

Curriculum Vitae

Name Tor Lattimore
Location London, UK
Email tor.lattimore@gmail.com

Education and Academic Positions

- 2017–present: Senior research scientist at DeepMind
- 2016–2017: Assistant professor at Indiana University
- 2014–2016: Postdoctoral fellow at the University of Alberta supervised by Csaba Szepesvári
- 2010–2013: PhD in Computer Science with Marcus Hutter at the Australian National University
- 2006–2009: Bachelor of Science (Advanced), 1st class honours in Mathematics at the Australian National University

Publications

- [1] Branislav Kveton, Csaba Szepesvári, Zheng Wen, Mohammad Ghavamzadeh, and Tor Lattimore. Garbage in, reward out: Bootstrapping exploration in multi-armed bandits. In *Proceedings of the 36th International Conference on Machine Learning*. 2019.
- [2] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press (draft), 2019.
- [3] Tor Lattimore and Csaba Szepesvári. Cleaning up the neighborhood: A full classification for adversarial partial monitoring. In Aurélien Garivier and Satyen Kale, editors, *Proceedings of the 30th International Conference on Algorithmic Learning Theory*, volume 98 of *Proceedings of Machine Learning Research*, pages 529–556. PMLR, 2019.
- [4] Tor Lattimore and Csaba Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. In *Conference on Learning Theory*. 2019.
- [5] Shuai Li, Tor Lattimore, and Csaba Szepesvári. Online learning to rank with features. In *Proceedings of the 36th International Conference on Machine Learning*. 2019.
- [6] Branislav Kveton, Chang Li, Tor Lattimore, Ilya Markov, Maarten de Rijke, Csaba Szepesvári, and Masrour Zoghi. Bubblerank: Safe online learning to rerank. *arXiv preprint*, 2018.
- [7] Tor Lattimore. Refining the confidence level for optimistic bandit strategies. *Journal of Machine Learning Research*, 19(20):1–32, 2018.
- [8] Tor Lattimore, Branislav Kveton, Shuai Li, and Csaba Szepesvári. Toprank: A practical algorithm for online stochastic ranking. In *Proceedings of the 31st Conference on Neural Information Processing Systems*. 2018.
- [9] Laurent Orseau, Levi Leis, Tor Lattimore, and Theophane Weber. Single-agent policy tree search with guarantees. In *Proceedings of the 31st Conference on Neural Information Processing Systems*. 2018.
- [10] Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying pac and regret: Uniform pac bounds for episodic reinforcement learning. In *Proceedings of the 30th Conference on Neural Information Processing Systems*. 2017.

- [11] Ruitong Huang, Tor Lattimore, András György, and Csaba Szepesvári. Following the leader and fast rates in online linear prediction: Curved constraint sets and other regularities. *Journal of Machine Learning Research*, 18(145):1–31, 2017.
- [12] Tor Lattimore. A scale free algorithm for stochastic bandits with bounded kurtosis. In *Proceedings of the 30th Conference on Neural Information Processing Systems*. 2017.
- [13] Tor Lattimore and Csaba Szepesvari. The End of Optimism? An Asymptotic Analysis of Finite-Armed Linear Bandits. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 728–737. PMLR, Fort Lauderdale, FL, USA, 20–22 Apr 2017.
- [14] Laurent Orseau, Tor Lattimore, and Shane Legg. Soft-bayes: Prod for mixtures of experts with log-loss. In *Proceedings of the 28th International Conference on Algorithmic Learning Theory*. 2017.
- [15] Joel Veness, Tor Lattimore, Avishkar Bhoopchand, Agnieszka Grabska-Barwinska, Christopher Mattern, and Peter Toth. Online learning with gated linear networks. Technical report, 2017.
- [16] 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.
- [17] Sébastien Gerchinovitz and Tor Lattimore. Refined lower bounds for adversarial bandits. In *Proceedings of the 29th Conference on Neural Information Processing Systems (NIPS)*. 2016.
- [18] 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.
- [19] 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.
- [20] Tor Lattimore. Regret analysis of the anytime optimally confident UCB algorithm. Technical report, 2016.
- [21] Tor Lattimore. Regret analysis of the finite-horizon Gittins index strategy for multi-armed bandits. In *Proceedings of Conference On Learning Theory (COLT)*. 2016.
- [22] 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.
- [23] Yifan Wu, Roshan Shariff, Tor Lattimore, and Csaba Szepesvári. Conservative bandits. In *Proceedings of the International Conference on Machine Learning (ICML)*. 2016.
- [24] Tor Lattimore. Optimally confident UCB : Improved regret for finite-armed bandits. Technical report, 2015.
- [25] Tor Lattimore. The pareto regret frontier for bandits. In *Proceedings of the 28th Conference on Neural Information Processing Systems (NIPS)*. 2015.
- [26] 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.
- [27] 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.
- [28] 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.

- [29] Tor Lattimore, András György, and Csaba Szepesvári. On learning the optimal waiting time. In *Proceedings of the 25th Conference on Algorithmic Learning Theory (ALT)*. 2014.
- [30] Tor Lattimore and Marcus Hutter. Asymptotics of continuous Bayes for non-i.i.d. sources. Technical report, 2014.
- [31] Tor Lattimore and Marcus Hutter. Bayesian reinforcement learning with exploration. In *Proceedings of the 25th Conference on Algorithmic Learning Theory (ALT)*. 2014.
- [32] Tor Lattimore and Marcus Hutter. General time consistent discounting. *Theoretical Computer Science*, 519(0):140 – 154, 2014.
- [33] Tor Lattimore and Marcus Hutter. On Martin-löf (non-)convergence of Solomonoff’s universal mixture. *Theoretical Computer Science*, 2014.
- [34] 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.
- [35] 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.
- [36] 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.
- [37] 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.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] 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.
- [43] 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.

Invited Talks

- SequeL machine learning group, Lille. 2018
- Data Learning and Inference (DALI). Auto-tuning stochastic bandits. 2016
- 4th Rutgers Applied Probability Conference. Presenting the Optimally Confident UCB. 2015
- Microsoft Research. Presenting Optimal Resource Allocation with Semi-Bandit Feedback. 2014

Prizes and Grants

- UAI 2016. Best student paper award [22] (Jan as student)
- ICML 2015. Reviewer Award
- UAI 2014. Runner-up best paper award [28]
- Google Australia Fellowship in Machine Learning (\$15 000), 2013
- Vice-Chancellor's HDR Travel Grant (\$1 000), 2012
- PASCAL Internal Visiting Programme (€1 000), 2012
- Boyapati Computer Science & Mathematical Honours Scholarship (\$5 000), 2009
- 3rd year Boyapati Computer Science & Mathematics Prize, 2008 (joint winner)
- 2st year Boyapati Computer Science & Mathematics Prize, 2007 (joint winner)
- 1st year Boyapati Computer Science & Mathematics Prize, 2006 (joint winner)
- 1st place, UNSW Programming Competition, 2005 (with Alex Osborne and Chris Pelling)
- 3rd place, UNSW Programming Competition, 2004 (with Alex Osborne and Chris Pelling)

Teaching

Lecturer

- B551: Introduction to AI, Indiana University 2016 (60 masters students)

Tutorials

- AAAI, 2018 (on bandits with Csaba Szepesvári)
- Oxford statistics group. Tutorial on bandit algorithms, 2018
- Reinforcement learning summer school (on bandit algorithms and exploration). Toronto, 2018
- European Workshop on Reinforcement Learning (bandits). Lille, 2018
- Winter School on Learning and Artificial intelligence (bandits). Trieste, 2018
- Gatsby Computational Neuroscience Unit (bandits), 2018
- Imperial College London machine learning tutorials (bandits), 2018

Teaching assistant

- MATH1014 (Introduction to Linear Algebra)
- MATH2306 (Partial Differential Equations and Applications)
- COMP6467 (Statistical Machine Learning)
- COMP4620/COMP8620 (Advanced Artificial Intelligence)

Reviewing

I have been on the program committee for the following conferences:

- Artificial Intelligence and Statistics Conference (AISTATS), 2017, 2018
- Algorithmic Learning Theory (ALT), 2017, 2019
- Neural Information Processing Systems (NIPS), 2016, 2017, 2018
- AAI Conference on Artificial Intelligence (AAAI), 2016
- International Joint Conference on Artificial Intelligence (IJCAI), 2015
- International Conference on Machine Learning (ICML), 2015
- Artificial General Intelligence (AGI), 2014

I have reviewed articles for a number of conferences and journals including:

- Conference on Learning Theory (COLT)
- Journal of Machine Learning Research (JMLR)
- Machine Learning (ML)
- Mathematics of Operations Research (MOR)
- Uncertainty in Artificial Intelligence (UAI)
- IEEE Transactions on Automatic Control (TAC)
- IEEE Transactions on Information Theory (TIT)
- Neural Information Processing Systems (NIPS)
- International Conference on Artificial Intelligence and Statistics (AISTATS)
- Theoretical Computer Science (TCS)
- European Conference on Machine Learning (ECML)
- Algorithmic Learning Theory (ALT)
- European Workshop on Reinforcement Learning (EWRL)
- Artificial General Intelligence (AGI)

Outreach

- “Lunch and Learn” presentation for year 11/12 students at the University of Alberta
- Co-supervising four high school students over two summers for the University of Alberta’s internship program

Other skills

- C/C++, Ruby, Python
- SQL
- Joint author of school timetabling application <https://prophetapp.com.au> (in use by most colleges in the Australian Capital Territory)

References

- Csaba Szepesvári (postdoctoral supervisor at the University of Alberta, szepesva@ualberta.ca)
- Marcus Hutter (PhD supervisor at the Australian National University, marcus.hutter@anu.edu.au)
- Rémi Munos (collaborator at Google DeepMind/INRIA Lille, remi.munos@inria.fr)
- András György (collaborator at University of Alberta, a.gyorgy@imperial.ac.uk)