NAN JIANG

CONTACT

Office: 3322 Siebel Center, 201 N Goodwin Ave, Urbana, IL 61801
Email: nanjiang@illinois.edu
Website: http://nanjiang.cs.illinois.edu

EMPLOYMENT

University of Illinois at Urbana-Champaign 2018 – present
Position: Assistant Professor Department of Computer Science

Microsoft Research, New York City 2017 – 2018
Position: Postdoctoral Researcher Machine Learning Group

EDUCATION

PhD, Computer Science and Engineering 2011 – 2017
University of Michigan, Ann Arbor, MI, USA
Division: Artificial Intelligence Advisor: Satinder Singh

Bachelor of Engineering, Department of Automation 2007 – 2011
Tsinghua University, Beijing, China
Division: Control Theory and Application Graduate with Distinction

MONOGRAPH

Reinforcement Learning: Theory and Algorithms. (*working draft*)
Alekh Agarwal, Nan Jiang, Sham Kakade, Wen Sun.

AWARDS AND GRANTS

Google Research Scholar Award Mar. 2024

Sloan Research Fellowship Feb. 2024

ICML 2022 Outstanding Paper Runner-Up July 2022

NSF CAREER Award Mar. 2022
Title: Theoretical Foundations of Offline Reinforcement Learning.

Engineering Council Outstanding Advisor Award Mar. 2022

Adobe Data Science Research Award Sept. 2021

NSF AI Institute for Future Edge Networks and Distributed Intelligence Oct. 2021
Role: Co-PI.

ARL IoBT Collaborative Research Alliance Feb. 2021
Role: Co-PI on Task "Enabling the Safe and Responsible Use of RL"

Rackham Predoctoral Fellowship Mar. 2016

AAMAS 2015 Best Paper Award May 2015

SERVICES

• Journal Editing
Editor: Foundations and Trends in Machine Learning (FnT in ML)
Action Editor: Journal of Machine Learning Research (JMLR)

• Senior Area Chair: ICLR'24

• Area Chair/Senior PC: ICML'19–24, NeurIPS'20–23, COLT'20, AAAI'19, AISTATS'19

• Reviewer: ICML'16–18, IJCAI'16, AAAI'17–18, AISTATS'17–18, NeurIPS'17–19, ALT'15'17, COLT'16, ICLR'18, JMLR, MathOR, JAIR, MLJ

• Grant Panelist: NSF 2019, 2022

• Event (Co-)Organization:
NeurIPS 2021 Workshop on Offline Reinforcement Learning Dec. 2021
Workshop on Advances in Theory and Algorithms for Deep Reinforcement Learning at ICERM Aug. 2021
RL Theory Virtual Seminar (co-host) 2021, 2023

• Internal Graduate Admissions 2021–2023
Graduate Study Committee 2019–
Subcommittee for Course Proposal Evaluation, College of Engineering 2020
Undergraduate Study Committee 2018–2019

INVITED TALKS

Conference on Information Sciences and Systems (CISS), 2024 Mar. 2024
Session: Modern Reinforcement Learning Host: Chi Jin

IISc Workshop: Reinforcement Learning: Recent Advances and Challenges Ahead Feb. 2024
Keynote Speaker Host: Gugan Thoppe

University of Michigan, EECS CSP seminar Nov. 2023
Host: Lei Ying

Carnegie Mellon University, SCS Special Seminar Mar. 2023
Host: Andrej Risteski

Tutorial Lectures, Workshop on Machine Learning and Its Applications at NUS Oct. 2022
Topic: Offline RL Theory

International Conference on Continuous Optimization (ICCOPT) 2022 July 2022
Session: Stochastic Algorithms Chair: Ashwin Pananjady

Microsoft Research MTL Seminar July 2022
Host: Geoff Gordon

New Models in Online Decision Making for Real-World Applications July 2022
Part of TTIC 2022 Summer Workshop Program

Reinforcement Learning and Decision Making (RLDM) June 2022
Workshop on RL as a Model of Agency

Conference on Information Sciences and Systems Mar. 2022
Session: Frontiers of Theoretical Reinforcement Learning Host: Chi Jin

Vector Institute Aug. 2021
Host: Amir-massoud Farahmand

ICERM Workshop on Advances in Theory and Algorithms for Deep RL Aug. 2021
Organizers: Nan Jiang, Sanjay Shakkottai, R. Srikant, Mengdi Wang

Learning Theory Seminar, Google Research NYC July 2021
Host: Christoph Dann

Microsoft Research Asia, Machine Learning Group June 2021
Host: Li Zhao

DeepMind, RL Team Mar. 2021
Host: Junhyuk Oh

Online Data Science Seminar, LSE Stats Feb. 2021
Host: Chengchun Shi

NeurIPS Workshop on Offline Reinforcement Learning Dec. 2020
Organizers: Aviral Kumar, Rishabh Agarwal, George Tucker, Lihong Li, Doina Precup

Simons Institute Workshop on RL from Batch Data and Simulation Nov. 2020
Organizers: Mengdi Wang, Emma Brunskill, Sean Meyn

RL Theory Virtual Seminars June 2020
Hosts: Gergely Neu, Ciara Pike-Burke, Csaba Szepesvári

Simons Institute Workshop: Emerging Challenges in Deep Learning Aug. 2019
Host: Matus Telgarsky

Annual ShanghaiTech Symposium on Information Science and Technology July 2019
Host: Ziyu Shao

2nd Machine Learning Theory Workshop at Peking University June 2019
Host: Liwei Wang

Workshop on Machine Learning for All-Inclusive Finance at ICML-19 Expo June 2019
Organizer: Ant Financial Service Group

MSR Talk Series, Microsoft Research Redmond May. 2019
Host: Alekh Agarwal

Statistics Research Colloquium, Purdue University Jan. 2019
Host: Guang Cheng

AI Seminar at CSE, University of Michigan Nov. 2018
Host: Satinder Singh

Seminar on Decision, Optimization, and Learning (DOL) at Caltech June 2018
Host: Yisong Yue

Machine Learning Seminar at UIUC Apr. 2018

Host: Girish Chowdhary

Conference on Information Sciences and Systems Mar. 2018
Session: Algorithmic Reinforcement Learning Host: Mengdi Wang

DARPA Workshop: *Diverse Ways of Inferring Missions* Oct. 2017
Invited to present the work on repeated inverse RL

Job Talk: *New results in statistical reinforcement learning* Feb. – Mar. 2017
University of Illinois Urbana-Champaign Host: Matus Telgarsky
University of California, Santa Barbara Host: Xifeng Yan
University of Maryland, College Park Host: Marine Carpuat
McGill University Host: Joelle Pineau
University of Massachusetts Amherst Host: Akshay Krishnamurthy
Toyota Technological Institute at Chicago Host: Matthew Walter
University of Minnesota Twin Cities Host: Arindam Banerjee

Microsoft Research, NYC Jan. 2017
Host: John Langford

Carnegie Mellon University, Machine Learning Department Nov. 2016
Host: Emma Brunskill

IBM Thomas J. Watson Research Center Aug. 2016
Host: Kartik Talamadupula

International Joint Conference on Artificial Intelligence July 2016
Invited to present the AAMAS 2015 paper at the Sister Conference Best Paper track.

International Conference on Robotics and Automation May 2015
Invited to present the AAAI 2015 paper at the 50th Anniversary of Shakey.

Carnegie Mellon University, Machine Learning Department Mar. 2015
Host: Emma Brunskill

PREPRINTS

Finite Sample Analysis of Minimax Offline Reinforcement Learning: Completeness, Fast Rates and First-Order Efficiency. (pdf)
Masatoshi Uehara, Masaaki Imaizumi, Nan Jiang, Nathan Kallus, Wen Sun, Tengyang Xie.

Model-free Representation Learning and Exploration in Low-rank MDPs. (pdf)
Aditya Modi, Jinglin Chen, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal.
(JMLR'23) *Journal of Machine Learning Research* (minor revision).

PUBLICATIONS

Harnessing Density Ratios for Online Reinforcement Learning. (pdf)
Philip Amortila, Dylan Foster, Nan Jiang, Ayush Sekhari, Tengyang Xie.
(ICLR'24) *The 11th International Conference on Learning Representations, with spotlight.*

Marginalized Importance Sampling for Off-Environment Policy Evaluation. (pdf)
Pulkit Katdare, Nan Jiang, Katherine Driggs-Campbell.
(CoRL'23) *7th Conference on Robot Learning.*

Future-Dependent Value-Based Off-Policy Evaluation in POMDPs. (pdf)
Masatoshi Uehara, Haruka Kiyohara, Andrew Bennett, Victor Chernozhukov, Nan Jiang, Nathan Kallus, Chengchun Shi, Wen Sun.
(NeurIPS'23) *37th Neural Information Processing Systems, with spotlight.*

Adversarial Model for Offline Reinforcement Learning. (pdf)
Tengyang Xie, Mohak Bhardwaj, Nan Jiang, Ching-An Cheng.
(NeurIPS'23) *37th Neural Information Processing Systems.*

Reinforcement Learning in Low-Rank MDPs with Density Features. (pdf)
Audrey Huang, Jinglin Chen, Nan Jiang.
(ICML'23) *40th International Conference on Machine Learning.*

The Optimal Approximation Factors in Misspecified Off-Policy Value Function Estimation. (pdf)
Philip Amortila, Nan Jiang, Csaba Szepesvari.
(ICML'23) *40th International Conference on Machine Learning.*

Offline Learning in Markov Games with General Function Approximation. (pdf)
Yuheng Zhang, Yu Bai, Nan Jiang.
(ICML'23) *40th International Conference on Machine Learning.*

The Role of Coverage in Online Reinforcement Learning. (pdf)
Tengyang Xie, Dylan Foster, Yu Bai, Nan Jiang, Sham Kakade.
(ICLR'23) *The 11th International Conference on Learning Representations*.

Explaining RL Decisions with Trajectories. (pdf)
Shripad Vilasrao Deshmukh, Arpan Dasgupta, Balaji Krishnamurthy, Nan Jiang, Chirag Agarwal, Georgios Theocharous, Jayakumar Subramanian.
(ICLR'23) *The 11th International Conference on Learning Representations*.

A Few Expert Queries Suffices for Sample-Efficient RL with Resets and Linear Value Approximation. (pdf)
Philip Amortila, Nan Jiang, Dhruv Madeka, Dean P. Foster.
(NeurIPS'22) *36th Neural Information Processing Systems.*

Interaction-Grounded Learning with Action-inclusive Feedback. (pdf)
Tengyang Xie, Akanksha Saran, Dylan J Foster, Lekan Molu, Ida Momennejad, Nan Jiang, Paul Mineiro, John Langford.
(NeurIPS'22) *36th Neural Information Processing Systems.*

On the Statistical Efficiency of Reward-Free Exploration in Non-Linear RL. (pdf)
Jinglin Chen, Aditya Modi, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal.
(NeurIPS'22) *36th Neural Information Processing Systems.*

Tiered Reinforcement Learning: Pessimism in the Face of Uncertainty and Constant Regret. (pdf)
Jiawei Huang, Li Zhao, Tao Qin, Wei Chen, Nan Jiang, Tie-Yan Liu.
(NeurIPS'22) *36th Neural Information Processing Systems.*

Beyond the Return: Off-policy Function Estimation under User-specified Error-measuring Distributions. (pdf)

Audrey Huang, Nan Jiang.
(NeurIPS'22) *36th Neural Information Processing Systems.*

Offline Reinforcement Learning Under Value and Density-Ratio Realizability: the Power of Gaps. (pdf)
Jinglin Chen, Nan Jiang.
(UAI'22) *38th Conference on Uncertainty in Artificial Intelligence.*

Offline Reinforcement Learning with Realizability and Single-policy Concentrability. (pdf)
Wenhao Zhan, Baihe Huang, Audrey Huang, Nan Jiang, Jason D. Lee.
(COLT'22) *35th Annual Conference on Learning Theory.*

Adversarially Trained Actor Critic for Offline Reinforcement Learning. (pdf)
Ching-An Cheng, Tengyang Xie, Nan Jiang, Alekh Agarwal.
(ICML'22 **Outstanding Paper Runner-up**) *39th International Conference on Machine Learning.*

A Minimax Learning Approach to Off-Policy Evaluation in Partially Observable Markov Decision Processes. (pdf) Chengchun Shi, Masatoshi Uehara, Jiawei Huang, Nan Jiang.
(ICML'22) *39th International Conference on Machine Learning.*

Towards Deployment-Efficient Reinforcement Learning: Lower Bound and Optimality. (pdf)
Jiawei Huang, Jinglin Chen, Li Zhao, Tao Qin, Nan Jiang, Tie-Yan Liu.
(ICLR'22) *The 10th International Conference on Learning Representations*.

On the Convergence Rate of Off-Policy Policy Optimization Methods with Density-Ratio Correction. (pdf)
Jiawei Huang, Nan Jiang.
(AISTATS'22) *The 25th International Conference on Artificial Intelligence and Statistics; also presented at Offline RL Workshop at NeurIPS-20*.

Bellman-consistent Pessimism for Offline Reinforcement Learning. (pdf)
Tengyang Xie, Ching-An Cheng, Nan Jiang, Paul Mineiro, Alekh Agarwal.
(NeurIPS'21) *35th Neural Information Processing Systems, with oral presentation (less than 1%)*.

Towards Hyperparameter-free Policy Selection for Offline Reinforcement Learning. (pdf)
Siyuan Zhang, Nan Jiang.
(NeurIPS'21) *35th Neural Information Processing Systems*.

Policy Finetuning: Bridging Sample-Efficient Offline and Online Reinforcement Learning. (pdf) Tengyang Xie, Nan Jiang, Huan Wang, Caiming Xiong, Yu Bai.
(NeurIPS'21) *35th Neural Information Processing Systems.*

Empirical Study of Off-Policy Policy Evaluation for Reinforcement Learning. (pdf)
Cameron Voloshin, Hoang Le, Nan Jiang, Yisong Yue.
(NeurIPS'21) *35th Neural Information Processing Systems, Datasets and Benchmarks Track*

On Query-efficient Planning in MDPs under Linear Realizability of the Optimal State-value Function. (pdf) Gellert Weisz, Philip Amortila, Barnabás Janzer, Yasin Abbasi-Yadkori, Nan Jiang, Csaba Szepesvári.
(COLT'21) *The 34th Annual Conference on Learning Theory.*

Batch Value-function Approximation with Only Realizability. (pdf)
Tengyang Xie, Nan Jiang.
(ICML'21) *38th International Conference on Machine Learning.*

Minimax Model Learning. (pdf)
Cameron Voloshin, Nan Jiang, Yisong Yue.
(AISTATS'21) 24th International Conference on Artificial Intelligence and Statistics.

Improved Worst-Case Regret Bounds for Randomized Least-Squares Value Iteration. (pdf) Priyank Agrawal, Jinglin Chen, Nan Jiang.
(AAAI'21) 35th AAAI Conference on Artificial Intelligence.

Minimax Value Interval for Off-Policy Evaluation and Policy Optimization. (pdf)
Nan Jiang, Jiawei Huang.
(NeurIPS'20) 34rd Neural Information Processing Systems.

Minimax Weight and Q-function Learning for Off-Policy Evaluation. (pdf)
Masatoshi Uehara, Jiawei Huang, Nan Jiang.
(ICML'20) *37th International Conference on Machine Learning.*

From Importance Sampling to Doubly Robust Policy Gradient. (pdf)
Jiawei Huang, Nan Jiang.
(ICML'20) *37th International Conference on Machine Learning.*

$Q^\star$ **Approximation Schemes for Batch RL: A Theoretical Comparison.** (pdf)
Tengyang Xie, Nan Jiang.
(UAI'20) *Conference on Uncertainty in Artificial Intelligence.*

Sample Complexity of RL using Linearly Combined Model Ensembles. (pdf)
Aditya Modi, Nan Jiang, Ambuj Tewari, Satinder Singh.
(AISTATS'20) *23rd International Conference on Artificial Intelligence and Statistics.*

Provably Efficient Q-Learning with Low Switching Cost. (pdf)
Yu Bai, Tengyang Xie, Nan Jiang, Yu-Xiang Wang.
(NeurIPS'19) *33rd Neural Information Processing Systems.*

Information-Theoretic Considerations in Batch Reinforcement Learning. (pdf)
Jinglin Chen, Nan Jiang.
(ICML'19) *36th International Conference on Machine Learning.*

Provably Efficient RL with Rich Observations via Latent State Decoding. (pdf)
Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudík, John Langford.
(ICML'19) *36th International Conference on Machine Learning.*

Model-based RL in Contextual Decision Processes: PAC bounds and Exponential Improvements over Model-free Approaches. (pdf)
Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford.
(COLT'19) *32nd Annual Conference on Learning Theory.*

On Oracle-Efficient PAC Reinforcement Learning with Rich Observations. (pdf)
Christoph Dann, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, Robert

E. Schapire.
(NeurIPS'18) *32nd Neural Information Processing Systems, with spotlight presentation; also presented at: 12th NYAS Machine Learning Symposium.*

Completing State Representations Using Spectral Learning. (pdf)
Nan Jiang, Alex Kulesza, Satinder Singh.
(NeurIPS'18) *32nd Neural Information Processing Systems.*

Open Problem: The Dependence of Sample Complexity Lower Bounds on Planning Horizon. (pdf)
Nan Jiang, Alekh Agarwal.
(COLT'18) *31st Annual Conference on Learning Theory.*

Hierarchical Imitation and Reinforcement Learning. (pdf)
Hoang M. Le, Nan Jiang, Alekh Agarwal, Miroslav Dudík, Yisong Yue, Hal Daumé III.
(ICML'18) *35th International Conference on Machine Learning.*

Markov Decision Processes with Continuous Side Information. (pdf)
Aditya Modi, Nan Jiang, Satinder Singh, Ambuj Tewari.
(ALT'18) *29th International Conference on Algorithmic Learning Theory.*

PAC Reinforcement Learning with an Imperfect Model. (pdf)
Nan Jiang.
(AAAI'18) *32nd AAAI Conference on Artificial Intelligence.*

Repeated Inverse Reinforcement Learning. (pdf)
Kareem Amin*, Nan Jiang*, Satinder Singh. (*Equal contribution.)
(NeurIPS'17) *31st Neural Information Processing Systems, with spotlight presentation; also presented at Reinforcement Learning and Decision Making 2017.*

Contextual Decision Processes with Low Bellman Rank are PAC-Learnable. (pdf)
Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, Robert E. Schapire.
(ICML'17) *34th International Conference on Machine Learning; also presented at 11th NYAS Machine Learning Symposium (with presentation award).*

Doubly Robust Off-policy Value Evaluation for Reinforcement Learning. (pdf)
Nan Jiang, Lihong Li.
(ICML'16) *33rd International Conference on Machine Learning; also presented in Workshop "Machine Learning for eCommerce" in Neural Information Processing Systems 2015.*

On Structural Properties of MDPs that Bound Loss due to Shallow Planning. (pdf)
Nan Jiang, Satinder Singh, Ambuj Tewari.
(IJCAI'16) *25th International Joint Conference on Artificial Intelligence.*

Improving Predictive State Representations via Gradient Descent. (pdf)
Nan Jiang, Alex Kulesza, Satinder Singh.
(AAAI'16) *30th AAAI Conference on Artificial Intelligence.*

Abstraction Selection in Model-based Reinforcement Learning. (pdf)
Nan Jiang, Alex Kulesza, Satinder Singh.
(ICML'15) *32nd International Conference on Machine Learning.*

The Dependence of Effective Planning Horizon on Model Accuracy. (pdf)
Nan Jiang, Alex Kulesza, Satinder Singh, Richard Lewis.
(AAMAS'15 **best paper**) *14th International Conference on Autonomous Agents and Multiagent Systems.*

Low-Rank Spectral Learning with Weighted Loss Functions. (pdf)
Alex Kulesza, Nan Jiang, Satinder Singh.
(AISTATS'15) *18th International Conference on Artificial Intelligence and Statistics.*

Spectral Learning of Predictive State Representations with Insufficient Statistics. (pdf)
Alex Kulesza, Nan Jiang, Satinder Singh.
(AAAI'15) *29th AAAI Conference on Artificial Intelligence*.

Improving UCT Planning via Approximate Homomorphisms. (pdf)
Nan Jiang, Satinder Singh, Richard Lewis.
(AAMAS'14) *13th International Conference on Autonomous Agents and Multiagent Systems*.

WORKSHOP PAPERS AND TECHNICAL NOTES

A Variant of the Wang-Foster-Kakade Lower Bound for the Discounted Setting. (pdf)
Philip Amortila, Nan Jiang, Tengyang Xie.

On Value Functions and the Agent-Environment Boundary. (pdf)
Nan Jiang.

Deterministic Bellman Residual Minimization.
Ehsan Saleh, Nan Jiang.
OptRL Workshop at NeurIPS'19.

PROFESSIONAL MEMBERSHIPS

Member of Association for the Advancement of Artificial Intelligence — Since 2018
Member of Association for Computing Machinery — Since 2018

TEACHING (UIUC)

CS 542 Statistical Reinforcement Learning (F21, F22, & F23 teaching excellence)
CS 443 Reinforcement Learning (S23 teaching excellence)
CS 598 Special Topics: Statistical Reinforcement Learning (F18 & S19 teaching excellence, F20 *outstanding*)
CS 498 Special Topics: Reinforcement Learning (F19; S21 teaching excellence)

STUDENT ADVISING

• PhD Theses Advised

Jinglin Chen. 2023.
Reinforcement Learning under General Function Approximation and Novel Interaction Settings.

Tengyang Xie. 2023.
Reinforcement Learning with Offline Data: Foundations and Algorithms.

• Current PhD Students

Philip Amortila, Audrey Huang, Yuheng Zhang, Wei Xiong

• **Master Theses Advised**

Priyank Agrawal. 2021.
Improved worst-case regret bounds for randomized least-squares value iteration.

Siyuan Zhang. 2021.
Batch value function tournament for offline policy selection in reinforcement learning.

• **Undergraduate Theses Advised**

Siyuan Zhang — Senior Thesis, 2019
Topic: Improving predictive state representations by optimizing transformation matrices.

• **PhD Thesis Committee**

Germano Gabbianelli (Universitat Pompeu Fabra)	Advisor: Gergely Neu
Ziping Xu (University of Michigan)	Advisor: Ambuj Tewari
Yuxuan Li (UIUC)	Advisor: Alfred Chong
Andrew Patterson (McGill)	Advisor: Martha White
Muhammad Aneeq Uz Zaman (UIUC)	Advisors: Tamer Basar, Geir Dullerud
Ehsan Saleh (UIUC)	Advisors: Timothy Bretl, Matthew West
Yikun Ban (UIUC)	Advisor: Jingrui He
Masatoshi Uehara (Cornell)	Advisor: Nathan Kallus
Unnat Jain (UIUC)	Advisors: Svetlana Lazebnik, Alexander Schwing
Belinda Tzen (UIUC)	Advisor: Maxim Raginsky
Aditya Modi (University of Michigan)	Advisors: Ambuj Tewari, Satinder Singh
Iou-Jen Liu (UIUC)	Advisor: Alexander Schwing

• **Visiting Students**

Shivangi Agarwal — Visiting PhD Student, 2024
Topic: AI for networks

Yash Nair — Summer REU, 2020
Topic: Off-policy evaluation in POMDPs.

Masatoshi Uehara — Visiting PhD Student, 2019
Topic: Marginalized importance sampling for off-policy evaluation.

Jiawei Huang — Visiting Student, 2019
Topic: Doubly robust policy gradient.

Kaiqian Han — Summer REU, 2019
Topic: Question & answer network for model-based RL.

Jiachen Hu — Summer REU, 2019
Topic: Concurrent exploration in RL.