

# Yingtong Dou

Palo Alto, CA | +1 312-785-5168 | ytongdou@gmail.com | [Personal Website](#) | [Google Scholar](#) | [GitHub](#)

---

## SUMMARY

Applied Scientist with experience in FinTech, GenAI, and Trust&Safety. Focused on data-driven research and building production-ready ML solutions. Industry experience includes developing Visa's foundation models for fraud detection and recommendation, LLM fine-tuning, and graph-based anomaly detection. Strong collaborator with a continual-learning mindset and a track record of impactful applied research.

---

## EXPERIENCE

### Visa Research, Foundational AI

Dec. 2022 – Present

*Research Scientist.*

*Foster City, CA*

- **Improved fraud detection performance by 22%** relative to the production model by designing, implementing, training, and optimizing **TransactionGPT**, Visa's proprietary foundation model for payment intelligence
- **Fine-tuned 0.1B to 7B LLMs on 5B+ transactions (480B tokens)** using LoRA and distributed training, improving key business metric by **2.5×** and beating the production model after optimizing data and training recipes
- Evaluating XGBoost, RNNs, Transformers, and frontier LLMs on Visa's internal benchmark to understand the pros and cons of different models and determine their best production fit
- Contributed to agentic coding and tooling projects inside Visa. Collaborated with university partners on efficient RAG, efficient LLM inference, and LLM safety research

### University of Illinois Chicago, Big Data and Social Computing Lab

Jun. 2019 – May 2022

*Research Assistant. Advisor: Philip S. Yu*

*Chicago, IL*

- Led research on **graph-based outlier/misinformation detection**, had publications at NeurIPS and KDD etc.
- Co-created **PyGOD** and **SafeGraph**, popular graph outlier detection libraries in academia and industry
- Worked with Tencent, Microsoft, Grab, and F5 on applied risk & safety projects (**fraud/spam** detection)

### Snap Research, User Modeling and Personalization

Jun. 2021 – Sep. 2021

*Research Intern.*

*Remote*

- Queried and analyzed the production graph-structured data for the Snapchat recommender system from GCP
- Investigated the robustness of **graph neural networks under concept drift** using different measurements

---

## EDUCATION

### University of Illinois Chicago - Ph.D. in Computer Science

2017 – 2022

*Dissertation: Robust Graph Learning for Misbehavior Detection [PDF]*

*Chicago, IL*

### Beijing University of Posts and Telecommunications - B.E. in IoT Engineering

2013 – 2017

*Graduated with Beijing Excellent Graduate Award*

*Beijing, China*

---

## RECENT PUBLICATIONS

**Y. Dou**, Y. Chen, et al. "TransactionGPT." *KDD*, 2026. [PDF]

H. Le, S. Zhong, Y. Lu, **Y. Dou**. et al. "FAFO: Lossy KV Cache Compression for Lossless Inference Acceleration via Draftless Fumble Decoding." *ICML*, 2026. [PDF]

Y. Jin, K. Sharma, V. Mohan, **Y. Dou**. et al. "SARA: Selective and Adaptive Retrieval-augmented Generation with Context Compression." *ACL*, 2026. [PDF]

K. Sharma, Y. Jin, V. Mohan, **Y. Dou**. et al. "Sysformer: Safeguarding Frozen Large Language Models with Adaptive System Prompts." *ICLR*, 2026. [PDF]

---

## SKILLS

**Machine Learning:** Tabular/Sequential Transformer, LLM Post Training & Evals, GNNs, Anomaly Detection

**GenAI Tools:** Hugging Face, PEFT, DDP/FSDP, Unsloth, vLLM, Claude Code, Cline

**Developer Tools:** PyTorch, TensorFlow, Spark, Hadoop, Ray, GCP, mlflow, Grafana, AWS, Docker