

# Hao Zheng



Forceless



Forceless



+86 182-1676-7170



wszh712811@gmail.com

## SUMMARY

Hao Zheng is a Master's student (Class of 2023) at the Institute of Software, Chinese Academy of Sciences, expected to graduate in June 2026. His research focuses on LLM Agents and Agentic RL, with full-stack capabilities in large model pretraining, fine-tuning, and RLHF, along with excellent engineering skills. His open-source project PPTAgent has gained 2,000+ Stars on GitHub, and he has published a first-author paper at EMNLP 2025.

**University of Chinese Academy of Sciences · M.S. in Software Engineering** Sep. 2023 – Jun. 2026

Institute of Software · Chinese Information Processing Laboratory  
Sun & Prof. Xianpei Han

Advisors: Prof. Le

**Zhejiang University of Technology · B.S. in Software Engineering** Sep. 2019 – Jun. 2023  
College of Computer Science

## WORK EXPERIENCE

**Personalized Large Language Model**, JD Explore Academy Jun. 2025 – Present

- **Business Context:** Building personalized LLMs with empathy and conversational expression capabilities to serve personalized user interaction scenarios
- **Phase 1 - Offline Fine-tuning Optimization:**
  - Developed a personalized dialogue data generation system: Based on PersonaHub seed data, designed a multi-stage data augmentation pipeline that significantly enhanced conversational expression and emotional empathy; produced 20K rounds of high-quality dialogue data covering core scenarios including knowledge QA, emotional companionship, and casual chat, tailored to the capabilities of different base models
  - External knowledge integration: Retrieved external information through search and tool-calling APIs, enabling comprehensive external knowledge retrieval and integration during both training and inference phases, improving model knowledge timeliness and accuracy
  - Preference learning and evaluation framework: Built multi-dimensional evaluation datasets and iteratively optimized LLM-as-Judge evaluators to guide model iteration; ensured response diversity through multiple sampling and deduplication strategies, constructing a 5K high-quality preference dataset
  - **Results:** Human evaluation showed model response pass rate reached 86%, **a 31 percentage point improvement** compared to open-source baseline model (55%)
- **Phase 2 - Online Reinforcement Learning** (In Progress):
  - PersonaGRM: To address increased hallucination and mechanical dialogue patterns caused by enhanced empathy in Phase 1 model, trained a generative reward model to evaluate response quality. Using Qwen3-14B as the base model, achieved 87% accuracy through GRPO training

## PROJECTS

**PPTAgent: Agent-based Presentation Generation**

Jun. 2024 – Mar. 2025

- Proposed a two-stage agent-based generation method, combining Code Action with LLMs to generate visually rich slides (PPTEval score 3.67 vs. baseline 3.00), achieving  $\geq 95\%$  generation success rate through self-correction mechanisms
- Designed rule-based rewards using ROUGE similarity and trained Qwen2.5-3B model with GRPO algorithm, improving first-attempt slide generation success rate from 37.4% to 99.3% (+61.9pp), surpassing Qwen2.5-72B (85.4%)
- Developed a slide semantic extraction module to extract functional categories and data patterns from reference slides, ensuring consistent content and layout generation
- Built the largest slide datasets to date: Zenodo10K and pptagent-code26k

- Project gained over **2,000 GitHub Stars** after open-sourcing, receiving widespread recognition from the community

## Foundation Model Training

Sep. 2023 – May 2024

- **Fine-grained CommonCrawl Data Processing Framework**: Developed an innovative deduplication algorithm combining distributed domain aggregation with n-gram matching, effectively achieving line-level deduplication of same-domain data, addressing the coarse-grained limitations of traditional document-level deduplication. In GPT-2 experiments, this method reduced PPL by 34%, building PB-scale high-quality training corpora that significantly improved model performance and knowledge timeliness on downstream tasks.
- Trained BERT classifiers using high-quality corpora from Wikipedia, Sina News, Baidu Baike, combined with randomly sampled low-quality data, achieving 91%+ accuracy in data quality classification to ensure high-quality pretraining data
- Integrated constructed 50K Chinese reasoning dataset (Cherry-LLM + Evol-Instruct) with 260K open-source dialogue/preference data; through model fine-tuning and RLHF alignment, achieved 21% improvement in Chinese comprehension (C-Eval) and 15% improvement in multi-turn dialogue quality (MT-Bench)

## PUBLICATIONS

- 
- [1] **Hao Zheng**, Xinyan Guan, et al. “PPTAgent: Generating and Evaluating Presentations Beyond Text-to-Slides”  
**EMNLP 2025 Main Conference** *First Author*