

XUELI PAN

AI ENGINEER | LLMs, RAG, KNOWLEDGE GRAPHS & ML SYSTEMS

Amsterdam, Netherlands | +31 06 831 71882 | x.pan791@gmail.com | xuelipan.github.io | linkedin.com/in/xueli-pan

PROFESSIONAL SUMMARY

AI engineer and applied AI researcher with 5+ years of experience designing end-to-end machine learning and data systems, with a focus on LLM applications, retrieval-augmented generation, semantic search, knowledge graphs, and model evaluation. Strong in Python, PyTorch, Hugging Face Transformers, retrieval pipelines, API integration, Docker, and scalable batch processing. Experienced in fine-tuning open-source LLMs, building reproducible inference and evaluation workflows, and translating ambiguous requirements into measurable AI solutions for technical and non-technical stakeholders.

CORE SKILLS

Applied AI & LLMs	LLM applications, RAG, prompt engineering, supervised fine-tuning, LoRA/QLoRA, embeddings, semantic search, recommendation systems, NLP, structured-output generation
Programming & APIs	Python, SQL, SPARQL, pandas, NumPy, PyTorch, Hugging Face Transformers, FastAPI, REST APIs, backend integration, asynchronous programming fundamentals
Retrieval & Data	LlamaIndex, ChromaDB, Sentence-BERT, RDF/knowledge graphs, QLever, Virtuoso, GraphDB, Neo4j, ETL, data validation, feature/data preprocessing
Engineering & Delivery	Git, Linux, Docker, HPC/SLURM, AWS basics, Databricks basics, scalable batch pipelines, experiment design, error analysis, reproducible workflows
AI-Assisted Development	Claude Code, Codex, modern AI-assisted coding and documentation workflows

PROFESSIONAL EXPERIENCE

PhD Researcher - AI Engineering & LLM Systems VU Amsterdam, Netherlands	Sep 2019 - Present
- Owned end-to-end development of applied AI systems for question answering, retrieval, and recommendation over scientific knowledge, covering problem definition, data preparation, model and retrieval design, evaluation, error analysis, and iterative improvement. - Developed FIRESPARQL, an LLM-based natural-language-to-SPARQL system; fine-tuned Llama 3 8B with LoRA using Hugging Face and LLaMA-Factory, built inference and evaluation pipelines, and achieved 0.85 relaxed exact match with 448 of 513 successful executions on SciQA. - Built ontology-grounded RAG workflows using GPT-4, embedding-based retrieval, and structured prompts to generate competency questions; designed evaluation criteria and validation processes across two domains. - Designed a property-guided subgraph retrieval framework combining semantic entity/property alignment, graph traversal, LLM-based filtering, and structured evidence selection for multi-hop knowledge-graph question answering. - Engineered reproducible batch pipelines to execute and evaluate 2,600+ SPARQL queries on million-scale RDF data using Python, QLever, Virtuoso, Docker, and HPC/SLURM; added validation, failure categorization, and comparative evaluation. - Built an end-to-end recommendation pipeline over 100k+ scientific papers, from parsing and preprocessing through embeddings, user-profile generation, ranking, and evaluation; introduced parallel processing to improve runtime. - Collaborated with supervisors, researchers, and external partners across multiple countries; presented applied AI work to academic and industry audiences, including Elsevier, and communicated technical trade-offs and outcomes clearly.	

EARLIER BUSINESS EXPERIENCE

Data Analyst - User Growth Didi Chuxing, Beijing, China	Sep 2018 - Jan 2019
- Prepared, cleaned, and segmented user-growth data for recurring performance analysis and growth decision support. - Clarified stakeholder questions, translated them into analysis-ready datasets and logic, and communicated findings in an actionable form.	
Data Analyst Wiley, Beijing, China	Mar 2018 - Sep 2018
- Prepared and analyzed publishing and operational datasets for performance reporting and ad hoc business analysis. - Improved data reliability through validation, cleaning, and repeatable, documented reporting workflows.	

EDUCATION

PhD in Artificial Intelligence Vrije Universiteit Amsterdam, Netherlands	2019 - 2026 (expected)
Research: LLMs, RAG, information retrieval, question answering, knowledge graphs, recommendation systems, and AI evaluation.	
MSc in Information Science Institute of Scientific and Technical Information of China, China	2016 - 2019
BSc in Library Science Northeast Normal University, China	2010 - 2014

LANGUAGES

English: C1 professional working proficiency | Mandarin: Native | Cantonese: Native | Dutch: A1, actively learning