

Hyeontae Kim

 hyeontae.me |  github.com/h383kim |  h383kim@uwaterloo.ca |  [LinkedIn](#)

EDUCATION

University of Waterloo

Data Science Co-op, Bachelor of Honours Mathematics

Waterloo, ON, Canada

GPA: 3.8/4.0

EXPERIENCE

Data Scientist @ Verily

Jan 2026 - Apr 2026

ex-Google Life Sciences (NASDAQ: GOOG)

Waterloo, ON

- Developed a protocol-driven **medication refill agent** with **LangGraph-based** orchestration for EHR eligibility review and safety screening to reduce clinician review time while preserving clinical safety. Achieved **95.1% decision accuracy** and a **91.1% mean rubric score** across end-to-end clinical workflows.
- Built an AI-powered **FHIR curation pipeline** converting raw patient records into LLM-ready longitudinal narratives, reducing **token usage by 97.3%** while improving clinical **QA accuracy by 3%**.
- Shipped a **GraphRAG clinical QA agent** using Text2Cypher retrieval over a Neo4j knowledge graph with dynamic result processing, achieving **92% accuracy** and **95% retrieval recall** as the team's production baseline.
- Designed **evaluation pipelines** with rubric-based LLM-as-a-Judge scoring and multi-agent rubric generation, achieving **>60% clinical intent alignment** and **0.966 AUROC**.

Data Scientist @ Rocket Innovation Studio

May 2025 - Aug 2025

Rocket Companies (NYSE: RKT)

Detroit, MI

- Researched **deep learning for tabular data** (TabNet, DCN, TabTransformer) as XGBoost alternatives for banker-lead allocation, projecting up to a **60% conversion-rate improvement**.
- Deployed a **banker-lead recommender system** on AWS SageMaker, integrating preprocessing, PyTorch models, and inference wrappers for low-latency production inference (**<80 ms/request**).
- Tracked **>200** experiments, artifacts, and schema versions with MLflow; reduced debugging/deployment turnaround by **40%** through schema validation and dependency logging.
- Designed online **A/B test plan** and ran rigorous offline backtests (time-based holdouts/shadow scoring) comparing XGBoost vs neural networks.

Data Scientist @ Canadian IP Office

May 2024 - Apr 2025

Government of Canada

Gatineau, QC

- Built a large-scale entity resolution pipeline for **CA (13M+)** and **global (330M+)** patent records on AWS.
- Combined SBERT (Siamese) embeddings with a gradient-boosting classifier and fuzzy/blocking rules, cutting duplicate entities by **35%** and improving match **F1 by 40%**.
- Optimized **batch ETL pipelines for classification tasks**, reducing end-to-end processing time by **82%** while handling large-scale data (**350M+** records).

PROJECTS

Smart Air Purifier Voice Control ([GitHub](#), [HuggingFace](#))

FastAPI, STT, vLLM

- Built an end-to-end **voice-control system** for an interactive **air-purifier robot**, converting speech into validated device actions with **SenseVoice STT**, a **LoRA-fine-tuned 4-bit Llama-3.2-1B** model, **grammar-constrained decoding**, and a stateful animated UI. Achieved **91.6%** audio-command accuracy within a **3 GB VRAM / 0.45 s** on-device target.

VLM Serving Benchmark & Routing Platform ([GitHub](#))

vLLM, llama.cpp, FastAPI/Pydantic

- Built a VLM benchmarking and routing platform for document extraction, running one-command sweeps across Transformers, vLLM, llama.cpp, and FP16→INT4 quantization to select the cheapest backend that meets accuracy and JSON-reliability SLAs.

TECHNICAL SKILLS

Programming Languages: Python, C++/C, SQL (PostgreSQL, Oracle), R, Go, Bash

Frameworks & Tools: LangGraph, LangChain, PyTorch, TensorFlow, MLflow, PySpark, Airflow, Docker, Kubernetes, W&B, FastAPI, AWS, GCP

Libraries: scikit-learn, Pandas, NumPy, OpenCV, Matplotlib, NLTK, spaCy, Selenium, BeautifulSoup