

Yunhao Tang

CONTACT INFORMATION	Email: yt2541@columbia.edu Address: 345 Mudd, 500 West 120th St., New York, NY, 10027 Homepage: https://robintyh1.github.io/ Github: https://github.com/robintyh1
EDUCATION	Columbia University in the City of New York , New York, NY, USA Ph.D., Operations Research Sep. 2017 - Present Advisor: Shipra Agrawal Columbia University in the City of New York , New York, NY, USA M.S., Financial Engineering Sep. 2016 - June 2017 Fudan University , Shanghai, China B.S., Physics Sep. 2012 - June 2016 University of Toronto , Toronto, Canada Exchange Student, Faculty of Art and Science Sep. 2014 - Dec. 2014
WORK EXPERIENCE	Research Intern Aug. 2019 - Dec. 2019, May 2021 - Aug. 2021 (expected) DeepMind, Paris, France Supervisor: Rémi Munos <ul style="list-style-type: none">• Studied both model-free and model-based RL algorithms, which led to improvements over large-scale distributed RL agents and two publications at ICML 2020.
RESEARCH INTEREST	My research focuses on various aspects of reinforcement learning (RL), for example, <ul style="list-style-type: none">• Scalable model-free and model-based RL algorithms (see papers [1,2,3])• Evolutionary strategies and blackbox optimization for RL (see papers [1,2,3])• Probabilistic inference and RL (see papers [1,2,3])• Accelerating integer programming solving with RL (see papers [1])
SELECTED PUBLICATIONS	<ol style="list-style-type: none">1. Y. Tang*, M. Rowland, R. Munos, M. Valko, Taylor Expansion of Discount Factors, <i>International Conference on Machine Learning (ICML)</i>, 20212. T. Kozuno*, Y. Tang*, M. Rowland, R. Munos, W. Dabney, S. Kaputrowski, D. Abel, M. Valko, Revisiting Peng’s $Q(\lambda)$ for Modern Reinforcement Learning, <i>International Conference on Machine Learning (ICML)</i>, 20213. Y. Tang, A. Kucukelbir, Hindsight Expectation Maximization, <i>International Conference on Artificial Intelligence and Statistics (AISTATS)</i>, 20214. Y. Tang, Self-imitation Learning via Generalized Lower Bound Q-learning, <i>Neural Information Process Systems (NeurIPS)</i>, 20205. Y. Tang, M. Valko, R. Munos, Taylor Expansion Policy Optimization, <i>International Conference on Machine Learning (ICML)</i>, 2020

6. J.B. Grill^{*}, F. Altche^{*}, **Y. Tang^{*}**, T. Hubert, M. Valko, I. Antonoglou, R. Munos, Monte Carlo Tree Search as Regularized Policy Optimization Algorithms, *International Conference on Machine Learning (ICML)*, 2020
7. **Y. Tang**, S. Agrawal and Y. Faenza, Reinforcement Learning for Integer Programming: Learning to Cut, *International Conference on Machine Learning (ICML)*, 2020

SELECTED
PREPRINTS

1. **Y. Tang**, M. Rowland, R. Munos, M. Valko, Marginalized Operators for Off-Policy Reinforcement Learning, *in submission*
2. **Y. Tang**, K. Chormanski, Online Hyper-parameter Tuning for Off-policy Learning via Evolutionary Strategies, *arxiv*
3. **Y. Tang**, M. Yin, M. Zhou, Augment-Reinforce-Merge Policy Gradient for Binary Stochastic Policy, *arxiv*

CONFERENCE
PUBLICATIONS

1. **Y. Tang**, Guiding Evolutionary Strategies with Off-policy Actor-Critic, *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2021
2. K. Choromanski^{*}, J. Park-Holder^{*}, A. Pacchiano^{*}, **Y. Tang^{*}**, A. Choromanska, M. Jordan, Learning to Score for Guided Policy Optimization, *International Conference on Machine Learning (ICML)*, 2020
3. Y. Yue, **Y. Tang**, M. Yin and M. Zhou, Discrete Action On-Policy Learning with Action-Value Critic, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020
4. **Y. Tang**, K. Choromanski and A. Kucukelbir, Variance Reduction for Evolutionary Strategies via Structured Control Variate, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020
5. K. Choromanski^{*}, J. Park-Holder^{*}, A. Pacchiano^{*}, **Y. Tang^{*}**, Practical Nonisotropic Monte Carlo Sampling in High Dimensions via Determinantal Point Processes, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020
6. **Y. Tang**, S. Agrawal, Discretizing Continuous Action Space for On Policy Optimization, *Association for the Advancement of Artificial Intelligence (AAAI)*, 2020
7. X. Song, W. Gao, Y. Yang, K. Choromanski, A. Pacchiano, **Y. Tang**, ES-MAML: Simple Hessian-Free Meta Learning *International Conference on Learning Representations (ICLR)*, 2020
8. K. Choromanski^{*}, J. Park-Holder^{*}, A. Pacchiano^{*}, **Y. Tang^{*}** Adaptive Sample-Efficient Blackbox Optimization via ES-active Subspaces, *Neural Information Processing Systems (NeurIPS)* 2019.
9. K. Choromanski, J. Park-Holder and A. Pacchiano, **Y. Yang**, D. Jain, Y. Yang, A. Iscen, J. Hsu, V. Sindhwani, Provably Robust Blackbox Optimization for Reinforcement Learning, *Conference on Robot Learning (CORL)*, 2019
10. K. Choromanski^{*}, A. Pacchiano^{*}, J. Pennington^{*}, **Y. Tang^{*}**, KAMA-NNs: low-dimensional rotation based neural networks, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019
11. J. Hron^{*}, M. Rowland^{*}, **Y. Tang^{*}**, K. Choromanski, T. Sarlos, A. Weller, Orthogonal Estimation of Wasserstein Distances, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019
12. **Y. Tang**, S. Agrawal, *Exploration by Distributional Reinforcement Learning*, *International Joint Conference on Artificial Intelligence (IJCAI)* 2018

WORKSHOP PUBLICATIONS AND PREPRINTS	<ol style="list-style-type: none"> 1. X. Song, K. Choromanski, J. Parker-Holder, Y. Tang, W. Gao, A. Pacchiano, T. Sarlos, D. Jain, Y. Yang, Reinforcement Learning with Chromatic Networks, <i>Workshop on Neural Architectural Search, International Conference on Learning Representations (ICLR), 2020</i> 2. Y. Tang, S. Agrawal, Boosting Trust Region Policy Optimization with Normalizing Flows Policy, <i>Deep Reinforcement Learning Workshop, Neural Information Processing Systems (NIPS) 2018</i>. 3. Y. Tang, X. Cao, Variational Auto-encoding Contexts for Control, <i>NIPS Workshop on Inference to Control, 2018</i> 4. Y. Tang, A. Kucukelbir, Variational Deep Q Network, <i>Workshop on Bayesian Deep Learning, Neural Information Processing Systems (NIPS) 2017</i>. 5. Y. Tang, S. Agrawal, Implicit Policy for Reinforcement Learning, <i>arXiv</i>
INVITED TALKS	<ol style="list-style-type: none"> 1. INFORMS, sessions on sequential decision making Virtual, Nov 2021 2. CORS (Canadian Operations Research Society), sessions on ML and optimization Virtual, May 2021 3. MSR (Microsoft Research) RL Day Virtual, Feb 2021 4. IPAM (Institute for Pure and Applied Mathematics) workshop on Deep Learning and Combinatorial Optimization Virtual, Feb 2021 5. INFORMS, sessions on ML and optimization Virtual, Nov 2020
TEACHING ASSISTANT EXPERIENCE	<p>IEOR 8100 Reinforcement Learning Fall 2020, Spring 2019, Spring 2018 Department of IEO, Columbia University, New York, USA</p> <p>IEOR 4525 Machine Learning for OR and FE Spring 2020, Fall2018 Department of IEO, Columbia University, New York, USA</p> <p>COMS6998 Probabilistic Programming Fall 2018 Department of Computer Science, Columbia University, New York, USA</p> <p>IEOR6711 Stochastic Modeling I Fall 2017 Department of IEO, Columbia University, New York, USA</p>
ACADEMIC SERVICE	Reviewer: ICML 2020'21, NeurIPS 2018'19'20, AISTATS 2019'20
HONORS AND AWARDS	<p>Class of '88 Fellowship, Department of IEO, Columbia University Dec 2020</p> <p>Data Science Institute (DSI) Outstanding Teaching Assistant, Data Science Institute, Columbia University June 2019</p> <p>Robert Gartland Scholarship for Academic Excellence (Top 2%), Department of IEO, Columbia University April 2017</p> <p>First-Class Scholarship, Fudan University June 2016, Oct. 2014</p>