

NAN JIANG

CONTACT

Email: nanjiang@illinois.edu.

Website: <http://nanjiang.cs.illinois.edu>.

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

EMPLOYMENT

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

Microsoft Research, New York City Summer 2016
Position: Research Intern. Mentor: John Langford

Microsoft Research, Redmond Summer 2015
Position: Research Intern. Mentor: Lihong Li

RESEARCH INTERESTS

Machine Learning, with focus on **Reinforcement Learning**.

PREPRINTS

Hierarchical Imitation and Reinforcement Learning. ([pdf](#))
Hoang M. Le, Nan Jiang, Alekh Agarwal, Miroslav Dudk, Yisong Yue, Hal Daum III.

On Polynomial Time PAC Reinforcement Learning with Rich Observations. ([pdf](#))
Christoph Dann, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, Robert E. Schapire.
presented at: 12th NYAS Machine Learning Symposium.

PUBLICATIONS

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.)

(NIPS'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*.

INVITED TALKS

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

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.

AWARDS

Rackham Predoctoral Fellowship Mar. 2016

CSE Honor Competition, 2nd Place Nov. 2015
Representing AI Division in CSE's annual research competition at UMich.

AAMAS 2015 Best Paper Award May 2015

REVIEWING

Reviewer: ICML'16-18, IJCAI'16, EWRL'16, AAAI'17-18, AISTATS'17-18, NIPS'17
Auxiliary Reviewer: ALT'15'17, COLT'16, ICLR'18

TEACHING

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