

Daniel R. Jiang

+1 765-409-9472
danielrjiang@gmail.com
<https://danielrjiang.github.io/>

Professional Experience

Meta (formerly Facebook), New York, NY and Menlo Park, CA.

Research Scientist, Applied Reinforcement Learning, June 2023 – Present, New York, NY

Co-creator of **Pearl**, an open source library for reinforcement learning (RL).

Research Scientist, Core Data Science, Jan. 2019 – June 2023, Menlo Park, CA

Co-creator of **BoTorch**, an open source library for Bayesian optimization (BO), built on PyTorch.

University of Pittsburgh, Pittsburgh, PA.

Adjunct Professor, Industrial Engineering, Sept. 2021 – Present.

Assistant Professor (tenure-track), Industrial Engineering, Sept. 2016 – Sept. 2021.

On leave-of-absence from Jan. 2019 – Sept. 2021.

Education

Princeton University

Ph.D. Operations Research & Financial Engineering, 2016.

Thesis: "Risk-Neutral and Risk-Averse Approximate Dynamic Programming Methods."

Advisor: Warren B. Powell

M.A. Operations Research & Financial Engineering, 2013.

Purdue University

B.S. Electrical and Computer Engineering, *With Highest Distinction*, 2011.

B.S. Mathematics, *With Highest Distinction*, 2011.

Research Overview

Keywords: Reinforcement learning, approximate dynamic programming, Bayesian optimization, adaptive large-scale online experimentation, operations management, energy systems, sharing economy

Summary: I am interested in methods for decision-making under uncertainty, drawing inspiration from both operations research (OR) and machine learning (ML). Two of my recent directions are:

1. **Leveraging structure in reinforcement learning.** How can problem structure be used as a type of "side information" to more effectively solve sequential decision problems? I work on reinforcement learning (RL) and approximate dynamic programming (ADP) algorithms that can exploit to problem structure, including temporal structure (e.g., fast and slow information dynamics), reward/value structure (e.g., convexity and monotonicity), or decomposability (e.g., weakly coupled structure).
2. **Acquisition of scarce and costly information.** Data can be rare and expensive to obtain. How can we design efficient methods to decide how (where and when) to collect such data? I've worked on methods based on Bayesian optimization (BO) for this purpose, with applications to adaptive online experimentation, tuning large-scale machine learning models, and fast simulations for augmented-reality hardware design. I'm particularly interested in using multi-step formulations to reason about the sequential nature of information collection.

Publications

Working Papers ([§]direct mentoring, *joint, †alphabetical)

1. Ethan Che, Daniel R. Jiang, Hongseok Namkoong, Jimmy Wang. Planning Batch Adaptive Experiments with Model-Predictive Control. Preprint, 2024.
2. Guojun Xiong, Shufan Wang, Daniel R. Jiang, Jian Li. Personalized Federated Reinforcement Learning with Shared Representations. Preprint, 2024.
3. Zheqing Zhu, Rodrigo de Salvo Braz, Jalaj Bhandari, Daniel R. Jiang, Yi Wan, Yonathan Efroni, Liyuan Wang, Ruiyang Xu, Hongbo Guo, Alex Nikulkov, Dmytro Korenkevych, Frank Cheng, Urun Dogan, Zheng Wu, Wanqiao Xu. Pearl: A Production-Ready Reinforcement Learning Agent. Preprint, 2024.
4. Yijia Wang[§], Daniel R. Jiang. Faster Approximate Dynamic Programming by Freezing Slow States. Major revision at **Management Science**, 2023.
5. Yijia Wang[§], Daniel R. Jiang. Structured Actor-Critic for Managing Public Health Points of Dispensing. Under revision for resubmission, 2022.

Peer-Reviewed Publications ([§]direct mentoring, *joint, †alphabetical)

6. Ibrahim El-Shar[§], Daniel R. Jiang. Weakly Coupled Deep Q-Networks. *Advances in Neural Information Processing Systems*, **NeurIPS 2023**.
7. Yijia Wang[§], Matthias Poloczek, Daniel R. Jiang. Dynamic Subgoal-based Exploration via Bayesian Optimization. **Transactions on Machine Learning Research**, 2023.
Preliminary version: *Task-Agnostic Reinforcement Learning Workshop* at **ICLR 2019**
8. Kevin Kuo, Pratiksha Thaker, Mikhail Khodak, John Nguyen, Daniel R. Jiang, Ameet Talwalkar, Virginia Smith. On Noisy Evaluation in Federated Hyperparameter Tuning. **MLSys 2023**.
9. Saif Benjafaar[†], Daniel R. Jiang[†], Xiang Li[†], Xiaobo Li[†]. Dynamic Inventory Repositioning in On-Demand Rental Networks. **Management Science**, 68(11), pp. 7793-8514, 2022.
10. Han Wu*, Sarah Tan*, Weiwei Li, Mia Garrard, Adam Obeng, Drew Dimmery, Shaun Singh, Hanson Wang, Daniel R. Jiang, Eytan Bakshy. Interpretable Personalized Experimentation. **KDD 2022**.
11. Raul Astudillo[§], Daniel R. Jiang, Max Balandat, Eytan Bakshy, Peter I. Frazier. Multi-Step Budgeted Bayesian Optimization with Unknown Costs. *Advances in Neural Information Processing Systems*, **NeurIPS 2021**.
12. Shali Jiang*, Daniel R. Jiang*, Max Balandat*, Brian Karrer, Jake Gardner, Roman Garnett. Efficient Nonmyopic Bayesian Optimization via One-Shot Multi-Step Trees. *Advances in Neural Information Processing Systems*, **NeurIPS 2020**.
13. Max Balandat, Brian Karrer, Daniel R. Jiang, Sam Daulton, Ben Letham, Andrew G. Wilson, Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. *Advances in Neural Information Processing Systems*, **NeurIPS 2020**.
14. Ibrahim El-Shar[§], Daniel R. Jiang. Lookahead-Bounded Q-Learning. *International Conference on Machine Learning*, **ICML 2020**.
15. Daniel R. Jiang, Lina Al-Kanj, Warren B. Powell. Optimistic Monte Carlo Tree Search with Sampled Information Relaxation Dual Bounds. **Operations Research**, 68(6), pp. 1678-1697, 2020.

16. Daniel R. Jiang, Emmanuel Ekwedike, Han Liu. [Feedback-Based Tree Search for Reinforcement Learning](#). *International Conference on Machine Learning, ICML 2018* (selected for *long talk*; 8.6% acceptance).
17. Daniel R. Jiang, Warren B. Powell. [Risk-Averse Approximate Dynamic Programming with Quantile-Based Risk Measures](#). **Mathematics of Operations Research**, 43(2), pp. 554-579, 2018.
Math of OR Editor's Pick for [INFORMS ICYMI](#) in December 2017.
18. Andrew L. Johnson, Daniel R. Jiang. [Shape Constraints in Economics and Operations Research](#). **Statistical Science**, 33(4), pp. 527-546, 2018.
19. Daniel R. Jiang, Warren B. Powell. [An Approximate Dynamic Programming Algorithm for Monotone Value Functions](#). **Operations Research**, 63(6), pp. 1489-1511, 2015.
20. Daniel R. Jiang, Warren B. Powell. [Optimal Hour-Ahead Bidding in the Real-Time Electricity Market with Battery Storage using Approximate Dynamic Programming](#). **INFORMS Journal on Computing**, 27(3), pp. 525-543, 2015.

Workshops and Long Abstracts

21. Jalaj Bhandari, Yonathan Efroni, Mohammad Ghavamzadeh, Daniel Jiang, Aldo Pacchiano, Yi Wan, Kelly Zhang, Angela Zhou. [Deployable RL: From Research to Practice](#). *Workshop Proposal at RLC 2024*.
22. Chu Wang, Yingfei Wang, Haipeng Luo, Daniel Jiang, Jinghai He, Zeyu Zheng. [2nd Workshop on Multi-Armed Bandits and Reinforcement Learning: Advancing Decision Making in E-Commerce and Beyond](#). *Workshop Proposal at KDD 2023*.
23. Udit Gupta, Daniel R. Jiang, Max Balandat, Carole-Jean Wu. [Towards Green, Accurate, and Efficient AI Models Through Multi-Objective Optimization](#). *Proposal Track: Tackling Climate Change with Machine Learning: Global Perspectives and Local Challenge Workshop at ICLR 2023*.
24. Daniel R. Jiang, Jelena Markovic, Shubhankar Ray, Adam Obeng, Max Balandat, Eytan Bakshy. [Adaptive Policies for Staggered Rollout of Large-scale Online Experiments](#). *Conference on Digital Experimentation @ MIT, CODE 2021*.
25. Han Wu, Sarah Tan, Weiwei Li, Mia Garrard, Hanson Wang, Daniel R. Jiang, Adam Obeng, Eytan Bakshy. [Learning Efficient Interpretable Policies on Experimental Data](#). *Conference on Digital Experimentation @ MIT, CODE 2021*.
26. Daniel R. Jiang[†], Haipeng Luo[†], Chu Wang[†], Yingfei Wang[†]. [Multi-Armed Bandits and Reinforcement Learning: Advancing Decision Making in E-Commerce and Beyond](#). *Workshop Proposal at KDD 2021*.
27. Yijia Wang[§], Brian Bell, Matthias Poloczek, Daniel R. Jiang. ["Exploration via Sample-Efficient Subgoal Design."](#) *Task-Agnostic Reinforcement Learning Workshop at ICLR 2019*.
28. Yijia Wang, Daniel R. Jiang. [Hierarchical Reinforcement Learning for Naloxone Procurement and Distribution](#). *Women in Machine Learning Workshop, WiML 2017*.
29. Daniel R. Jiang, Thuy V. Pham, Warren B. Powell, Daniel F. Salas, Warren R. Scott. [A Comparison of Approximate Dynamic Programming Techniques on Benchmark Energy Storage Problems: Does Anything Work?](#) *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning, ADPRL 2014*.

Recent Awards & Honors

1. First Prize Winner (1 out of 300 submissions), Kaggle + NFL Big Data Bowl, February 2024.
Uncovering Tackle Opportunities and Missed Opportunities

Authors: Matthew Chang, Katherine Dai, Daniel Jiang, and Harvey Cheng

\$25,000 prize and invited to present to NFL personnel and media at the 2024 NFL Scouting Combine.

2. *Transportation Science Meritorious Service Award, 2022.*

These awards recognize associate editors and reviewers who have offered exceptional service in the review process.

3. National Science Foundation ECCS Award 1807536, \$350,892. 2018-2022.

Dynamic Risk-Averse Optimization of Distributed Energy Resource Aggregators.

PI: Daniel Jiang, co-PI Jeff Kharoufeh

4. Tencent AI Lab Faculty Award, \$50,000. 2017-2018.

“Advanced Reinforcement Learning Methods for Gameplay AI.”

5. Central Research Development Fund, University of Pittsburgh, \$16,000. 2017-2019.

“Naloxone Procurement and Distribution in the Presence of High Drug Prices.”

6. Pittsburgh Supercomputing Center Grant, 2017-2018.

7. Wu Prize for Excellence (Sir Gordon Y.S. Wu '58), Princeton University, 2015.

Mentoring Experience

Intern Supervision at Meta/Facebook

Jiayue Wan, Cornell University, 2022-2023 (CDS Research Intern).

Raul Astudillo, Cornell University, 2020-2021 (CDS Research Intern).

Student Supervision at University of Pittsburgh

Ibrahim Elshar, Ph.D., 2017-2022 (first position: researcher at Hitachi AI Lab).

Thesis: *Exploiting Structure and Relaxations in RL and Stochastic Optimal Control.*

Yijia Wang, Ph.D., 2016-2022 (first position: applied scientist at Amazon SCOT).

Thesis: *Structured Strategies for Learning and Exploration in Sequential Decision Making.*

Qingshi Sun, B.S., 2021-2022 (now Ph.D student at USC ISyE).

Dissertation Committee or Reader

Liang Xu (committee), University of Pittsburgh, 2022.

Seyed-Danial Mohseni-Taheri (committee), University of Illinois Chicago, 2021.

Joe Durante (reader), Princeton University, 2020.

Emmanuel Ekwedike (reader), Princeton University, 2020.

David Abdul-Malak (committee), University of Pittsburgh, 2018.

Felipe Hernandez Cruz (committee), University of Pittsburgh, 2018.

Arnab Bhattacharya (committee), University of Pittsburgh, 2017.

Onur Tavaslioglu (committee), University of Pittsburgh, 2017.

Bolong Cheng (reader), Princeton University, 2017.

Other Undergraduate Mentoring at University of Pittsburgh

Mackenzie Cavanaugh, Trevor Dycio, Carolyn Pontari, and Julie Shields, Fall 2018.

“Optimizing the Production of Meds2Home Prescriptions for Grane Rx Pharmacy.”

First Place IE Senior Design Project, Swanson School of Engineering Design EXPO.

Nick Kelly, Undergraduate Intern, Fall 2018.

Mohamed Kashkoush, Undergraduate Intern, 2017-2018.

Service

Co-organizer of RLC Workshop “Deployable RL: From Research to Practice,” August 9, 2024.

Co-organizer of KDD Workshop “MAB and RL: Advancing Decision Making in E-Commerce and Beyond,” (August 15, 2021, August 7, 2023); see (26):

Co-organizer of Facebook Adaptive Experimentation Workshop, February 13, 2020.

Associate Editor, *Operations Research Letters*, February 2018–April 2021.

Reviewer for

Journals: *Operations Research* (2017-24), *Management Science* (2019-22), *Manufacturing & Service Operations Management* (2017-20), *INFORMS Journal on Computing* (2020-21), *INFORMS Journal on Data Science* (2021-23), *Transportation Science* (2021), *IEEE Transactions on Automatic Control* (2019-20), *Transactions on Machine Learning Research* (2022-24), *SIAM Journal on Optimization* (2022-23), *European Journal of Operational Research* (2018, 2021), *Optimization Letters* (2023), *Probability in the Engineering and Informational Sciences* (2019-20), *Naval Research Logistics* (2018-19), *Annals of Operations Research* (2018-19), *IEEE Transactions on Sustainable Energy* (2017), *IEEE Transactions on Smart Grid* (2016), *IEEE Transactions on Power Systems* (2015-16), *IIEE Transactions* (2018-20), *Production and Operations Management* (2019), *Omega* (2018).

Conferences: *NeurIPS* (2021-23), *ICLR* (2022-23), *AISTATS* (2023), *ICML* (2023-24), *AutoML* (2023).

Workshops: *ML and the Physical Sciences Workshop*, *NeurIPS* (2019), *Algorithmic Fairness through the Lens of Causality and Privacy*, *NeurIPS* (2022).

Mentor, Women in Machine Learning (WiML) Workshop at NeurIPS 2017.

Teaching Experience

@ University of Pittsburgh

Decision Models (newly developed), IE 1086/2086, Fall 2016-Fall 2021 (5 times)

Advanced undergraduate/master’s level; data-driven decision making under uncertainty with a broad set of applications in energy, inventory, revenue management, healthcare, and finance.

Approximate Dynamic Programming (newly developed), IE 3186, Spring 2017, Fall 2018.

Ph.D. level; fundamental theory of approximate dynamic programming/reinforcement learning.

Reinforcement Learning (newly developed), IE 2186, Summer 2018.

Master’s level; covers policy evaluation, control, and applications of reinforcement learning.

Innovate: Global Issues in Engineering and Business, March 2017, 2018 (one-week abroad).

Undergraduate level; taught by Prof. Jayant Rajgopal and covers global engineering and business. The week-long study abroad portion visits Beijing, Suzhou, and Shanghai, China.

@ Princeton University

Probability and Stochastic Systems, ORF 309, Fall 2015.

Junior level course on probability with weekly precepts.

Senior Thesis Writing Group Leader, ORFE Department, 2013–2016.

Mentoring group for undergraduate seniors working on their senior theses.

Operations and Information Engineering, ORF 411, Fall 2013, 2014.

Senior level course on decision-making in OR.

Optimal Learning, ORF 418, Spring 2013.

Senior level course on sequential information collection.

Selected Invited Seminars & Presentations

Keynote Speaker, IISE Pittsburgh Chapter Awards Dinner, Pittsburgh, November 2023.

Machine Learning and Inference Research, Netflix, June 2023.

Industrial Engineering Seminar, University of Pittsburgh, December 2022.

Amazon Science Seminar, Amazon, August 2022.

ISOM Seminar, Hong Kong University of Science and Technology Business School, May 2022.

Decision Sciences Seminar, Duke Fuqua Business School, April 2022.

Industrial Engineering Seminar, University of Pittsburgh, April 2021.

OM Seminar, Lee Kong Chian School of Business, Singapore Management University, April 2021.

Institute of OR and Analytics Seminar, National University of Singapore, December 2020.

Pitt SWE Code Day (for middle schoolers), Pittsburgh, PA, November 2020.

Industrial and Systems Engineering Seminar, University of Tennessee, March 2020.

Information and Decision Sciences Seminar, University of Illinois at Chicago Business, Sept 2019.

ICML Conference (Long Talk), Stockholm, Sweden, July 2018.

Facebook, Menlo Park, CA, June 2018.

Systems and Industrial Engineering Seminar, University of Arizona, April 2018.

Commodity and Energy Markets Annual Meeting, Oxford, UK, June 2017.

4th International Conference on Computational Sustainability, Cornell University, July 2016

Industrial Engineering Seminar, University of Pittsburgh, January 2016.

OM Seminar, Tepper School of Business, Carnegie Mellon University, January 2016.

Patents

1. Warren B. Powell, Daniel R. Jiang. "Method for Bidding Battery Storage into Hour-Ahead Energy Markets," **U.S. Patent No. 9,965,802**, issued May 8, 2018.

Professional Activities

Industrial Mathematics Workshop and Clinic, IMA at University of Minnesota, July-Aug 2017.

IISE New Faculty Colloquium, Pittsburgh, PA, May 2017.

INFORMS New Faculty Colloquium, Nashville, TN, November 2016.

Member, INFORMS, 2013-Present.

Executive Board, Princeton Graduate Engineering Council, 2013-2015.

Other Awards & Honors

First Year Fellowship, Princeton University, 2011-2012.

Eta Kappa Nu Outstanding Junior in Electrical Engineering Award, Purdue University, Fall 2009.

Awarded to 1 out of approximately 200 undergraduate ECE students.

First Place, Purdue Mathematics Department Problem of the Week Competition, Spring 2008.

MathCounts Scholarship (full tuition, room, and board), Purdue University, 2007-2011.

Valedictorian, West Lafayette High School, May 2007.

Qualifier, USA Mathematical Olympiad (USAMO), April 2006.

Top 0.2% (approximately 400/200,000) of participating US high school students.

Undergraduate Research & Internship Experience

Undergraduate Researcher, Computer Science, **Stanford University**, June-August 2010.

Advised by Prof. Monica Lam. Developed an algorithm to use email & social networking data to generate graph-based representations of social affiliations to improve sharing granularity on social networks.

Engineering Intern, **Boston Scientific**, St. Paul, MN, May 2009-August 2009.

Undergraduate Researcher, Electrical Engineering, **Purdue University**, June-August 2007.

Advised by Prof. Hong Tan. Developed a virtual Chinese calligraphy brush using a haptic device.

Media

2024 Big Data Bowl Winners. (National Football League, 3/4/2024)

Facebook open-sources Ax and BoTorch to simplify AI model optimization. (Techcrunch, 5/1/2019)

Pitt researcher uses video games to unlock new levels of AI. (Swanson Newsroom, 11/5/2018)

Swanson School professor using video games to test AI. (University Times, 12/12/2018)

The Story Behind "What Would I Say?" (The New Yorker, 11/12/2013)

Personal

3 on 3 Intramural Basketball Champions, Princeton University Dillon Intramural Sports, Fall 2013.

Princeton University Science Action Most Viral Video Award for *Charged*, May 2013.

Co-written, directed, shot, edited with Omer Malik, Harvey Cheng, Matt Chang, Yasmin Afsar.

Co-creator of social media webapp *What Would I Say?*, Hack Princeton, November 2013.

17 million pageviews, 9 million unique visitors, winner of "Best Facebook Integrated Hack," and the subject of articles in The New Yorker, CNN, The Telegraph, and Time.