

References

[1] Barnes, C. P., Filippi, S., Stumpf, M. P., and Thorne, T. (2012). Considerate approaches to constructing summary statistics for abc model selection. *Statistics and Computing*, pages 1–17.

[2] Bertsekas, D. P. and Tsitsiklis, J. N. (1996). *Neuro-Dynamic Programming*. Athena Scientific.

[3] Boucheron, S., Lugosi, G., and Bousquet, O. (2003). Concentration inequalities. In *Advanced Lectures in Machine Learning*, pages 208–240. Springer, London, UK.

[4] Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning and Games*. Cambridge University press, Cambridge, UK.

[5] DeGroot, M. H. (1970). *Optimal Statistical Decisions*. John Wiley & Sons.

[6] Dimitrakakis, C., Nelson, B., Zhang, Z., Mitrokotsa, A., and Rubinstein, B. I. P. (2017). Differential privacy for Bayesian inference through posterior sampling. *Journal of Machine Learning Research*, 18(11):1–39.

[7] Dimitrakakis, C. and Tziortziotis, N. (2013). ABC reinforcement learning. In *ICML 2013*, volume 28(3) of *JMLR W & CP*, pages 684–692. See also arXiv:1303.6977.

[8] Duchi, J. C., Jordan, M. I., and Wainwright, M. J. (2013). Local privacy and statistical minimax rates. Technical Report 1302.3203, arXiv.

[9] Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). *Calibrating Noise to Sensitivity in Private Data Analysis*, pages 265–284. Springer Berlin Heidelberg, Berlin, Heidelberg.

[10] Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.

[11] Puterman, M. L. (1994). *Markov Decision Processes : Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New Jersey, US.

[12] Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press.

[13] Tossou, A. C. Y. and Dimitrakakis, C. (2016). Algorithms for differentially private multi-armed bandits. In *13th International Conference on Artificial Intelligence (AAAI 2016)*.

[14] Wilkinson, R. D. (2013). Approximate bayesian computation (ABC) gives exact results under the assumption of model error. *Statistical applications in genetics and molecular biology*, 12(2):129–141.