

Parsa Miraghaei

✉ Parsa.miraghaei@gmail.com ☎ +358 466 450 848 📄 in/parsa-miraghaei 🌐 parsamiraghaei.com

As a Machine Learning Engineer with a passion for continuous learning and innovation, I bring a strong foundation of academic excellence and professional experience.

Currently finalizing my Master's degree in Signal Processing and Machine Learning at Tampere University, I have developed a deep understanding of machine learning models, including their limitations, edge cases, and vulnerabilities, enabling me to design reliable and robust systems for real-world applications.

My hands-on expertise spans the full spectrum of AI development. I specialize in end-to-end development, from designing infrastructure with Infrastructure as Code (IaC) and architecting systems for CI/CD pipelines to building scalable APIs, managing large-scale data infrastructures, and implementing advanced machine learning workflows. This expertise extends to crafting intuitive user interfaces and delivering optimized front-end endpoints, ensuring seamless integration and functionality across the entire stack.

My holistic approach enables rapid development, debugging, and deployment of robust, scalable solutions tailored to user and business needs. I have worked extensively with big data and cloud services, leveraging my skills in web development and machine learning to deliver scalable solutions. With a comprehensive understanding of the ML ecosystem—including configuration, data collection, resource management, process management, training infrastructure, serving infrastructure, and monitoring—I excel at implementing seamless MLOps operations and ensuring scalable, reliable deployment pipelines.

Collaboration and effective communication are integral to my approach, as I believe they drive system improvements and foster personal and professional growth. Thriving in dynamic, innovation-driven environments, I am eager to contribute my expertise and enthusiasm to forward-thinking teams, delivering impactful machine learning solutions.

Reaserch Assistance

Tampere University

March 2024 - Present

- Conducting research under the supervision of **David Hästbacka** and **Sergio Moreschini**, focusing on language models and agnetic operations especially with small language models on edge devices. The work integrates MLOps best practices, incorporating advanced features such as synthetic data management, prompt and embedding management, agent and chain management, pipeline orchestration, automated testing, autonomous retrieval-augmented generation (RAG), and transfer learning techniques, including distillation, pruning, and quantization. Additionally, it involves developing specialized model evaluations, tailored A/B testing for different agent chains and prompt versions, governance and safety checks, and creating a custom playground for experimentation and development.
- Built and managed full-stack DataOps pipelines across ingestion, transformation, storage, and analytics—integrating OLTP, APIs, logs, and event streams into data lakes and warehouses using batch and stream processing frameworks. Ensured quality, observability, and governance through metadata management, data contracts, and automated validation.
- Developed ML systems with end-to-end lifecycle support—including feature platforms, experiment tracking, orchestration, model registry, CI/CD pipelines, deployment, and vector DB integration—ensuring scalable serving, automated validation, versioned model handovers, and robust monitoring across training and inference stages.

ML Engineer Trainee in Innovation Team

Elisa Polystar

May 2023 - November 2024, Finland

- Designed and implemented a full-stack solution integrating Large Language Models (LLMs) to enhance enterprise infrastructure, including a copilot feature seamlessly embedded into existing telecom software systems. co-Led the Generative AI engineering team in scaling proof-of-concept (PoC) projects into reliable, production-grade solutions.
- Engineered full-stack architectures to support Automated Machine Learning (AutoML) pipelines tailored for ad hoc data scenarios, ensuring scalability, maintainability, and adaptability to telecom-specific workflows.
- Worked extensively with big data in a telecom environment, building systems capable of handling large-scale, high-throughput data efficiently. Designed data processing pipelines and services to ensure performance and reliability at scale.
- Reported directly to the Chief Technology Officer (CTO), providing technical insights and advancing strategic initiatives. Delivered several presentations on the solution and proof of concept to effectively convince non-technical stakeholders.

AI Scientist & Python Developer

Magister Solution Ltd

June 2022 - March 2023, Finland

- Designed and implemented sophisticated attacks on Machine Learning-based models / library, focusing on identifying and exploiting vulnerabilities in advanced AI systems.
- Co-Authored and published a comprehensive paper at the 2022 IEEE Conference on Standards for Communications and Networking (CSCN), [On Assessing Vulnerabilities of the 5G Networks to Adversarial Examples](#).
- Conducted research [On Attacking Future 5G Networks with Adversarial Examples](#), paving the way for enhanced security protocols in next-generation telecommunication networks.

Research Assistant

Shahid Rajaei Teacher Training University

September 2020 -December 2021

- Conducted time series forecasting research, focusing on a Forex trading data case study. Utilized advanced methods including MLP, CNN, LSTM, RNN, GRU, ResNet, signal processing, fuzzy systems, hybrid models (e.g., CNN+LSTM, associated LSTM), RBF Neural Networks, Random Forest, Variational Autoencoders (VAE), ensemble learning (boosting), and multimodal machine learning techniques.
- Developed expertise in Generative Adversarial Networks (GANs) for generating synthetic data to enhance cognitive modeling research under Prof. Reza Ebrahimpour. Created GAN-based datasets that improved model training efficiency and robustness by 25%.

R&D Engineer

Mehregan Pahnesh

June 2017 - December 2018

- Assisted in designing and conducting PCB layouts and feasibility studies for industrial projects.
- Gained hands-on experience in Embedded Systems development and implementation for various industrial applications.

Educations

Bachelor of Science (B.Sc) in Electrical engineering

Minor in Control Systems • K.N. Toosi University of Technology

Master of Science (M.Sc.) in Signal Processing and Machine Learning

Tampere University • Finland • 2025

Skills

Front-End: HTML, CSS, JavaScript (JS), TypeScript (TS), React, Next.js; experienced with design tools like Figma.

Back-End:

- Python (Flask, FastAPI), TypeScript (Hono, Express)
- Protocols: OAuth, GraphQL, SSE, WebSocket, REST, gRPC, tRPC, WebRTC, Pub/Sub, AMQP, MQTT.

Database :

- Clickhouse, PostgreSQL, MySQL, DuckDB, MongoDB, BigQuery, Redis (caching), Neo4j.
- ORMs : SQLAlchemy (Python) , Prisma (TypeScript)

DevOps :

- Linux/Unix, Git, Docker, Azure DevOps (pipelines), Github actions, Terraform, Grafana, Prometheus, Sentry
- Kubernetes : Core Kubernetes Concepts, Helm, Kustomize, Access Control , OpenTelemetry, Cert-Manager, Argo CD, Istio

Cloud :

- Azure: Virtual Machines, Blob Storage, App Service, Load Balancer, Machine Learning, Autoscale Functions
- AWS: EC2, S3, RDS, DynamoDB, Lambda, Step function, VPC, Route 53, IAM, Amazon Sage Maker
- Google Cloud: Compute Engine, Google Cloud Storage (GCS), App Engine, Load Balancing, IAM, IAP, Cloud DNS, GKE, BigQuery, Vertex AI, Cloud Run, Artifact Registry

MLOPS:

- Orchestration : KubeFlow, ZenML, Flyte
- ML Lifecycle : MLflow, W&B, Comet.
- Monitoring : Evidently, Arize
- Model Serve : BentoML, Seldon Core, KFServe, Torchserve, NVIDIA Triton,
- Feature Engineering : Feast
- AutoML : FLAML, evalml, autogluon, TPOT
- End-to-End : Amazon Sage Maker, Vertex AI, databricks, familiar with Azure Machine learning, ClearML, Lightning AI , Qwak
- Machine learning: TensorFlow, Keras, PyTorch, PyTorch Lightning, Scikit-learn, statsmodels, Pandas, Polars, pyspark, XGBoost, LightGBM, Plotly, Matplotlib, Statsmodels, Transformers (Hugging face), nltk, Spacy, Librosa , prophet, PyCaret, familiar with JAX

DataOps:

- Data Versioning & Lineage: DVC, LakeFS
- Metadata & Cataloging: DataHub, Unity Catalog, Nessie, Hive Metastore
- Data Modeling & Transformation: dbt
- Search Engines: ElasticSearch
- ELT : Airbyte, dlt
- Batch Processing: Apache Spark, Dask, Ray
- Stream: Apache Kafka, RabbitMQ
- Workflow Orchestration: Flyte, MageAI, Kestra, Airflow
- Data Storage (Lake & Warehouse): Delta Lake, Apache Iceberg
- File Formats: Parquet, ORC, Avro
- Object Storage: S3, GCS, HDFS
- Ad Hoc Query & Analytics Engines: Trino, Dremio
- Observability & Quality: Great Expectations, Evidently
- Application Output & Dashboards: Looker, Superset, Metabase, Kibana, Plotly, Streamlit

LLM:

- Chains: LangChain(LangGraph, LangMem), LlamaIndex, Haystack, adk-python (Google), DSPY, Promptflow, Mem0, E2B (code sandbox)
- Agents : openai-agents-python, agents (Cloudflare), Autogen(Microsoft), beeai_framework (IBM), Agent Development Kit (google), Crewai, Agno
- Middleware : LiteLLM
- Fine-tuning: Unsloth, Axolotl, PEFT, MergeKit, Distillable
- Embeddings : Sentence Transformers, FastEmbed
- Observability: Phoenix, Opik, Helicone, OpenLLMerty, Mlflow, Langfuse
- Data labeling : Argilla
- Evaluation : Ragas, DeepEval
- VectorDB : ChromaDB, Qdrant , Clickhouse, Weaviate
- Inference : vLLM, llama.cpp, Accelerate, CTranslate2, deepspeed, TensorRT-LLM, LitServe
- Protocols : MCP , A2A

Languages: Proficient in Python and TypeScript, with experience in C++ (for robotics), C# (Unity projects), Go (backend services), C, and Java.