Chuyan Zhou

chyzhou24@gmail.com · Github · Homepage

Education

ShanghaiTech University, Shanghai, China 2022.9 – 2026.7 (expected)

Bachelor of Engineering in Computer Science
Advised by Prof. Kewei Tu
Overall GPA 3.85/4.0, ranked 3/162 (top 2.5%), *Major GPA* 4.0/4.0

University of California, Berkeley, California, USA 2024.8 – 2025.1

GLOBE Program in Computer Science major, College of Engineering
Overall GPA 4.0/4.0

Research Experience & Papers under Review

VDI Center ShanghaiTech University 2023.10 – Present

Research Intern Advised by Prof. Kewei Tu

Training-free Process Reward for LLM Reinforcement Finetuning via FOL & MCTS In Progress
Contributor

- In modern RL-based LLM finetuning, ORMs are a common type of rewards which predict the reward based on the final outcome of the task. Though simple, ORMs may suffer reward sparsity issues making LLMs fail to generate useful thinking process. PRMs (Process Reward (Models)), in comparison, are proposed to predict the reward by intermediate steps of the LLMs, providing more fine-grained and informative feedback. However, existing PRMs either (1) evaluate with sentence-level segmentation but need human annotations & offline or online training; (2) are training-free based on statistical heuristics or relying completely on an ORM, and only segment the generation by fixed number of tokens not semantically complete units.
- We designed a neuro-symbolic framework, mainly a training-free process reward paradigm based on First-order Logic (FOL) reasoning and Monte Carlo Tree Search (MCTS):
  - We force LLMs to generating reasoning steps in the $p \wedge q \rightarrow r$ form, thus enabling semantically complete segmantation of the reasoning process.
  - We implement a training-free process reward by integrating a FOL solver (Z3); and utilize an LLM teacher to translate natural language reasoning into FOL for ground-truth correctness verification. We can easily combine this reward with outcome reward from an ORM.
  - For RL finetuning, we plan to leverage MCTS or its variants to (1) find good multiple generations for group sampling; (2) backpropagate the process rewards to previous steps to find a more explainable action-state (Q) value estimation. We aim to develop entropy-guided sampling by first trying to sample some, and then finding the most uncertain step in a trajectory and branching out from there.
- We plan to base the experiments on logical reasoning tasks such as LogiQA, Reclor and AR-LSAT, and evaluate and compare the performance by accuracy improvement after RL finetuning with our proposed process reward.

Augmenting Transformer Language Models with Dependency Graphs ACL 2026 Main Conference
Author

- In this project, we proposed a novel method to augment Transformer-based syntactic language models with dependency graphs, especially with word-synchronous semantic dependency graphs, which is to model the dependencies and tokens jointly in a unified framework with a dependency-predicting head so that the original word sequences are kept intact.
- Since the use of syntactic information in past works is limited to recursive structures i.e. trees, and graphs are more general structures which we consider have more potential information to be utilized, based on a codebase of Transformer-XL, we implemented the model architecture which operates on the dependency graphs and word sequences jointly.

- As a contributor, I conducted baseline experiments on metrics such as perplexity (PTB 1987-89 WSJ Treebanks), SG and BLiMP tests.

Other Contributions

- (2025.9-Present, as a primary contributor) Based on our group's previous research on the Jacobi Iteration method for Latent CoT, I am currently progressing on KV cache expansion on looped transformers with exit gates.
- (2023.12-2025.12, as a primary contributor) Developed a binary classification and regression model predicting packaging efficiency leveraging ESM Encoders, for synthesis (impainting) of gene sequence insertions into AAV2 caspid proteins with the motif for gene therapy vectors.
- (2024.9-2025.2, as a primary contributor) Investigated the feasibility of a Jacobi decoding method in a continuous fashion, as a prototype of parallel continuous chain-of-thought generation (PCCoT).

Miscellaneous Projects & Competitions

LLM-powered Lecture Generation Application UC Berkeley, LLM Agent Hackathon 2024.8 – 2024.12
Primary Contributor Teamed Project

- Independently developed the backend framework of the lecture generation pipeline based on FastAPI as a deployable web service. The backend service involves asynchronous task running using multithreading, task managing by API powered by Redis databases, and a metadata system managing the generated data.
- Worked as a main developer and integrated the respective model components, which run on the above backend framework.
- Developed an additional LLM-powered QA agent based on the built backend, which involves interacting with LLMs using a long context of generated lectures with a RAG system dynamically indexing the sources (e.g. textbooks) for LLM grounding.

Multi-step ASR Pipeline Bengali.AI Speech Recognition Challenge, Kaggle 2023.7 – 2023.10
Primary Contributor Teamed Competition

- We developed an automatic speech recognition (ASR) pipeline for a low-resource language (Bangla/Bengali) in a team of 2. With supervised finetuning from augmented low-resource speech datasets, the inference pipeline mainly consists of denoising → inference main ASR decoder model using finetuned Wav2Vec2 whose decoder probabilities are fused with those of an n-gram draft language model → post-processing by rearranging the punctuation from plain words using a BERT-based punctuation prediction model.
- I worked mainly on the n-gram & punctuation enhancing and the finetuning of the Wav2Vec2 model.

Activities

SI 140A Course (Probability and Statistics for EECS) ShanghaiTech 2025.2 – 2025.7
Head Teaching Assistant

- Held discussion classes in a weekly basis for 40 students, covering the course materials, assignments and providing additional insights. Mainly responsible for creating & grading assignments and exams.

Honors and Awards

- ***Outstanding Student Scholarship***, ShanghaiTech University 2025
- ***Silver Medal*** (24-th place), Award on Bengali.AI Speech Recognition Challenge on Kaggle 2023.10