

Xiaopeng Li

✉ eelxpeng@gmail.com | 📞 +1 425-324-6638 | [in /in/eelxpeng](https://www.linkedin.com/in/eelxpeng) | [eelxpeng.github.io](https://github.com/eelxpeng)

Research Interest

I am primarily interested in Large Language Models training, Reinforcement Learning from Human Feedback, Natural Language/Code Generation, Agent Orchestration, Retrieval Augmented Generation, etc.

Education

The Hong Kong University of Science and Technology

2014 - 2019

Ph.D. in Computer Science and Engineering

Thesis: Harnessing the Synergy between Neural and Probabilistic Machine Learning

Experience

Amazon Q CodeGen / AWS AI Labs

Oct. 2022 - Present

Senior Applied Scientist

- Initiator and science lead for Amazon Q CodeGen Chat project. Developed initial chat model training and **created the first Q IDE Chat demo for the first Q IDE Chat in Feb 2023**. Persuaded the team, leadership and product to buy in the project. Helped scoped the Q IDE Chat product, build team and led the science strategy and planning timeline.
- Led the team to develop chat model training recipe, create high-quality human annotation data and human alignment data, develop automatic chat evaluation pipeline. Delivered key benchmarks to help promote the product launch.
- Initiated and led the Reinforcement Learning from Human Feedback (RLHF) group in Q CodeGen team. Led the team to develop distributed and optimized RLHF training pipeline, and led code SFT data and preference data curation and PPO training. Collaborated with Titan team, **we achieved 10x speedup in RLHF training, and better human preference for the Titan model**. Led the integration of multi-task learning loss to PPO to maintain performance on general benchmarks. Successfully launched the RLHF trained Titan model in Bedrock.
- Led the distributed training of **Direct Preference Optimization (DPO)** in the team and supervised the DPO training for code generation model that significantly improves coding benchmark.
- Proposed a novel training framework for improving self-debugging on code generation, including synthetic data generation pipeline, supervised finetuning, and Reinforcement Learning with explanation and execution reward.
- Led the work on finetuning internal model with hundred billion size model for Amazon Q, and helped launch the model to production.
- Led on Q CodeGen Chat **LLM Agent** project to orchestrate the search tool using for project-level context, customer database and web/documentation search with agentic method.

Amazon CodeWhisperer / AWS AI Labs

Oct. 2019 - Oct. 2022

Applied Scientist

- Initiated the code LLM project in 2020 and pretrained the very initial code GPT-2 model, being one of the **earliest GPT LLM advocates**. Created proof-of-concept code completion demo with the LLM in 2020, and it led to the founding of CodeWhisperer team and product.
- Developed and maintained distributed training infrastructure for pretraining and finetuning that are scalable to thousands of GPUs for LLMs with tens of billion parameters since 2021. Identified and resolved several bugs in open source framework and contributed to DeepSpeed and Pytorch Lightning.
- Led the GPT model **pretraining of 6.7B, 13B and 26B model parameters** since 2021, a lot earlier than most companies and teams, and **pushed internal code LLMs to state-of-the-art**. Successfully launched Amazon CodeWhisperer IDE code completion model and product.
- Led the investigation of science strategy, and architecture choices. Investigated GPT, BART and semantic parsing models, and delivered key benchmarks to finalize the decision of going towards GPT model.
- Led the key hyperparameter investigation work including learning rate scheduling, large batch size investigation, sparse attention, etc. Designed the model architectures for 600M, 2.7B, 6.7B, 13B, 26B.
- Initiated and led the finetuning pipeline work post pretraining to finetune model on multiple task datasets simultaneously including AWS API, unit test generation, programming problem, etc. Developed **experience-replay finetuning strategy that prevent catastrophic forgetting**, and is the most important training stage post pretraining.
- Led the investigation of **Retrieval-Augmented Generation** for code generation, and proposed a RAG model structure for fast encoding of retrievals.
- Developed multilingual code generation benchmarks **MBXP**, Multi-lingual HumanEval and MathQA-X, and pretrain and investigate the performance of multilingual code LLMs on code generation, insertion, robustness, etc.

Software Engineering Intern

- Worked on AutoML Recommendations with Google Brain and Cloud AI team, mainly focusing on sequence modeling with Convolutional Neural Networks, Recurrent Neural Networks and Attention Networks for recommendation.

Selected Publications

Training LLMs to Better Self-Debug and Explain Code. Nan Jiang, [Xiaopeng Li](#), Shiqi Wang, Qiang Zhou, Soneya Binta Hossain, Baishakhi Ray, Varun Kumar, Xiaofei Ma, Anoop Deoras. Under submission 2024.

Multi-lingual Evaluation of Code Generation Models. Ben Athiwaratkun, Sanjay Krishna Gouda, Zijian Wang, [Xiaopeng Li](#), Yuchen Tian, Ming Tan, Wasi Uddin Ahmad, Shiqi Wang, Qing Sun, Mingyue Shang, Sujan Kumar Gonugondla, Hantian Ding, Varun Kumar, Nathan Fulton, Arash Farahani, Siddhartha Jain, Robert Giaquinto, Haifeng Qian, Murali Krishna Ramanathan, Ramesh Nallapati, Baishakhi Ray, Parminder Bhatia, Sudipta Sengupta, Dan Roth, Bing Xiang. **ICLR 2023**

ContraCLM: Contrastive Learning For Causal Language Model. Nihal Jain, Dejiao Zhang, Wasi Uddin Ahmad, Zijian Wang, Feng Nan, [Xiaopeng Li](#), Ming Tan, Ramesh Nallapati, Baishakhi Ray, Parminder Bhatia, Xiaofei Ma, Bing Xiang. **ACL 2023**

Exploring Continual Learning for Code Generation Models. Prateek Yadav, Qing Sun, Hantian Ding, [Xiaopeng Li](#), Dejiao Zhang, Ming Tan, Xiaofei Ma, Parminder Bhatia, Ramesh Nallapati, Murali Krishna Ramanathan, Mohit Bansal, Bing Xiang. **ACL 2023**

Not All Attention is Needed: Gated Attention Network for Sequence Data. Lanqing Xue, [Xiaopeng Li](#), Nevin L. Zhang. **AAAI 2020**

Learning Latent Superstructures in Variational Autoencoders for Deep Multidimensional Clustering. [Xiaopeng Li](#), Zhouong Chen, Leonard K. M. Poon and Nevin L. Zhang. **ICLR 2019**

Learning to Abstract for Memory-augmented Conversational Response Generation. Zhiliang Tian, Wei Bi, [Xiaopeng Li](#) and Nevin L. Zhang. **ACL 2019**

Building Sparse Deep Feedforward Networks using Tree Receptive Fields. [Xiaopeng Li](#), Zhouong Chen and Nevin L. Zhang. **IJCAI 2018**

Collaborative Variational Autoencoder for Recommender Systems. [Xiaopeng Li](#) and James She. **KDD 2017**

Program Committees

Served as a reviewer for ICLR, ICML, NeurIPS, ACL, AAI, IJCAI.

Project Links