

Ching-An Cheng

Principal Researcher
AI Frontiers, Microsoft Research

chinganc@microsoft.com
www.chinganc.com

Building 99
14850 NE 36th St
Redmond, WA 98052

Research Interests

Theory, algorithms, and applications of learning in sequential decision making, agents, and robotics.

Learning from Feedback · Generative Optimization · Reinforcement Learning · Imitation Learning · Online Learning · LLM Agent · Embodied Agent · Manipulation

Work Experiences

2024–
2020–2024 | **Principal Researcher, Microsoft Research, Redmond, USA**
Senior Researcher, Microsoft Research, Redmond, USA

Education

2015–2019 | Ph.D. in Robotics, Georgia Institute of Technology, USA
Thesis: *Efficient and Principled Robot Learning: Theory and Algorithms*
Advisor: Byron Boots
Committee: Seth Hutchinson, Geoff Gordon, Evangelos A. Theodorou, Karen Liu

2011–2013 | M.S. in Mechanical Engineering, National Taiwan University, Taiwan
Thesis: *Robot Dynamics Learning and Human-Robot Interaction*
Advisor: Han-Pang Huang

2007–2011 | B.S. in Mechanical Engineering, National Taiwan University, Taiwan
B.S. in Electrical Engineering, National Taiwan University, Taiwan

Awards and Honors

2022 | **Outstanding Paper Award (Runner-Up), ICML 2022**

2019 | **Best Paper Award, OptRL Workshop @ NeurIPS 2019**

2019 | **Best Student Paper Award, RSS 2019**

2019 | **Best Systems Paper Award (Finalist), RSS 2019**

2019 | **Google PhD Fellowship, Machine Learning** (declined)

2019 | **Nvidia PhD Fellowship**

2018 | **Best Systems Paper Award (Finalist), RSS 2018**

2018 | **Best Paper Award, AISTATS 2018**

2015 | Government Scholarship to Study Abroad, Ministry of Education, Taiwan

2014 | Excellent Project Award, Industrial Technology Research Institute, Taiwan

2010 | Cheng-Tai Scholarship, Taiwan

Publications

- 2024 C.-A. Cheng*, A. Nie*, and A. Swaminathan*. Trace is the next autodiff: Generative optimization with rich feedback, execution traces, and llms. *Advances in Neural Information Processing Systems*, 2024 (**Open-sourced Project. Featured at NeurIPS 2024 Expo**)
- 2024 Y. Fan, J. Li, A. Swaminathan, A. Modi, and C.-A. Cheng. How to solve contextual goal-oriented problems with offline datasets? *Advances in Neural Information Processing Systems*, 2024
- 2024 R. Zheng, C.-A. Cheng, H. Daumé III, F. Huang, and A. Kolobov. Prise: Llm-style sequence compression for learning temporal action abstractions in control. *International Conference on Machine Learning*, 2024 (**Oral (<2%)**)
- 2024 S. Geng, A. Pacchiano, A. Kolobov, and C.-A. Cheng. Improving offline rl by blending heuristics. *International Conference on Learning Representations*, 2024 (**Spotlight (5%)**)
- 2023 A. Li, D. Misra, A. Kolobov, and C.-A. Cheng. Survival instinct in offline reinforcement learning. *Conference on Neural Information Processing Systems*, 2023 (**Spotlight (<3%)**)
- 2023 M. Bhardwaj, T. Xie, B. Boots, N. Jiang, and C.-A. Cheng. Adversarial model for offline reinforcement learning. *Conference on Neural Information Processing Systems*, 2023
- 2023 G. Thomas, C.-A. Cheng, R. Loynd, V. Vineet, M. Jalobeanu, and A. Kolobov. Plex: Making the most of the available data for robotic manipulation pretraining. *Conference on Robot Learning*, 2023
- 2023 V. Myers, A. He, K. Fang, H. Walke, P. Hansen-Estruch, C.-A. Cheng, M. Jalobeanu, A. Kolobov, A. Dragan, and S. Levine. Goal representations for instruction following: A semi-supervised language interface to control. *Conference on Robot Learning*, 2023
- 2023 A. Li, B. Boots, and C.-A. Cheng. Mahalo: Unifying offline reinforcement learning and imitation learning from observations. *International Conference on Machine Learning*, 2023
- 2023 H.-A. Nguyen and C.-A. Cheng. Provable reset-free reinforcement learning by no-regret reduction. *International Conference on Machine Learning*, 2023
- 2023 S. R. Sinclair, F. Frujeri, C.-A. Cheng, L. Marshall, H. Barbalho, J. Li, J. Neville, I. Menache, and A. Swaminathan. Hindsight learning for mdps with exogenous inputs. *International Conference on Machine Learning*, 2023
- 2023 S. Amani, L. F. Yang, and C.-A. Cheng. Provably efficient lifelong reinforcement learning with linear function representation. *International Conference on Learning Representations*, 2023
- 2022 N. Wagener, A. Kolobov, F. V. Frujeri, R. Loynd, C.-A. Cheng, and M. Hausknecht. Mocapact: A multi-task dataset for simulated humanoid control. *Neural Information Processing Systems Datasets and Benchmarks Track*, 2022
- 2022 C.-A. Cheng*, T. Xie*, N. Jiang, and A. Agarwal. Adversarially trained actor critic for offline reinforcement learning. *International Conference on Machine Learning*, 2022 (**Outstanding Paper Award, Runner-Up**) (*equal contribution)
- 2021 T. Xie, C.-A. Cheng, N. Jiang, P. Mineiro, and A. Agarwal. Bellman-consistent pessimism for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 34:6683–6694, 2021 (**Oral Presentation (<1%)**) (*equal contribution)
- 2021 C.-A. Cheng, A. Kolobov, and A. Swaminathan. Heuristic-guided reinforcement learning. *Advances in Neural Information Processing Systems*, 34:13550–13563, 2021
- 2021 N. C. Wagener, B. Boots, and C.-A. Cheng. Safe reinforcement learning using advantage-based intervention. In *International Conference on Machine Learning*, pages 10630–10640. PMLR, 2021

- 2021 A. Zanette, C.-A. Cheng, and A. Agarwal. Cautiously optimistic policy optimization and exploration with linear function approximation. In *Conference on Learning Theory*, pages 4473–4525. PMLR, 2021
- 2021 A. Li*, C.-A. Cheng*, M. A. Rana, M. Xie, K. Van Wyk, N. Ratliff, and B. Boots. RMP²: A structured composable policy class for robot learning. *Robotics: Science and Systems*, 2021 (*equal contribution)
- 2021 X. Yan, B. Boots, and C.-A. Cheng. Explaining fast improvement in online imitation learning. In *Uncertainty in Artificial Intelligence*, pages 1874–1884. PMLR, 2021
- 2021 C.-A. Cheng, M. Mukadam, J. Issac, S. Birchfield, D. Fox, B. Boots, and N. Ratliff. RMPflow: A geometric framework for generation of multitask motion policies. *IEEE Transactions on Automation Science and Engineering*, 18(3):968–987, 2021
- 2020 C.-A. Cheng, A. Kolobov, and A. Agarwal. Policy improvement via imitation of multiple oracles. *Advances in Neural Information Processing Systems*, 33:5587–5598, 2020 (**Spotlight Talk (<3%)**)
- 2020 A. Rahimi*, A. Shaban*, C.-A. Cheng*, B. Boots, and R. Hartley. Intra order-preserving functions for calibration of multi-class neural networks. *Advances in Neural Information Processing Systems*, 33:13456–13467, 2020 (*equal contribution)
- 2020 C.-A. Cheng, R. T. Combes, B. Boots, and G. Gordon. A reduction from reinforcement learning to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3514–3524. PMLR, 2020
- 2020 C.-A. Cheng*, J. Lee*, K. Goldberg, and B. Boots. Online learning with continuous variations: Dynamic regret and reductions. In *International Conference on Artificial Intelligence and Statistics*, pages 2218–2228. PMLR, 2020 (*equal contribution)
- 2020 B. Wingo, C.-A. Cheng, M. Murtaza, M. Zafar, and S. Hutchinson. Extending Riemannian motion policies to a class of underactuated wheeled-inverted-pendulum robots. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3967–3973. IEEE, 2020
- 2019 Y. Pan, C.-A. Cheng, K. Saigol, K. Lee, X. Yan, E. A. Theodorou, and B. Boots. Imitation learning for agile autonomous driving. *The International Journal of Robotics Research*, 39(2-3):286–302, 2020
- 2019 C.-A. Cheng, X. Yan, and B. Boots. Trajectory-wise control variates for variance reduction in policy gradient methods. In *Conference on Robot Learning*, pages 1379–1394. PMLR, 2020 (*equal contribution)
- 2019 M. Mukadam, C.-A. Cheng, D. Fox, B. Boots, and N. Ratliff. Riemannian motion policy fusion through learnable lyapunov function reshaping. In *Conference on robot learning*, pages 204–219. PMLR, 2020
- 2019 A. Li, C.-A. Cheng, B. Boots, and M. Egerstedt. Stable, concurrent controller composition for multi-objective robotic tasks. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 1144–1151. IEEE, 2019
- 2019 N. Wagener*, C.-A. Cheng*, J. Sacks, and B. Boots. An online learning approach to model predictive control. *Robotics: Science and Systems*, 2019 (*equal contribution) **Best Student Paper Award; Best Systems Paper Award, Finalist**
- 2019 C.-A. Cheng, X. Yan, N. Ratliff, and B. Boots. Predictor-corrector policy optimization. In *International Conference on Machine Learning*, pages 1151–1161. PMLR, 2019 (**Long Talk (<5%)**)
- 2019 C.-A. Cheng, X. Yan, E. Theodorou, and B. Boots. Accelerating imitation learning with predictive models. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3187–3196. PMLR, 2019
- 2019 A. Shaban*, C.-A. Cheng*, N. Hatch, and B. Boots. Truncated back-propagation for bilevel optimization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1723–1732. PMLR, 2019 (*equal contribution)
- 2019 Z.-H. Kang, C.-A. Cheng, and H.-P. Huang. A singularity handling algorithm based on operational space control for six-degree-of-freedom anthropomorphic manipulators. *International Journal of Advanced Robotic Systems*, 16(3), 2019

- 2018 C.-A. Cheng, M. Mukadam, J. Issac, S. Birchfield, D. Fox, B. Boots, and N. Ratliff. RMPflow: A computational graph for automatic motion policy generation. In *International Workshop on the Algorithmic Foundations of Robotics*, pages 441–457. Springer, 2018
- 2018 H. Salimbeni*, C.-A. Cheng*, B. Boots, and M. Deisenroth. Orthogonally decoupled variational Gaussian processes. *Conference on Neural Information Processing Systems*, 2018 (*equal contribution)
- 2018 C.-A. Cheng, X. Yan, N. Wagener, and B. Boots. Fast policy learning using imitation and reinforcement. *Conference on Uncertainty in Artificial Intelligence*, 2018 (**Plenary Presentation (<9%)**)
- 2018 Y. Pan, C.-A. Cheng, K. Saigol, K. Lee, X. Yan, E. Theodorou, and B. Boots. Agile off-road autonomous driving using end-to-end deep imitation learning. *Robotics: Science and Systems*, 2018 **Best Systems Paper Award, Finalist**
- 2018 C.-A. Cheng and B. Boots. Convergence of value aggregation for imitation learning. In *International Conference on Artificial Intelligence and Statistics*, volume 84, pages 1801–1809, 2018 **Best Paper Award**
- 2018 J. L. Molnar, C.-A. Cheng, L. O. Tiziani, B. Boots, and F. L. Hammond. Optical sensing and control methods for soft pneumatically actuated robotic manipulators. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3355–3362. IEEE, 2018
- 2017 C.-A. Cheng and B. Boots. Variational inference for Gaussian process models with linear complexity. In *Advances in Neural Information Processing Systems*, 2017
- 2017 M. Mukadam, C.-A. Cheng, X. Yan, and B. Boots. Approximately optimal continuous-time motion planning and control via probabilistic inference. In *IEEE International Conference on Robotics and Automation*, pages 664–671, 2017
- 2016 C.-A. Cheng and H.-P. Huang. Learn the Lagrangian: A vector-valued RKHS approach to identifying Lagrangian systems. *IEEE Transactions on Cybernetics*, 46(12):3247–3258, 2016
- 2016 S.-Y. Lo, C.-A. Cheng, and H.-P. Huang. Virtual impedance control for safe human-robot interaction. *Journal of Intelligent & Robotic Systems*, 82(1):3, 2016
- 2016 C.-A. Cheng, H.-P. Huang, H.-K. Hsu, W.-Z. Lai, and C.-C. Cheng. Learning the inverse dynamics of robotic manipulators in structured reproducing kernel Hilbert space. *IEEE Transactions on Cybernetics*, 46(7):1691–1703, 2016
- 2016 C.-A. Cheng and B. Boots. Incremental variational sparse Gaussian process regression. In *Advances in Neural Information Processing Systems*, pages 4410–4418, 2016
- 2015 H.-P. Huang, Y.-H. Liu, W.-Z. Lin, Z.-H. Kang, C.-A. Cheng, and T.-H. Huang. Development of a p300 bci and design of an elastic mechanism for a rehabilitation robot. *International Journal of Automation and Smart Technology*, 5(2):91–100, 2015
- 2015 C.-H. Chang, H.-P. Huang, H.-K. Hsu, and C.-A. Cheng. Humanoid robot push-recovery strategy based on cmp criterion and angular momentum regulation. In *IEEE International Conference on Advanced Intelligent Mechatronics*, pages 761–766. IEEE, 2015
- 2015 M.-B. Huang, H.-P. Huang, C.-C. Cheng, and C.-A. Cheng. Efficient grasp synthesis and control strategy for robot hand-arm system. In *IEEE International Conference on Automation Science and Engineering*, pages 1256–1257. IEEE, 2015
- 2013 C.-A. Cheng, H.-P. Huang, H.-K. Hsu, W.-Z. Lai, C.-C. Cheng, and Y.-C. Li. Identification of the inverse dynamics of robot manipulators with the structured kernel. In *International Automatic Control Conference*, pages 266–271. IEEE, 2013
- 2013 T.-H. Huang, C.-A. Cheng, and H.-P. Huang. Self-learning assistive exoskeleton with sliding mode admittance control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 698–703. IEEE, 2013
- 2013 C.-A. Cheng, T.-H. Huang, and H.-P. Huang. Bayesian human intention estimator for exoskeleton system. In *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, pages 465–470. IEEE, 2013

- 2012 T.-H. Huang, H.-P. Huang, C.-A. Cheng, J.-Y. Kuan, P.-T. Lee, and S.-Y. Huang. Design of a new hybrid control and knee orthosis for human walking and rehabilitation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3653–3658. IEEE, 2012
- 2012 H.-P. Huang, Y.-H. Liu, T.-H. Huang, Z.-H. Kang, W.-Z. Lin, W. Ching-Ping, and C.-A. Cheng. Development of a brain-machine interface for motor imagination task. In *International Conference on Automation Technology*, 2012
- 2011 Y.-H. Liu, C.-A. Cheng, and H.-P. Huang. Novel feature of the EEG based motor imagery BCI system: Degree of imagery. In *International Conference on System Science and Engineering*, pages 515–520. IEEE, 2011
- 2010 C.-A. Cheng, Y.-H. Liu, and H.-P. Huang. Motor imagery recognition for brain-computer interfaces using Hilbert-Huang transform and effective event-related-desynchronization features. In *CSME National Conference*, 2010

Workshop

- 2024 C.-A. Cheng*, A. Nie*, and A. Swaminathan*. Trace is the new autodiff—unlocking efficient optimization of computational workflows. *ICML 2024 AutoRL Workshop*, 2024
- 2023 A. Nie, C.-A. Cheng, A. Kolobov, and A. Swaminathan. Importance of directional feedback for llm-based optimizers. *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023
- 2023 Y. Fan, J. Li, A. Swaminathan, A. Modi, and C.-A. Cheng. Simple data sharing for multi-task goal-oriented problems. *Goal-Conditioned Reinforcement Learning Workshop at NeurIPS 2023*, 2023
- 2023 H. Liu, A. Chen, Y. Zhu, A. Swaminathan, A. Kolobov, and C.-A. Cheng. Interactive robot learning from verbal corrections. *Workshop on Language and Robot Learning at CoRL 2023*, 2023
- 2023 R. Zheng, C.-A. Cheng, H. Furong, and A. Kolobov. Learning multi-task action abstractions as sequence compression problem. *CoRL 2023 Workshop on Pre-training for Robot Learning*, 2023 **Spotlight**
- 2023 A. Li, D. Misra, A. Kolobov, and C.-A. Cheng. Survival instinct in offline reinforcement learning and implicit human bias in data. *Interactive Learning with Implicit Human Feedback Workshop at ICML 2023*, 2023 **Oral Presentation**
- 2023 H.-A. Nguyen and C.-A. Cheng. Provable reset-free reinforcement learning by no-regret reduction. *AAAI 2023 Reinforcement Learning Ready for Production Workshop*, 2023
- 2022 T. Xie, M. Bhardwaj, N. Jiang, and C.-A. Cheng. Armor: A model-based framework for improving arbitrary baseline policies with offline data. *NeurIPS 2022 Offline RL Workshop*, 2022
- 2022 G. Thomas, A. Kolobov, C.-A. Cheng, V. Vineet, and M. Jalobeanu. Heetr: Pre-training for robotic manipulation on heteromodal data. In *CoRL 2022 Workshop on Pre-training Robot Learning*, 2022
- 2020 A. Li*, C.-A. Cheng*, M. A. Rana, N. Ratliff, and B. Boots. RMP²: A differentiable policy class for robotic systems with control-theoretic guarantees. *NeurIPS 2020 3rd Robot Learning Workshop*, 2020
- 2019 J. Lee*, C.-A. Cheng*, K. Goldberg, and B. Boots. Continuous online learning and new insights to online imitation learning. *NeurIPS 2019 Optimization Foundations of Reinforcement Learning Workshop*, 2019 (*equal contribution) **Best Paper Award**
- 2019 C.-A. Cheng*, X. Yan*, and B. Boots. Trajectory-wise control variates for variance reduction in policy gradient method. *NeurIPS 2019 Optimization Foundations of Reinforcement Learning Workshop*, 2019
- 2018 C.-A. Cheng, X. Yan, N. Ratliff, and B. Boots. Predictor-corrector policy optimization. *Deep Reinforcement Learning Workshop NeurIPS*, 2018

- 2017 Y. Pan, C.-A. Cheng, K. Saigol, K. Lee, X. Yan, E. Theodorou, and B. Boots. Learning deep neural network control policies for agile off-road autonomous driving. *The NIPS Deep Reinforcement Learning Symposium*, 2017
- 2017 C.-A. Cheng and B. Boots. Convergence of value aggregation for imitation learning. In *The NIPS Deep Reinforcement Learning Symposium*, 2017
- 2016 C.-A. Cheng and B. Boots. Incremental variational sparse Gaussian process regression. In *NIPS Workshop on Adaptive and Scalable Nonparametric Methods in Machine Learning*, 2016

Invited Talks

- 2024 “Trace”-ing the Path to Self-adapting AI Agents,” Robotics Colloquium, University of Washington (host: Abhishek Gupta)
- 2024 “Trace”-ing the Path to Self-adapting AI Agents,” AutoGen community, (host: Qingyun Wu)
- 2024 “Trace”-ing the Path to Self-adapting AI Agents,” Cortex AI Research Talk Series, (host: Ade Famoti)
- 2023 “ARMOR: A Model-based Framework for Improving Arbitrary Baseline Policies with Offline Data,” RL Theory Seminar (host: Gergely Neu)
- 2022 “A Game Theoretic Approach to Offline Reinforcement Learning,” Workshop on Structure of Constraints in Sequential Decision-Making, Simons Institute, USA
- 2020 “Efficient Policy Optimization by Online Imitation Learning,” MSR AI Seminar, Microsoft Research, USA
- 2020 “Trajectory-wise Control Variates for Policy Gradient Methods,” GTC, Nvidia, USA
- 2019 “RMPflow: A Geometric Framework for Policy Fusion,” Fanuc, USA (host: Hsien-Chung Lin)
- 2019 “Learning from Past Mistakes and Future Predictions for Sequential Decision Making,” RIKEN, Osaka, Japan (host: Yoshinobu Kawahara)
- 2019 “Learning to Optimize,” NTU, Taipei, Taiwan (host: Han-Pang Huang)
- 2018 “Policy Optimization as Predictable Online Learning Problems,” Microsoft Research, Montreal, Canada (host: Geoffrey Gordon)
- 2018 “Policy Optimization as Predictable Online Learning Problems: Imitation Learning and Beyond,” Microsoft Research, Redmond, USA (host: Andrey Kolobov)

Service

- Reviewer:
JMLR · NATURE · JAIR · IEEE RA-L · IEEE TPAMI · IEEE SPL · JINT · MACH · IEEE Trans SMC Systems · NeurIPS · ICML · COLT · ICLR · AAAI · RSS · ICRA · AIM · IROS · WAFR
- Area Chair
ICLR
- Organizer:
Microsoft Research Summit (RL track) 2021
MSR Reinforcement Learning Day 2021
ICML 2019 Workshop: Real-World Sequential Decision Making: Reinforcement Learning and Beyond

Mentorship

- Interns: Wanqiao Xu (PhD Student, Stanford) Kaiwen Wang (PhD Student, Cornell Tech) Huihan Liu (PhD Student, University of Texas at Austin) Ruijie Zheng (PhD Student, University of Maryland) Allen Nie (PhD Student, Stanford) Sumedh A Sontakke (PhD Student, University of Southern

California) Ying Fan (PhD Student, UW-Madison) Sinong Geng (PhD Student, Princeton) Hoai-An Nguyen (Undergraduate Student, Rutgers) Garrett Thomas (PhD Student, Stanford) Sean Sinclair (PhD Student, Cornell) Nolan Wagener (PhD Student, Georgia Tech) Tengyang Xie (PhD Student, UIUC) Andrea Zanette (PhD Student, Stanford)

- Collaborators: Aditya Soni (Research Fellow, Microsoft) Anqi Li (PhD Student, UW) Sanae Amani (PhD Student, UCLA) Mohak Bhardwaj (PhD Student, UW) Tengyang Xie (PhD Student, UIUC) Jonathan Lee (Undergraduate Student, UC Berkeley) Bruce Wingo (PhD Student, Georgia Tech)

Teaching Experience

- | | |
|-------------|---|
| 2018 | Teaching Assistant in Statistical Techniques in Robotics (CS8803, Byron Boots), Georgia Institute of Technology, USA (lectured “Gaussian Process Regression”) |
| 2017 | Teaching Assistant in Statistical Techniques in Robotics (CS8803, Byron Boots), Georgia Institute of Technology, USA (lectured “Inference of Graphical Model”) |
| 2013 | Lectured “Introduction to Machine Learning” in Intelligent Control (ME7144, Han-Pang Huang), National Taiwan University, Taiwan |
| 2012 & 2013 | Lectured induction training of Robotics Laboratory in linear algebra, matrix theory, convex optimization, machine learning, Matlab, and C++, National Taiwan University, Taiwan |

Skills

Python, Pytorch, Tensorflow, LLM, \LaTeX

Languages

Chinese (native), English (fluent), Japanese (basic)