

Kamyar Azizzadenesheli

Nvidia Corporation
Santa Clara, California

Homepage: [kamyar.page](#)

Research Interests

My research interest is mainly in the area of Machine Learning, from theory to practice. I work on topics including but not limited to Reinforcement Learning, Bandits, Applied Mathematics, Deep Learning, Neural Operators, Control Theory, Optimization, High Dimensional Statistics, Risk Assessment, Online learning, Domain adaptation, Active Learning, Robotics, Safety, Adversarial Attacks, and Generative models through both learning theory and core scientific lenses. I am actively working on trustworthy and domain specific AI solutions to real-world problems.

Current Positions

Senior Research Scientist, Nvidia Corporation
Santa Clara, CA, USA

Since July 2022

Education

Ph.D. & M.Sc. in Electrical Engineering & Computer Science
University of California, Irvine, CA, USA.

2014-2019
Advised by Prof. Anima Anandkumar.

B.Sc. in Electrical Engineering
Sharif University of Technology, Tehran, Iran.

2010-2014
Advised by Prof. F. Ashtiani & Prof. F. Marvasti,

Positions

- Senior Research Scientist at **Nvidia**, Santa Clara, CA, USA. *(July 2022-Present)*
- Assistant Professor at **Purdue University**, West Lafayette, IN, USA. *(August 2020-July 2022)*
- Visiting Faculty at **Caltech**, Pasadena, CA, USA. *(August 2021)*
At the Geological and Planetary Sciences Division,
Host: Prof. Zachary E. Ross.
At the Computing + Mathematical Sciences Department,
Host: Prof. Animashree Anandkumar.
- Postdoctoral Scholar at **Caltech**, Pasadena, CA, USA. *(July 2019-August 2020)*
Hosts: Prof. Animashree Anandkumar & Prof. Yisong Yue
- Special Student at **Caltech**, Pasadena, CA, USA. *(April 2019-June 2019)*
Hosts: Prof. Animashree Anandkumar & Prof. Yisong Yue
- Visiting Researcher at **Caltech**, Pasadena, CA, USA. *(June 2018-April 2019)*
Hosts: Prof. Animashree Anandkumar & Prof. Yisong Yue
- Visiting Researcher at **Stanford University**, Palo Alto, CA, USA. *(October 2017-June 2018)*
Host: Prof. Emma Brunskill

- Long term visiting researcher at **Simons Institute, UC Berkeley, CA.** *(Jan 2017-May 2017)*
- Guest Researcher at **INRIA, France.** *(August 2016-November 2016)*
Host: Dr. Alessandro Lazaric
- Short-term Visitor at **MSR, New York City, NY, USA.** *(2016)*
Host: Prof. Animashree Anandkumar
- Short-term Visitor at **MSR, New England, MA, USA.** *(2016)*
Host: Prof. Animashree Anandkumar

Awards and Honors

- **EECS Department Fellowship**, University of California, Irvine, CA, USA, 2014.
- **Silver Medal** in International Olympiad in Astronomy & Astrophysics, Beijing, China, 2010.
- **Gold Medal** in National Olympiad in Astronomy & Astrophysics, Tehran, Iran, 2009.
- **First rank elite student** in National Elite Foundation, Iran, 2009-2014.
- **Second prize** in Sharif RoboCup Competition, Machine vision section, Tehran, Iran, 2012.

Books

- **Deep Learning - The Straight Dope**, an online Deep Learning book on Amazon Mxnet Library. Zachary C. Lipton, Mu Li, Alex Smola, Sheng Zha, Aston Zhang, Joshua Z. Zhang, Eric Junyuan Xie, Kamyar Azizzadenesheli, Jean Kossaifi, Stephan Rabanser, [\[link\]](#)

Papers

Conference:

1. Taylan Kargin, Sahin Lale, Kamyar Azizzadenesheli, Anima Anandkumar, Babak Hassibi. Thompson Sampling Achieves $\tilde{O}(\sqrt{T})$ Regret in Linear Quadratic Control
Appeared at Conference on Learning Theory (COLT) 2022.
[paper]
2. Pan Xu, Hongkai Zheng, Eric Mazumdar, Kamyar Azizzadenesheli, Anima Anandkumar. Langevin Monte Carlo for Contextual Bandits
Appeared at International Conference on Machine Learning (ICML) 2022.
[paper]
3. Liu Leqi, Audrey Huang, Zachary C. Lipton, Kamyar Azizzadenesheli. Supervised Learning with General Risk Functionals
Appeared at International Conference on Machine Learning (ICML) 2022.
[paper]
4. Audrey Huang, Liu Leqi, Zachary C. Lipton, Kamyar Azizzadenesheli. Off-Policy Risk Assessment for Markov Decision Processes
Appeared at International Conference on Artificial Intelligence and Statistics (AISTATS) 2022.
[paper]
5. Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, Anima Anandkumar. Reinforcement Learning with Fast Stabilization in Linear Dynamical Systems
Appeared at International Conference on Artificial Intelligence and Statistics (AISTATS) 2022.
[paper]

6. Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, Anima Anandkumar. Model Learning Predictive Control in Nonlinear Dynamical Systems
Appeared at 60th IEEE Conference on Decision and Control(CDC) 2021.
[paper]
7. Audrey Huang, Leqi Liu, Zachary C. Lipton, Kamyar Azizzadenesheli. Off-Policy Risk Assessment in Contextual Bandits
Appeared at Neural Information Processing Systems (NeurIPS) 2021.
[paper]
8. Guanya Shi, Kamyar Azizzadenesheli, Soon-Jo Chung, Yisong Yue. Meta-Adaptive Nonlinear Control: Theory and Algorithms
Appeared at Neural Information Processing Systems (NeurIPS) 2021.
[paper]
9. Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, Anima Anandkumar. Finite-time System Identification and Adaptive Control in Autoregressive Exogenous Systems.
Appeared at Annual Conference on Learning for Dynamics and Control (L4DC), 2021.
[paper]
10. Soheil Esmailzadeh, Chiyu Max Jiang, Kamyar Azizzadenesheli, Karthik Kashinath, Mustafa Mustafa, Hamdi A Tchelepi, Philip Marcus, Mr Prabhat, Anima Anandkumar. A Deep Learning Based Physics Informed Continuous Spatio Temporal Super-Resolution Framework.
Bulletin of the American Physical Society, 2020.
[paper]
11. Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar. Fourier Neural Operator for Parametric Partial Differential Equations,
Appeared at International Conference on Learning Representations (ICLR) 2021.
[paper] [Code] [Blog Post]
12. Akella Ravi Tej, Kamyar Azizzadenesheli, Mohammad Ghavamzadeh, Yisong Yue, Anima Anandkumar. Deep Bayesian Quadrature Policy Optimization. *Appeared at AAAI Conference on Artificial Intelligence 2021 (AAAI-21) .*
[paper] [Code]
13. Manish Prajapat, Kamyar Azizzadenesheli, Alexander Liniger, Yisong Yue, Anima Anandkumar. Competitive Policy Optimization, 2020.
Appeared at The Conference on Uncertainty in Artificial Intelligence (UAI) 2021 .
[paper] [Code] [Package] [Page] [Blog]
14. Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar. Multipole Graph Neural Operator for Parametric Partial Differential Equations, 2020.
Appeared at Neural Information Processing Systems (NeurIPS) 2020.
[paper] [Code]
15. Chiyu (Max) Jiang, Soheil Esmailzadeh, Kamyar Azizzadenesheli, Karthik Kashinath, Mustafa Mustafa, Hamdi A. Tchelepi, Philip Marcus, Prabhat, Anima Anandkumar. MeshfreeFlowNet: A Physics-Constrained Deep Continuous Space-Time Super-Resolution Framework.
Appeared at SC20, The International Conference for High Performance Computing, Networking, Storage, and Analysis. Best Paper Finalist.
[paper] [Package] [Talk] [Page] [paper] [Code] [Colab]
16. Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, Animashree Anandkumar. Logarithmic Regret Bound in Partially Observable Linear Dynamical Systems, 2020.
Appeared at Neural Information Processing Systems (NeurIPS) 2020.
[paper]

17. Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, Animashree Anandkumar. Adaptive Control and Regret Minimization in Linear Quadratic Gaussian (LQG) Setting, 2021.
Appeared at American Control Conference ACC 2021 .
[paper]
18. Amy Zhang, Zachary C Lipton, Luis Pineda, Kamyar Azizzadenesheli, Animashree Anandkumar, Laurent Itti, Joelle Pineau, Tommaso Furlanello. Learning Causal State Representations of Partially Observable Environments,
Appeared at the Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM) 2019.
[paper]
19. Kamyar Azizzadenesheli. Maybe a few considerations in Reinforcement Learning Research?, 2019.
[paper]
20. Kamyar Azizzadenesheli, Anqi Liu, Fanny Yang, Animashree Anandkumar. Regularized Learning for Domain Adaptation under Label Shifts,
Appeared at International Conference on Learning Representations (ICLR) 2019.
[paper]
21. Jeremy Bernstein, Jiawei Zhao, Kamyar Azizzadenesheli, Anima Anandkumar. signSGD with Majority Vote is Communication Efficient and Fault Tolerant,
Appeared at International Conference on Learning Representations (ICLR) 2019.
[paper]
22. Guanya Shi, Xichen Shi, Michael O’Connell1, Rose Yu, Kamyar Azizzadenesheli, Animashree Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural Lander: Stable Drone Landing Control using Learned Dynamics,
Appeared at International Conference on Robotics and Automation (ICRA) 2019 .
[paper] [video]
23. Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, Anima Anandkumar. signSGD: Compressed Optimisation for Non-Convex Problems,
Appeared at International Conference on Machine Learning (ICML) 2018.
[paper]
24. Guneet S. Dhillon, Kamyar Azizzadenesheli, Jeremy D. Bernstein, Jean Kossaifi, Aran Khanna, Zachary C. Lipton, Animashree Anandkumar. Stochastic activation pruning for robust adversarial defense,
Appeared at International Conference on Learning Representations (ICLR) 2017.
[paper]
25. Kamyar Azizzadenesheli, Animashree Anandkumar. Efficient Exploration through Bayesian Deep Q-Networks, *Appeared at Information Theory and Application (ITA) 2018*.
[paper] [talk]
26. Kamyar Azizzadenesheli, Alessandro Lazaric, Anima Anandkumar. Reinforcement Learning in Rich Observation MDPs using Spectral Methods,
Appeared at Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM) 2017.
[paper]
27. Kamyar Azizzadenesheli, Alessandro Lazaric, Anima Anandkumar. Open Problem: Approximate Planning of POMDPs in the class of Memoryless Policies,
Appeared at Conference on Learning Theory (COLT) 2016 .
[paper] [talk]
28. Kamyar Azizzadenesheli, Alessandro Lazaric, Anima Anandkumar. Reinforcement Learning of POMDPs using Spectral Methods,

Appeared at Conference on Learning Theory (COLT) 2016 .
[paper] [talk]

Journal:

29. Mridul Agarwal, Vaneet Aggarwal, Kamyar Azizzadenesheli. Multi-Agent Multi-Armed Bandits with Limited Communication,
Appeared at Journal of Machine Learning Research (JMLR), 2022. [paper]
30. Michael O'Connell, Guanya Shi, Xichen Shi, Kamyar Azizzadenesheli, Anima Anandkumar, Yisong Yue, Soon-Jo Chung. Neural-Fly enables rapid learning for agile flight in strong winds,
Appeared at Science Robotics, 2022. [paper] [media] [demo] [package]
31. Gege Wen, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, Sally M. Benson. U-FNO - an enhanced Fourier neural operator based-deep learning model for multiphase flow,
Appeared at Advances in Water Resources, Elsevier, 2022. [paper] [software] [blog post] [media]
32. Yan Yang, Angela F. Gao, Jorge C. Castellanos, Zachary E. Ross, Kamyar Azizzadenesheli, Robert W. Clayton. Seismic wave propagation and inversion with Neural Operators
Appeared at The Seismic Record Journal, 2021.
[paper]
33. Burigede Liu, Nikola Kovachki, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, Andrew Stuart, Kaushik Bhattacharya. A learning-based multiscale method and its application to inelastic impact problems,
Appeared at Journal of the Mechanics and Physics of Solids, 2021.
[paper]
34. Kamyar Azizzadenesheli. Importance Weight Estimation and Generalization in Domain Adaptation under Label Shift.
Appeared at the IEEE Transactions on Pattern Analysis and Machine Intelligence, 2021
[paper]
35. K. Kashinath, M. Mustafa, A. Albert, J-L. Wu, C. Jiang, S. Esmailzadeh, Kamyar. Azizzadenesheli, R. Wang, A. Chattopadhyay, A. Singh, A. Manepalli, D. Chirila, R. Yu, R. Walters, B. White, H. Xiao, H. A. Tchelepi, P. Marcus, A. Anandkumar, P. Hassanzadeh and Prabhat. Physics-informed machine learning: case studies for weather and climate modelling,
Appeared at The Royal Society, 2021.
[paper]
36. Jonathan D. Smith, Zachary E. Ross, Kamyar Azizzadenesheli, Jack B. Muir. HypoSVI: Hypocenter inversion with Stein variational inference and Physics Informed Neural Networks.
Appeared at Geophysical Journal International, 2021.
37. Zachary E Ross, Daniel T Trugman, Kamyar Azizzadenesheli, Animashree Anandkumar. Directivity Modes of Earthquake Populations with Unsupervised Learning,
Appeared at the Journal of Geophysical Research 2020,
[paper]

Workshop:

38. William Wong, Liu Leqi, Audrey Huang, Kamyar Azizzadenesheli, Zachary C. Lipton. RiskyZoo: A Library for Risk-Sensitive Supervised Learning,
Appeared in the Responsible Decision Making in Dynamic Environments Decision Workshop at ICML 2022.
[paper]

39. Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, Anima Anandkumar. Physics-Informed Neural Operator for Learning Partial Differential Equations, 2021
Appeared in the Continuous Time Methods Workshop at ICML 2022.
[paper]
40. Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, Pedram Hasanzadeh, Karthik Kashinath, Animashree Anandkumar. FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators,
Appeared in the AI for Earth and Space Science at ICLR 2022.
[paper] [news]
41. Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar. Neural Operator: Graph Kernel Network for Partial Differential Equations,
Appeared at International Conference on Learning Representations (ICLR) 2020 Workshop.
[paper] [Code] [Blog Post]
42. Kamyar Azizzadenesheli, Brandon Yang, Weitang Liu, Zachary C Lipton, Animashree Anandkumar. Surprising Negative Results for Generative Adversarial Tree Search,
Appeared at International Conference on Machine Learning (ICML) 2018 workshop.
[paper]
43. Jeremy Bernstein, Kamyar Azizzadenesheli, Yu-Xiang Wang, Anima Anandkumar. Compression by the signs: distributed learning is a two-way street,
Appeared at International Conference on Learning Representations (ICLR) 2018 Workshop.
[paper]
44. Zachary C. Lipton, Kamyar Azizzadenesheli, Abhishek Kumar, Lihong Li, Jianfeng Gao, Li Deng. Combating Reinforcement Learning’s Sisyphian Curse with Intrinsic Fear,
Appeared at Neural Information Processing Systems (NeurIPS) 2016 Workshop.
[paper]
45. Kamyar Azizzadenesheli, Alessandro Lazaric, Anima Anandkumar. Experimental paper: Reinforcement Learning of POMDPs using Spectral Methods,
Appeared at Neural Information Processing Systems (NeurIPS) 2016 Workshop.
[paper]

Under review:

46. Victor D. Dorobantu, Kamyar Azizzadenesheli, and Yisong Yue. Compactly Restrictable Metric Policy Optimization Problems, 2022
[paper]
47. Sahin Lale, Yuanyuan Shi, Guannan Qu, Kamyar Azizzadenesheli, Adam Wierman, Anima Anandkumar. KCRL: Krasovskii-Constrained Reinforcement Learning with Guaranteed Stability in Nonlinear Dynamical Systems, 2022
[paper]
48. Qian Zhang, Anuran Makur, Kamyar Azizzadenesheli. Functional Linear Regression of CDFs, 2022
[paper]
49. Abhijeet Vyas, Kamyar Azizzadenesheli. Competitive Gradient Optimization, 2022
[paper] [code]
50. Md Ashiqur Rahman, Manuel A. Florez, Anima Anandkumar, Zachary E. Ross, Kamyar Azizzadenesheli. Generative Adversarial Neural Operators, 2022
[paper] [code]

51. Md Ashiqur Rahman, Zachary E. Ross, Kamyar Azizzadenesheli. U-NO: U-shaped Neural Operators, 2022
[paper] [code] [colab]
52. Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, Anima Anandkumar. Physics-Informed Neural Operator for Learning Partial Differential Equations, 2021
[paper]
53. Nikola Kovachki, Zongyi Li, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar. Neural Operator: Learning Maps Between Function Spaces, 2021
[paper]
54. Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar. Markov Neural Operators for Learning Chaotic Systems, 2021
[paper]
55. Jafar Abbaszadeh Chekan, Kamyar Azizzadenesheli, Cedric Langbort. Joint Stabilization and Regret Minimization through Switching in Systems with Actuator Redundancy, 2021
[paper]
56. Audrey Huang, Liu Leqi, Zachary C. Lipton, Kamyar Azizzadenesheli. On the Convergence and Optimality of Policy Gradient for Coherent Risk , 2021
[paper]
57. Kamyar Azizzadenesheli, Yisong Yue, Animashree Anandkumar. Policy Gradient in Partially Observable Environments: Approximation and Convergence, 2020.
[paper]
58. Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, Animashree Anandkumar. Stochastic Linear Bandits with Hidden Low Rank Structure, *Early version appeared at Neural Information Processing Systems (NeurIPS) 2019 Workshop.*
[paper]
59. Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, Animashree Anandkumar. Regret Minimization in Partially Observable Linear Quadratic Control, *Early version appeared at Neural Information Processing Systems (NeurIPS) 2019 Workshop.*
[paper]

Thesis

- * Kamyar Azizzadenesheli. Reinforcement Learning in Structured and Partially Observable Environments, 2019. [Thesis]

Talks and Interviews

*some contains links

- **Invited Talk:** “Neural Operators, Learn to Solve Partial Differential Equations”, Alan Turing Institute (*March, 2022*)
— UK’s national data science & AI centre
- **Invited Talk:** “Neural Operators, Learn to Solve Partial Differential Equations”, The University of Sydney (*March, 2022*)
- **Invited Talk:** “Neural Operators, Learn to Solve Partial Differential Equations”, Baidu Research (*March, 2022*)

- **Podcast:** This Week in Machine Learning AI (TWIML AI) by Sam Charrington (*Feb, 2022*)
- **Invited Talk:** “Neural Operators, Learn to Solve Partial Differential Equations”, Gypsum Seminar (*Nov, 2021*)
— The GeodesY and geoPhysics Seminar of the Upper Midwest.
- **Invited Talk:** “A Crash Course on Neural Operators”, British Mathematical Colloquium (BMC) (*April, 2021*)
— The British Mathematical Colloquium (BMC) is the largest pure mathematical conference to be held annually in the UK.
- **Podcast:** TalkRL by Robin R Chauhan (*October, 2019*)
- **Invited Talk:** “Statistical Learning Theory in Practice”, UC. San Diego, CA, USA (*March, 2019*)
— The THEORY SEMINAR at UC San Diego, department of Computer Science and Engineering, mainly organized by Shachar Lovett.
- **Invited Talk:** Reinforcement Learning in Structured Environments, UT Austin, Tx, USA (*October, 2018*)
— An invited talk hosted by Alexandros G Dimakis of UT Austin,
- **Podcast:** This Week in Machine Learning AI (TWIML AI) by Sam Charrington (*August, 2018*)
- **Talk:** “Exploration in Reinforcement Learning workshop”, International Conference on Machine Learning (ICML) 2018, Stockholm, Sweden (*July, 2018*) — A talk on my recent work on Exploration in Reinforcement Learning
- **Invited Talk:** “Bayesian deep Reinforcement Learning”, Stanford, CA, USA (*May, 2018*)
— An invited talk hosted by Machine Learning group at Stanford
- **Invited Talk:** “Efficient Exploration Through Bayesian Deep Q-Networks ”, Information Theory and Applications Workshop, (ITA), San Diego, CA, USA (*February, 2018*)
- **Invited Talk:** “Reinforcement Learning in rich observable environment”, Stanford, CA, USA (*February, 2018*)
— An invited talk hosted by Machine Learning group at Stanford
- **Talk:** at MLTrainon workshop, NeurIPS 2017, Long Beach, CA, USA (*December, 2017*)
A talk on my recent work on Exploration in Reinforcement Learning
- **Invited talk:** at Amazon AWS, Palo Alto, CA, USA (*April, 2017*)
- **Invited Talk:** at “Interactive Learning” workshop, Simons Institute, UC Berkeley, CA, USA (*Feb, 2017*)
- **Invited Talk:** “RL of Partially Observable Environment”, Caltech, CA, USA (*Jan, 2017*)
— An invited talk hosted by Yisong Yue of Caltech,
- **Talk:** at ”Open Problem” , *Colombia University, (COLT), (June, 2016)*
- **Talk:** at ”Bandit and Reinforcement Learning”, *Colombia University, (COLT), (June, 2016)*

Certifications

- **Certified:** in ”Topological Data Analysis” *NSF-CBMS, (June, 2016)*

Media Coverage

- *The Year in Math and Computer Science* – by "Quanta Magazine", (Dec, 2021)
- *Latest Neural Nets Solve World's Hardest Equations Faster Than Ever Before* – by "Quanta Magazine", (April, 2021)
- *AI has cracked a key mathematical puzzle for understanding our world* – by "MIT Technology Review", (October, 2020)
- *Caltech Open-Sources FNO: A Deep Learning Method For Solving PDEs (Partial differential equations)* – by "MarkTechPost", (January, 2021)
- *Caltech Open-Sources AI for Solving Partial Differential Equations* – by "InfoQ", (Dec, 2020)
- *"Neural Lander" Uses AI to Land Drones Smoothly* – by "Caltech news", (May, 2019)
- *Drone Uses AI to Overcome Turbulence* – by "dronebelow", (May, 2019)
- *Neural Network Provides Stable Drone Landing* – by "dronebelow", (November, 2018)

Current Students

Ph.D. students:

- Abhijeet Vyas, *Topic: Online Meta Learning*
- Mirza Masfiquir Rahman, *Topic: Neural Operator*
- Md Ashiqur Rahman, *Topic: Domain Adaptation*
- Qifan Zhang, *Topic: Risk Assessment*

Masters students:

- Vivek Gupta, *Topic: Bayesian Decision Making*

Undergraduate students:

- David Kim, *Topic: Non-Convex problems in Linear Algebra*
- Shatayu Kulkarni, *Topic: Algorithmic Games*

Graduated Students

Graduate students:

-

Undergraduate students:

-

Service on M.Sc./Ph.D. Thesis Committees

Masters:

- Xia Tian, Purdue University, *(To be held)*
- Siddharth Divi, Purdue University, *(June, 2021)*
- Manish Prajapat, ETH Zürich, *(Dec, 2020)*

Ph.D.:

- Jingya Wang, Purdue University, *(To be held)*
- Fnu Chandra Mouli Sekar, Purdue University, *(To be held)*
- Xu Quiling, Purdue University, *(To be held)*
- Nan Jiang, Purdue University, *(To be held)*
- Yuantong Li, Purdue University, *(To be held)*
- Xin Cheng, Purdue University, *(To be held)*
- Tinghan Yang, Purdue University, *(To be held)*
- Victor Dorobantu, Caltech, *(To be held)*

Teaching Assignments

- CS59300-RL Spring 2021 Reinforcement Learning
- CS47100-AI Fall 2021 Artificial Intelligence
- CS59000-AI Spring 2021 Artificial Intelligence (Newly developed course)
- CS59000-RL Fall 2020 Reinforcement Learning (Newly developed course)

Administration Services at the Purdue University

- Updating the graduate student admissions website to emphasize on TOEFL requirement of 80 instead of 100.
- Leading the attempt to make the graduate Artificial Intelligence as a core course
- Leading the effort to pass an important campus wide admission rule that removes the requirement for official transcripts form graduate student admission requirements
- Participate in faculty hiring in 2020
- Member of faculty hiring team in 2021
- Leading the effort to change the name of the artificial intelligence cohort and its descriptions
- Leading the effort to update the area study of our Master program.

Professional Services

Reviewer (Program Chair member):

- Journal of Machine Learning Research (JMLR)
- Association for Computing Machinery (ACM) SIGMETRICS
- International Journal of Computational Linguistics and Applications (IJCLA)
- Association for the Advancement of Artificial Intelligence (AAAI)
- Conference on Learning Theory (COLT)
- International Conference on Learning Representations (ICLR)
- International Conference on Machine Learning (ICML)
- International Conference on Machine Learning, Optimization, and Data Science (LOD)
- Conference on Neural Information Processing Systems (NeurIPS)
- Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM)
- Institute of Electrical and Electronics Engineers (IEEE)
- Springer Nature
- Artificial Intelligence and Statistics (AISTAT)
- Conference on Decision and Control (CDC)
- The Conference on Uncertainty in Artificial Intelligence (UAI)
- Journal of Computational Physics

Area Chair:

- “Area Chair” Conference on Neural Information Processing Systems (NeurIPS)
- “Senior Program Committee” Association for the Advancement of Artificial Intelligence (AAAI)

Organizer:

- Co-Organizer: “Challenges in Data-Driven, Optimization-Based Control” at Control Systems Society (CDC) 2022 an Invited Session with Cedric Langbort

Mentees

Graduate:

- Zongyi Li (Caltech), Manish Kumar Bera (ETH), Sahin Lale (Caltech), Jeremy Bernstein (Caltech)

Undergraduate:

- Chase Blagden (Caltech), Damon Lin (Caltech), Videh Raj Nema (IIT Madras), Guneet S Dhillon (UT Austin, Amazon), Weitang Liu (UC, Davis), Brandon Yang (Stanford, Google Brain), Jiawei Zhao (Nanjing University, Caltech), Hongjie Chen (Shanghai Jiao Tong University, Caltech), Zongyi Li (Washington University in St. Louis, Caltech), Vivek Bharadwaj (Caltech), Akshay R. Vegesna (Caltech), Ravi Tej Akella (IIT Roorkee), Manish Prajapat (ETH), Albert Zhai (Caltech), Abhijeet Vyas (IIT), Jihwan Bae (Gwangju Institute of Science and Technology), Sirui Li (MIT), Saturnin J. Pugnet (Caltech)