

# Zhihong Shao

E-mail: [zhshaothu@gmail.com](mailto:zhshaothu@gmail.com)

Phone: +86 13121259158

Web: <https://ZhihongShao.github.io>

---

## RESEARCH INTERESTS

My interests are in natural language processing and deep learning. I am particularly interested in how we can build a robust and scalable AI system that can leverage diverse skills (e.g., tool use and reasoning) to aggregate possibly-heterogeneous information and answer natural language questions precisely regardless of their complexity.

## EDUCATION

**Tsinghua University**, Beijing, China

September 2019 - Present

*Ph.D. Student*, Computer Science and Technology

*Advisor*: [Minlie Huang](#)

**Beihang University**, Beijing, China

September 2015 - July 2019

*B.E.*, Computer Science and Technology

*GPA*: 3.86/4, *Rank*: 2/213

## RESEARCH HIGHLIGHTS

LLM Multi-Step Reasoning & Tool Augmentation

- Train LLMs to Integrate Tool Use into Generation: ToRA [1] (ToRA-34B is the first open-source TOOL-AUGMENTED LLM scoring over 50% on the competition-level MATH dataset, close to GPT-4 solving problems with code, with 770+ github stars);
- Improve Math Reasoning in LLMs via Math Training and RL: DeepSeekMath [10] (DeepSeekMath 7B is the first open-source LLM scoring over 50% WITHOUT RELYING ON TOOLS on the competition-level MATH dataset, close to GPT-4 and Gemini Ultra, with 440+ github stars);
- Optimize LLM-Tool Interaction and Adapt Tools for LLMs based on the Interaction Data: ITER-RETGEN [2] with generation-augmented retrieval and generation-augmented retrieval adaptation;
- Inference-Time Optimization: (i) Prompt Optimization: Synthetic Prompting [4] for automatically synthesizing high-quality CoT demonstrations for self-improvement; (ii) Self-Correction based on Feedback from Tools: CRITIC [3] (Basically the first paper to show that current LLMs struggle with intrinsic self-correction and propose tool-aided correction for more stable improvements).

## PUBLICATIONS

- [1] [ToRA: A Tool-Integrated Reasoning Agent for Mathematical Problem Solving](#)  
**Zhihong Shao\***, Zhibin Gou\*, Yeyun Gong, Yelong Shen, Yujiu Yang, Minlie Huang, Nan Duan, Weizhu Chen  
*International Conference on Learning Representations (ICLR)*, 2024.
- [2] [Enhancing Retrieval-Augmented Large Language Models with Iterative Retrieval-Generation Synergy](#)  
**Zhihong Shao**, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, Weizhu Chen  
*Findings of Empirical Methods in Natural Language Processing (Findings of EMNLP)*, 2023.
- [3] [CRITIC: Large Language Models Can Self-Correct with Tool-Interactive Critiquing](#)  
Zhibin Gou, **Zhihong Shao**, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, Weizhu Chen  
*International Conference on Learning Representations (ICLR)*, 2024.

- [4] [Synthetic Prompting: Generating Chain-of-Thought Demonstrations for Large Language Models](#)  
**Zhihong Shao**, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen  
*International Conference on Machine Learning (ICML)*, 2023.
- [5] [Chaining Simultaneous Thoughts for Numerical Reasoning](#)  
**Zhihong Shao**, Fei Huang, and Minlie Huang  
*Findings of Empirical Methods in Natural Language Processing (Findings of EMNLP)*, 2022.
- [6] [Answering Open-Domain Multi-Answer Questions via a Recall-then-Verify Framework](#)  
**Zhihong Shao**, and Minlie Huang  
*Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022.  
(Best QA system on the [AmbigNQ leaderboard](#))
- [7] [AdvExpander: Generating Natural Language Adversarial Examples by Expanding Text](#)  
**Zhihong Shao**, Zhongqin Wu, and Minlie Huang  
*IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*, vol. 30, pp. 1184-1196, 2022.
- [8] [A Mutual Information Maximization Approach for the Spurious Solution Problem in Weakly Supervised Question Answering](#)  
**Zhihong Shao**, Lifeng Shang, Qun Liu, and Minlie Huang  
*Annual Meeting of the Association for Computational Linguistics (ACL)*, 2021.
- [9] [Long and Diverse Text Generation with Planning-based Hierarchical Variational Model](#)  
**Zhihong Shao**, Minlie Huang, Jiangtao Wen, Wenfei Xu, and Xiaoyan Zhu  
*Empirical Methods in Natural Language Processing (EMNLP)*, 2019.

## PREPRINT

- [10] [DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models](#)  
**Zhihong Shao**, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y.K. Li, Y. Wu, Daya Guo  
*Arxiv abs/2402.03300*, 2024.
- [11] [DeepSeek LLM: Scaling Open-Source Language Models with Longtermism](#)  
DeepSeek-AI  
*Arxiv abs/2401.02954*, 2024.
- [12] [Math-Shepherd: A Label-Free Step-by-Step Verifier for LLMs in Mathematical Reasoning](#)  
Peiyi Wang, Lei Li, **Zhihong Shao**, R.X. Xu, Damai Dai, Yifei Li, Deli Chen, Y.Wu, Zhifang Sui  
*Arxiv abs/2312.08935*, 2023.
- [13] [CoTK: An Open-Source Toolkit for Fast Development and Fair Evaluation of Text Generation](#)  
Fei Huang, Dazhen Wan, **Zhihong Shao**, Pei Ke, Jian Guan, Yilin Niu, Xiaoyan Zhu, and Minlie Huang  
*Arxiv abs/2002.00583*, 2020.

## RESEARCH EXPERIENCE

### Microsoft Research Asia

Sep 2022 - Nov 2023, Beijing, China

Research Intern (Supervisors: *Yeyun Gong, Nan Duan, Yelong Shen, Weizhu Chen*)

- [Knowledge-Grounded Generation]: **ITER-RETGEN** [2], which synergizes retrieval and generation iteratively, is a strong method that enables large language models to leverage intrinsic and extrinsic knowledge flexibly;

- [Tool-Augmented Reasoning]: (i) **Synthetic Prompting** [4] elicits better reasoning in large language models with model-synthesized chain-of-thought demonstrations, achieving a new state-of-the-art on numerical reasoning, symbolic reasoning, and algorithmic reasoning tasks; (ii) **ToRA** [1] integrates natural language reasoning with program-based tool use for mathematical reasoning; ToRA-34B is the first open-source model to attain an accuracy over 50% on the competition-level MATH dataset; (iii) **CRITIC** [3] teaches large language models to correct themselves via interactions with tools.

## CoAI Lab, Tsinghua University

Sep 2019 - Present, Beijing, China

*Ph.D. Student (Supervisor: Minlie Huang)*

- [Knowledge-Grounded Generation]: **RECTIFY** [6], a Recall-then-Verify framework that exploits retrieved knowledge comprehensively, tops AmbigNQ;
- [Tool-Augmented Reasoning]: (i) **CANTOR** [5], the first non-autoregressive numerical reasoner, outperforms  $174\times$  larger PaLM 62B on grade school math; (ii) **MIMAX** [8], a weakly-supervised training algorithm that is applicable to various neuro-symbolic reasoning models;
- [Robust Natural Language Understanding]: **AdvExpander** [7], a novel insertion based textual adversarial attack, reveals new robustness issues.
- [Text Generation and Evaluation]: (i) **PHVM** [9], a data-to-text generation model; (ii) **CoTK** [13], a toolkit for fair evaluation.

## AWARDS

<b>Lenovo Scholarship</b> , Tsinghua University	2023
<b>1st Prize</b> , Comprehensive Scholarship, Tsinghua University	2022
<b>2nd Prize</b> , Comprehensive Scholarship, Tsinghua University	2021
<b>3rd Prize</b> , the National Final of "LAN QIAO CUP" C/C++ Group	2018
<b>1st Prize</b> , National College Students Mathematics Competition (non-math-major)	2016
<b>China National Scholarship</b>	2016, 2017, 2018

## SERVICES

**Reviewer/Program Committee:** ACL, EMNLP, NLPCC, ARR

## TEACHING ASSISTANT

<b>Artificial Neural Network</b>	Fall 2019 - 2022
<i>Instructor: Minlie Huang</i>	

<b>Object-Oriented Programming</b>	Spring 2020 - 2023
<i>Instructor: Minlie Huang</i>	
<i>Also gave guest lectures and made assignments</i>	