Siteng Kang

siteng.kang@gmail.com | sitengkang.github.io

Education

University of Illinois Chicago — *PhD in Computer Science — 2025*
Brown University — *MS in Computer Science — 2020*
Pennsylvania State University — *BS in Computer Science; BS in Mathematics; Minor in Physics — 2018*

Industry Experience

Amazon Jul 2025 – Present

Applied Scientist, Alexa — Large-Scale Customer Intent Classification

- Developed multilingual intent classification systems for large-scale customer language understanding applications.
- Designed LLM-assisted annotation and model adaptation workflows to improve data quality and support scalable expansion across categories and locales.

Amazon May 2022 – Sep 2024

Applied Scientist Intern (3x)

- Developed machine learning models for customer behavior prediction, including churn forecasting and behavioral signal discovery.
- Built predictive models for subscription conversion and customer lifecycle analysis across digital products.
- Applied sequential modeling to trust and safety tasks, including bad-actor detection from user activity patterns.

Accenture Jun 2018 – Aug 2018

Data Scientist Intern

- Developed OCR-based document understanding methods for automated receipt text extraction workflows.

Research Experience

Adaptive Data Harvesting for Universal-Constraint Learning (under review) **IJCAI 2026**

- Proposed a framework that formulates collocation selection as a sequential decision-making problem and trains an RL policy to adaptively propose samples based on evolving learning signals.
- Validated the framework under Lyapunov NN and PINN settings, where task-specific policies improve convergence speed, training stability, and constraint satisfaction compared to fixed sampling heuristics.

Data Poisoning for Offline-to-Online Reinforcement Learning **UAI 2024**

- Proposed a bi-level optimization data poisoning attack that remains *stealthy* during offline training but induces severe performance degradation during online fine-tuning by triggering overestimation and distribution shift.
- Validated the attack across four environments and multiple algorithms, causing >20% return drop under realistic threat models (limited poisoning budget and no white-box access).

Poisoning Generative Replay in Continual Learning **ICML 2023**

- Developed a dirty-label, input-aware backdoor attack that exploits generative models' inability to capture input-dependent triggers, promoting forgetting while preserving current-task accuracy.
- Validated the attack on four datasets; effective against strong defenders, driving prior-task accuracy to $< 10\%$.

Data-Independent Memory Hard Functions **CRYPTO 2019**

- Analyzed space–time complexity of Argon2 variants under parallel/generic attacks, contributing to stronger constructions for memory-hard functions.

Selected Publications

- **Kang S.**, Zhang X. *Adaptive Data Harvesting for Efficient Neural Network Learning with Universal Constraints.* Under review at **IJCAI 2026**.
- Yu Z.*, **Kang S.***, Zhang X. *Offline Reward Perturbation Boosts Distributional Shift in Online RL.* **UAI 2024**.
- **Kang S.**, Shi Z., Zhang X. *Poisoning Generative Replay in Continual Learning to Promote Forgetting.* **ICML 2023**.
- Blocki J., Harsha B., **Kang S.**, Lee S., Xing L., Zhou S. *Data-Independent Memory Hard Functions: New Attacks and Stronger Constructions.* **CRYPTO 2019**.