

Bing Liu

v.bingliu@gmail.com
+1 (650) 215-8735
<http://bingliu.me>

Summary

I am a tech lead manager at Meta GenAI, building generative AI technologies for conversational agents. I lead a machine learning team to build and productionize large language models (LLM) at scale. Previously, I was at Meta Reality Labs, led and grew the NLU team to build cutting-edge AI technologies for the voice Assistant for AR, VR, and Metaverse. I am an expert in Machine Learning, NLP, and Conversational AI. I have experience in leading large projects that involve 30+ engineers and cross-functional partners from conception to completion, in ambiguous environments and demanding schedules.

Professional Experience

Meta, Generative AI Engineering Lead - Large Language Models (LLMs) 2023 - present

- Building the foundation Multilingual Llama model from pretraining to fine-tuning to productionization. Leading a team of senior engineers and scientists on the internalization of Meta AI Assistant.
- Built a ML team from the ground up focusing on LLMs for business AI agents. Launched the LLM powered Business Messaging agent which helps businesses better engage with their customers and scale the business.

Meta, Reality Labs Tech Lead Manager - NLP/AI 2018 - 2022

- Led the 0-to-1 effort of building the self-learning system for Assistant that personalized the voice interaction experiences. Launched it to Assistant on VR Headsets, Smart Displays, AR Glasses to support millions of users with top-line product quality and user engagement impact.
- Served as the engineering lead for the v-team across ML, Platform, Data, Analytics, and Product Experience pillars in Meta Assistant in driving the end-to-end product quality program.
- Led and grew the NLU team. Supported a team of ML engineers and scientists to develop cutting-edge NLP technologies for Assistant across multiple product verticals, languages, and device lines.

Google Research Software Engineer Intern - NLP/AI 2017

- Developed state-of-the-art deep learning and reinforcement learning technologies for end-to-end spoken dialog systems at Google Research.

Capio.ai (Acq. by Twilio) Software Engineer - Full Stack 2014 - 2015

- Built the 1st version of Capio's cloud infrastructure and Speech (ASR) APIs. Developed contextual language modeling and language understanding systems power Capio's voice intelligence system.

Oracle Member of Technical Staff - Big Data 2011 - 2013

- Worked with client, pre-sale consultants, and technical staff to deliver the 1st Oracle Big Data Appli-
ance solution in Oracle APAC to a top-tier telecommunications company.

Education

Carnegie Mellon University 2013 - 2018

- Ph.D., Electrical and Computer Engineering
- Research area: natural language processing, spoken dialog system, machine learning, AI
- Thesis: *Learning Task-Oriented Dialog with Neural Network Methods*

Nanyang Technological University, Singapore 2007 - 2011

- B.Eng., Electrical and Electronic Engineering; Minor Degree in Business
- First Class Honors (GPA 4.9/5.0); Dean's List; NTU President Research Scholar

KTH Royal Institute of Technology, Sweden 2010

- Exchange program, Information and Communication Technology

Professional Service

Served as Chair, Senior PC, Editor, and Reviewer for top NLP and AI conferences and journals.

- Area Chair, Large Language Models track @ ACL 2023.
- Senior Program Committee: EMNLP 2023, AACL 2023, ICASSP 2023.
- General Chair, 4th Workshop on NLP for Conversational AI @ ACL 2022.
- Journal Reviewer (past and/or present): IEEE/ACM Transactions on Audio, Speech, and Language Processing; IEEE Signal Processing Letters; Neural Networks Journal; Information Systems Journal;
- Program Committee (past and/or present):
 - NLP Conferences: ACL, EMNLP, NAACL, COLING, SIGDIAL, EACL, AACL
 - Speech Conferences: ICASSP, Interspeech, ASRU, SLT
 - Machine Learning Conferences: AACL, KDD

Invited Talk

- Seminar in NLP, "ChatGPT and Large Language Models", University of California, Santa Cruz, May 2023.
- Panel speaker at the AI & NLP Career Event, University of California, Santa Cruz, Feb 2023.
- Seminar in Natural Language Processing, "Advancing Spoken Dialog System with Deep Neural Networks", University of California, Santa Cruz, Nov 2020
- RE•WORK AI Assistant Summit, "Towards Interactive Learning of Spoken Dialog Systems", San Francisco, Jan 2020
- Rasa Developer Summit, "Interactive Learning of Task-Oriented Dialog Systems", San Francisco, Sep 2019
- GPU Technology Conference (GTC), "Reinforcement Learning for End-to-End Trainable Neural Dialog Models", Santa Clara, Mar 2018

Selected Publications

Published 30+ papers in top NLP, Speech, and AI conferences. See [Google Scholar](#) for full publications.

1. Derek Xu, Shuyan Dong, Changhan Wang, Suyoun Kim, Zhaojiang Lin, Akshat Shrivastava, Shang-Wen Li, Liang-Hsuan Tseng, Alexei Baevski, Guan-Ting Lin, Bing Liu, Hung-yi Lee, Yizhou Sun and Wei Wang, "**Introducing Semantics into Speech Encoders**", in ACL, 2023
2. Zhiyu Chen, Bing Liu, Seungwhan Moon, Chinnadhurai Sankar, Paul Crook, William Yang Wang, "**KETOD: Knowledge-Enriched Task-Oriented Dialogue**", in NAACL, 2022
3. Zhaojiang Lin, Bing Liu, Andrea Madotto, Seungwhan Moon, Zhenpeng Zhou, Paul Crook, Zhiguang Wang, Zhou Yu, Eunjoon Cho, Rajen Subba, Pascale Fung, "**Zero-Shot Dialogue State Tracking via Cross-Task Transfer**", in EMNLP, 2021
4. Andrea Madotto, Zhaojiang Lin, Zhenpeng Zhou, Seungwhan Moon, Paul Crook, Bing Liu, Zhou Yu, Eunjoon Cho, Zhiguang Wang, "**Continual Learning in Task-Oriented Dialogue Systems**", in EMNLP, 2021
5. Zhiyu Chen, Honglei Liu, Hu Xu, Seungwhan Moon, Hao Zhou, Bing Liu, "**NUANCED: Natural Utterance Annotation for Nuanced Conversation with Estimated Distributions**", in EMNLP, 2021
6. Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho, Rajen Subba, "**Leveraging Slot Descriptions for Zero-Shot Cross-Domain Dialogue State Tracking**", in NAACL, 2021
7. Kai Sun, Seungwhan Moon, Paul Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho, and Claire Cardie, "**Adding Chit-Chats to Enhance Task-Oriented Dialogues**", in NAACL, 2021
8. Tianxing He, Jun Liu, Kyunghyun Cho, Myle Ott, Bing Liu, James Glass, Fuchun Peng, "**Analyzing the Forgetting Problem in the Pretrain-Finetuning of Dialogue Response Models**", in EACL, 2021
9. Hu Xu, Seungwhan Moon, Honglei Liu, Bing Liu, Pararth Shah, Philip S Yu, "**User memory reasoning for conversational recommendation**", in COLING, 2020
10. Bing Liu, and Ian Lane, "**Adversarial Learning of Task-Oriented Neural Dialog Models**", in SIGDIAL, 2018
11. Bing Liu, Gokhan Tur, Dilek Hakkani-Tur, Pararth Shah, and Larry Heck, "**Dialogue Learning with Human Teaching and Feedback in End-To-End Trainable Task-Oriented Dialogue Systems**", in NAACL, 2018
12. Pararth Shah, Dilek Hakkani-Tur, Bing Liu, and Gokhan Tur, "**Bootstrapping a Neural Conversational Agent with Dialogue Self-Play, Crowdsourcing and On-Line Reinforcement Learning**", in NAACL, 2018
13. Bing Liu, Tong Yu, Ian Lane, and Ole Mengshoel, "**Customized Nonlinear Bandits for Online Response Selection in Neural Conversation Models**", in AAAI, 2018
14. Bing Liu, Gokhan Tur, Dilek Hakkani-Tur, Pararth Shah, and Larry Heck, "**End-to-End Optimization of Task-Oriented Dialogue Model with Deep Reinforcement Learning**", in NIPS Workshop, 2017
15. Bing Liu, and Ian Lane, "**Multi-Domain Adversarial Learning for Slot Filling in Spoken Language Understanding**", in NIPS Workshop, 2017
16. Bing Liu, and Ian Lane, "**Iterative Policy Learning in End-To-End Trainable Task-Oriented Neural Dialog Models**", in IEEE ASRU, 2017

17. Bing Liu, and Ian Lane, "**An End-to-End Trainable Neural Network Model with Belief Tracking for Task-Oriented Dialog**", in INTERSPEECH, 2017
18. Bing Liu, and Ian Lane, "**Dialog Context Language Modeling with Recurrent Neural Networks**", in IEEE ICASSP, 2017
19. Bing Liu, and Ian Lane, "**Joint Online Spoken Language Understanding and Language Modeling With Recurrent Neural Networks**", in SIGDIAL, 2016
20. Bing Liu, and Ian Lane, "**Attention-Based Recurrent Neural Network Models for Joint Intent Detection and Slot Filling**", in INTERSPEECH, 2016
21. Bing Liu, and Ian Lane, "**Recurrent Neural Network Structured Output Prediction for Spoken Language Understanding**", in NIPS Workshop, 2015

Patents

1. Honglei Liu, Hao Zhou, Seungwhan Moon, Bing Liu, Yulong Qiu, Daniel Chai, Pararth Paresh Shah, Xiaolei Li, Rajen Subba, Hu Xu, "**Personalized conversational recommendations by assistant systems**", US Patent, April 2023
2. Paul Anthony Crook, Baiyang Liu, Pararth Paresh Shah, Bing Liu, "**Dialog Session Override Policies for Assistant Systems**", US Patent, May 2022
3. Zhiyu Chen, LIU Honglei, Hu Xu, Seungwhan Moon, Hao Zhou, Bing Liu, Zihang Meng, Amy Lawson Bearman, Ahmed Magdy Hamed Mohamed, Eric Robert Northup, Jihang Li, Jinsong Yu, Jeremy Gillmor Kahn, Lloyd Hilaiel, Justin Denney, Christophe Chaland, "**Systems and Methods for Implementing Smart Assistant Systems**", US Patent, April 2022