Davronbek Malikov

LLM & RAG Engineer | Agentic AI Engineer | ML & Deep Learning Engineer | Computer Vision | Ph.D.
Seoul, South Korea
davronbekmalikov96@gmail.com | +82-10-2474-2901 | LinkedIn | Portfolio

Summary

Ph.D. in AI Convergence Engineering and AI Engineer & Team Lead at TripleH, leading the team that builds on-premise multi-agent LLM systems for enterprise clients. Post-trained LLMs with SFT (LoRA, QLoRA), DPO, and RLHF on limited GPU resources (free Kaggle T4 GPUs and a laptop CPU), raising a Qwen2.5 model's test score from 15/68 to 44/68; built RAG pipelines and LangGraph agents monitored with LangSmith and, for on-premise deployments, Arize Phoenix.

Builds end-to-end prediction models: feature engineering, leakage-safe validation, XGBoost, LightGBM, and stacking ensembles, Optuna tuning, and SHAP explainability. Author of 5 first-author papers (SCIE), with published models improving prediction accuracy by up to 17%, deployed as live apps.

Technical Skills

- **LLM & RAG:** SFT, LoRA/QLoRA, DPO, RLHF (PPO, RLOO), vLLM, Ollama, llama.cpp, GGUF/AWQ/GPTQ, hybrid RAG (BM25 + dense, RRF), cross-encoder reranking, GraphRAG, pgvector, Qdrant, Milvus
- **Agents & Observability:** LangGraph, LangChain, MCP, A2A, LangSmith, Arize Phoenix, OpenTelemetry, Ragas, DeepEval, LLM-as-a-Judge
- **Machine Learning:** Scikit-learn, XGBoost, LightGBM, stacking ensembles, AutoML, Optuna, SHAP, Pandas, NumPy
- **Deep Learning & Vision:** PyTorch, TensorFlow/Keras, Hugging Face (Transformers, PEFT), CNNs, Transformers, YOLO, OpenCV
- **MLOps & Engineering:** Python, SQL, FastAPI, Docker, GitHub Actions, MLflow, Weights & Biases, PostgreSQL, Redis, Kafka, Next.js

Experience

AI Engineer & Team Lead, TripleH, Seoul — Dec 2025–Present

- Lead the AI engineering team building production autonomous agent systems for enterprise clients, from research and prototyping to on-premise and cloud deployment; set engineering standards, run code reviews, and mentor engineers.
- **Multi-Agent Orchestration Platform:** central hub that plans, routes, and monitors specialized sub-agents (automated market reports, a financial-document RAG assistant, an enterprise-messenger chatbot).
- **Legal Research Assistant:** multi-LLM ensemble that queries several models in parallel, aggregates their answers, and verifies every cited regional legal precedent before it reaches the user.
- **On-Premise AI Platform:** every agent runs locally on open-weight LLMs (Llama 4, Qwen3) with no external API calls, so no data ever leaves the client. Also delivered a meeting-intelligence agent (on-premise speech-to-text) and per-employee AI assistants. Stack: SFT/LoRA and DPO, hybrid RAG (pgvector, Qdrant), LangGraph, LangSmith and Arize Phoenix monitoring.

Independent Researcher (Sports Analytics) & AI Educator, Self-directed and online — 2020–2025

- Researched ML, deep learning, and LLMs for soccer analytics; taught machine learning, deep learning, computer vision, and LLM courses and mentored students on end-to-end AI projects.

Graduate Researcher, Gyeongsang National University, South Korea — 2018–2025

- Big Data Systems Software Lab (2022–2025): ML for soccer analytics, covering salary prediction with traditional ML and AutoML, a dual expected/actual-goals performance model, and injury-type prediction from engineered features. Distributed Systems Lab (2018–2021): mathematics for distributed systems. 5 first-author papers (see Publications).

Mathematics Instructor, Fergana, Uzbekistan — 2014–2017

- Prepared students for university entrance exams, the SAT, and international mathematics exams.

Selected Projects

LLM & Agentic Systems

ERA Academy: AI EdTech platform, multi-agent engine, permission-aware RAG — [GitHub] · **[Live]**

- One provider-agnostic agent engine behind three role-scoped assistants (tutor, teacher, admin): Claude, OpenAI, and local Ollama adapters, SSE streaming, usage tracking, and cost and rate limits.
- RAG over uploaded course files with pgvector; access rules enforced inside the retrieval SQL query, cited answers, prompt-injection defense; 206 automated tests, all passing.

Touchline: AI scouting copilot (Claude, LangGraph, hybrid RAG, LangSmith) — [GitHub]

- Claude agent on LangGraph over 7 schema-validated tools with human-in-the-loop shortlist approval; hybrid RAG (dense + BM25, RRF) in Qdrant with cited reports; DuckDB warehouse of 50,149 players plus StatsBomb event data.
- Every run traced in LangSmith; 120-question golden set run as experiments (filter accuracy, precision/recall@10, numeric grounding, invented-player detection); 102 automated tests.

TeachHelp: Post-training an LLM into an Uzbek and English teaching assistant [GitHub] · **[Live]**

- Qwen2.5 through SFT → DPO → RLHF, 14 models on one 68-question held-out test: SFT raised Qwen2.5-1.5B from 15/68 to 44/68; DPO from scratch, with v1's length bias diagnosed and fixed in v2 (44/68, refusals 8/8).
- RLHF with Bradley-Terry reward models and PPO/RLOO across 4 judges, with reward-hacking checks.

Portiq: Multi-agent job hub (LangGraph, Kafka, A2A, pgvector) **[Live]**

- LangGraph orchestrator routes 11 agents over Kafka and A2A, with a Judge agent, pgvector job matching, and OpenTelemetry and Jaeger tracing.

Machine Learning & Deep Learning

LLM from Scratch: GPT-style Transformer in PyTorch *(pretraining in progress)* [GitHub]

- Implemented a decoder-only Transformer from first principles: multi-head causal self-attention, MLP blocks, weight tying; verified with shape, initial-loss, overfit-one-batch, and causality tests.
- Full pretraining pipeline: cleaning, filtering, exact and MinHash near-duplicate detection, BPE tokenizer, binary token datasets; warmup + cosine LR, gradient accumulation and clipping, bf16/fp16 mixed precision, resumable checkpoints; perplexity evaluation and temperature/top-k/top-p sampling.

Soccer Player Salary Predictor: end-to-end supervised ML [GitHub] · **[Live]**

- Leakage-safe pipeline on 8 seasons of Bundesliga, La Liga, and Serie A data: 21 engineered features, baseline → gradient boosting (XGBoost, LightGBM) → stacking and voting ensembles, Optuna Bayesian tuning.
- SHAP global and local explanations; Gradio demo on Render with uptime monitoring; 17% accuracy gain (SCIE 2025).

TacticalDNA: unsupervised player-archetype discovery [GitHub] · **[Live]**

- Clustered 6,314 player-seasons on per-90 features; KMeans, Agglomerative, DBSCAN, and Gaussian Mixture compared on Silhouette, Davies-Bouldin, and Calinski-Harabasz; a blind stat line matched prime Messi at 98% cosine similarity.

More on the portfolio: TACTIQ (evidence-first football RAG, in development), and RAG-powered apps Palov, Hamyon, and Kotib, at davronbek-portfolio.vercel.app.

Publications (first author)

- *Predicting Soccer Player Salaries with Both Traditional and Automated Machine Learning Approaches.* Applied Sciences (SCIE), 2025. [DOI]
- *Beyond xG: A Dual Prediction Model for Analyzing Player Performance Through Expected and Actual Goals in European Soccer Leagues.* Applied Sciences (SCIE), 2024. [DOI]
- *Comparative Evaluation of Machine Learning Models for Predicting Soccer Injury Types.* KSCIE, 2024. [DOI]
- *The Application of Machine Learning on the Injury Prediction of Soccer Players.* ISE@APSEC, 2022. [PDF]
- *On the Construction and Properties of Lattice-Group Structure in Cartesian Product Spaces.* Journal of Computer Science, 2020. [DOI]

Education

Ph.D. in AI Convergence Engineering, Gyeongsang National University, Republic of Korea 2020–2025
Thesis: *Advancements in Soccer Analytics: Players' Goals and Salary Prediction Models with Machine Learning*
Master of Informatics, Gyeongsang National University, Republic of Korea 2018–2020
Bachelor of Mathematics, Fergana State University, Uzbekistan 2014–2018

Languages

English (Fluent), Uzbek (Native), Turkish (Advanced), Russian (Intermediate), Korean (Conversational)