Availability: Immediate start – Dec 2026 (Long-term). **Schedule:** Full-time during semester breaks, Part-time during academic terms

Chunxiao (Elin) Ren

elin.chunxiao@gmail.com | (+ 65) 82616564

EDUCATION

National University of Singapore (NUS) -- *Singapore* | Expected: 2025.08-2026.12

- *MSc Master of computing (Specialization in Computer Science)*

Lappeenranta-Lahti University of Technology (LUT) -- *Finland* | 2022.09-2025.06

- *BSc Software & System Engineering (Double Degree) |Overall: GPA 4.87/5 (Top 10%)*

Hebei University of Technology (211) -- *China* | 2021.09-2025.06

- *BSc Computer Science (Double Degree)*

RESEARCH EXPERIENCE

Policy Evidence Knowledge Graph Platform | *Research Assistant* 2026.01-present

- Designed an anti-hallucination control framework for LLM Agents: agent accesses Neo4j exclusively via a KG API, with layered validation including citation whitelist checks, directional conflict detection, unit normalization, and NLI / LLM-as-judge post-hoc review — ensuring full answer traceability
- Built a GraphRAG-style retrieval pipeline with 5-way parallel retrieval and query decomposition; applied driver-level Cypher splitting with deduplication-merge to improve multi-hop query recall and citation diversity
- Developed a dual-track confidence scoring system modeling evidence quality (source credibility, coverage) and answer quality (consistency, inference distance) separately; distinguished Extract vs. Extrapolate claims; persisted user feedback and confidence snapshots to SQLite to drive OLS calibration and feature weight hot-updates
- Built an end-to-end paper ingestion and evaluation pipeline covering PDF parsing, finding extraction, snippet-claim alignment, and Neo4j ingestion; designed an LLM-as-judge auto-evaluation module scoring Agent responses on accuracy, traceability, and consistency across three dimensions to support iterative system improvement

INTERN EXPERIENCE

Beijing Seeyon Internet Software Corp. | *LLM Machine Learning Engineer Intern* Beijing, China | 2025.05 – 2025.08

- Built a vertical Code-Generation Agent for OA workflow automation on Seeyon V5 PaaS, enabling non-technical users to generate executable Python business logic from natural language. Designed the full pipeline: requirement parsing → template retrieval → code generation → deployment, reducing manual configuration time by 20%.
- Designed a model evaluation and prompt optimization framework to benchmark Qwen, GLM, and DeepSeek on reasoning and code synthesis tasks. Developed structured system prompts and few-shot constraints to stabilize output format and improve controllability for the internal CoMi agent.
- Improved agent reliability via Supervised Fine-Tuning (SFT) on code generation tasks. Built a custom training pipeline integrating AST-based loss and semantic consistency loss to enforce syntactic correctness and logic alignment.

PROJECT EXPERIENCE

Automated Valuation Model for Singapore Public Housing 2025.09-11

- Integrated 5 categories of geospatial facility data (~774 POIs including MRT stations, schools, and shopping malls) and engineered ~20 proximity features per sample using sklearn BallTree with Haversine metric — covering dual-radius density counts, nearest distances, and facility-tier flags — achieving ~10× speedup over brute-force search
- Designed a compound sample-weighting scheme combining exponential time decay (to emphasize recent market trends) and high-price boosting, with monotonicity constraints on floor area and remaining lease to prevent counter-intuitive predictions
- Built a CatBoost + LightGBM + XGBoost stacking ensemble with 5-Fold OOF meta-features and a linear meta-learner; further applied two-stage refinement for top-10% high-value properties and 3-seed averaging, reducing validation RMSE (log space) from 0.061 to 0.050 (~18% improvement)

TECHNICAL SKILLS

- Languages & Data Tools: Python, SQL, pandas, NumPy, scikit-learn, Jupyter
- Machine Learning: CatBoost, LightGBM, XGBoost, Stacking Ensemble, PyTorch
- LLM & NLP: RAG / GraphRAG, NLP / NLI, Embeddings, LangChain, LangGraph, Prompt Engineering, SFT / PEFT / LoRA
- Model Deployment: FastAPI (REST API + Streaming), Vercel, Railway
- Databases & Storage: Neo4j, PostgreSQL, SQLite, Redis, Vector DB