

Dr. Oliver Köhn

Senior Data & AI Engineer

Consultant for Data, Automation and AI

kontakt@oliverkoehn.info

+49 152 06348816

oliverkoehn.com

linkedin.com/in/droliverkoehn

Saarbrücken (Germany) · nationwide & remote

PROFILE

Senior Data & AI Engineer and PhD physicist with more than ten years of experience in data engineering, data science, cloud platforms, process automation, and production AI systems.

I develop complete solutions from analysis and architecture through data pipelines, backend and AI components to user interface, deployment, and production operation. My work combines technical depth with a clear focus on business processes, user needs, and economic value.

My focus areas include modern data platforms, automated document and communication workflows, knowledge assistants, predictive analytics, as well as local and hybrid AI architectures. Depending on requirements I integrate existing commercial systems or develop sovereign, self-hostable solutions with open and replaceable components.

Positioning: Data Engineering · Automation · Applied AI · Solution Architecture

AREAS OF EXPERTISE

Data Engineering & Platforms	Design and delivery of data pipelines, data models, data warehouses, lakehouses, and cloud data platforms. Experience with data migration, data quality, orchestration, monitoring, and both batch and event-based processing.
Analytics & Data Science	Development of statistical models, predictive-analytics solutions, management dashboards, and explanatory analyses. Translating domain questions into measurable metrics, data models, and decision support.
Automation & Integration	Automation of recurring business processes via APIs, workflows, events, and system integrations. Connecting e-mail, file storage, task management, knowledge systems, line-of-business applications, and communication.
Applied AI	Development of production AI applications using language models, retrieval-augmented generation, agent systems, document analysis, semantic search, source display, and structured model outputs.
Architecture & Operations	Building operable solutions for cloud, European infrastructure, and local environments. Containerisation, CI/CD, monitoring, model gateways, and integration of commercial and self-hosted components.

PROFESSIONAL EXPERIENCE

Consultant for Data, Automation and AI

Self-employed · 12/2024–present

Development of individual data and AI applications, automated business processes, and sovereign, self-hostable system architectures. Consulting and technical implementation from a single source — from process analysis and pilot to production integration.

Project: Automated processing of incoming e-mails

CONTEXT: For a recurring e-mail workflow, an application was built that analyses free-form messages and turns them into structured cases. The solution detects, among other things, requests, involved parties, references, deadlines, priorities, and possible responsibilities. From the results, tasks, follow-up processes, and notifications can be created automatically.

IMPLEMENTATION: The application includes a user interface for the traceable review of extracted information, a FastAPI backend, structured model outputs, and workflow integration. Depending on operational requirements, commercial model APIs, European services, or locally deployed models can be used.

TECHNOLOGIES: *Python · FastAPI · PostgreSQL · LLMs · structured outputs · n8n · REST APIs · Docker · Ollama · LiteLLM*

Project: Automated document processing across multiple systems

CONTEXT: Development of an end-to-end workflow in which incoming PDFs and other documents are automatically detected, analysed, and further processed. Extracted information is stored in a structured form and then handed over to task management, documentation, and team communication.

IMPLEMENTATION: The solution connects file storage, AI-driven content analysis, workflow orchestration, kanban tasks, and knowledge documentation. This reduces media breaks and provides traceable process documentation.

TECHNOLOGIES: *Nextcloud Files · Nextcloud Deck · Nextcloud Talk · Docmost · n8n · Python · LLMs · Docker · REST APIs*

Project: Internal knowledge assistant with verifiable sources

CONTEXT: Development of a RAG application that lets employees ask questions about internal contracts, policies, and working documents. The system searches relevant passages, generates a concise answer, and shows the underlying documents and citations directly.

IMPLEMENTATION: In addition to document ingestion and semantic search, permissions, source display, model routing, and local operation options were implemented. Models can be accessed locally via Ollama or flexibly through LiteLLM across different internal and external providers.

TECHNOLOGIES: *RAG · LangChain · LangGraph · Hugging Face · vLLM · Ollama · LiteLLM · PostgreSQL · pgvector · FastAPI*

Project: AI-supported management and reporting application

CONTEXT: Implementation of an agent system that aggregates data from multiple sources, analyses developments, and prepares structured management reports. The system helps contextualize notable metrics and suggests possible next analytical steps.

IMPLEMENTATION: The architecture combines data pipelines, analytical models, dashboards, and generative AI. Particular attention was paid to traceable intermediate results, controlled tool calls, and a clear separation between computed metrics and generated explanations.

TECHNOLOGIES: *Python · SQL · LangGraph · Agentic AI · Power BI · Apache Superset · Grafana · REST APIs · Docker*

Project: Voice assistant with corporate knowledge

CONTEXT: Development of a voice-RAG prototype for recurring internal or external requests. Spoken questions are transcribed, checked against approved corporate documents, and returned as spoken answers.

IMPLEMENTATION: In addition to speech recognition and speech synthesis, source display, conversation logging, and local model options were integrated. Sensitive documents and model components can be operated fully on controlled infrastructure.

TECHNOLOGIES: *Speech-to-Text · Text-to-Speech · RAG · Ollama · LiteLLM · FastAPI · WebSockets · Docker*

Senior Data & AI Engineer

Public Cloud Group · 08/2024–11/2024

Development of production AI solutions and data-driven processes in a cloud and enterprise context. Focus on scalable platform components, document analysis, and moving machine learning and generative-AI use cases into robust architectures.

Project: Enterprise-wide Generative-AI platform for content workflows

CONTEXT: Contribution to the design and development of a central platform for AI-supported content processes. The architecture provided reusable capabilities for model access, prompt management, structured outputs, and integration of different data sources.

IMPLEMENTATION: This enabled business units to embed generative-AI functionality into existing content and approval workflows without building a completely new infrastructure for every use case. In addition to the application layer, scaling, access control, monitoring, and cloud deployment were also addressed.

TECHNOLOGIES: *AWS · Python · FastAPI · Docker · LLM APIs · RAG · REST APIs · CI/CD · Monitoring*

Project: Automated document analysis for reporting workflows

CONTEXT: Development of a solution for analysing extensive and partly unstructured documents. Relevant content was extracted, classified, and transformed into a structured form suitable for downstream reporting and review processes.

IMPLEMENTATION: A key focus was on reliable structured outputs, validation rules, and traceable mapping of results back to the original document locations. The solution was integrated into an existing Azure environment and into pre-existing data workflows.

TECHNOLOGIES: *Azure · Python · LLMs · document processing · FastAPI · structured outputs · data validation*

Project: Predictive analytics for industrial machine data

CONTEXT: Development respectively enhancement of a forecasting system for the early detection of possible machine failures. Operating, condition, and sensor data were aggregated from SAP-adjacent and further data sources, prepared, and made usable for forecasting models.

IMPLEMENTATION: Results were delivered in a way that allowed subject-matter experts to prioritise risks and plan maintenance activities more precisely. Beyond the model itself, the data pipeline, result interpretation, and integration into operational workflows were central to the project.

TECHNOLOGIES: *Azure · Python · SAP · Predictive Analytics · Machine Learning · ETL · Monitoring*

Design and implementation of cloud data platforms, analytics solutions, and machine-learning applications in the Microsoft Azure environment. Advising business units and technical teams and delivering production-grade data and analytics processes.

Project: Predictive maintenance system for industrial plants

CONTEXT: Development of a forecasting system for detecting imminent machine failures. Historical operating data, sensor data, and maintenance records were aggregated, cleaned, and prepared for machine-learning models.

IMPLEMENTATION: The solution covered not only the prediction model but the full processing chain from the data source through feature engineering and model evaluation to result delivery. Predictions were integrated into analysis and reporting interfaces for subject-matter experts.

TECHNOLOGIES: *Python · Azure Synapse Analytics · Azure Data Factory · Azure Databricks · Spark · Machine Learning · Power BI*

Project: Azure analytics platform with automated data pipelines

CONTEXT: Design and build-out of a cloud-based analytics platform for centrally processing and providing business data. Data from multiple operational systems was harmonised through automated ETL and ELT processes and unified in a consistent data model.

IMPLEMENTATION: On the platform, dashboards and self-service analytics solutions were created for various business units. Data quality, pipeline reusability, and traceable operation were central pillars of the architecture.

TECHNOLOGIES: *Azure Data Factory · Azure Synapse · Databricks · SQL · Spark · Power BI · Tableau · CI/CD*

Project: Migration of legacy data and modernisation of reporting

CONTEXT: Replacement respectively modernisation of historically grown data and reporting structures. Existing data sources, transformation logic, and reports were analysed, cleaned, and gradually migrated to an Azure-based platform.

IMPLEMENTATION: Beyond the technical migration, data models were harmonised, dependencies documented, and queries optimised. This produced a more resilient basis for reporting, analytics, and later AI use cases.

TECHNOLOGIES: *Azure · SQL · Python · ETL · data modelling · cloud migration · Power BI*

Project: NLP-based analysis of customer feedback

CONTEXT: Development of an application for the automated analysis of customer comments and free-text feedback. Texts were prepared, classified thematically, and examined for sentiment and recurring topics.

IMPLEMENTATION: Results were visualised in a dashboard so that business units could identify trends and relevant problem areas more quickly. The solution combined NLP processing, a data pipeline, and business-intelligence components.

TECHNOLOGIES: *Azure Databricks · Python · NLP · Sentiment Analysis · Power BI · Spark*

Development of data-driven analytical, risk, and simulation solutions in the financial services environment. Combining quantitative models with stable data and application architectures and evaluating novel computational methods.

Project: Platform for risk and scenario analyses

CONTEXT: Design of a technical architecture for running data-intensive risk and scenario computations. Different market, position, and reference data were prepared and made available for reproducible simulation runs.

IMPLEMENTATION: The platform supported the execution of various scenarios, structured storage of results, and their subsequent analysis by subject-matter experts. Special emphasis was placed on traceability, data quality, and reproducible calculation states.

TECHNOLOGIES: *Python · SQL · quantitative models · simulation · data pipelines · REST APIs · containerisation*

Project: Data-driven decision support in financial services

CONTEXT: Development of an analytical and simulation solution for evaluating different business options. Complex computational results were translated into understandable metrics and visualisations.

IMPLEMENTATION: The work included structuring the underlying question, technical modelling, processing large amounts of data, and delivering results to decision makers.

TECHNOLOGIES: *Python · SQL · statistical modelling · scenario analysis · reporting · data visualisation*

Project: Proof of concept for Quantum Computing

CONTEXT: Investigation of whether selected optimisation or simulation problems can be tackled with quantum-based respectively quantum-inspired methods. Classical approaches were compared with alternative computational paradigms and evaluated for technical feasibility and potential value.

IMPLEMENTATION: The focus was not only on the algorithm but also on data preparation, interfaces, reproducible testing, and a realistic assessment of the technology for enterprise use.

TECHNOLOGIES: *Python · Quantum Computing · optimisation · simulation · classical baseline methods*

Development of analytical solutions for compliance, stress tests, data quality, and decision support. Automation of recurring analytical and reporting processes in financial services.

Project: Automated compliance and control evaluations

CONTEXT: Development of a data-based solution for auditing large data sets against defined compliance and control criteria. Data from different sources was harmonised, validated, and evaluated against domain-specific rules.

IMPLEMENTATION: Results were prepared for subject-matter experts in a traceable form and integrated into recurring reporting and audit processes. Manual Excel and database work could thereby be transformed into reproducible Python and SQL workflows.

TECHNOLOGIES: *Python · SQL · data quality · rule engines · reporting · automation*

Project: Stress-test and scenario analyses

CONTEXT: Implementation of quantitative calculations for evaluating the impact of different market and risk scenarios. Input data was prepared, models were implemented, and results were validated.

IMPLEMENTATION: Results were made available for further analyses and management reports. In addition to modelling, data reconciliation, reproducibility, and domain documentation were essential parts of the work.

TECHNOLOGIES: *Python · SQL · quantitative models · simulation · stress tests · data visualisation*

Project: Sentiment and free-text analysis

CONTEXT: Development of an analytical prototype for evaluating unstructured texts. Relevant content was prepared, categorised, and examined for sentiment respectively topics.

IMPLEMENTATION: The work combined statistical analysis, NLP techniques, and a clear presentation of results for business users.

TECHNOLOGIES: *Python · NLP · Sentiment Analysis · statistical modelling · reporting*

DOCTORAL RESEARCH

2017 – 2020

Ph.D. in Physics (Dr. rer. nat.)

Saarland University and University of Wisconsin–Madison

Grade: very good

Thesis: *Statistical analysis of bacterial locomotion*

During my doctorate I developed statistical methods and reproducible analysis pipelines for complex motion and trajectory data. Large experimental datasets were processed, quality-checked, and compared against mathematical and physical models. The work combined scientific software development, time-series analysis, probabilistic modelling, and the visualisation of complex results. It also involved interdisciplinary collaboration with researchers from physics and biology as well as the presentation and documentation of scientific results.

EDUCATION

2023 – 2024

Master of Science in Data Science

IU International University of Applied Sciences

Thesis: *Data-Science solution with a business approach*

Combining technical data-science methods with economic evaluation, feasibility, and concrete business value.

2014 – 2015

Master of Science in Physics

Saarland University

Thesis: *Application of Machine Learning on the characterization of two-level systems*

Application of machine-learning methods for the analysis and characterisation of physical systems.

2010 – 2014

Bachelor of Science in Physics

Saarland University

Foundations in theoretical and experimental physics, mathematics, statistics, modelling, and scientific data analysis.

CERTIFICATIONS

- Microsoft Certified: Fabric Analytics Engineer Associate

- Microsoft Certified: Azure Data Engineer Associate
- Microsoft Certified: Azure Data Scientist Associate

TECHNOLOGIES AND TOOLS

DATA ENGINEERING & PLATFORMS

Python · SQL · PostgreSQL · Spark · PySpark · Databricks · Azure Data Factory · Azure Synapse Analytics · Microsoft Fabric · AWS · GCP · ETL/ELT · Data warehouses · Lakehouses · Data modelling · Data quality

APPLICATIONS & INTEGRATION

FastAPI · REST APIs · Docker · n8n · Nextcloud · Docmost · Webhooks · Event workflows · Microservices · OAuth/OIDC · PostgreSQL · Redis

AI & MACHINE LEARNING

LLMs · RAG · LangChain · LangGraph · Hugging Face · vLLM · Ollama · LiteLLM · scikit-learn · PyTorch · TensorFlow · XGBoost · NLP · Predictive Analytics · Agentic AI · Semantic search

ANALYTICS & OPERATIONS

Power BI · Tableau · Apache Superset · Grafana · Monitoring · CI/CD · Cloud · EU infrastructure · On-premises · Containerisation

WORKING STYLE

- ✓ Direct collaboration instead of long handover chains
- ✓ Technology selection appropriate to the problem
- ✓ Understandable communication with both domain and engineering teams
- ✓ Working results instead of pure concept papers
- ✓ Traceable and operable solutions
- ✓ Deliberate balance between commercial, open, and self-hosted components