

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

Position: Assistant Professor

2018 – present
Department of Computer Science

Microsoft Research, New York City

Position: Postdoctoral Researcher

2017 – 2018
Machine Learning Group

EDUCATION

PhD, Computer Science and Engineering

University of Michigan, Ann Arbor, MI, USA

Division: Artificial Intelligence

2011 – 2017
Advisor: Satinder Singh

Bachelor of Engineering, Department of Automation

Tsinghua University, Beijing, China

Division: Control Theory and Application

2007 – 2011
Graduate with Distinction

MONOGRAPH

Reinforcement Learning: Theory and Algorithms. (*working draft*)

Alekh Agarwal, Nan Jiang, Sham Kakade, Wen Sun.

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.

PUBLICATIONS

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*.

The Caltech Off-Policy Policy Evaluation Benchmarking Suite. ([pdf](#))

(Previously: Empirical Study of Off-Policy Policy Evaluation for Reinforcement Learning)

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^* 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*; also presented in Workshop "From Bad Models to Good Policies" in Neural Information Processing Systems 2014.

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

On the Convergence Rate of Off-Policy Policy Optimization Methods with Density-Ratio Correction. ([pdf](#))

Jiawei Huang, Nan Jiang.

Offline RL Workshop at NeurIPS-20.

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.

INVITED TALKS

- | | |
|--|-----------|
| Vector Institute
Host: Amir-massoud Farahmand | Aug. 2021 |
| ICERM Workshop on Advances in Theory and Algorithms for Deep RL
Organizers: Nan Jiang, Sanjay Shakkottai, R. Srikant, Mengdi Wang | Aug. 2021 |
| Learning Theory Seminar, Google Research NYC
Host: Christoph Dann | July 2021 |
| Microsoft Research Asia, Machine Learning Group
Host: Li Zhao | June 2021 |
| DeepMind, RL Team
Host: Junhyuk Oh | Mar. 2021 |
| Online Data Science Seminar, LSE Stats
Host: Chengchun Shi | Feb. 2021 |
| NeurIPS Workshop on Offline Reinforcement Learning
Organizers: Aviral Kumar, Rishabh Agarwal, George Tucker, Lihong Li, Doina Precup | Dec. 2020 |
| Simons Institute Workshop on RL from Batch Data and Simulation
Organizers: Mengdi Wang, Emma Brunskill, Sean Meyn | Nov. 2020 |
| RL Theory Virtual Seminars
Hosts: Gergely Neu, Ciara Pike-Burke, Csaba Szepesvári | June 2020 |
| Simons Institute Workshop: Emerging Challenges in Deep Learning
Host: Matus Telgarsky | Aug. 2019 |
| Annual ShanghaiTech Symposium on Information Science and Technology
Host: Ziyu Shao | July 2019 |
| 2nd Machine Learning Theory Workshop at Peking University
Host: Liwei Wang | June 2019 |
| Workshop on Machine Learning for All-Inclusive Finance at ICML-19 Expo
Organizer: Ant Financial Service Group | June 2019 |

MSR Talk Series, Microsoft Research Redmond Host: Alekh Agarwal	May. 2019
Statistics Research Colloquium, Purdue University Host: Guang Cheng	Jan. 2019
AI Seminar at CSE, University of Michigan Host: Satinder Singh	Nov. 2018
Seminar on Decision, Optimization, and Learning (DOL) at Caltech Host: Yisong Yue	June 2018
Machine Learning Seminar at UIUC Host: Girish Chowdhary	Apr. 2018
Conference on Information Sciences and Systems Session: Algorithmic Reinforcement Learning	Mar. 2018 Host: Mengdi Wang
DARPA Workshop: <i>Diverse Ways of Inferring Missions</i> Invited to present the work on repeated inverse RL	Oct. 2017
Job Talk: <i>New results in statistical reinforcement learning</i> University of Illinois Urbana-Champaign University of California, Santa Barbara University of Maryland, College Park McGill University University of Massachusetts Amherst Toyota Technological Institute at Chicago University of Minnesota Twin Cities	Feb. – Mar. 2017 Host: Matus Telgarsky Host: Xifeng Yan Host: Marine Carpuat Host: Joelle Pineau Host: Akshay Krishnamurthy Host: Matthew Walter Host: Arindam Banerjee
Microsoft Research, NYC Host: John Langford	Jan. 2017
Carnegie Mellon University, Machine Learning Department Host: Emma Brunskill	Nov. 2016
IBM Thomas J. Watson Research Center Host: Kartik Talamadupula	Aug. 2016
International Joint Conference on Artificial Intelligence Invited to present the AAMAS 2015 paper at the Sister Conference Best Paper track.	July 2016
International Conference on Robotics and Automation Invited to present the AAI 2015 paper at the 50th Anniversary of Shakey.	May 2015
Carnegie Mellon University, Machine Learning Department Host: Emma Brunskill	Mar. 2015

AWARDS AND GRANTS

Adobe Data Science Research Award	Sept. 2021
NSF AI Institute for Future Edge Networks and Distributed Intelligence Role: Co-PI.	Oct. 2021
ARL IoBT Collaborative Research Alliance Role: Co-PI on Task “Enabling the Safe and Responsible Use of RL”	Feb. 2021
Rackham Predoctoral Fellowship	Mar. 2016
AAMAS 2015 Best Paper Award	May 2015

SERVICES

• 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	2021
• Meta-Reviewer (Area Chair/Senior PC) AAAI’19, AISTATS’19, ICML’19–21, COLT’20, NeurIPS’20–21	
• 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 Small 2019	
• Internal Graduate Admissions	2021–
Graduate Study Committee	2019–2021
Subcommittee for Course Proposal Evaluation, College of Engineering	2020
Undergraduate Study Committee	2018–2019

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) CS 598 Special Topics: Statistical Reinforcement Learning (F18 & S19 teaching excellence , F20 outstanding—highest honor given by Center of Innovation in Teaching and Learning) CS 498 Special Topics: Reinforcement Learning (F19, S21 teaching excellence)	
---	--

- MSR NYC
COMS E6998.001 Bandits and RL @ Columbia University (F17) Project Advisor
- U Michigan
STATS 710 Sequential Decision Making with mHealth Applications (F16) Guest Lecturer
EECS 545 Machine Learning (W14) Graduate Student Instructor
EECS 574 Computational Complexity (F13) Graduate Student Instructor

STUDENT ADVISING

• Visiting Students

- 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.

• Undergraduate Theses

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

• Master Theses

- Priyank Agrawal, MS. Graduation: 2021
Thesis: *Improved worst-case regret bounds for randomized least-squares value iteration.*
- Siyuan Zhang, MS. Graduation: 2021
Thesis: *Batch value function tournament for offline policy selection in reinforcement learning.*

• PhD Thesis Committee

- Unnat Jain Advisors: Svetlana Lazebnik, Alexander Schwing
- Belinda Tzen Advisor: Maxim Raginsky
- Aditya Modi Advisors: Ambuj Tewari, Satinder Singh
- Iou-Jen Liu Advisor: Alexander Schwing

• Current PhD Students

- Jinglin Chen, Tengyang Xie, Philip Amortila, Jiawei Huang, Audrey Huang