

# ANDREI-IOAN BLEAHU

## AI Engineer & Data Scientist

@andreibleahu@gmail.com

+372 53603943

Tartu, Estonia

### PROFILE

MSc Data Science graduate (Distinction, Dean's List, GPA 18.1/20) from the University of St Andrews with hands-on experience building LLM-based systems, including retrieval-augmented generation pipelines, FastAPI services, and evaluation workflows. **Background in computational imaging and optical physics research**, with strong quantitative foundations and experience developing applied data and ML systems across startup and research environments.

### EDUCATION

MSc Data Science (Distinction (GPA 18.1/20) – Dean's List)

University of St Andrews

Sept 2023 – Sept 2025

Remote

- **Dissertation:** Implemented **Bayesian Neural Networks** in the language of **Julia** from scratch for efficient optimisation; achieved up to **90% parameter reduction** in MLP architectures during training, including significant parameter reduction in Feed-Forward Networks layers embedded into Transformer networks, while preserving predictive performance.

MPhys Physics (2:1 Hons)

University of Southampton

Oct 2014 – Jun 2019

Southampton, UK

- Final Year Project (First Class, 71/100): trade-off analysis of hollow-core fibres in quantum repeater architectures

### WORK EXPERIENCE

AI Engineer (Contributor)

Maxibook (Early-Stage Startup, Contributor)

Feb 2026 – Present

Remote

- Designed and deployed **6 independent production microservices** on AWS EC2 with systemd for a knowledge extraction AI engine, including an end-to-end runner for advanced testing.
- Implemented retrieval-augmented generation (**RAG**) pipelines with **Tavily search integration** and LLM-as-a-judge evaluation workflows, to improve the quality of generated content and add validation guardrails before final cleanup
- Developed a **hybrid semantic clustering approach** combining agglomerative clustering with LLM-generated semantic labelling, prioritising deterministic grouping behaviour and lower gpt-inference cost compared to fully prompt-driven clustering strategies.
- Delivered full **technical documentation** and a deployment operations guide for each service.

Data Analyst

Car Rental Gateway

Sept 2025 – Jan 2026

Tartu, Estonia

- Designed **Python + SQL reporting workflows** for daily reservation monitoring, built to be reproducible and auditable, effectively improving the data correctness and validation process.
- Conducted a **cross-system financial reconciliation** between Stripe/Braintree transaction records and the BI tool, identifying unreported cost items (platform taxes, fees) and delivering concrete recommendations for incorporating these fees in internal reporting.

Junior Research Fellow – Computational Imaging

University of Tartu

Aug 2022 – Jul 2023

Tartu, Estonia

- Built end-to-end **computational imaging pipelines** for optical data acquisition and reconstruction, focusing on inverse modelling, signal quality optimisation, and cost-efficient imaging architectures.
- Contributed to the development of the **SWITSH imaging technique**, achieving up to a **10× improvement in optical resolution** under experimental conditions.
- First author on **two peer-reviewed research publications**; presented research at the international **SPIE conference** and co-organised the “**Ciphering with CIPHR**” hackathon.

---

## Data Scientist (Contract)

### Gauss & Neumann

📅 Jul 2023 – Sept 2023

📍 Remote

- Delivered end-to-end **digital marketing campaign analysis**: data extraction, cleaning, performance modelling, and investment recommendations

## CORE SKILLS

---

### Languages:

Python

SQL

Julia

### AI & ML:

LLMs

RAG

LangChain

PyTorch

Bayesian ML

Prompt Engineering

### Backend & APIs:

FastAPI

REST APIs

Pydantic

Docker

### Cloud & DevOps:

AWS EC2

Systemd

Linux / Ubuntu

Git

### Data & Analysis:

Pandas

NumPy

Jupyter

Statistical Modelling