

Yifei Ma

 Palo Alto, CA, USA
  mayf.mark@gmail.com
  (412) 425-2659
  yma.io
  Google Scholar

Summary

8+ years of industry experience designing, training, and deploying large-scale recommender systems and domain-adapted language models. Research background in reinforcement learning, Bayesian modeling, counterfactual reasoning, and temporal models, with publications at top-tier venues. Contributor to problem framing and technical evaluation across projects involving scientific ambiguity.

Experience

- Amazon**, Senior Applied Scientist, Sponsored Products (Ads) Palo Alto, CA, USA
Sept 2025 – present
6 months
- Built LLM-based ad retrieval systems leveraging near-real-time shopper intent prediction
 - Led multi-objective reinforcement fine-tuning using distant rewards and online A/B feedback to maximize ad relevance and conversion
 - Delivered statistically significant revenue lifts on search and homepage widgets
- Amazon**, Senior Applied Scientist, Stores Foundational AI Palo Alto, CA, USA
Jan 2025 – Sept 2025
9 months
- Led a cross-team effort (5 scientists, 1 engineer) to improve a personalized query builder widget on the Amazon search navigation bar using fine-tuned LLM rerankers; delivered ~20% offline recall@1 lift and advanced to online experimentation
 - Led synthetic data generation and instruction-tuning for a domain-adapted 7B LLM; improved IFEval performance by ~15% while maintaining parity on other benchmarks
 - Self-organized a team of 1 scientist and 3 engineers to develop MCP servers for monitoring and debugging internal ML workflows; won a GenAI hackathon hosted by the Rufus team
- Amazon**, Applied Scientist → Senior Applied Scientist, AWS Personalize East Palo Alto & Santa Clara, CA, USA
Sept 2017 – Dec 2024
7 years 4 months
- Early work on counterfactual optimization for personalization ([AISTATS 2019 paper](#))
 - Developed scalable hybrid sequential recommender systems for real-time personalization in AWS AI; lead author of [KDD 2020 Best Paper](#) (Applied Data Science Track)
 - Extended item-to-user recommendation to user-to-item segmentation using a novel temporal point process framework ([ICLR 2022](#))
 - Scaled recommender systems to 5M+ unique items via noise-contrastive loss and approximate nearest-neighbor search, addressing large-catalog customer needs ([AWS blog](#))
 - Led development of an LLM-based content generator to enhance thematic recommendation quality using human feedback ([WSDM 2023 demo](#))

Education

- Ph.D.** **Carnegie Mellon University**, Machine Learning Aug 2011 – May 2017
- Advisor Jeff Schneider
 - Thesis Active Search with Complex Actions and Rewards
- B.S.** **Tsinghua University**, Automation and Mathematics (Dual Major) Aug 2007 – July 2011

Publications

- Optimal Design for Human Preference Elicitation** 2024
 Subhojyoti Mukherjee, Anusha Lalitha, Kousha Kalantari, Aniket Anand Deshmukh, Ge Liu, Yifei Ma, Branislav Kveton
 (NeurIPS 2024)

Logic-Scaffolding: Personalized Aspect-Instructed Recommendation Explanation Generation Using LLMs	2024
Behnam Rahdari, Hao Ding, Ziwei Fan, Yifei Ma, Zhuotong Chen, Anoop Deoras, Branislav Kveton (Proceedings of the 17th ACM International Conference on Web Search and Data Mining)	
Trending Now: Modeling Trend Recommendations	2023
Hao Ding, Branislav Kveton, Yifei Ma, Youngsuk Park, Venkataramana Kini, Yupeng Gu, Ravi Divvela, Fei Wang, Anoop Deoras, Hao Wang (Proceedings of the 17th ACM Conference on Recommender Systems)	
Fixed-Budget Best-Arm Identification with Heterogeneous Reward Variances	2023
Anusha Lalitha, Kousha Kalantari, Yifei Ma, Anoop Deoras, Branislav Kveton (Uncertainty in Artificial Intelligence)	
Context Uncertainty in Contextual Bandits with Applications to Recommender Systems	2022
Hao Wang, Yifei Ma, Hao Ding, Yuyang Wang (Proceedings of the AAAI Conference on Artificial Intelligence)	
Bridging Recommendation and Marketing via Recurrent Intensity Modeling	2022
Yifei Ma, Ge Liu, Anoop Deoras (International Conference on Learning Representations)	
Temporal-Contextual Recommendation in Real-Time	2020
Yifei Ma, Balakrishnan Narayanaswamy, Haibin Lin, Hao Ding (Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining)	
Imitation-Regularized Offline Learning	2019
Yifei Ma, Yu-Xiang Wang, Balakrishnan Narayanaswamy (The 22nd International Conference on Artificial Intelligence and Statistics)	
Towards Optimal Off-Policy Evaluation for Reinforcement Learning with Marginalized Importance Sampling	2019
Tengyang Xie, Yifei Ma, Yu-Xiang Wang (Advances in Neural Information Processing Systems 32)	

Additional

Service: Reviewer for NeurIPS, ICML, KDD, UAI, AAAI (multiple years)