Tuan Dung (Daniel) Ngo

AI/ML Research Scientist, New York City, NY
tuandungngo207@gmail.com — (717) 713-3090
www.linkedin.com/in/ngod207 — dtngo.com — Google Scholar

Ph.D. in CS (UMN). Researcher in LLM personalization, synthetic personas and watermarking at J.P. Morgan AI Research.

RESEARCH INTERESTS

LLM Personalization, Watermarking LLMs, Machine Learning Theory, Algorithmic Game Theory, Incentive-Aware Machine Learning, Causal Inference, Uncertainty Quantification.

SKILLS

Python, PyTorch, Transformers, LLM fine-tuning, reinforcement learning (RLHF), evaluation, statistics and causal inference.

EDUCATION

University of Minnesota, Minneapolis, MN — August 2019 – December 2024
Computer Science Ph.D.
Advisors: Steven Wu and Maria Gini
Thesis: Incentive-Aware Machine Learning for Social Welfare Maximization under Uncertainty
GAGE Fellowship (August 2019, May 2020)

Dickinson College, Carlisle, PA — August 2015 – May 2019
Bachelor of Science: Computer Science *(summa cum laude)*
Bachelor of Science: Mathematics *(summa cum laude)*

EMPLOYMENT

J.P. Morgan AI Research — New York, NY
AI Research Senior Associate — January 2025 – Present

- Research end-to-end synthetic persona generation and evaluation models using LLMs, covering persona construction, simulation of user behavior, and assessment of persona fidelity.
- Conduct research on LLM personalization, including context management for long-running interactions, modeling how user preferences evolve over time, and evaluating the quality of personalized responses.
- Study the robustness and security of watermarking for LLM outputs and generative tabular data, including attacks on distortion-free watermarks (NeurIPS 2025 Lock-LLM Workshop).

J.P. Morgan AI Research — New York, NY
AI Research Summer Associate — Summer 2024

- Designed an adaptive and robust watermarking approach for synthetic tabular data generated by machine learning models (UAI 2026).
- Researched incentive-aware algorithms for two-sided matching markets, where participants act strategically (NeurIPS 2024 Workshop on Bayesian Decision-making and Uncertainty).

University of Minnesota — Minneapolis, MN
Research Assistant — August 2019 – December 2024

Carnegie Mellon University — Pittsburgh, PA
Visiting Scholar — August 2021 – December 2024

- Conducted research on incentive-aware machine learning, spanning incentivized exploration, strategic causal inference and differential privacy (TMLR; NeurIPS 2022; ICML 2021, 2022).
- Studied reliability and fairness topics in machine learning, including multicalibration and model multiplicity (ICLR 2025; NeurIPS 2025).

PUBLICATIONS

Journal:

1. **Daniel Ngo***, Keegan Harris*, Anish Agarwal, Vasilis Syrgkanis, and Zhiwei Steven Wu. *Incentive-Aware Synthetic Control: Accurate Counterfactual Estimation via Incentivized Exploration.* **Transactions on Machine Learning Research (TMLR)**.

*Denotes equal contribution and/or alphabetical ordering

Conference Proceedings:

1. **Daniel Ngo***, Archan Ray*, Akshay Seshadri*, Daniel Scott, Saheed Obitayo, Niraj Kumar, Vamsi K. Potluru, Marco Pistoia, and Manuela Veloso. *Adaptive and Robust Watermark for Generative Tabular Data.* **Conference on Uncertainty in Artificial Intelligence (UAI)**, 2026.
2. Hongyi Henry Jin, Zijun Ding, **Daniel Ngo**, and Zhiwei Steven Wu. *Discretization-free Multicalibration through Loss Minimization over Tree Ensembles.* **Conference on Neural Information Processing Systems (NeurIPS)**, 2025.
3. Ally Yalei Du*, **Daniel Ngo***, and Zhiwei Steven Wu. *Reconciling Model Multiplicity for Downstream Decision Making.* **International Conference on Learning Representations (ICLR)**, 2025.
4. Xinyan Hu*, **Daniel Ngo***, Zhiwei Steven Wu, and Aleksandrs Slivkins. *Incentivizing Combinatorial Bandit Exploration.* **Conference on Neural Information Processing Systems (NeurIPS)**, 2022.
5. **Daniel Ngo***, Giuseppe Vietri*, and Zhiwei Steven Wu. *Improved Regret for Differentially Private Exploration in Linear MDP.* **International Conference on Machine Learning (ICML)**, 2022.
6. Keegan Harris, **Daniel Ngo***, Logan Stapleton*, Hoda Heidari, and Zhiwei Steven Wu. *Strategic Instrumental Variable Regression: Recovering Causal Relationships From Strategic Responses.* **International Conference on Machine Learning (ICML)**, 2022.
7. **Daniel Ngo***, Logan Stapleton*, Vasilis Syrgkanis, and Zhiwei Steven Wu. *Incentivizing Exploration with Algorithmic Instruments.* **International Conference on Machine Learning (ICML)**, 2021.

Workshop Papers and Preprints:

1. Shayleen Reynolds, Hengzhi He, **Daniel Ngo**, et al. *Breaking Distortion-free Watermarks in Large Language Models.* **NeurIPS 2025 Lock-LLM Workshop**. arXiv preprint, 2025. `https://arxiv.org/abs/2502.18608`
2. Jane Hsieh, Angie Zhang, Erik Chou, Mialy Rasetarinera, **Daniel Ngo**, et al. *Supporting Worker Needs and Advancing Policy: Exploring Gig Worker-Centered Data-Sharing Alternatives.* arXiv preprint, 2024. `https://arxiv.org/abs/2412.02973`
3. **Daniel Ngo**, Vamsi K. Potluru, and Manuela Veloso. *Incentivized Exploration in Two-sided Matching Markets.* **NeurIPS 2024 Workshop on Bayesian Decision-making and Uncertainty**.
4. Shengyuan Hu*, **Daniel Ngo***, Shuran Zheng*, Virginia Smith, and Zhiwei Steven Wu. *Federated Learning as a Network Effects Game.* arXiv preprint, 2023. `https://arxiv.org/abs/2302.08533`

TALKS

1. *"Incentivizing Combinatorial Bandit Exploration"*

   INFORMS Annual Meeting (October 2024), invited talk

2. *"Incentive-Aware Synthetic Control: Accurate Counterfactual Estimation via Incentivized Exploration"*

   INFORMS Annual Meeting (October 2023), invited talk

3. *"Incentivizing Exploration with Algorithmic Instruments"*

   ICML 2021 (July 2021, Short Presentation)

SERVICE

Conference Reviewer

- Artificial Intelligence and Statistics Conference (AISTATS) 2026, 2025
- Causal Learning and Reasoning (CLeaR) 2023
- International Conference on Learning Representations (ICLR) 2025, 2024
- International Conference on Machine Learning (ICML) 2025, 2024, 2023, 2022, 2021
- Conference on Neural Information Processing Systems (NeurIPS) 2026, 2025, 2024, 2023, 2022, 2021
- Conference on Uncertainty in Artificial Intelligence (UAI) 2026

TEACHING

- Guest Lecturer: Human-AI Interaction (Fall 2023). Instructor: Haiyi Zhu
- Teaching Assistant: Algorithms and Data Structures (Fall 2020). Instructor: Carl Sturtivant

*Denotes equal contribution and/or alphabetical ordering