

Vikas Ramaswamy

Data Engineer · Analytics Engineer · AI Engineer

Emden, Germany · +49 1635483706 · vikasramaswamyj@gmail.com · [LinkedIn](#) · [github.com/vi6120](#) · Available immediately

PROFESSIONAL SUMMARY

Data and AI Engineer with experience across India and Germany, building production data pipelines at Capgemini, delivering analytical infrastructure at Mercedes-Benz and Mutares, and shipping two live AI applications with real external users. Strong across the full stack: dbt, Airflow, DuckDB, SQL on the engineering side; LangGraph, RAG, and FastAPI on the AI/LLM side. M.Eng. in Business Intelligence & Data Analytics from Hochschule Emden-Leer, graduating with 1.91 overall and a thesis graded 1.30 (sehr gut). Available immediately for full-time roles in Germany.

PROFESSIONAL EXPERIENCE

Analytics Engineer (Working Student) | Mutares SE & Co. KGaA, Munich

07/2024 – 08/2025

- Built end-to-end data pipelines across 5+ post-merger portfolio companies, unifying heterogeneous source systems into dbt medallion layers, automated JSON processing cut transformation time by 50%, and reduced manual reporting effort by 40%.
- Implemented Git-versioned dbt transformations with 28 automated data quality tests, eliminating reconciliation errors across 5+ reporting streams and establishing a reliable single source of truth.
- Simulated a full SAP MM/SD/WM pipeline: 4,300+ daily stock movements, Apache Airflow orchestration, 16-model dbt medallion architecture with 28 DQ tests, DuckDB OLAP warehouse with sub-50ms queries on 11 Parquet tables.
- Deployed a FastAPI + LLaMA 3.3 70B AI advisor (live context from 6 gold-layer tables) and a React dashboard (8 pages: inventory KPIs, vendor scorecards, demand signals, supply risk).

Business Analyst (Intern) | Mercedes-Benz Group AG, Stuttgart

10/2023 – 03/2024

- Built Power BI dashboards and SQL datasets, giving management direct visibility into Azure Cloud infrastructure cost allocation, enabling a 15% reduction in identified wasted spend.
- Standardised KPI frameworks for the Platform Service Desk and product teams, accelerating decision cycles by 25% and cutting ad-hoc reporting workload across multiple stakeholder groups.

Data Engineering & Analytics | Capgemini Technology Services India, Bengaluru

09/2019 – 08/2022

- Maintained production ETL pipelines with CI/CD practices across hybrid cloud architectures, achieving 98.5% SLA adherence over a 3-year tenure across all production workloads.
- Partnered with data scientists to deliver demand forecasting models that improved forecast accuracy by 40%, and engineered executive KPI frameworks that cut manual reporting dependencies by 65%.
- Designed OLAP query optimisations and semantic layers that cut analytical query latency by 60%+ across Power BI dashboards used by portfolio company leadership for executive decisions.
- Engineered an end-to-end live data pipeline from ERP to live dashboards using SAP BW, SAP SLT, and Kafka for data streaming, which drastically reduced the Latency by more than 50%.

EDUCATION

M.Eng. Business Intelligence and Data Analytics

09/2022 – 08/2025

Hochschule Emden-Leer, Emden, Germany | Overall Grade: 1.91 | Thesis Grade: 1.30 (sehr gut)

- Thesis: "Business Management Framework for Integrating Predictive Analytics into Financial Operations Using Power BI"
- Notable grades: Digitalization & Automation 1.0, Strategic Management 1.3, Business Management 1.0, Communication & Culture 1.7
- Deployed a public GitHub App using LangGraph to orchestrate 3 parallel LLM agents (code quality, security, performance) for automated PR reviews, installed on 3 external repos with zero-configuration onboarding and zero webhook failures in production.
- Integrated Supabase pgvector + HuggingFace embeddings for persistent per-repo review memory, plus a multi-turn @mention chat interface grounded in the full diff, changed files, and comment thread.

B.Eng. Mechanical Engineering

06/2015 – 06/2019

Visvesvaraya Technological University, Belagavi, Karnataka, India

TECHNICAL SKILLS

Data Engineering: Python, SQL (Advanced), dbt, Apache Airflow, Spark, PostgreSQL, Neo4j, SAP (MM/SD/WM), CI/CD.

Cloud & Platforms: Azure Data Factory, Microsoft Fabric, Power BI, Databricks, SAP HANA, Vercel, Railway, Supabase

AI & LLM Stack: LangChain, LangGraph, RAG, Vector Databases (pgvector, FAISS), Embeddings, NLP, Groq, Azure OpenAI

Frameworks: FastAPI, React (Vite, Tailwind, Recharts), PyTorch, Hugging Face, MCP

Dev & Operations: Git, Docker, GitHub Actions, GitHub Apps (JWT, Webhooks), REST APIs, Agile / Scrum

CERTIFICATIONS & PERSONAL DETAILS

Oracle Cloud Infrastructure 2025 Certified AI Foundations Associate Valid until 09/2027

Date of Birth: 11.06.1997, KR Pet, Karnataka, India · **Nationality:** Indian **Languages:** English (C1) | German (A2)