

# HAIKANG DENG

+1 (919) 260-9698 | frankdenghaikang@gmail.com | Google Scholar | LinkedIn

## EDUCATION

---

**University of California, Los Angeles** Sep 2024 - Present  
Ph.D. in Computer Science, Artificial Intelligence  
Advisor: Nanyun Peng

**University of North Carolina, Chapel Hill** Aug 2019 - May 2023  
B.S. in Computer Science & Statistics GPA: 3.96/4.0  
Advisor: Colin Raffel, Guanting Chen

## PUBLICATIONS

---

- [1] Reward-Augmented Decoding: Efficient Controlled Text Generation With a Unidirectional Reward Model  
**Haikang Deng** and Colin Raffel  
*Empirical Methods in Natural Language Processing (EMNLP)*, 2023
- [2] Large Language Models Struggle to Learn Long-Tail Knowledge  
Nikhil Kandpal, **Haikang Deng**, Adam Roberts, Eric Wallace, and Colin Raffel  
*International Conference on Machine Learning (ICML)*, 2023
- [3] Decoupling Task-Solving and Output Formatting in LLM Generation  
**Haikang Deng**, Po-Nien Kung, Nanyun Peng  
*arXiv*, under review, 2025
- [4] DialectGen: Benchmarking and Improving Dialect Robustness in Multimodal Generation  
Yu Zhou, Sohyun An, **Haikang Deng**, Da Yin, Clark Peng, Cho-Jui Hsieh, Kai-Wei Chang, Nanyun Peng  
*arXiv*, under review, 2025

## RESEARCH & PROFESSIONAL EXPERIENCE

---

**University of North Carolina, Chapel Hill** Aug 2022 - Mar 2024  
*Research Assistant* Advised by Prof. Colin Raffel

- Conducted research on human feedback alignment, controlled text generation, and retrieval augmentation
- Trained expert model routing and tagging in a contrastive manner with signals from a reward model

**Amazon** Dec 2023 - Jul 2024  
*Software Dev Engineer* Bellevue, WA

- Designed automation for product pre-compute based on given marketplace, merchant, and store utilization score
- Composed a product lead time update tool that supports bulk updates (up to 1 million entries)
- Built a Local Landing Page to display local products available for pick up

## PROJECTS

---

**Power Prover Models with Curriculum Learning** Jul 2025 - Present

- Benchmark prover model's performance on textbook problems
- Design a curriculum learning framework for prover models using a mixture of SFT and RLVR
- Collect Lean conjectures through generating instantiations

**Neurally Decomposed Oracle for Preference Optimization** Jul 2024 - Nov 2024

- Derived training objectives from NADO and adapt to preference data
- Analyzed model-generated response as compared to SOTA methods such as DPO
- Interpolated and extrapolated policy model to explore the potential of expedited training

**Embedding-Based Routing and Tagging of PEFT Modules** Oct 2023 - Mar 2024

- Trained a router that routes a query to a PEFT module using input embeddings
- Examined performance gap between held-in and held-out data routing
- Enabled zero-shot PEFT module addition to the pool of modules

**Tracing Model-Generated Wrong Answers** Nov 2022 - Feb 2023

- Explored correlation between QA accuracy and the number of relevant documents that contain negative co-occurrences
- Identified that incorrect answers are not driven by an increased association between the wrong answer and query entities