

# Ming Yin

✉ my0049@princeton.edu | 🌐 | 📧 | 🐦 | 📱 +1 (805)-637-8752

## EDUCATION & APPOINTMENT

---

Princeton University, Postdoctoral Associate in Electrical and Computer Engineering	2023-
University of California, Santa Barbara, Ph.D. in Computer Science	2019-23
University of California, Santa Barbara, Ph.D. in Statistics and Applied Probability	2016-22
University of Science and Technology of China, B.S. in Applied Mathematics	2012-16

## RESEARCH INTERESTS

---

My research spans theory, methodology, and applications of reinforcement learning (RL) and generative artificial intelligence (GenAI). Concretely, I am interested in:

- Machine learning theory, with a particular interest in understanding the fundamental limits of reinforcement learning and bandits.
- Developing algorithms for GenAI models that enhance efficiency, robustness, and adaptivity.
- AI applications in biotechnology, healthcare, and real-world operational problems.

## PREPRINTS AND WORKSHOPS [ \* DENOTES EQUAL CONTRIBUTION ]

---

### **MATH-Perturb: Benchmarking LLMs' Math Reasoning Abilities against Hard Perturbations**

Kaixuan Huang, Jiacheng Guo, Zihao Li, Xiang Ji, Jiawei Ge, Wenzhe Li, Yingqing Guo, Tianle Cai, Hui Yuan, Runzhe Wang, Yue Wu, **Ming Yin**, Shange Tang, Yangsibo Huang, Chi Jin, Xinyun Chen, Chiyuan Zhang, Mengdi Wangg (2025)

*Under Review.*

### **On the Statistical Complexity for Offline and Low-Adaptive Reinforcement Learning with Structures**

**Ming Yin**, Mengdi Wang, and Yu-Xiang Wang (2025)

*Invited Review Article to Statistical Science Journal (To Appear).*

### **CRISPR-GPT: LLM Agents for Automated Design of Gene-Editing Experiments**

Yuanhao Qu\*, Kaixuan Huang\*, **Ming Yin**, Kanghong Zhan, Dyllan Liu, Di Yin, William A Johnson, Xiaotong Wang, Denny Zhou, Russ Altman, Mengdi Wang\*, and Le Cong\* (2024)

*Under Review by Nature Biomedical Engineering*

### **On Langevin Posterior Sampling for Offline Reinforcement Learning**

Thanh Nguyen-Tang, **Ming Yin**, Masatoshi Uehara, Yu-Xiang Wang, Mengdi Wang, Raman Arora (2024)

*Under Review.*

### **Why Quantization Improves Generalization: NTK of Binary Weight Neural Networks**

Kaiqi Zhang, **Ming Yin**, Yu-Xiang Wang (2023)

*In ICML workshop in Neural Compression, Honolulu, HI, USA.*

### **Offline Policy Evaluation for Reinforcement Learning with Adaptively Collected Data**

Sunil Madhow, Dan Qiao, **Ming Yin**, Yu-Xiang Wang (2022)

*In NeurIPS workshop in Offline RL (2022), New Orleans, LA, USA.*

### **A Theoretical Perspective for Speculative Decoding Algorithm**

Ming Yin, Minshuo Chen, Kaixuan Huang, Mengdi Wang (2024)

In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, Vancouver, Canada.

### **NetworkGym: Reinforcement Learning Environments for Multi-Access Traffic Management in Network Simulation**

Momin Haider, Ming Yin, Menglei Zhang, Arpit Gupta, Jing Zhu, Yu-Xiang Wang (2024)

In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS Dataset and Benchmark Track 2024)*, Vancouver, Canada.

### **Fast Best-of-N Decoding via Speculative Rejection**

Ruiqi Zhang\*, Momin Haider\*, Ming Yin, Jiahao Qiu, Mengdi Wang, Peter Bartlett, Andrea Zanette (2024)

In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, Vancouver, Canada.

### **Transfer $Q^*$ : Principled Decoding for LLM Alignment**

Souradip Chakraborty\*, Soumya Suvra Ghosal\*, Ming Yin, Dinesh Manocha, Mengdi Wang, Amrit Singh Bedi, Furong Huang (2024)

In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, Vancouver, Canada.

### **Offline Multitask Representation Learning for Reinforcement Learning**

Haque Ishfaq\*, Thanh Nguyen-Tang, Songtao Feng, Raman Arora, Mengdi Wang, Ming Yin\*, Doina Precup\* (2024)

In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, Vancouver, Canada.

### **Towards General Function Approximation in Nonstationary Reinforcement Learning**

Songtao Feng, Ming Yin, Ruiquan Huang, Yu-Xiang Wang, Jing Yang, Yingbin Liang (2024)

In *IEEE Journal on Selected Areas in Information Theory (JSAIT)*.

### **Learning the Target Network in Function Space**

Ming Yin\*, Kavosh Asadi\*, Yao Liu\*, Shoham Sabach\*, Rasool Fakoor (2024)

In *Proceedings of the 40th International Conference on Machine Learning (ICML 2024)*, Vienna, Austria.

### **Improving Sample Efficiency of Model-Free Algorithms for Zero-Sum Markov Games**

Songtao Feng, Ming Yin, Yu-Xiang Wang, Jing Yang, Yingbin Liang (2024)

In *Proceedings of the 40th International Conference on Machine Learning (ICML 2024)*, Vienna, Austria.

### **Logarithmic Switching Cost in Reinforcement Learning beyond Linear MDPs**

Dan Qiao, Ming Yin, Yu-Xiang Wang (2024)

In *IEEE International Symposium on Information Theory (ISIT 2024)*, Athens, Greece.

### **MMMU: A Massive Multi-discipline Multimodal Understanding and Reasoning Benchmark for Expert AGI**

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhao Chen (2024)

**Best Paper Finalist & Oral presentation**<sup>1</sup>

The *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2024)*, Seattle, WA, USA.

---

<sup>1</sup>Total of 24 award candidate papers selected out of 11532 submissions and 2719 accepted papers (about 0.2% of total).

## **Posterior Sampling with Delayed Feedback for Reinforcement Learning with Linear Function Approximation**

**Ming Yin\***, Nikki Kuang\*, Mengdi Wang, Yu-Xiang Wang, Yian Ma (2023)

In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS 2023)*, New Orleans, LA, USA.

## **TheoremQA: A Theorem-driven Question Answering dataset**

Wenhu Chen, **Ming Yin**, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, Tony Xia (2023)

*Conference on Empirical Methods in Natural Language Processing [Main] (EMNLP 2023)*, Singapore, Asia.

## **No-Regret Linear Bandits beyond Realizability**

Chong Liu, **Ming Yin**, Yu-Xiang Wang (2023)

In *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence (UAI 2023)*, Pittsburgh, PA, USA.

## **Non-stationary Reinforcement Learning under General Function Approximation**

Songtao Feng, **Ming Yin**, Ruiquan Huang, Yu-Xiang Wang, Jing Yang, Yingbin Liang (2023)

In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)*, Honolulu, HI, USA.

## **Offline Reinforcement Learning with Closed-Form Policy Improvement Operators**

Jiachen Li, Edwin Zhang, **Ming Yin**, Qinxun Bai, Yu-Xiang Wang, William Yang Wang (2023)

In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)*, Honolulu, HI, USA.

*Short version accepted to NeurIPS workshop in Offline RL (2022), New Orleans, LA, USA.*

## **Offline Reinforcement Learning with Differentiable Function Approximation is Provably Efficient**

**Ming Yin**, Mengdi Wang, Yu-Xiang Wang (2023)

In *Proceedings of the 10th International Conference on Learning Representations (ICLR 2023)*, Kigali Rwanda, Africa.

## **On Instance-Dependent Bounds for Offline Reinforcement Learning with Linear Function Approximation**

Thanh Nguyen-Tang, **Ming Yin**, Sunil Gupta, Svetha Venkatesh, Raman Arora (2023)

In *Proceedings of Association for the Advancement of Artificial Intelligence (AAAI 2023)*, Washington, DC, USA.

## **Offline Stochastic Shortest Path: Learning, Evaluation and Towards Optimality**

**Ming Yin\***, Wenjing Chen\*, Mengdi Wang, Yu-Xiang Wang (2022)

In *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence (UAI 2022)*, Eindhoven, Netherlands.

## **Sample-Efficient Reinforcement Learning with $\log\log(T)$ Switching Cost**

Dan Qiao, **Ming Yin**, Ming Min, Yu-Xiang Wang (2022)

In *Proceedings of the 39th International Conference on Machine Learning (ICML 2022)*, Baltimore, MD, USA.

## **Near-optimal Offline Reinforcement Learning with Linear Representation: Leveraging Variance Information with Pessimism**

**Ming Yin**, Yaqi Duan, Mengdi Wang, Yu-Xiang Wang (2022)

In *Proceedings of the 10th International Conference on Learning Representations (ICLR 2022)*, Virtual.

## **Towards Instance-Optimal Offline Reinforcement Learning with Pessimism**

**Ming Yin**, Yu-Xiang Wang (2021)

In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, Vancouver, Canada.

## **Optimal Uniform OPE and Model-based Offline Reinforcement Learning in Time Homogeneous, Reward-Free and Task-Agnostic Settings**

**Ming Yin**, Yu-Xiang Wang (2021)

In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, Vancouver, Canada.

*Short version accepted to ICML 2021 Reinforcement Learning Theory Workshop*

## **Near-Optimal Offline Reinforcement Learning via Double Variance Reduction**

**Ming Yin**, Yu Bai, Yu-Xiang Wang (2021)

In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, Vancouver, Canada.

*Short version accepted to ICML 2021 Reinforcement Learning Theory Workshop*

## **Near-Optimal Provable Uniform Convergence in Offline Policy Evaluation for Reinforcement Learning**

**Ming Yin**, Yu Bai, Yu-Xiang Wang (2021) **Oral presentation**<sup>2</sup>

In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS 2021)*, Virtual.

*Short version accepted to NeurIPS 2020 Offline Reinforcement Learning Workshop.*

## **Asymptotically Efficient Off-Policy Evaluation for Tabular Reinforcement Learning**

**Ming Yin**, Yu-Xiang Wang (2020)

In *Proceedings of the 23th International Conference on Artificial Intelligence and Statistics (AISTATS 2020)*, Sicily, Italy.

## ACADEMIC SERVICES

---

### Senior Academic Service

- (NeurIPS) Area Chair, 2025
- (ICML) Area Chair, 2025
- (AISTATS) Area Chair, 2025
- (ICCOPT) International Conference on Continuous Optimization, Session Creator & Chair, 2025.
  - Optimization for Large Language Models and Kernels
- (UMAP) Workshop & Tutorial Co-chair for ACM Conference on User Modeling, Adaptation, and Personalization, 2025.
- (NeurIPS) Area Chair, 2024
- (INFORMS) 2024 INFORMS Annual Meeting, Invited Session Creator & Chair, 2024.
  - Integrating Generative AI with Sequential Decision-Making: Theory and Applications
- (ISMP) 25th International Symposium on Mathematical Programming, Session Creator & Chair, 2024.
  - Statistical Learning, Optimization and Stochastic Programming for Reinforcement Learning
- (CISS) 58th Annual Conference on Information Sciences and Systems, Session Chair, 2024.
- (NeurIPS) Area Chair, 2023

---

<sup>2</sup>Total of 48 oral presentations selected out of 1527 submissions and 455 accepted papers (about 3.1% of total).

## Conference Reviewers

- (ICML) International Conference on Machine Learning, 2020,2021,2022, 2023, 2024
- (AISTATS) International Conference on Artificial Intelligence and Statistics, 2021,2022,2023,2024
- (NeurIPS) Conference on Neural Information Processing Systems, 2021,2022
- (ICLR) International Conference on Learning Representations, 2022,2023, 2024
  - ICLR Blogpost Reviewer, 2024
- (AAAI) AAAI Conference on Artificial Intelligence, 2023, 2024
- (UAI) Conference on Uncertainty in Artificial Intelligence, 2023
- (EMNLP) Conference on Empirical Methods in Natural Language Processing, 2023
- (COLM) The first & second Conference on Language Modeling, 2024, 2025
- (COLT) Conference on Learning Theory, 2024, 2025
- (L4DC) Learning for Dynamics and Control, 2025

## Journal Reviewers

- (JASA) Journal of the American Statistical Association
- (JMLR) Journal of Machine Learning Research
- (JAIR) Journal of Artificial Intelligence Research
- (SIMODS) SIAM Journal on Mathematics of Data Science
- (TMLR) Transactions on Machine Learning Research
- (Ann. Stat.) Annals of Statistics
- (MACH) Machine Learning, Journal by Springer
- (JDS) ACM/IMS Journal of Data Science (3-year appointment)

## Workshop Services

- Program Committee of *the 19th Workshop on Women in Machine Learning (NeurIPS 24)*,
- Program Committee of *the 3rd Workshop on Mathematical Reasoning and AI (NeurIPS 23)*,
- Program Committee of *the AI for Scientific Discovery Workshop (NeurIPS 23)*,
- Program Committee of *the Interactive Learning with Implicit Human Feedback (ICML 23)*,
- Program Committee of *the Neural Compression Workshop (ICML 23)*,
- Program Committee of *the 3rd (Launchpad) Offline RL Workshop (NeurIPS 22)*,
- Program Committee of *the 2nd Offline RL workshop (NeurIPS 21)*,
- Program Committee of *RL theory workshop (ICML 21)*

## SELECTED TALKS

---

**8th Workshop on Cognition & Control in Complex Systems** Gainesville, FL, *Mar., 2025*

KEYNOTE SPEAKER: On Foundation and Applications of Offline Reinforcement Learning

**NSF Cyber Physical Systems PI Meeting** Nashville, TN, *Mar., 2025*

POSTER TITLE: Reinforcement Learning in the Era of Generative AI

**AAAI-25 Tutorial on Offline RL** Philadelphia, PA, *Feb., 2025*

TUTORIAL TITLE: Advancing Offline Reinforcement Learning: Essential Theories and Techniques for Algorithm Developers

CO-SPEAKERS: Fengdi Che, Chenjun Xiao, Csaba Szepesvari.

**Contributed talk at 2024 Informs Annual Meeting** Seattle, WA, *Oct., 2024*

TALK: On the evaluation benchmarks for Multimodal Generative AI

**The 2024 Young Researchers Workshop at Cornell University** Ithaca, NY, *Oct., 2024*  
TALK: Understanding Q\*: From sequential decision-making to Large Language Model Alignment

**The 25th International Symposium on Math Programming** Montreal, Canada, *July, 2024*  
TALK: Learning the Target Network in Function Space

**The First CMLR Seminar at Peking University** Beijing (Virtual), China, *April, 2024*  
TALK: Recent Advances in Offline Reinforcement Learning

**Princeton Postdoc Council Seminar** Princeton, NJ, *Mar., 2024*  
TALK: A Probabilistic Characterization of Speculative Decoding for Large Language Models

**Contributed talk at 2023 Informs Annual Meeting** Phoenix, AZ, *Oct., 2023*  
SESSION CHAIRS: Asuman Ozdaglar, Kaiqing Zhang, Jiawei Zhang  
TALK: Offline Reinforcement Learning with Differentiable Function Approximation is Provably Efficient

**The 2023 Young Researchers Workshop at Cornell University** Ithaca, NY, *Oct., 2023*  
ATTENDANCE: Poster Presentation  
TITLE: Principled Posterior Sampling for Batch RL with Neural Network Approximation

**The 68th TrustML Young Scientist Seminar at RIKEN** Tokyo (Virtual), Japan, *April, 2023*  
HOST: Masashi Sugiyama  
TALK: Towards Sample-Optimal Offline Reinforcement Learning

**Information Theory and Applications Workshop** San Diego, CA, *Feb., 2023*  
AWARD: Graduation Day Award  
TALK: “Towards Instance-dependent and Optimal Offline Reinforcement Learning”

**Institute for Foundations of Data Science at Yale** New Haven (Virtual), CT, *Jan., 2023*  
HOST: Daniel Spielman  
TALK: “Towards Instance-dependent and Optimal Offline Reinforcement Learning”

**Amazon AWS AIRE Intern Summit** Sunnyvale, CA, *Aug, 2022*  
TALK: “Practical off-policy evaluation by leveraging the power of more data”  
VENUE: Amazon AWS AI Research & Education Intern Summit

**Big Data and Machine Learning Seminar at UCLA** Los Angeles, CA, *April, 2022*  
HOST: Quanquan Gu  
TALK: “Towards better Instance-dependent offline Reinforcement Learning”

**The Ohio State University, ECE Department** Columbus, OH (Virtual), *Dec, 2021*  
HOST: Yingbin Liang  
TALK: “Optimal Uniform OPE and Model-based Offline Reinforcement Learning in Time-Homogeneous, Reward-Free and Task-Agnostic Settings”

**Princeton University, ECE Department** Princeton, NJ (Virtual), *July, 2021*  
TALK: “Uniform OPE, Adaptive, Optimal offline RL and Beyond”

**AISTATS 2021 Oral Presentation** Virtual, *April, 2021*  
HOST: Amin Karbasi  
TALK: “Near-Optimal Provable Uniform Convergence in Offline Policy Evaluation for Reinforcement Learning”.

## AWARDS

---

**Rising Star Award at KAUST AI Symposium, 2025**

**Rising Star Award at Conference on Parsimony and Learning, 2025**

**CVPR Best Paper Finalist, 2024**

**Young Researcher (Rising Star) Workshop at Cornell ORIE, 2024**

**Graduation Day Award, *Information Theory and Applications Workshop (ITA)*, 2023**

**UCSB Graduate Division Dissertation Fellowship, 2023 [Declined for other options]**

**UCSB CS Academic Excellence Fellowship, 2021**

## PROFESSIONAL EXPERIENCE

---

**Amazon AWS AI Research & Education, *Applied Scientist Intern*** *CA, Summer, 2023*

**Amazon AWS AI Team, *Applied Scientist Intern*** *CA, Summer, 2022*

## TEACHING EXPERIENCE

---

**AAAI Tutorial on Offline RL** *AAAI, 2025*

**Princeton Postdoc Council Seminar, *Moderator*** *Princeton, 2024*

**Research Methods in Computer Science, *CS 110, Course Instructor*** *UCSB, 2020-2021*

**Probability Theory and Stochastic Processes, *PSTAT 213, TA*** *UCSB, 2018*

**Applied Stochastic Processes, *PSTAT 160, TA*** *UCSB, 2017,2018*

## OUTREACH AND GENERAL SERVICES

---

**Mentor of Doctoral Consortium at AAAI, 2025**

Participating in roundtable discussions, attending poster sessions in the same room, and joining students for lunch as a mentor.

**Program Committee for the WiML workshop, 2024**

Reviewed abstract submissions for the Women in Machine Learning workshop.

**Princeton Research Day, *Judge, 2024***

Served in the Judge panel at Princeton Research Day. Responsible for award assignments.

**Long term UNICEF doner, 2022-present**

Help the children globally who need it most.

**Mentor for the *Early Research Scholar Program (ERSP)* at UCSB, 2020 - 2021**

Instructed underrepresented minority students doing research.

## STUDENT MENTEES

---

**Nguyen Anh Duc, *NUS undergraduate, recipient of 2024 INFORMS Undergraduate OR Prize.***

**Jiahao Shi, *Tsinghua Undergraduate, now Princeton PhD***

**Jiahao Qiu, *Princeton PhD***

## OTHER TALKS OF MY WORKS

---

### **Simons Institute for the Theory of Computing**

Berkeley, CA (Virtual), *Dec, 2020*

PRESENTER: Yu-Xiang Wang

TALK: “Near-Optimal Provable Uniform Convergence in Offline Policy Evaluation for Reinforcement Learning”

VENUE: Reinforcement Learning from Batch Data and Simulation workshop

### **Virtual RL Theory Seminars**

Virtual, *Oct, 2020*

HOST: Gergely Neu, Ciara Pike-Burke, Csaba Szepesvári

PRESENTER: Yu-Xiang Wang

TALK: “Near-Optimal Provable Uniform Convergence in Offline Policy Evaluation for Reinforcement Learning”