APOSTOLOS NTELOPOULOS

(+45) 52640486 ◇ Denmark
ntelopoulosa@gmail.com ◇ linkedin.com/in/antelopoulos/ ◇ https://github.com/antelopoul

OBJECTIVE

Machine Learning Engineer with experience applying predictive modeling and deep learning techniques to real-world systems. Strong focus on time-series forecasting, feature engineering, and model optimization using XGBoost and deep learning methods. Delivered production ML systems through Kubeflow pipelines and Kubernetes-based serving with KServe. Experienced in designing REST APIs and scalable architectures that integrate ML models into operational systems with monitoring tools Grafana. Interested in building applied ML with measurable impact on forecasting, optimization, and developing intelligent agentic systems with MCP servers & tools .

EDUCATION

MSc Computer Engineering / Artificial Intelligence, Vision & Sound, Aalborg University 2022–2024
Specialized in deep learning methods for computer vision and AI systems in engineering contexts.
Master's thesis focused on video anomaly detection using vision-language models, accepted at IEEE IPTA.

Thesis: "CALLM: Cascading Autoencoder and Large Language Model for Video Anomaly Detection"

BEng Electronic Engineering, ASPETE 2014 - 2017 Courses: Digital Signal Analysis & Processing, Broadcasting Systems, Data Acquisition, Microcomputers, Computer Networks

Thesis: "Sound Classification using Convolutional Neural Networks"

Skills: Python, Tensorflow and handling large datasets including Google's AudioSet to fine-tune a lightweight CNN (i.e. MobileNet v2) for classifying environmental sounds.

SKILLS

Technical Skills	Machine Learning: PyTorch, TensorFlow, Scikit-learn Data/ML Systems: Kubeflow, Kafka, PostgreSQL, S3 Backend & Infra: Python, Docker, Kubernetes, REST APIs Data Processing: Pandas, NumPy, OpenCV Optimization: OR-Tools (CP-SAT) MLOps: CI/CD , model versioning, deployment pipelines
ML Focus Areas	End-to-end ML systems (training → deployment → monitoring), time-series forecasting, regression modeling, distributed ML pipelines, applied optimization models, VLMs

EXPERIENCE

Software Developer — Machine Learning Specialist Jan 2025 -
Veovo *Aalborg, Denmark*

- Built and operationalized Kubeflow pipelines to support end-to-end ML experimentation at scale, including automated data extraction from PostgreSQL and S3 data lake, feature engineering, model training, hyperparameter tuning, and rigorous validation with baseline comparisons and deployment of regression and time series models.
- Built and maintained Bamboo-based CI/CD pipelines for ML and software services, including automated training validation, unit/integration testing, artifact versioning, and deployment to Kubernetes environments.
- Designed and deployed cloud-native microservices on Kubernetes, integrating REST APIs and Kafka-based messaging to support scalable data processing, model serving, and distributed ML workflows.
- Collaborated with distributed cross-functional teams across engineering and product functions to define requirements and deliver production solutions.
- Designed and implemented optimization models using OR-Tools CP-SAT, translating business requirements into mathematical constraints and objective functions to generate scalable, high-quality scheduling solutions.

Presales Engineer (Internship) Jul 2020 - Jan 2021
OTE SA *Athens, Greece*

- Contributed to the architecture and design of large-scale distributed network infrastructure, including a $3.4M IP/MPLS VPN system (NETVIS) connecting 84 global sites, applying technical feasibility analysis, system design principles, and cross-functional collaboration to define scalable, reliable infrastructure under real-world constraints.

PROJECTS

Diffusion Models and GANs for Super-Resolution. Creating synthetic low-resolution and low-light images using Stable Diffusion. Trained it on a multi-gpu cluster and evaluated it on a super-resolution (SR) framework, whereas the results showed a significant enhancement in light & image quality.

Implicit Neural Speech and Audio Compression. Focus on compressing audio signals using Implicit Neural Representation. Optimization was achieved through techniques including quantization and pruning. Engaged also with the research on meta-learning.

Leveraging LLMs for Effective Anomaly Detection. Worked on a multimodal Visual Inspection Model, which provided the localization result in a 3x3 grid and textual explanation using the LLM after fine-tuned with LoRA (training a small set of weights with matrix decomposition). The application was investigated in the Industrial Anomaly Detection domain on public datasets including MV-Tec and Visa.

CALLM: Cascading Autoencoder and Large Language Model for Video Anomaly Detection. Published MSc research in a peer-reviewed conference, designing and evaluating a cascaded 3D Autoencoder–Visual Language Model framework for abnormal scene detection and mitigation of out-of-distribution failures in surveillance video. Integrated and fine-tuned the Video-LLaMA foundation model using PEFT/QLoRA, performed large-scale distributed training and inference on SLURM-managed GPU clusters, and benchmarked the proposed approach against baseline methods.

Code: https://github.com/antelopoul/CaLLM

Paper: https:/vbn.aau.dk/ws/portalfiles/portal/740054157/author.pdf

EXTRA-CURRICULAR ACTIVITIES

- Developed a vector database system using the Unsplash dataset to study embedding-based similarity search and vector indexing mechanisms for large-scale retrieval applications.

LANGUAGES

English (Full professional proficiency), Danish (Professional working proficiency, daily use), Greek (native)

ACTIVITIES

- Sailing
- Member of a wakeboard club, supporting park operations and participating in community activities.