

Junjie Li

☎ +86 16623824868 • ✉ lij850601@gmail.com • 🌐 github.com/Picrew • 🔗 LinkedIn

Research Focus: Agent Harness Engineering, Agentic RL, RAG/Retrieval, Long-Context LLMs, Evaluation

EDUCATION

University of Chinese Academy of Sciences

M.S. in Computer Science and Technology

Research direction: LLM agents, agent harness engineering, retrieval systems, and reinforcement learning.

Beijing, China

Sep 2024 – Jun 2027

Dalian Maritime University (211)

B.Eng. in Electronic Information Engineering

Dalian, China

Sep 2020 – Jun 2024

RESEARCH & PUBLICATIONS

Agent Harness Engineering: A Survey

First / Equal-Contribution Author • [OpenReview Archive](#), May 2026 • [Paper](#) • [Code](#) • [Website](#)

- Defined **agent harness engineering** as the production reliability layer around LLM agents; proposed **ETCLOVG** and mapped **170+** open-source harness systems.

Lost in Stories: Consistency Bugs in Long Story Generation by LLMs

First Author • [ACL 2026 \(Accepted\)](#) • [Paper](#) • [Code](#) • [Website](#)

- Proposed **ConStory-Bench** (2,000 prompts, 5-category / 19-subtype error taxonomy) and **ConStory-Checker**, an evidence-grounded contradiction-detection pipeline for long-form narrative consistency evaluation.

Good Teachers, Better Students: A Survey of Reward Models for LLMs

Co-First Author • [TechRxiv preprint](#), Dec 2025 • [DOI](#)

A hyperparameter-fusion neural network for deposition prediction

Co-Author • [Engineering Applications of AI \(JCR Q1\)](#), 2025, Vol. 162 • [DOI](#)

INTERNSHIP

Microsoft | LLM Research Intern

Beijing, China • Apr 2025 – Apr 2026

- **Led the Agentic Deep Research harness:** reproduced *Search-R1* as baseline and upgraded single-shot QA into an Agentic-RL-trained **multi-turn evidence-seeking agent** with search tools, evidence context, rollout lifecycle, and verification feedback for multi-hop reasoning.
- **Resolved reward sparsity in multi-hop GRPO:** used three-tier staged reward shaping – format compliance, evidence grounding, and final correctness – to add dense early-phase signal and stabilize convergence.
- **Fixed tool-use exploration collapse:** combined forced exploration with retrieved-token masking so on-policy GRPO did not reinforce the no-search shortcut after stochastic correct rollouts.
- **Engineered Qwen3-30B-scale GRPO post-training:** built async rollout/training, vLLM-accelerated inference, gradient checkpointing, and mixed precision; final model learned stable evidence-first reasoning beyond prompting and SFT baselines.

Shanghai AI Laboratory | LLM Algorithm Intern

Shanghai, China • Dec 2024 – Apr 2025

- **Owned end-to-end delivery of Text2Brep** on the InternAI **Wing-Wing** platform: designed a 7-sub-model language+diffusion system that generates industrial-grade B-Rep CAD from natural-language specifications.
- **Addressed multi-model error propagation:** used staged freeze-then-joint fine-tuning plus interface-level geometric consistency loss to control upstream drift into invalid downstream topology.

- **Built industrial STEP/CAD data infrastructure** for format normalization, geometric repair, and topology completion, enabling million-scale training/retrieval with **300%** throughput improvement.
- **Applied RL reasoning-policy optimization** on DeepSeek-R1 / Qwen-7B / Llama-8B for autonomous task decomposition and geometric verification: **~50%** faster inference, **+15%** accuracy, and **95%+** precision.

NetEase  | LLM Algorithm Intern

Beijing, China • Aug 2024 – Dec 2024

- **Led RAG core-model training upgrades** for open-source QAnything ([GitHub](#) / [qanything.ai](#)), shifting from prompt tuning to retrieval-model fine-tuning after diagnosing generic embeddings as the bottleneck on Youdao verticals.
- **Fine-tuned a domain embedding model**: mined production query-document pairs and designed BM25 hard negatives plus ratio-tunable easy negatives, materially outperforming random-negative baselines on domain recall.
- **Trained a cross-encoder reranker** from online bad cases where the correct document was retrieved but ranked low, using combined pointwise + pairwise loss to improve top-K precision.
- **Built a closed-loop RAG quality-governance system**: offline benchmarks, online probes, bad-case harvesting, layered fault attribution, regression gates, GraphRAG multi-hop retrieval, answer verification, and parsing fixes.

PROJECTS

Seek-Writer • AI Long-Form Fiction Studio ([seekcopywriter.com](#))

- Independently designed, developed, and launched a writer-first platform with split-view drafting, structured character/chapter state, continue / rewrite / generate modes, and an integrated Seek Agent for chapter planning and consistency repair.
- Focused on long-context agent optimization: key-fact memory compaction, chapter-aware model routing, and persistent local story state for stable cross-session generation.

StoryForge • Open-Source Long-Form Writing CLI

- Developed an open-source CLI compiling structured story state (character sheets, timeline, foreshadowing) with chapter-level planning, token-budget control, and consistency self-check hooks to reduce multi-chapter continuation drift.

HONOURS & COMPETITIONS

Outstanding Student Scholarship	Dec 2023
National Undergraduate Embedded Chip and System Design Contest, National First Prize	Aug 2023
National Undergraduate Electronic Design Contest (TI Cup), Provincial First Prize	Aug 2023
National Undergraduate Intelligent Vehicle Competition, National Second Prize	Aug 2023
China Robot and Artificial Intelligence Competition, National First Prize	Jun 2023
China Collegiate Intelligent Robot Creativity Competition, National Second Prize	Aug 2022

TECHNICAL SKILLS

- **Agent Harness Architecture**: ETCLOVG taxonomy across execution, tooling, context, lifecycle, observability, verification, and governance; MCP-style tool contracts, sandbox / computer-use runtimes, and permission models.
- **Evaluation & Reliability**: trace-based failure diagnosis, regression gates, LLM-as-judge review, fault attribution, benchmark construction, and evidence-grounded checker design.
- **Post-Training & Retrieval**: SFT, LoRA/QLoRA, DPO-style preference optimization, GRPO/PPO, reward shaping, forced exploration, retrieved-token masking, embedding and cross-encoder reranker training.
- **Infrastructure & Engineering**: Python, PyTorch, TensorFlow, vLLM, async RL rollout, 30B-scale gradient checkpointing + mixed precision, SQL, Shell, Git, Docker.