

Yash Chandak

+1-413-345-3192
✉ ychandak@stanford.edu
📁 yashchandak.github.io/

Research Interests

My interests are at the intersection of foundation models, decision making, and formal reasoning. I have also worked on strategic data collection under unobserved confounding, reward design, and different aspects of continual and reinforcement learning.

Education

- 2022–Current **Postdoctoral Scholar**, *Stanford University, USA*
Advised by Prof. Emma Brunskill
- 2017–2022 **PhD, Computer Science**, *University of Massachusetts Amherst, USA*
Advised by Prof. Philip S. Thomas, Autonomous Learning Lab.
- 2017–2020 **MS in Computer Science**, *University of Massachusetts Amherst, USA*
Advised by Prof. Philip S. Thomas, Autonomous Learning Lab.
- 2013–2017 **B.Tech, Computer Science & Engineering**, *VIT University, Chennai, India*
Advised by Prof. Nithya Darisini P.S.

Awards & Honors

- 2022 Google Ph.D. Fellowship. (Declined)
- 2022 *RLDM*: Best Paper.
- 2021 Outstanding Reviewer for NeurIPS (8%) and Top Reviewer for ICML (10%).
- 2020 *AAAI*: Outstanding student paper honorable mention.
- 2020 Graduated with Distinction, MS in Computer Science.
- 2013 Played at the national level for the Basketball Federation of India (BFI).

Internships

- May–Sept '22 **Google DeepMind, London**, under *Diana Borsa, Mohammad Azar, and Prof. Remi Munos*
- Jun–Aug '19 **Adobe Research, San Jose**, under *Georgios Theodorou and Prof. Sridhar Mahadevan*
- Jun–Aug '18 **Adobe Research, San Jose**, under *Georgios Theodorou*
- Jan–Jun '17 **Indian Institute of Technology Madras, India**, under *Prof. Balaraman Ravindran*
- Jun–Jul '16 **University of Technology Troyes, France**, under *Prof. Babiga Birregah*
- Feb–Mar '16 **Defence Research & Development Organisation (IRDE, DRDO)**, under *Sh. Jai Prakash Singh*
- Feb–May '15 **The Aspiring Researcher Challenge**, under *Prof. James Davis, UCSC and Rajan Vaish, Stanford*

Teaching

- Fall '20 Teaching Assistant for CS 687: Reinforcement Learning, UMass.
- Fall '19 Guest Lecture *On Designing Reward Signals*, CS 687: Reinforcement Learning, UMass.
- Spring '18 Teaching Assistant for CS 240: Reasoning Under Uncertainty, UMass.
- Fall '17 Teaching Assistant for CS 383: Artificial Intelligence, UMass.

Service

- Reviewer Journals: JMLR (2021-24), MLJ (2022-23), AIJ(2023).
Conferences: ICML (2019-23), NeurIPS (2019-2023), ICLR (2022), TMLR (2022-23), CoLLAs (2022-23).

Outreach and Mentorship

- Fall '20 - PhD mentor, in collaboration with Microsoft New England, for CS696DS.
Spring '21 Guided 3 Master's students for the project on optimizing interventions in shared autonomy.
- Spring '20 PhD mentor, in collaboration with Microsoft New England, for CS696DS.
Guided 4 Master's students towards developing human-agent cooperative RL systems.
- Spring '19 PhD mentor, in collaboration with Microsoft Research Montreal, for CS696DS.
Guided 4 Master's students towards developing RL algorithms for text-based games.

Invited Talks

- 2023 Stanford Causal Inference Seminar
Topic: Adaptive Instrument Design for Indirect Experiments.
- 2022 Safe Reinforcement Learning workshop at IJCAI.
Topic: Going Beyond Expected Cost/Return Metrics and Stationarity Assumptions.
- 2022 University of New Hampshire
Topic: Reinforcement Learning for Non-stationary Environments.

Journal Publications

- Under review **Data-Efficient Policy Evaluation Through Behavior Policy Search**
Josiah P. Hanna, **Yash Chandak**, Philip S. Thomas, Martha White, Peter Stone, Scott Niekum
- Frontiers 2022 **Scaling Graph Propagation Kernels for Predictive Learning**
Priyesh Vijayan, **Yash Chandak**, Mitesh M Khapra, Srinivasan Parthasarathy, Balaraman Ravindran
Frontiers in Big Data, section Data Mining and Management.
- UMass 2022 **Reinforcement Learning for Non-Stationary Problems**
Yash Chandak
Ph.D. Thesis, University of Massachusetts Amherst

Preprints

- In Preparation **Neural Theorem Proving with Information Directed Search**
Yash Chandak, Jonathan Lee, Emma Brunskill
- Under Review **Adaptive Instrument Design for Indirect Experiments**
Yash Chandak, Shiv Shankar, Vasilis Syrgkanis, Emma Brunskill
- Under Review **OPERA: Offline Policy Evaluation with Re-weighted Aggregates of Multiple Estimators**
Allen Nie, **Yash Chandak**, Christina J. Yuan, Anirudhan Badrinath, Yannis Flet-Berliac, Emma Brunskill
- Under Review **A/B testing under Interference with Partial Network Information**
Shiv Shankar, Ritwik Sinha, **Yash Chandak**, Saayan Mitra, Madalina Fiterau

Conference Publications

- NeurIPS 2023 **Behavior Alignment via Reward Function Optimization**
Dhawal Gupta*, **Yash Chandak***, Scott M. Jordan, Bruno Castro da Silva, Philip S. Thomas
Neural Information Processing Systems.
- NeurIPS 2023 **In-Context Decision-Making from Supervised Pretraining**
Jonathan Lee, Annie Xie, Aldo Pacchiano, **Yash Chandak**, Chelsea Finn, Ofir Nachum, Emma Brunskill
Neural Information Processing Systems.

- ICML **Representations and Exploration for Deep Reinforcement Learning using Singular Value Decomposition**
2023 **Yash Chandak**, Shantanu Thakoor, Zhaohan Daniel Guo, Yunhao Tang, Remi Munos, Will Dabney, Diana Borsa
International Conference on Machine Learning
- ICML **Understanding Self-Predictive Learning for Reinforcement Learning**
2023 Yunhao Tang, Zhaohan Daniel Guo, Pierre Harvey Richemond, Bernardo Avila Pires, **Yash Chandak**, Remi Munos, Mark Rowland, Mohammad Gheshlaghi Azar, Charline Le Lan, Clare Lyle, Andras Gyorgy, Shantanu Thakoor, Will Dabney, Bilal Piot, Daniele Calandriello, Michal Valko
International Conference on Machine Learning
- AISTATS **SSOPE: Asymptotically Unbiased Off-Policy Policy Evaluation when Reusing Old Data in Nonstationary Environments**
2023 Vincent Liu, **Yash Chandak**, Philip S. Thomas, Martha White
International Conference on Artificial Intelligence and Statistics
- NeurIPS **(Off) Policy Evaluation for Action-Dependent Non-stationary Environments**
2022 **Yash Chandak**, Shiv Shankar, Nate Bastian, Bruno Castro da Silva, Emma Brunskill, Philip S. Thomas
Neural Information Processing Systems.
- NeurIPS **Factored DRO: Factored Distributionally Robust Policies for Contextual Bandits**
2022 Tong Mu, **Yash Chandak**, Tatsunori Hashimoto, Emma Brunskill
Neural Information Processing Systems.
- AAAI **On Optimizing Interventions in Shared Autonomy**
2022 Weihao Tan*, David Koleczek*, Siddhant Pradhan*, Nicholas Perello, Vivek Chettiar, Nan Ma, Aaslesha Rajaram, Vishal Rohra, Soundararajan Srinivasan, H M Sajjad Hossain[†], **Yash Chandak**[†]
Association for the Advancement of Artificial Intelligence. *Equal contribution, [†]Equal advising.
- NeurIPS **Universal Off-Policy Evaluation**
2021 **Yash Chandak**, Scott Niekum, Bruno Castro da Silva, Erik Learned-Miller, Emma Brunskill, Philip S. Thomas
Neural Information Processing Systems.
- NeurIPS **SOPE: Spectrum of Off-Policy Estimators**
2021 Christina Yuan, **Yash Chandak**, Stephen Giguere, Philip S. Thomas, Scott Niekum
Neural Information Processing Systems.
- ICML **High Confidence Generalization for Reinforcement Learning**
2021 James Kostas, **Yash Chandak**, Scott Jordan, Philip S. Thomas
International Conference on Machine Learning.
- AAAI **High-Confidence Off-Policy (or Counterfactual) Variance Estimation**
2021 **Yash Chandak**, Shiv Shankar, Philip S. Thomas
Association for the Advancement of Artificial Intelligence.
- NeurIPS **Towards Safe Policy Improvement for Non-Stationary MDPs**
2020 **Yash Chandak**, Scott Jordan, Georgios Theocharous, Martha White, Philip S. Thomas
Neural Information Processing Systems.
- ICML **Optimizing for the Future in Non-Stationary MDPs**
2020 **Yash Chandak**, Georgios Theocharous, Shiv Shankar, Martha White, Sridhar Mahadevan, Philip S. Thomas
International Conference on Machine Learning.
- ICML **Evaluating the Performance of Reinforcement Learning Algorithms**
2020 Scott Jordan, **Yash Chandak**, Daniel Cohen, Mengxue Zhang, Philip S. Thomas
International Conference on Machine Learning.

AAAI **Lifelong Learning with a Changing Action Set**
2020 **Yash Chandak**, Georgios Theodorou, Chris Nota, Philip S. Thomas
Association for the Advancement of Artificial Intelligence.

AAAI **Reinforcement Learning When All Actions are Not Always Available**
2020 **Yash Chandak**, Georgios Theodorou, Blossom Metevier, Philip S. Thomas
Association for the Advancement of Artificial Intelligence.

ICML **Learning Action Representations for Reinforcement Learning**
2019 **Yash Chandak**, Georgios Theodorou, James Kostas, Scott Jordan, Philip S. Thomas
International Conference on Machine Learning.

Workshop Publications and Others

OPT **Optimization using Parallel Gradient Evaluations on Multiple Parameters**
(NeurIPS) **Yash Chandak**, Shiv Shankar, Venkata Gandikota, Philip S. Thomas, Arya Mazumdar
2022 OPT workshop at Neural Information Processing Systems.

RLDM **A Generalized Learning Rule for Asynchronous Coagent Networks**
2022 James Kostas, Scott Jordan, **Yash Chandak**, Georgios Theodorou, Dhawal Gupta, Philip S. Thomas
Multidisciplinary Conference on Reinforcement Learning and Decision Making.

SafeRL **Behavior Policy Search for Risk Estimators in Reinforcement Learning**
(NeurIPS) Elita Lobo, **Yash Chandak**, Dharmashankar Subramanian, Josiah Hanna, Marek Petrik
2021 Workshop on Safe and Robust Control of Uncertain Systems at NeurIPS.

HumanAI **Intervention Aware Shared Autonomy**
(ICML) Weihao Tan*, David Koleczek*, Siddhant Pradhan*, Nicholas Perello, Vivek Chettiar, Nan Ma,
2021 Aaslesha Rajaram, Vishal Rohra, Soundararajan Srinivasan, H M Sajjad Hossain[†], **Yash Chandak**[†]
Human-AI Collaboration in Sequential Decision-Making, ICML. *Equal contribution, [†]Equal advising.

arXiv **Reinforcement Learning for Strategic Recommendations**
2020 Georgios Theodorou, **Yash Chandak**, Philip S. Thomas, Frits de Nijs
Arxiv 2009.07346.

RLDM **Improving Generalization over Large Action Sets**
2019 **Yash Chandak**, Georgios Theodorou, James Kostas, Scott Jordan, Philip S. Thomas
Multidisciplinary Conference on Reinforcement Learning and Decision Making.

RLDM **Evaluating RL Algorithms Using Cumulative Distributions of Performance**
2019 Scott Jordan, **Yash Chandak**, Mengxue Zhang, Daniel Cohen, Philip S. Thomas
Multidisciplinary Conference on Reinforcement Learning and Decision Making.

CL **Reinforcement Learning with a Dynamic Action Set**
(NeurIPS) **Yash Chandak**, Georgios Theodorou, James Kostas, Philip Thomas
2018 Continual Learning workshop at Neural Information Processing Systems.

arXiv **Classical Policy Gradient: Preserving Bellman's Principle of Optimality**
2018 Philip S. Thomas, Scott Jordan, **Yash Chandak**, Chris Nota, James Kostas
Arxiv 1906.03063.

STAR-AI **HOPF: Higher Order Propagation Framework for Deep Collective Classification**
(IJCAI) Priyesh Vijayan, **Yash Chandak**, Mitesh M Khapra, Srinivasan Parthasarathy, Balaraman Ravindran
2018 International Workshop on Statistical Relational AI, IJCAI.

MLG **Fusion Graph Convolutional Networks**
(KDD) Priyesh Vijayan, **Yash Chandak**, Mitesh M Khapra, Srinivasan Parthasarathy, Balaraman Ravindran
2018 International Workshop on Machine Learning with Graphs, KDD.

HCOMP **On Optimizing Human-Machine Task Assignments**
(AAAI) Andreas Veit, Michael Wilber, Rajan Vaish, Serge Belongie, James Davis and **others**
2015 AAAI Conference on Human Computation and Crowdsourcing, work-in-progress.

Patents

- 2023 **Reinforcement learning with a stochastic action set**
Yash Chandak, Georgios Theocharous
US Patent 11615293.
- 2022 **Forecasting and learning accurate and efficient target policy parameters for dynamic processes in non-stationary environments**
Yash Chandak, Georgios Theocharous, Sridhar Mahadevan
US Patent 17072868.
- 2022 **Lifelong learning with a changing action set**
Yash Chandak, Georgios Theocharous
US Patent 11501207.
- 2020 **Generating and providing proposed digital actions in high-dimensional action spaces using reinforcement learning models**
Yash Chandak, Georgios Theocharous
US Patent 16261092.