

Publications of Tor Lattimore

- [1] Aurélien Garivier, Emilie Kaufmann, and Tor Lattimore. On explore-then-commit strategies. In *Proceedings of the 29th Conference on Neural Information Processing Systems (NIPS)*. 2016.
- [2] Sébastien Gerchinovitz and Tor Lattimore. Refined lower bounds for adversarial bandits. In *Proceedings of the 29th Conference on Neural Information Processing Systems (NIPS)*. 2016.
- [3] Ruitong Huang, Tor Lattimore, András Gyögy, and Csaba Szepesvári. Following the leader and fast rates in linear prediction: Curved constraint sets and other regularities. In *Proceedings of the 29th Conference on Neural Information Processing Systems (NIPS)*. 2016.
- [4] Finnian Lattimore, Tor Lattimore, and Mark Reid. Causal bandits: Learning good interventions via causal inference. In *Proceedings of the 29th Conference on Neural Information Processing Systems (NIPS)*. 2016.
- [5] Tor Lattimore. Regret analysis of the anytime optimally confident UCB algorithm. Technical report, 2016.
- [6] Tor Lattimore. Regret analysis of the finite-horizon Gittins index strategy for multi-armed bandits. In *Proceedings of Conference On Learning Theory (COLT)*. 2016.
- [7] Tor Lattimore and Csaba Szepesvári. The end of optimism? An asymptotic analysis of finite-armed linear bandits. Technical report, 2016.
- [8] Jan Leike, Tor Lattimore, Laurent Orseau, and Marcus Hutter. Thompson sampling is asymptotically optimal in general environments. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*. 2016.
- [9] Yifan Wu, Roshan Shariff, Tor Lattimore, and Csaba Szepesvári. Conservative bandits. In *Proceedings of the International Conference on Machine Learning (ICML)*. 2016.
- [10] Tor Lattimore. Optimally confident UCB : Improved regret for finite-armed bandits. Technical report, 2015.
- [11] Tor Lattimore. The pareto regret frontier for bandits. In *Proceedings of the 28th Conference on Neural Information Processing Systems (NIPS)*. 2015.
- [12] Tor Lattimore, Koby Crammer, and Csaba Szepesvári. Linear multi-resource allocation with semi-bandit feedback. In *Proceedings of the 28th Conference on Neural Information Processing Systems (NIPS)*. 2015.
- [13] Tom Everitt, Tor Lattimore, and Marcus Hutter. Free lunch for optimisation under the universal distribution. In *Proceedings of IEEE Congress on Evolutionary Computing (CEC)*. 2014.
- [14] Tor Lattimore, Koby Crammer, and Csaba Szepesvári. Optimal resource allocation with semi-bandit feedback. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence (UAI)*. 2014.
- [15] Tor Lattimore, András Gyögy, and Csaba Szepesvári. On learning the optimal waiting time. In *Proceedings of the 25th Conference on Algorithmic Learning Theory (ALT)*. 2014.
- [16] Tor Lattimore and Marcus Hutter. Asymptotics of continuous Bayes for non-i.i.d. sources. Technical report, 2014.
- [17] Tor Lattimore and Marcus Hutter. Bayesian reinforcement learning with exploration. In *Proceedings of the 25th Conference on Algorithmic Learning Theory (ALT)*. 2014.
- [18] Tor Lattimore and Marcus Hutter. General time consistent discounting. *Theoretical Computer Science*, 519(0):140 – 154, 2014.
- [19] Tor Lattimore and Marcus Hutter. On Martin-löf (non-)convergence of Solomonoff’s universal mixture. *Theoretical Computer Science*, 2014.

- [20] Tor Lattimore and Rémi Munos. Bounded regret for finite-armed structured bandits. In *Proceedings of the 27th Conference on Neural Information Processing Systems (NIPS)*. 2014.
- [21] Tor Lattimore and Marcus Hutter. No free lunch versus Occam's razor in supervised learning. In David Dowe, editor, *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence*, volume 7070 of *Lecture Notes in Computer Science*, pages 223–235. Springer Berlin Heidelberg, 2013.
- [22] Tor Lattimore and Marcus Hutter. On Martin-Löf convergence of Solomonoff's mixture. In T-H. Hubert Chan, LapChi Lau, and Luca Trevisan, editors, *Theory and Applications of Models of Computation*, volume 7876 of *Lecture Notes in Computer Science*, pages 212–223. Springer Berlin Heidelberg, 2013.
- [23] Tor Lattimore, Marcus Hutter, and Peter Sunehag. Concentration and confidence for discrete bayesian sequence predictors. In Sanjay Jain, Rémi Munos, Frank Stephan, and Thomas Zeugmann, editors, *Proceedings of the 24th International Conference on Algorithmic Learning Theory*, pages 324–338. Springer, 2013.
- [24] Tor Lattimore, Marcus Hutter, and Peter Sunehag. The sample-complexity of general reinforcement learning. In *Proceedings of the 30th International Conference on Machine Learning*. 2013.
- [25] Laurent Orseau, Tor Lattimore, and Marcus Hutter. Universal knowledge-seeking agents for stochastic environments. In Sanjay Jain, Rémi Munos, Frank Stephan, and Thomas Zeugmann, editors, *Proceedings of the 24th International Conference on Algorithmic Learning Theory*, volume 8139 of *Lecture Notes in Computer Science*, pages 158–172. Springer Berlin Heidelberg, 2013.
- [26] Tor Lattimore and Marcus Hutter. PAC bounds for discounted MDPs. In Nader Bshouty, Gilles Stoltz, Nicolas Vayatis, and Thomas Zeugmann, editors, *Proceedings of the 23th International Conference on Algorithmic Learning Theory*, volume 7568 of *Lecture Notes in Computer Science*, pages 320–334. Springer Berlin / Heidelberg, 2012.
- [27] Tor Lattimore and Marcus Hutter. Asymptotically optimal agents. In Jyrki Kivinen, Csaba Szepesvári, Esko Ukkonen, and Thomas Zeugmann, editors, *Proceedings of the 22nd International Conference on Algorithmic Learning Theory*, volume 6925 of *Lecture Notes in Computer Science*, pages 368–382. Springer Berlin / Heidelberg, 2011.
- [28] Tor Lattimore and Marcus Hutter. Time consistent discounting. In Jyrki Kivinen, Csaba Szepesvári, Esko Ukkonen, and Thomas Zeugmann, editors, *Proceedings of the 22nd International Conference on Algorithmic Learning Theory*, volume 6925 of *Lecture Notes in Computer Science*, pages 383–397. Springer Berlin / Heidelberg, 2011.
- [29] Tor Lattimore, Marcus Hutter, and Vaibhav Gavane. Universal prediction of selected bits. In Jyrki Kivinen, Csaba Szepesvári, Esko Ukkonen, and Thomas Zeugmann, editors, *Proceedings of the 22nd International Conference on Algorithmic Learning Theory*, volume 6925 of *Lecture Notes in Computer Science*, pages 262–276. Springer Berlin / Heidelberg, 2011.