

Dongwei Jiang

✉ jiangdongwei0@gmail.com ☎ 443-453-7610 🌐 Some-random 🎓 Google Scholar 🌐 Personal Website

With six years of industry and research experience in speech processing and self-supervised speech models, my focus is now shifting to LLM, with a particular interest in reasoning and self-improvement.

EDUCATION

JOHNS HOPKINS UNIVERSITY	M.Eng. in Computer Science	Aug 2023 - Till Now
PEKING UNIVERSITY	B.S. in Geographic Information System	Sep 2011 - June 2016

RESEARCH EXPERIENCE

Mentor: Prof. Greg Durrett	UT Austin, Jul 2024 - Till Now
<i>Research Intern. Conducted comprehensive analysis of Chain-of-Thought (CoT) prompting effectiveness across 100+ papers and 20 datasets, demonstrating that CoT benefits are primarily concentrated in math and symbolic reasoning tasks [1].</i>	
Mentors: Prof. Daniel Khashabi, Benjamin Van Durme	JHU, Sep 2023 - Till Now
<i>Research Assistant. (1) Understanding the underlying reason that prevents LLMs from effective self-improvement [3] (2) Distilling rationales from unlabelled data to provide process feedback for various reasoning problems [2] (3) Formulating a consistent and theoretically grounded method for annotating decompositional entailment datasets [5].</i>	
Mentor: Prof. Shay B. Cohen	EDINBURGH, Apr 2023 - Dec 2023
<i>Research Intern. Leveraging theorem prover Lean to solve complex natural language logical reasoning problems using symbolic reasoning methods [4].</i>	
Mentor: Dr. Xiangang Li	DiDi, Jan 2018 - Dec 2020
<i>Research Scientist. Advanced self-supervised speech processing with MPC and Speech-SimCLR [6, 7, 8].</i>	

SELECTED PUBLICATIONS

-
- [1] Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, **Dongwei Jiang**, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, Greg Durrett. *To CoT or not to CoT? Chain-of-thought helps mainly on math and symbolic reasoning*, in submission to ICLR (score 6, 8, 8), 2025
- [2] **Dongwei Jiang**, Guoxuan Wang, Yining Lu, Andrew Wang, Jingyu Zhang, Chuyu Liu, Benjamin Van Durme, Daniel Khashabi. *RATIONALYST: Pre-training Process-Supervision for Improving Reasoning*, in submission to ARR Dec
- [3] **Dongwei Jiang**, Jingyu Zhang, Orion Weller, Nathaniel Weir, Benjamin Van Durme, Daniel Khashabi. *SELF-[IN]CORRECT: LLMs Struggle with Refining Self-Generated Responses*, in AAAI, 2025. ★ Featured in the EMNLP 2024 tutorial *Enhancing LLM Capabilities Beyond Scaling Up*
- [4] **Dongwei Jiang**, Marcio Fonseca, Shay B. Cohen. *LeanReasoner: Boosting Complex Logical Reasoning with Lean*, in NAACL, 2024
- [5] Nathaniel Weir, Kate Sanders, Orion Weller, Shreya Sharma, **Dongwei Jiang**, Zhengping Jiang, Bhavana Dalvi Mishra, Oyvind Tafjord, Peter Jansen, Peter Clark, Benjamin Van Durme. *Enhancing Systematic Decompositional Natural Language Inference Using Informal Logic*, in EMNLP, 2024
- [6] **Dongwei Jiang**, Wubo Li, Miao Cao, Wei Zou, Xiangang Li. *Speech simclr: Combining contrastive and reconstruction objective for self-supervised speech representation learning*, in InterSpeech 2021
- [7] **Dongwei Jiang**, Wubo Li, Ruixiong Zhang, Miao Cao, Ne Luo, Yang Han, Wei Zou, Kun Han, Xiangang Li. *A further study of unsupervised pretraining for transformer based speech recognition*, in ICASSP 2021
- [8] **Dongwei Jiang**, Xiaoning Lei, Wubo Li, Ne Luo, Yuxuan Hu, Wei Zou, Xiangang Li. *Improving transformer-based speech recognition using unsupervised pre-training*, arxiv, 2019

WORK EXPERIENCE

SHOPEE	Beijing, Nov 2021 - Nov 2022
<i>Led a team of four. Responsible for low-resource ASR and its application in products</i>	
<ul style="list-style-type: none">Low Resource ASR optimization: Used Youtube data crawling and filtration to accumulate weakly supervised training data for the base model. Used self-training to accumulate target domain training dataLow Resource ASR annotation: Established best practices and helped training of local agents	

- Product Support: Supported live stream clipping and live stream ongoing content search on various markets. The model was used as an extra source of supply for product introductory videos, which increased 23,000 orders each day

YUANFUDAO

Beijing, Dec 2020 - Oct 2021

Led a team of six. Responsible for the application of End-to-End ASR and Talking Face Generation

- Education ASR Optimization: Implemented relevant techniques including code switch and children speech recognition
- Product Support: Supported quality assurance and content understanding for various applications including children's short videos, live lesson streaming, and voice communication, leading to a monthly saving of 900,000 RMB in costs

DIDICHUXING TECHNOLOGY

Beijing, Jan 2018 - Dec 2020

Led a team of five. Responsible for the development and deployment of ASR and TTS

- End-to-End ASR: Implemented LAS and Speech Transformer with Attention + CTC multi-task loss. Implemented AED + CTC joint decoding and word-level WFST with contextual-graph based hotword-fix. Being one of the first to bring end-to-end ASR model online in China, Speech Transformer replaced all hybrid/CTC models for internal products
- End-to-End Streaming ASR: Investigated and implemented prevalent end-to-end streaming speech recognition solutions including neural transducer, RNN-T, and Mocha
- End-to-End TTS: Implemented Tacotron2 and FastSpeech end-to-end Speech Synthesis model with multi-speaker TTS and GST for personalized TTS

JD.COM INC

Beijing, Jan 2017 - Dec 2017

- Read through Kaldi source code. Became proficient in the concept and implementation of hybrid ASR system. Followed early CTC work from Google and experimented with different network structures, modeling units, and subsampling rate

OUTREACH AND PROFESSIONAL DEVELOPMENT

- Open-source Contribution
 - Led the effort to refactor DiDi Global's internal speech framework code base and released it on Github (project name Athena). Athena contains custom easy-to-use feature extractor, industrial engine for ASR/TTS/ASV/VC, various unsupervised pre-training algorithms and WFST construction/decoding pipeline with end-to-end framework. Currently, the project has 900+ stars and 200+ PRs
- Presentations
 - Summer 2021 Serving the underserved using low-resource speech technology. Company-wide sharing, Shopee
 - Fall 2020 An Introduction to Athena and its application at DiDi. Invited Talk, Linux Foundation AI & Data
 - Fall 2019 Extending BERT to speech recognition. Invited Talk, JiQiZhiXing (leading Chinese AI media platform)
- Research Advising and Mentoring
 - 2024-2025 Alvin Zhang at JHU
 - 2020-2021 Qixin Xie at YuanFuDao
 - 2018-2020 Wubo Li at DiDi Global
 - 2018-2020 Ne Luo at DiDi Global
 - 2018-2020 Ruixiong Zhang at DiDi Global