

Dennis Netzer

Senior MLOps & GenAI Engineer · Co-Founder

Frankfurt a. M., Germany | Remote / Frankfurt area on-site

[linkedin.com/in/dennis-netzer-4a782417b](https://www.linkedin.com/in/dennis-netzer-4a782417b) | [\[Website URL\]](#) | [\[Email\]](#) | [\[Phone\]](#)



Availability: from July 2026 for freelance engagements (up to 3 days/week) | **Day rate:** on request

Profile

Senior engineer with 6+ years of experience at the intersection of Machine Learning, MLOps and Generative AI. Hands-on across the full ML production lifecycle — from classical anomaly detection and time-series modelling to LLM-based multi-agent architectures with RAG, fine-tuning and MCP. Currently co-founding an EXIST-funded startup building intelligent anomaly detection and root-cause analysis for Oracle databases. Previously 2.5 years at Capgemini as Lead Data Scientist with focus on GenAI solutions for industrial clients, preceded by Python development at ORDIX.

Additional industrial background through a dual Bachelor in Industrial Engineering with hands-on production planning experience — a useful asset on industrial and manufacturing engagements.

Focus Areas

- **GenAI & LLM Engineering** — RAG architectures, fine-tuning (PEFT/LoRA), multi-provider gateways, MCP tool integration, LLM observability & guardrails
- **MLOps & ML Lifecycle** — experiment tracking, model versioning, workflow orchestration, reproducibility
- **Anomaly Detection & Predictive Analytics** — classical ML, deep-learning-based time-series models, root-cause analysis pipelines
- **Database-adjacent ML/AI** — Oracle performance data, SQL governance

Skills

GenAI & LLM Engineering

LLM APIs (various providers) · LangChain · Hugging Face Transformers · RAG architectures · Pinecone · Qdrant · Fine-tuning (PEFT, LoRA, QLoRA) · Model Context Protocol (MCP) · Function calling · LiteLLM (multi-provider gateway) · Langfuse · LangSmith · Prompt engineering · Guardrails · LLM evaluation · n8n

Machine Learning & Deep Learning

TensorFlow · Keras · PyTorch · scikit-learn · Hugging Face · classical ML (regression, classification, clustering, tree-based, SVM) · Deep Learning (CNNs, RNNs, LSTMs, autoencoders, GANs, transformers) · Anomaly detection (Isolation Forest, One-Class SVM, ARIMA, Prophet, LSTM-autoencoder, transformer-based) · Reinforcement Learning (research background) · Computer Vision (YOLO)

MLOps & ML Lifecycle

MLflow · Kubeflow · Airflow · Flyte (evaluated) · Model versioning · Experiment tracking · Model Registry · Reproducibility · Evidently (basics)

Backend & APIs

Python (senior level) · FastAPI · REST · GraphQL · pytest · async / await

Databases

Oracle · PostgreSQL · MySQL · Neo4j (Cypher, graph modelling) · Vector stores (Pinecone, Qdrant)

Cloud & Infrastructure

Microsoft Azure (Azure OpenAI, App Services) · Google Cloud Platform (Compute, AI Training) · Linux

Frontend (supporting)

Vue.js · React (touchpoints)

Methods & Domain

Agile / Scrum · Solution engineering · Vendor evaluation & make-or-buy · Stakeholder communication · Workshops & training · Conference speaking · Production planning (SAP) · Automotive procurement · Bill-of-materials & variant management

Languages

German (native) · English (fluent, C1) · Spanish (basics, B1)

Experience

Co-Founder & Head of AI/ML

aprevis, Darmstadt | February 2025 – present

Intelligent anomaly detection and root-cause analysis platform for Oracle databases. EXIST-funded (Sept 2025 – Jun 2026), preceded by HessenIdeen scholarship.

- Designed and built the ML/DL pipeline for anomaly detection on Oracle system metrics (CPU, memory, IO, wait events)
- Developed, trained and validated anomaly detection models — transformer-based architectures, LSTM autoencoders, Isolation Forest, One-Class SVM
- Architected and implemented the root-cause analysis pipeline with LLM-driven hypothesis generation and automated action-plan creation
- Fine-tuned European open-source models (Mistral 7B) and reasoning models (DeepSeek) using PEFT/LoRA for domain-specific database diagnostics
- Built a multi-provider LLM gateway (LiteLLM) giving customers free choice between OpenAI, Mistral and locally hosted models — on-premise deployment for enterprise clients
- RAG integration with Pinecone and Qdrant for domain-specific knowledge
- MCP tool integration: YouTrack (issue tracking), Slack (notifications), vector-DB queries
- LLM observability via Langfuse, plus prompt management and guardrail strategies
- Backend architecture in Python/FastAPI, frontend contributions in Vue.js
- Communication with pilot customers (CLAAS, Team PB, SE Padersoft), conference talks, strategic product development
- Co-responsibility for sales, marketing and finance topics within the founding team

Stack: Python · FastAPI · TensorFlow · PyTorch · Hugging Face · MLflow · Kubeflow · Airflow · LangChain · LiteLLM · Langfuse · LangSmith · Pinecone · Qdrant · n8n · Mistral · DeepSeek · OpenAI · Vue.js · Oracle

Lead Data Scientist (from April 2024) / Consultant Data Science

Capgemini, Frankfurt | October 2022 – February 2025

Lead Python Developer & Data Scientist | Premium Automotive Manufacturer

March 2024 – February 2025 (~11 months)

Modernization and digitalization of procurement at a leading German premium automotive manufacturer. Embedded in a large programme of 50–70 people spanning multiple workstreams.

- Technical lead of a 4-person development team (1 lead, 2 junior developers, business steering by client lead) building a platform for automated component evaluation and quotation creation
- Backend architecture in Python/FastAPI with PostgreSQL as application database
- Integration of internal component data via GraphQL, automated component classification using Azure OpenAI (GPT-4) with embedding-based matching
- Modelled the bill of materials in a Neo4j graph database (>1,000 nodes) for variant management, supply-chain transparency and impact analysis — Cypher queries for multi-level traversal
- Advised the procurement department on data science topics: price prediction models, exploratory data analysis with Jupyter notebooks
- Delivered internal Python training sessions
- Project mandate extended at the end of 2024; consistently positive feedback from business users
- Promoted to Lead Data Scientist in April 2024 during the engagement

Stack: Python · FastAPI · PostgreSQL · Azure OpenAI (GPT-4) · GraphQL · Neo4j · Cypher · React (touchpoints) · Jupyter

MLOps Framework Evaluation (internal)

January 2024 – February 2024

- Strategic evaluation and comparison of MLOps frameworks (MLflow, Kubeflow, Flyte) for internal Capgemini standardization
- Built proof-of-concepts and an evaluation matrix to derive use-case-specific recommendations

Solution Engineer | Insurance Group

July 2023 – December 2023 (6 months)

Conception of a new visitor platform for the headquarters of a German insurance group (registration, hotel booking, meeting management, leisure activities).

- Worked alongside a senior solution engineer on feasibility and tech-stack options
- Translated user journeys defined by the business into technical architecture options
- Conducted vendor scouting and evaluation conversations with platform providers
- Provided make-or-buy advisory input for the client's decision board

Generative AI R&D + Hannover Messe 2023

October 2022 – June 2023

- Early exploration of GPT-3 and GPT-4 with focus on industrial use cases
- Co-developed the “Trustworthy GPT-4 Advisor” — a multimodal agent (GPT-4 + computer vision + ElevenLabs speech) that evaluated physically built Lego structures against rule sets
- Vision pipeline using YOLO-based colour and structural recognition through a stereo camera setup
- Backend in Python/FastAPI, single-turn interaction model with configurable personality modes (kind / harsh)

- Co-speaker at Hannover Messe (April 2023): *“Generative AI in Intelligent Industry — Capgemini Trustworthy GPT-4 Advisor”*
- Contributed to internal bid responses, including a public-sector engagement with elevated security clearance requirements

Stack: Python · FastAPI · OpenAI GPT-4 · YOLO · Computer Vision · ElevenLabs · multi-sensor setup

Junior Consultant

ORDIX AG | April 2022 – September 2022

Two parallel engagements covering AIOps and database governance.

Anomaly Detection for IT Infrastructure (internal AIOps PoC)

- Evaluated and implemented anomaly detection algorithms on system metrics (CPU, memory, IO, wait events) for database infrastructures
- Comparative study of Isolation Forest, One-Class SVM, ARIMA, Prophet and LSTM-based architectures
- Co-speaker at DOAG 2022 (September 2022): *“From Performance Signature to AIOps”*

SQL Governance Platform | Banking

- Junior Python developer in a team building a SQL governance application for a bank
- Tooling for management, versioning and approval of SQL statements used in production
- Compliance- and audit-driven use case in a regulated environment

Stack: Python · Oracle · scikit-learn · TensorFlow · classical time-series methods

Working Student / Master Intern — Data Science

ORDIX AG (as part of the dual master programme) | October 2019 – March 2022

- Practical phases as part of the dual master programme in Data Science
- In-depth Linux and SQL training, followed by collaboration on database-adjacent analytics topics
- Designed and delivered the workshop *“Inside the Black Box of Machine Learning”* at IT-Tage 2020

Production Planner

BOS GmbH & Co. KG (automotive supplier) | October 2018 – June 2019

- Production and logistics planning for an automotive supplier specialized in customized equipment
- Make-or-buy analyses, planning of production tooling and resources
- Contribution to a new factory ramp-up in China in close coordination with sales and engineering
- Self-directed upskilling in Python, SQL and Data Science (Udemy, Coursera) in preparation for the master programme

Dual Bachelor Student / Working Student

BOS GmbH & Co. KG | October 2015 – September 2018

- Practical phases during the dual bachelor programme; focus on production planning and logistics
- International assignments at plants in Hungary and Poland

Education

Master of Science — Data Science

Darmstadt University of Applied Sciences (h_da), dual programme with ORDIX AG | October 2019 – December 2021 | **Final grade: 1.2 (German scale, equivalent to high distinction)**

- Research paper (grade 1.3): *“Faulty reward functions — How to ensure that reward maximization aligns with developer intent”* (Reinforcement Learning, applied to the h_da foosball project)
- Master thesis (grade 1.3): *“How can state-of-the-art generating algorithms be used to create interactive Artificial Intelligence that generates electronic music”* — own neural network architectures (autoencoders, GANs) in TensorFlow, training on Google Cloud

Bachelor of Engineering — Industrial Engineering (focus: production & logistics)

Baden-Württemberg Cooperative State University (DHBW) Stuttgart, dual programme with BOS GmbH & Co. KG
| October 2015 – September 2018 | **Final grade: 1.9 (German scale, equivalent to upper second class)**

- Exchange semester at **Edinburgh Napier University**, Scotland (5th semester)
- Bachelor thesis (grade 1.6): optimization of production planning using SAP Long-Term Planning

Talks & Publications

- **DOAG 2022** (Sep 2022) — *“From Performance Signature to AIOps”* (co-speaker)
- **Hannover Messe 2023** (Apr 2023) — *“Generative AI in Intelligent Industry — Capgemini Trustworthy GPT-4 Advisor”* (co-speaker)
- **IT-Tage 2021** (Dec 2021) — *“Artificial Intelligence as Artist — How modern algorithms generate music”* (solo speaker)
- **IT-Tage 2020** (Dec 2020) — Workshop *“Inside the Black Box of Machine Learning”* (co-speaker, team of 3)

Certifications & Continued Education

- **TensorFlow Developer Certificate** — Google (valid until March 2027)
- **IBM Data Science Professional Certificate** — Coursera Specialization
- **Complete Python Bootcamp: Go from zero to hero in Python 3** — Udemy
- **The Ultimate MySQL Bootcamp: Go from SQL Beginner to Expert** — Udemy

Grants & Recognition

- **EXIST Founder Scholarship** (German Federal Ministry / ESF) — September 2025 – June 2026
- **HessenIdeen Scholarship** — March 2025 – August 2025