Mohamed Abdelhamid

Munich, Germany | +49 152 3169 6069 | M.Abdelhamid@campus.lmu.de | GitHub | LinkedIn

Summary

M.Sc. Data Science student at LMU Munich and Machine Learning Engineer focused on reliable AI, model evaluation, multimodal/multilingual RAG, time-series forecasting, and computer vision. Experienced in building end-to-end ML pipelines, evaluating model failure modes, training deep learning models on GPU infrastructure, and productionizing systems with PyTorch, Docker, Airflow, MLflow, ONNX, TensorRT, and Triton.

Experience

Linde plc Apr 2025 – Present
Machine Learning Engineer, Working Student *Munich, Germany*

- Implemented an end-to-end multimodal document-intelligence pipeline that converts PDFs, images, PPTX, DOCX, XLSX, HTML, and other formats into structured Markdown using GLM-OCR for text/layout extraction and Gemma 4 for explaining and transcribing embedded images, figures, and visual content.
- Evaluating an enterprise RAG system across retrieval, reranking, OCR/document understanding, multimodal embeddings, and multilingual embeddings; designing experiments to compare quality, robustness, and failure modes.
- Improved satellite object detection accuracy from 84.5% to 92.0% mAP@50 by replacing a YOLOv8 baseline with RF-DETR Large and training on aerial imagery using 4 A100 GPUs.
- Automated training, evaluation, and deployment with Airflow and MLflow; productionized inference with ONNX, TensorRT, and Triton while optimizing latency/throughput trade-offs.

Upwork 2023 – 2025
Machine Learning Engineer *Remote*

- Served as the sole ML engineer for MknoonAI, a healthcare AI startup, owning the full ML lifecycle from dataset curation and PyTorch modeling to deployment and web integration.
- Developed medical-imaging models including a chest X-ray classifier with 95.5% F1 and a dual-view mammography model with 85% F1 and tumor-localization support; achieved Top Rated status with 100% job success rate.

Education

Ludwig Maximilian University of Munich Expected Apr 2027
M.Sc. Data Science, Elite Graduate Program *Munich, Germany*

- Recipient of the **Konrad Zuse School of Excellence in Reliable AI Scholarship**.

Innopolis University Jul 2024
B.Sc. Computer Science, GPA: 4.8 / 5.00 *Innopolis, Russia*

Selected Projects

Foundation Models for Industrial Time-Series Forecasting *LMU / **BASF SE***

- Evaluated six time-series foundation models – TimesFM, Chronos-2, MOMENT, Lag-Llama, Toto, and Moirai – against AutoARIMA, AutoSARIMA, and AutoSARIMAX baselines on 1,118 monthly industrial time series from BASF SE.
- Compared univariate and multivariate forecasting under zero-shot and supervised fine-tuning settings using rolling-origin evaluation and RMSSE at a three-month forecast horizon.
- Analyzed when foundation models outperform classical statistical baselines by relating forecasting performance to trend strength, seasonality strength, spectral entropy, sparsity, and coefficient of variation.

Membership Inference Attacks on Medical Language Models *AI Security / Privacy*

- Evaluated privacy risks in synthetic clinical language models using black-box membership-inference attacks based on likelihood signals and model-output behavior.
- Analyzed attack performance across patient subgroups and data conditions to understand when medical language models are more likely to leak training-data membership information.
- Connected empirical privacy-risk evaluation with broader questions in trustworthy AI, especially the deployment of language models in sensitive healthcare settings.

Research, Publications, and Open Source

Publication: B. Rasheed, M. Abdelhamid, A. Khan, I. Menezes, and A. M. Khatak, "Exploring the Impact of Conceptual Bottlenecks on Adversarial Robustness of Deep Neural Networks," *IEEE Access*, 2024.

Open Source: Contributed to the Ivy open-source library by adding `log1p` support to the PyTorch frontend in a merged pull request.

Skills

Programming: Python, C++
ML / AI: PyTorch, Transformers, OpenCV, RF-DETR, YOLO, Foundation models, LangChain, model evaluation, multimodal AI, document understanding
Infrastructure: Docker, Airflow, MLflow, FastAPI, Linux, Git, GCP, AWS, ONNX, TensorRT, Triton, Multi-GPU Training