

[Download Latest PDF Resume](#)

Profile Summary

Senior Data Scientist with 8+ years of experience transforming large-scale data into actionable insights for e-commerce and enterprise SaaS. Specialized in building and deploying scalable ML/AI solutions—including recommendation engines, RAG/LLM pipelines, and time-series forecasting—in cloud environments to drive revenue growth, automation, and operational efficiency.

Education

- Advance Machine Learning Specialisation | Higher School of Economics (2020 - 2022)
- Data Analysis/Statistics | Nackademin (2015 - 2017)
- International Relations | Stockholms University (2010 - 2015)

Courses & Specialisations

- Advance Machine Learning Specialisation | Higher School of Economics (2020 - 2022)
- [Natural Language Processing Specialisation \(2023\)](#)

Publications, Projects & Podcasts

1. Building Baseline Models with Design Patterns in ML (Part 1). [Publication](#)
2. Exploring Stable Diffusion in Google Colab using CUDA: A Step-by-Step Tutorial. [Publication](#)
3. Hyperdimensional Computing: Taking AI to the Next Level by Emulating the Brain. [Publication](#)
4. Data Ethics and Responsible AI. [Podcast](#)
5. Building a Comprehensive RAG Application: Frontend and Backend Development with Azure OpenAI, Flask & LangChain. [Publication](#)

More of my articles can be found on my blog: [Link to my blog](#)

Key Skills

- **Programming Languages & Libraries:** Python, SQL, R, Pandas, NumPy, TensorFlow, PyTorch, Scikit-Learn, Lifetimes Python Package, MLFlow, Streamlit
- **NLP and generative AI:** Hugging Face, NLTK, SpaCy, LangChain, LangGraph, AutoGen
- **Machine Learning & Data Science:** Reinforcement Learning, Deep Learning, Transfer Learning, Federated Learning, Recommendation Systems, Model Optimization, Model Tracking, Time Series Modeling, Feature Engineering, EDA, Bayesian Statistics, Dynamic Pricing

- **Cloud Services:** Snowflake, AWS, Azure, GCP
- **Development Tools & Version Control:** GIT, GitLab, JIRA, Trello, Asana
- **Visualization Tools:** Tableau, Google Data Studio, Power BI, Looker, D3.js
- **MLOps and Deployment:** CI/CD Pipelines, Kubernetes, Docker
- **Big Data Technologies:** Apache Spark, Hadoop
- **Data Governance and Compliance:** GDPR, CCPA
- **Experimentation and Testing:** A/B Testing, Experimentation Frameworks

Soft Skills

- Effective mentor to junior data scientists, interns, and university students, helping them develop skills in Python, ML fundamentals, and production workflows.
- Strong stakeholder communication, translating complex technical concepts into actionable business insights.
- Leadership experience in guiding project direction, reviewing code, and ensuring best practices across teams.
- Adaptability in fast-paced environments with changing priorities and emerging technologies.

Work Experience

Lead Data Scientist - AI Initiatives | Stora Enso (*January 2025 - Present*)

- Act as a full-time AI expert advising across multiple Stora Enso projects, guiding teams from problem definition through PoC development and production deployment.
- Lead technical design and model selection for AI initiatives, ensuring solutions are robust, scalable, and aligned with business objectives.
- Collaborate closely with engineers, data scientists, and engineering leadership to translate complex business challenges into practical AI solutions.
- Mentor data scientists and engineers through code reviews, architectural guidance, and best-practice recommendations, accelerating team delivery and technical maturity.
- Evaluate tools, frameworks, and vendors for AI and ML use cases, enabling informed and sustainable technology decisions.
- Establish AI standards and best practices across projects, improving consistency, reliability, and maintainability of AI solutions.
- Support the transition of AI initiatives from early exploration to validated PoCs and production-ready systems, reducing delivery risk and time-to-value.

Senior Data Scientist | Stora Enso (*June 2022 - 2025*)

AI Negotiation Agent for Stora Enso (*2024- Present*)

Building a production multi-agent AI system that powers B2B pricing negotiations for a global enterprise (~€10B revenue). The system orchestrates 5+ specialized LLM agents — each with domain-specific tools, data sources, and fact-checking — to deliver data-driven negotiation strategies in real time.

Key technical work: - Custom multi-agent orchestrator extending Microsoft's Magentic-One framework - Hallucination detection pipeline that verifies agent claims against source data - RAG system with PGVector for semantic search over financial reports, news, and customer documents - Snowflake → PostgreSQL data pipeline serving profitability analytics to agent tools - Evaluation framework with LLM judges, RAGAS, and reinforcement-learning reward models - Enterprise deployment on Azure (OpenAI, Key Vault, Managed Identity, Docker)

Stack: Python, FastAPI, AutoGen, PostgreSQL/PGVector, Snowflake, Azure, Docker

AI Acceleration with Generative AI Project (January 2024 - 2025):

- Built and deployed company-wide Q&A chatbots using LangChain, Hugging Face, Azure OpenAI, and RAG models.
- Led end-to-end development of scalable AI solutions with Flask, Docker, PostgreSQL, and Azure cloud.
- Researched LLM reasoning and multi-agent systems, applying reinforcement learning and LangGraph for innovation.
- Synthesized cutting-edge research into practical, production-ready generative AI applications.

Links : [Link to repository with demo examples of the implementation](#)
(Note that this project is still ongoing)

Topic Modelling for Safety notifications(January 2024 - December 2024): - Managed extensive volumes of safety notifications in textual format, utilizing topic modelling techniques to identify themes and generate clusters. - Employed models such as LDA, BERT, K-means, and hybrid variations to extract insights from the data efficiently.

Precision Forestry Project (August 2023 - December 2023):

- Optimized code structure using design patterns.
- Refactored regression models and implemented Azure AutoML SDK.
- Streamlined development processes, enhancing overall codebase efficiency.

Selfly Store Project (June 2022 - June 2023):

Leveraged AI to transform vending machines into AI-driven systems. Developed a patentable AI solution, currently in the process of being filed in the US. Project Stages:

A. Analytics/Engineering Stage:

- Managed and processed extensive cloud-stored data.
- Designed a preprocessing pipeline for product inventory.
- Improved product categorization using a logic-based approach and TF-IDF-based classifier.

- Created a cabinet lifetime value model for performance assessment.

B. Artificial Intelligence Stage:

- Enhanced cabinet intelligence through forecasting using LSTM models and ensemble methods.
- Developed deep neural network-based models for propensity score.
- Built collaborative filtering and hybrid recommender systems.
- Utilized Bayesian statistics for dynamic pricing to optimize revenue and customer satisfaction.

Tech Stack: Databricks, MLFlow, Hugging Face (Open-source NLP library), Lanchain, Streamlit, Flask, Docker Containers, Azure Open AI, PostgreSQL, Azure AutoML SDK, Scikit-Learn, Python, SQL, Snowflake, GIT, Azure Cloud, MongoDB

Data Scientist | TUI Global (March 2021 - June 2022)

Customer Life Time Value Project (Mar 2021 - Nov 2021):

- Goal: Develop a Customer Life Time Value model to inform business decisions.
- Solution: Implemented probabilistic and Neural Network models using Lifetimes Python package.
- Results stored in Snowflake; shared via Tableau dashboard.

Recommender Systems Project (Mar 2021 - Oct 2021):

- Goal: Personalize destination and hotel recommendations at the user level.
- Solution: Built item-based collaborative filtering and user-based probabilistic models.
- Utilized XGBoost for classifier model; recommendations sent via email.
- Results stored in Snowflake; ongoing work on API deployment.

Hotel Search Ranking Model Project (Mar 2021 - Jan 2022):

- Goal: Develop a hotel ranking solution using booking and clickstream data.
- Solution: Stacked model combining booking and clickstream data.
- Utilized RandomForest and XGBoost classifiers; results stored in MLFlow.
- Model to be integrated into website's search section.

ML Pipeline Automation Package POC (Mar 2021 - June 2022):

- Goal: Automate daily machine learning workflows to free up time for other project tasks.
- Solution: Developed flexible Python functions/classes covering all ML pipeline stages.
- Output stored in MLFlow and local folders for easy access and analysis.

Tech Stack: Lifetimes Python Package (Probabilistic modeling), Snowflake, Tableau, XGBoost, MLFlow (Model tracking), Python, SQL, Google Cloud, Amazon Web Services, GitLab

Data Science & Data Analytics | TUI Nordics (April 2019 - March 2021)

Machine Learning & Data Science:

- Built attribution models using Markov chains for revenue allocation and optimization.
- Conducted customer segmentation via RFM key metrics: recency, frequency, and monetary value.
- Implemented classification and regression algorithms such as KNN, Random Forest, and XGBoost.
- Engaged in error monitoring, feature engineering, and exploratory data analysis.
- Developed recommendation engines and evaluated recommender systems.
- Applied time series modeling techniques including LSTM, XGBoost, and ARIMA variations.

Data/Web Analysis:

- Provided weekly reporting and analysis to stakeholders for website optimization.
- Utilized SQL (BigQuery) and Python for data analysis.
- Extracted specific data from databases using SQL and analyzed further using Python, Jupyter Notebooks, and PyCharm.
- Created data monitoring dashboards using Google Data Studio or Tableau based on task objectives.

Tech Stack: Snowflake, Tableau, Python, SQL, Google Cloud, Amazon Web Services, GitLab, Google Cloud, Google Data Studio

Guest Lecturer (Data Analytics/Pandas Programming) | Hyper Island (December 2020)

I had the pleasure to be invited as a guest lecturer for one day to teach data analytics students of Hyper Island about the most modern techniques within data analysis with Pandas and Python. This was an incredible experience considering that the whole lecture was online due to the pandemic. I would love to have more opportunities to keep teaching students about the most high in demand skills within data analytics since I also have the opportunity to learn a lot.

Tech Stack: Python, Pandas, Jupyter Notebook

Data Analyst - E-commerce | Off The Wall (OTW) (June 2017 - April 2019)

- Conducted web/marketing analysis, optimizing website performance and content.
- Provided analytical support for business development and editorial teams.
- Implemented A/B testing and conversion optimization strategies.
- Managed Google Adwords campaigns, conducted keyword analysis, and developed reports from Google Analytics.
- Utilized Google Tag Manager for tag management and analyzed

SEO performance.

- Applied advanced statistical techniques for predictive analysis.
- Collaborated with social media team, managing ads, optimizing campaigns, and analyzing data using Facebook Analytics and Facebook pixel.

Contact info:

- **Email:** tabers77@gmail.com
- **GitHub:** [GitHub Profile](#)
- **Data Science Portfolio:** [GitHub Repository](#)
- **LinkedIn:** [LinkedIn Connect](#)
- **Digital Resume & Portfolio:** [[Digital Resume & Portfolio](#)]