

Farrukh Nauman

Principal Consultant | Enterprise Data Platforms · Coding Agents · Time Series ML | PhD

farrukh.nauman@inertialrange.com | (+46) 0702984959 | fnauman.com | inertialrange.com

Swedish citizen | LinkedIn: [fnauman](#) | Github: [fnauman](#)

SUMMARY

Principal consultant helping enterprise data teams de-risk platform migrations, cut warehouse cost, and automate analytics workflows with coding agents. I lead and ship in production-scale environments—owning roadmaps, managing stakeholders, and writing code, not slide decks. Selected results:

- **~1 week:** Cross-platform migration validator from zero to live dashboard (Databricks → Snowflake).
- **5–10× faster, 50%+ cheaper:** SQL pipeline optimization with row-level correctness validation.
- **10–100× faster:** Time series classification for industrial telemetry and edge deployment.
- **Team Lead:** Interim lead & project manager for 4–7 person DS and advanced analytics team.

CURRENT FOCUS

- Cross-platform migration validation and warehouse modernization (Snowflake / Databricks).
- Coding-agent workflows for migration, validation, analytics automation, and large code changes.
- Time series / telemetry ML for industrial production and edge deployment.
- Interim technical leadership: roadmap ownership, stakeholder steering, delivery discipline.

SKILLS & TECH STACK

Data Platforms	Snowflake, Azure Databricks, Spark, Snowpark, Azure Data Factory, SQL (advanced)
AI & Agents	AI Coding Agents (Cortex Code, Codex CLI, Claude Code), CLI-based Agentic Automation, LLM Agents (LangChain, OpenAI SDK), RAG, Prompt Engineering, Synthetic Data
ML & Analytics	PyTorch, Time Series, Predictive Modeling, Anomaly Detection, Computer Vision, Transformers, Weights & Biases
Engineering	Python (Expert, 8+ yrs), C/C++ (Proficient), MATLAB, Git, Docker, CI/CD, REST APIs, High Performance Computing
Business	Stakeholder Management, Project Scoping, Solution Architecture, Technical Leadership, ROI Analysis
Languages	English (Fluent), Swedish (SFI C2), Urdu (Native)

EXPERIENCE

InertialRange Labs AB - Self-employed

Principal AI Consultant

Linköping, Sweden

Aug 2025 - Present

Engagement: Interim Team Lead – Data Science & Advanced Analytics (Sep 2025 - Jun 2026):

Enterprise Data Platform Engagement

- Lead a 4–7 person DS/analytics team; own roadmap, stakeholder steering, and technical direction across telemetry analytics, ML, and GenAI automation.
- Coordinate transition planning/execution for **Databricks** → **Snowflake** while maintaining daily Spark/BI workloads with several terabytes of sensor data.
- Built a **cross-platform migration validation system** using AI coding agents: automated comparison of schemas, row counts, key distributions, and date ranges across dozens of tables—from zero to live dashboard in ~1 week. ([Case study](#))
- Connected **Databricks and Snowflake CLIs** through an autonomous coding agent (Codex CLI / Cortex Code CLI) to orchestrate migration, validation, and reconciliation pipelines end-to-end—shrinking multi-day manual runbooks into hands-off workflows.
- Redesigned a critical daily pipeline from full recompute to **validated incremental processing**: 5–10× speedup with row-level correctness proof, projected to cut compute costs by **50%+**. Killed a faster but fundamentally broken optimization using multi-scale benchmarks and EXCEPT-based validation. ([Case study](#))
- **Tech:** Snowflake, Databricks, Snowpark, AI Coding Agents, LLM Agents (LangChain), Streamlit, Model Monitoring.

Independent projects (separate from client engagement):

- Open-sourced [ts-agents](#) (PoC): a framework for autonomous time series agents; applied it to **machine activity recognition** from sensor data, demonstrating that classifier-heavy activity recognition workflows can be replaced by agent-driven analysis.

- Built a **Text-to-SQL pipeline against a production ERP database** (hundreds of tables, real data, not a tutorial schema): schema-aware retrieval to fit the context budget, a semantic layer mapping business vocabulary to columns and join paths, and result-level evaluation—a query that runs cleanly against a null column still returns zeros and gets logged as “success”.
- Developed a practitioner framework for **large code migrations with coding agents**: evidence-preservation methodology, behavioral oracles, and AGENTS .md-based quality contracts. ([Case study](#))

RISE Research Institutes of Sweden AB

AI Researcher & Consultant

Linköping, Sweden

Jul 2021 - Aug 2025

Project Lead: Sustainable Fashion AI Automation (2022-2025: 24 months): Leading two major initiatives: [Vinnova: AI for Circular Fashion](#) (Project Lead, ~ 9M SEK) and [CISUTAC](#) (AI Lead, ~ 2M SEK).

- **Challenge:** Manual quality inspection bottleneck in circular fashion supply chain, 30% inconsistency and 25% cost overhead.
- **Solution:** End-to-end computer vision system for automated attribute detection.
- **Impact:** 40% reduction in processing time, 50%+ reduction in data collection costs through synthetic data.
- **Tech:** PyTorch, Vision Transformers, CLIP, Gradio, Docker, Synthetic Data Generation.
- **Recognition:** 1 of 5 projects at [EU sustainable AI](#) (2023).

Low-Energy IoT Solutions (2022: 4 months):

- Identified miniROCKET-based time series methods for industrial edge deployment, enabling real-time analysis with ~90% lower hardware cost than heavier alternatives.

Additional:

- Established AI mentorship program for Master’s thesis students; supported drone edge AI, air-traffic automation, traffic analysis, and anomaly detection projects.

2MNordic IT Consulting AB

Data Scientist & Data Engineer

Gothenburg, Sweden

Dec 2019 - Jun 2021

Project: Early Warning System for Student Performance (2020: 6 months):

- Developed predictive analytics identifying at-risk students in 6th grade, enabling early intervention for 10% of the population (3,000+ students impacted).
- **Tech:** Azure DevOps, Azure Functions, Data Factory, Python, SQL, Power BI.

Project: Mathematics Assessment Optimization (2021: 4 months):

- Analysis of digital test results across 8 schools; findings informed district-wide grading policy adjustments.
- **Tech:** Scikit-learn, Statistical Analysis, Python, Azure Notebooks.

Previous Research Positions

2009–2019

- **Research Fellow, Chalmers University of Technology:** Gothenburg, Sweden
Complex systems modeling, large-scale data analysis 2018–2019
- **Research Scientist, Niels Bohr Institute:** Copenhagen, Denmark
Simulation, forecasting, computational modeling 2015–2018
- **Research Assistant/PhD Student, Univ. of Rochester:** New York, USA
Data analysis, predictive modeling 2009–2015

EDUCATION & CERTIFICATIONS

Microsoft Certified

Azure Data Engineer Certificate

Azure
2020

University of Rochester

PhD in Physics and Astronomy

Rochester, New York (USA)

Oct 2015

Focus: Complex Systems Modeling, Data Analysis, Computational Fluid Dynamics, High Performance Computing, C/C++

AWARDS & ACHIEVEMENTS

- Horton fellowship from Laboratory for Laser Energetics - full research funding award. 2010-2015
- Susumu Okubo Prize for highest performance on graduate comprehensive exam and excellence in coursework. 2011