Guus Bouwens

(253) 409-7155 • guus@bouwens.nl LinkedIn • Portfolio

EDUCATION

Texas Tech University, Rawls College of Business

Lubbock, TX, USA

Master of Science in Data Science

Expected May 2026

- Recipient of two academic scholarships; GPA: 4.0/4.0
- Relevant Coursework: Big Data Strategy and Security, Business Intelligence, Simulation & Optimization, Applied AI

Vrije Universiteit Amsterdam, School of Business and Economics

Amsterdam, NL

Bachelor of Science in Econometrics and Data Science, Minor in Big Data

Sep 2020 - Mar 2025

- Led three capstone projects that earned the highest marks
- Relevant Coursework: Multivariate Statistics, Linear Algebra, Analysis, Time Series, Machine Learning, Databases

PROFESSIONAL EXPERIENCE

Texas Tech University, Graduate School

Lubbock, TX, USA

Sep 2025 - Current

Statistical Consultant

• Advised graduate students, postdocs, and faculty on all stages of the quantitative research lifecycle, from experimental design and data cleaning to model selection and interpretation of results.

McDermott International, Ltd

The Hague, NL

AI Engineering Intern

Oct 2024 - Apr 2025

- Boosted model accuracy 2x and speed 5x by architecting an advanced RAG system with contextual chunking, hybrid retrieval, and reranking, deployed on Azure using Python, LlamaIndex, Chroma and CI/CD pipelines.
- Architected a security-focused **evaluation framework with in-text citations** to guarantee response factuality, partnering with PhD researchers and project owners in an Agile workflow.

Invest-NL Amsterdam, NL

AI Engineering Intern

Apr 2024 - Sep 2024

- Collaborated with Deep Tech and Energy investment teams to fine-tune a Llama-3-70B model, boosting accuracy from 60% to 95% for company classification while cutting API costs through output optimization.
- Automated the MLOps lifecycle from data ingestion to deployment using Python, Power BI, and Azure CI/CD, enabling continuous monitoring and dashboard updating.

Talpa eCommerceAmsterdam, NLData Science InternMar 2023 - Oct 2023

• **Lifted recommendation engine performance, increasing activity by 10%,** by adding **cosine similarity** measures, informed by customer insights uncovered via **Fourier analysis** of multi-year behavioral data.

• **Engineered a scalable deployment** by containerizing the model with **Docker** and deploying on **AWS SageMaker** with **Jenkins** for CI/CD, within a **Jira**-managed Agile workflow.

PROJECTS

- Applied signal processing and wavelet-based denoising to pulse data from 1,500 individuals, developed a Random Forest model with SHAP feature attribution, resulting in a heart rate predictor with MAE of 8 bpm.
- Trained a **Gradient Boosting classifier** on Mendeley Women's Health Dataset from temperature and sleep data to **predict menstrual cycle phases**, achieving **74% cross-validation accuracy**.
- Implemented **mechanistic interpretability** methods—including activation patching and causal tracing in **PyTorch**—to visualize neural patterns in the DeepSeek-R1 **transformer** model across 15 cognitive tasks.
- Converted FoodSeg103 to YOLO-seg polygons and **fine-tuned YOLO11x-seg** on GPU, enabling **accurate multi-ingredient food segmentation**.

CERTIFCATES

AWS Fundamentals | Microsoft Azure Data Scientist Associate | AI in Healthcare (Stanford)
Agentic AI & RAG (IBM) | Deep Learning and NLP (DeepLearning.AI) | Machine Learning (UW & ICL)

SKILLS_

AI & DS: Python (PyTorch, TensorFlow, Scikit-learn, Keras, Pandas), NLP (LLM Fine-Tuning, RAG, Transformers), Time-Series Analysis, Statistical Modeling, Monte Carlo

MLOps & Data: AWS (SageMaker), Azure (CI/CD), GCP, Docker, MLFlow, SQL, PostgreSQL, MongoDB, ChromaDB **Development & Visualization**: C++, VS Code, FastAPI, Git, GitHub Actions, Power BI

Languages: Dutch (Native), English (Fluent, 110/120 TOEFL).