PO-NIEN KUNG

Email Website Google Scholar GitHub

EDUCATION

University of California, Los Angeles *September 2022 – Present*
Ph.D. in Computer Science, expected 2027 Advisor: Nanyun (Violet) Peng

National Taiwan University *September 2017 – June 2021*
B.S. in Computer Science and Information Engineering

RESEARCH INTERESTS

Reliable and controllable language-model reasoning, reinforcement learning, agentic formal mathematics, instruction following, and constrained generation.

PUBLICATIONS AND PREPRINTS

LEAP: Supercharging LLMs for Formal Mathematics with Agentic Frameworks **Preprint, 2026**
Po-Nien Kung, Linfeng Song, Dawsen Hwang, Jinsung Yoon, Chun-Liang Li, Simone Severini, Mirek Olšák, Edward Lockhart, Quoc V. Le, Burak Gokturk, Thang Luong, Tomas Pfister, and Nanyun Peng

Learning Structured Reasoning via Tractable Trajectory Control **ICML 2026 Spotlight**
Po-Nien Kung, Zhen Yang, Jeffrey Luo, Cheng-Fu Yang, Haikang Deng, Zi-Yi Dou, Yinfei Yang, Nanyun Peng, Zhe Gan, and Kai-Wei Chang

Adaptable Logical Control for Large Language Models **NeurIPS 2024**
Honghua Zhang*, **Po-Nien Kung***, Masahiro Yoshida, Guy Van den Broeck, and Nanyun Peng

Decoupling Task-Solving and Output Formatting in LLM Generation **ACL 2026**
Haikang Deng, **Po-Nien Kung**, and Nanyun Peng

Active Instruction Tuning: Improving Cross-Task Generalization by Training on Prompt Sensitive Tasks **EMNLP 2023**
Po-Nien Kung, Fan Yin, Di Wu, Kai-Wei Chang, and Nanyun Peng

Do Models Really Learn to Follow Instructions? An Empirical Study of Instruction Tuning **ACL 2023**
Po-Nien Kung and Nanyun Peng

Improving Event Definition Following for Zero-Shot Event Detection **ACL 2024**
Zefan Cai*, **Po-Nien Kung***, Ashima Suvarna, Mingyu Derek Ma, Hritik Bansal, Baobao Chang, P. Jeffrey Brantingham, Wei Wang, and Nanyun Peng

STAR: Boosting Low-Resource Information Extraction by Structure-to-Text Data Generation with Large Language Models **AAAI 2024**
Mingyu Derek Ma, Xiaoxuan Wang, **Po-Nien Kung**, P. Jeffrey Brantingham, Nanyun Peng, and Wei Wang

Efficient Multi-Task Auxiliary Learning: Selecting Auxiliary Data by Feature Similarity **EMNLP 2021**
Po-Nien Kung, Sheng-Siang Yin, Yi-Cheng Chen, Tse-Hsuan Yang, and Yun-Nung Chen

Zero-Shot Rationalization by Multi-Task Transfer Learning from Question Answering **Findings of EMNLP 2020**
Po-Nien Kung, Tse-Hsuan Yang, Yi-Cheng Chen, Sheng-Siang Yin, and Yun-Nung Chen

MIDDAG: Where Does Our News Go? Investigating Information Diffusion via Community-Level Information Pathways **AAAI 2024**
Mingyu Derek Ma, Alexander K. Taylor, Nuan Wen, Yanchen Liu, **Po-Nien Kung**, Wenna Qin, Shicheng Wen, Azure Zhou, Diyi Yang, Xuezhe Ma, Nanyun Peng, and Wei Wang

GenEARL: A Training-Free Generative Framework for Multimodal Event Argument Role Labeling **Preprint, 2024**
Hritik Bansal, **Po-Nien Kung**, P. Jeffrey Brantingham, Kai-Wei Chang, and Nanyun Peng

LLM-REVal: Can We Trust LLM Reviewers Yet? **Preprint, 2025**
Rui Li, Jia-Chen Gu, **Po-Nien Kung**, Heming Xia, Junfeng Liu, Xiangwen Kong, Zhifang Sui, and Nanyun Peng

* Equal contribution.

RESEARCH EXPERIENCE

Google *January 2026 – September 2026*
Student Researcher

- Focused on developing the agentic reasoning and search components of **LEAP**, including informal proof planning, subgoal decomposition, iterative refinement, and Lean compiler verification.
- LEAP solved all 12 Putnam 2025 problems and achieved a 70% formal solve rate on Lean-IMO-Bench.

UCLA PlusLab *September 2022 – Present*
Graduate Researcher

- Conduct research on reliable and controllable language-model reasoning, including structured exploration, constrained decoding, instruction following, and formal mathematics.
- Developed **Ctrl-R**, which uses tractable HMM-DFA guidance and importance-sampled reinforcement learning; also contributed to Deco-G and adaptable logical control methods for constrained generation.

Zoom *June 2024 – August 2024*
Research Intern

- Investigated compositional instruction following and evaluated the controllability of large language models under complex task specifications.

MediaTek Research *October 2021 – January 2022*
Research Intern

- Developed a retrieval-augmented Transformer with end-to-end multi-hop retrieval; the resulting work was published at the FEVER Workshop at EACL 2023.

TECHNICAL SKILLS

Languages	Python, C, Lean, JavaScript
ML and systems	PyTorch, Hugging Face Transformers, vLLM, VERL, CUDA, Linux, Git
Research methods	Reinforcement learning, constrained decoding, distributed training, theorem proving

SELECTED HONORS

Principal Award, National Taiwan University Undergraduate Thesis Award
Second Place, DSTC9-SIMMC Response Generation