Tomas Pop

Machine Learning Engineer | MSc. Graduate

✉ pop.tomas@outlook.com 📞 +420 723 006 468

📍 Liberec, Czechia (open to relocation)   in linkedin.com/in/tomas-pop

 github.com/poptomas   📄 Citizenship: Czech (EU)

WORK EXPERIENCE

Machine Learning Engineer   08/2022 – 02/2026
Cisco Systems, Inc.   Prague, Czechia
- Phishing Detection: Contributed to an email anti-phishing engine protecting hundreds of millions of emails weekly.
- LLM Pipeline: Built a modular library for rapid text-classification deployment, processing 550k+ weekly detections.
- Optimization: Reduced ML-based models' inference latency by 90%+ via PyTorch, ONNX, and CUDA on AWS.
- Research: Contributed to a collaborative research paper (arXiv) and helped bridge the gap between experiments and production components.

EDUCATION

Master's Degree (Mgr., equivalent to MSc.) - Software and Data Engineering   10/2023 – 02/2026
Charles University, Faculty of Mathematics and Physics ⊘   Prague, Czechia
- Specialization: Big Data Processing
- Graduated with Honors (Grade: 1.0 - best possible)

Graduate Exchange Semester   02/2025 – 07/2025
University of Basel, Department of Mathematics and Computer Science ⊘   Basel, Switzerland

Bachelor's Degree (Bc., equivalent to BSc.) - Programming and Software Development   10/2019 – 09/2023
Charles University in Prague, Faculty of Mathematics and Physics ⊘   Prague, Czechia

LANGUAGES

English	**German**	**Czech**
C1 - fluent (Certificate of Advanced English)	C1 - professional working proficiency	Native speaker

KEY PROJECTS

Master Thesis (done in collaboration with Cisco) ⊘   09/2024 – 02/2026
Phishing Detection Using Knowledge Distilled From Large Language Models
- Available in the Charles University Digital Repository ⊘

Swiss Public Transport Visualizer using SBB API ⊘   04/2025 – 05/2025
- Co-led the development of an interactive geospatial tool on the SBB Open Transport API, visualizing real-time data from 2,500+ stations with 100% coverage of SBB regions.
- Designed scalable architecture using Docker for containerization and implemented data streaming principles, with the system explicitly prepared for Apache Spark deployment to handle Big Data volumes.
- Developed a responsive UI with Folium and Leaflet.js, enabling fast station-level insights and map-based queries.
- Technologies: Python (Streamlit, Folium), Docker, Apache Spark, NetworkX

TECHNICAL SKILLS

Machine Learning
- Transformers
- LLMs
- PyTorch, Hugging Face
- Knowledge Distillation

Cloud
- AWS (EC2, S3, Sagemaker)
- Docker & Containerization

Programming Languages
- Python (Primary)
- C++, C#, Java, JavaScript

Model Optimization & Production ML
- ONNX, TensorRT, CUDA
- GPU Acceleration
- Model Optimization & Deployment
- High-Throughput Inference Pipelines

Data Science & Analytics
- Exploratory Data Analysis, Statistical Modelling, Data Cleaning
- SQL & NoSQL databases (Redshift, RavenDB)
- Data Visualization
- Apache Spark
- Data Streaming & Real-Time Analytics

PUBLICATIONS

A Modular and Adaptive System for Business Email Compromise Detection 🔗   07/2023
- Contributed to a collaborative industry research paper (arXiv) and helped bridge the gap between experiments and production components.