Doğaç Eldenk

Evanston, IL
dogacel@gmail.com

Linkedin GitHub Website

I am currently doing my **Master's in Computer Science at Northwestern University**, working on simulating, running and optimizing AI models on distributed settings and edge devices. I have past research experience on software engineering, machine learning and generative AI. I am passionate **open-source** projects and contributed to over 40+.

Previously, I worked over 4 years of expertise in building, designing and optimizing **distributed systems**, **cloud infrastructure**, scalable **high-performance web applications** and developer tooling.

Technical Skills

Languages	Python, Kotlin, Java, Scala, TypeScript, JavaScript, Go, C++
Technologies	SQL, REST, CUDA, gRPC, Protobuf, Terraform, PostgreSQL, GitHub Actions, Envoy, Apache Airflow
Frameworks	PyTorch, Spring Boot, Armeria, React, Koin, Exposed
Tools	AWS, SLURM, Conda, Git, Gradle, Grafana, Linux

Experience

Northwestern University
Evanston, IL

Graduate Researcher — *June 2025 - Current*

Intelligent Mobile and Embedded Computing (IMEC) Lab

- **Design and train custom LMs** that are used for speculative decoding to accelerate LLM inference (EAGLE).
- **Conducting research on AI systems** with a focus on AI efficiency and distributed AI.
- **Built a discrete-event simulation suite** to allow running multiple AI models in a simulation to profile performance of distributed algorithms. Conducted large-scale experiments and numerical analysis to derive scaling laws.
- **Extended state-of-the-art AI models**, such as Llama 3, CLIP ViT, with support for parallelization techniques such as Tensor Parallelism, Model Parallelism, and other state-of-the art methods based on recent research papers.
- **Developing and prototyping distributed and efficient inference techniques**, including communication–computation overlap, speculative decoding, and early exiting, aimed at accelerating transformer inference.

Carbon Health (U.S.)
Remote - Chicago, IL

Software Engineer - Senior — *November 2023 - August 2025*

- **Re-designed and implemented** a new generic distributed fine-grained authorization solution compatible with the microservice architecture and also migrated over 1000+ endpoints with **zero down-time**.
- **Lead architecture and management** of the service infrastructure, including **observability**, monitoring, tracing, **scaling**, **networking**, security, and deployment pipelines, ensuring **high reliability and efficiency**.
- **Architected** the pipeline for an **LLM-based AI** to automatically respond to patient messages, schedule appointments, assign tasks and later implemented together with the Data Science team.

Software Engineer - Mid Level — *January 2022 - November 2023*

- Migrated **50+ tables** and **100+ gRPC** endpoints with **0 downtime** from monolith to microservices, remodelled top 3 biggest tables in the company with over **200M+ rows** with challenges such as **query performance optimizations**.
- Developed developer tools, such as **job schedulers**, remote **load testing** pipelines, **protobuf code generators**.
- **Full-stack development** with React-Native with Redux, GraphQL with Apollo and Netflix DGS.

Software Engineer - Junior — *December 2020 - January 2022*

Redgate Software (U.K.)
July 2019 - September 2019

Software Engineer - Intern — *Cambridge, UK*

Education

Northwestern University — Evanston, IL
M.Sc in Computer Science, GPA 4.0 / 4.0 — *September 2025 - September 2026 (Est.)*

Bilkent University – (Comprehensive scholarship) — Ankara, Turkey
B.Sc. in Computer Science, GPA 3.75 / 4.0 - Summa Cum Laude — *Sep 2017 - Jan 2022*

Research and Publications

Uniference: Using Discrete Event Simulation for Developing Distributed AI Models — FMSys 2026
Doğaç Eldenk and Stephen Xia

EarSleeve: Transforming Everyday Earphones into a 12-Lead ECG Sensing Platform — SenSys 2026
Junxi Xia, Doğaç Eldenk, Honjun Xu, Yang Liu and Stephen Xia

Incidents During Microservice Decomposition: A Case Study — EASE 2025
Doğaç Eldenk and Alperen Çetin

Best Paper Award – A case study analyzing 107 incidents during Carbon Health's microservice decomposition over five

Learning Portrait Drawing with Unsupervised Parts — International Journal of Computer Vision
Burak Taşdemir, M. Güdükbay, Doğaç Eldenk, Adil Meriç, Ayşegül Dündar

An image translation architecture combining high-level semantic understanding with unsupervised parts and identity preservation. Proposed a novel asymmetric pose-based cycle consistency loss that mitigates reconstruction constraints.

Key Projects and Contributions

Universal DeepSeek OCR / DeepSeek OCR 2 — 🤗 Dogacel/Universal-DeepSeek-OCR-2
CPU & MPS support to DeepSeekOCR, over 20k+ downloads on huggingface.

Large-Scale LLM Fine-tuning for Kubernetes Configurations — 🤗 Dogacel/Qwen3-Coder...Kubernetes

- **Fine-tuned** Qwen3-Coder-30B on Kubernetes Q&A data, achieving performance competitive with SOTA models.
- Implemented distributed training pipelines using 8×H100 GPUs, including data processing, training and evaluation.

Armeria - Microservice Framework — line/armeria
Contributor

- Contributed over 25+ issues and 15+ PRs, with over **12,000+ lines of code**. Mainly worked on **gRPC HTTP/JSON transcoding**, custom API development platform, and Kotlin support.

Kotlinx Protobuf Gen — Dogacel/kotlinx-protobuf-gen

- An open-source project for generating Kotlin data classes that use **kotlinx.serialization** by parsing protobuf files, used by several other open-source projects in production.

Defenchess, UCI Chess Engine — cetincan0/Defenchess
Co-author

- Modern chess engine, **ranked top 10 in the world** with 3285 ELO among other chess engines on CCRL.

Other Achievements and Participation

- 680th out of 2 million in National University Placement Exam
- Speaker - PgDay Chicago 2025: Integration testing with PostgreSQL