

Ted Moskovitz

✉ ted@gatsby.ucl.ac.uk
📄 tedmoskovitz.github.io

Education

- 2019–present **Gatsby Computational Neuroscience Unit, University College London**,
PhD, Machine Learning and Theoretical Neuroscience
Advisors: Maneesh Sahani (Primary) and Matt Botvinick (Secondary).
- 2017–2019 **Columbia University**,
MS, Computer Science (Machine Learning Track; 3.98/4.00).
Advisor: Larry Abbott.
- 2013–2017 **Princeton University**,
AB with Honors, Neuroscience; minors in Computer Science and Linguistics
Advisor: Jonathan Pillow.

Experience

- 2022 **DeepMind**
Research Scientist Intern, Discovery Team; Host: Tom Zahavy
Stable optimization for constrained Markov decision processes. Contributed to [Optax](#).
- 2019 **Uber AI Labs**
Research Intern, Horizons Team; Hosts: Aditya Rawal and Kenneth Stanley
Optimization algorithms for deep neural networks.
- 2019 **Columbia University Center for Theoretical Neuroscience**
Research Assistant; Supervisors: Larry Abbott and Ken Miller
Biologically-plausible algorithms and architectures for deep learning.
- 2016–2018 **Princeton Neuroscience Institute**
Research Assistant; Supervisor: Jonathan Pillow
Deep learning and linear-nonlinear Poisson spiking models of neural encoding.

Select Publications & Original Work (See [Google Scholar](#) for Complete List)

- 2023 **T. Moskovitz**, A. Singh, D. Strouse, T. Sandholm, R. Salakhutdinov, A. Dragan, S. McAleer.
Confronting Reward Model Overoptimization with Constrained RLHF. *ICLR 2024* (To appear).
***Spotlight, Top 5% of Submissions**
- T. Moskovitz**, S. Hromadka, A. Touati, D. Borsa, M. Sahani. A State Representation for Diminishing Rewards. *NeurIPS 2023*.
- A. Singh, S. Chan, **T. Moskovitz**, E. Grant, A. Saxe, F. Hill. The Transient Nature of Emergent In-context Learning in Transformers. *NeurIPS 2023*.
- T. Moskovitz**, B. O'Donoghue, V. Veeriah, S. Flennerhag, S. Singh, T. Zahavy. ReLOAD: Reinforcement Learning with Optimistic Ascent-Descent for Last-Iterate Convergence in Constrained MDPs. *ICML 2023*.
- T. Moskovitz**, T. Kao, M. Sahani, M. Botvinick. Minimum Description Length Control. *ICLR 2023*.
- 2022 **T. Moskovitz**, K. Miller, M. Sahani, M. Botvinick. A Unified Theory of Dual Process Control. [Preprint](#)
- T. Moskovitz**, M. Arbel, J. Parker-Holder, A. Pacchiano. Towards an Understanding of Default Policies in Multitask Policy Optimization. *AISTATS 2022*. ***Best Paper Award Honorable Mention**

- T. Moskovitz**, S.R. Wilson, M. Sahani A First-Occupancy Representation for Reinforcement Learning. *ICLR 2022*.
- 2021 **T. Moskovitz**, J. Parker-Holder, A. Pacchiano, M. Arbel, M.I. Jordan. Tactical Optimism and Pessimism for Deep Reinforcement Learning. *NeurIPS 2021*.
- T. Moskovitz***, M. Arbel*, F. Huszar, A. Gretton. Efficient Wasserstein Natural Gradients for Reinforcement Learning. *ICLR 2021*. *Equal contribution.
- 2020 W. K. Li, **T. Moskovitz**, H. Kanagawa, M. Sahani. Amortised Learning by Wake-Sleep. *ICML 2020*.
- 2019 **T. Moskovitz**, R. Wang, J. Lan, S. Kapoor, J. Yosinski, A. Rawal. Learned First-Order Preconditioning. *Beyond First Order Methods in ML Workshop, NeurIPS 2019*.
- 2018 **T. Moskovitz**, A. Litwin-Kumar, L. Abbott. Feedback alignment in deep convolutional networks. [Preprint](#).

Invited Talks

- 2023 Google DeepMind London, Integration Team
Meta Fundamental AI Research (FAIR) Paris, Reinforcement Learning Group
- 2022 Microsoft Research NYC, Reinforcement Learning Group
COSYNE 2022 Workshop, “The what, how and when of learning”
- 2020 UCLA Mathematics Department, Seminar on Optimal Transport and Mean Field Games

Awards and Honors

- 2024 Spotlight Paper, ICLR 2024 (Top 5% of Conference Submissions)
- 2022 Best Paper Award Honorable Mention, AISTATS 2022
- 2021 Outstanding Reviewer Award, NeurIPS (Top 8% of Reviewers)
- 2016 Highly Commended Entrant in Linguistics: The Global Undergraduate Awards

Service

- 2019-present Reviewer, NeurIPS, ICLR, ICML, AISTATS, CCN, TMLR
- 2020 Mentor, In2Science UK — Mentored high school students from disadvantaged backgrounds interested in careers in STEM.

Teaching

- Fall 2020 Gatsby Unit, UCL, *Teaching Assistant*
Probabilistic and Unsupervised Learning, Systems and Theoretical Neuroscience
- Spring 2018 Columbia University Department of Computer Science, *Teaching Assistant*
COMS 4995 Deep Learning

Technical Skills

- Programming Python (advanced), C (intermediate), C++ (intermediate), CUDA (intermediate), experience with distributed training on GPUs & TPUs
- ML Jax, PyTorch, TensorFlow, Scikit-Learn

Other

- 2013–2017 Princeton Varsity Heavyweight Rowing
- 2014–2017 Volunteer, Special Olympics of New Jersey