CHARALAMBOS THEODOROU

4, Dumfries Street, Luton, BEDFORDSHIRE. LU1 5FT
Tel. 07493237732
pampos20101995@hotmail.com

PROJECTS

ANTICIP8 — Human Behaviour Prediction Foundation Model - https://anticip8.ai/
Foundational model predicting *actions* (behaviour) rather than next-word text, enabling proactive decision-making across domains.

- Built **end-to-end from scratch**: research - data design - model training - evaluation - optimization - cloud deployment.
- Developed transformer-based behaviour prediction pipelines (e.g., **BERT-style encoder and downstream action heads**) and calibrated decision thresholds for reliable operational use.
- Designed benchmarking, failure-mode tests to ensure robust behaviour predictions across contexts.

MELELEM — Meta-Learning Enterprise AI (Privacy-Preserving Optimization) - https://melelem.ai/
Real-time evolution of enterprise AI via meta-learning with a Matrix Scoring System for dynamic optimization and reliable outputs.

- Built **0→1**: architecture design, implementation, testing, and deployment for mission-critical use cases.
- Implemented continuous evaluation and scoring loop to optimize responses while maintaining privacy and consistency requirements.

MAV — Multimodal Deepfake Detection (Voice, Text, Video) - Part of https://fountech.ai/

- Led development of a **multimodal detection system** combining audio, textual, and visual signals for robust deepfake classification.
- Led a **32-person cross-functional team**; reviewed research, set technical direction, and drove critical implementation decisions across model and pipeline components.

HUMANIX — Distributed Human–Machine Synergy Platform - Coming soon
Culturally-aware AI interactions using demographic targeting and graph-based reasoning.

- Built **from scratch**, designing a distributed AI workflow combining demographic/context features with **graph neural networks (GNNs)** for nuanced interaction policies.

SPRYT (NHS) — No-Show Prediction System (NSP) - https://spryt.com/

- Built a **no-show / show prediction** solution end-to-end for healthcare operations.
- Implemented a **DQN reinforcement learning** approach to optimize intervention strategies (e.g., reminders, timing, channel selection) and improve attendance outcomes.

SOFFOS — AI Training Solution Builder - Part of MELELM now into 1 unified platform
Platform enabling non-technical teams to create AI training solutions.

- Collaborated on the **AI/ML components**, contributing to model selection, evaluation, and integration into the product workflow.

AERO — AI Firefighting Drones (Thermal + Multi-Agent RL) - https://www.raidsai.ai/

- Built **0→1**: thermal-based detection and **multi-agent reinforcement learning** for real-time coordination and suppression optimization.
- Developed simulation-to-deployment workflow for robust decision-making under uncertainty.

RAIDS — Real-Time AI Safety Monitoring
AI safety monitoring layer that detects rogue behaviour early; starting with LLMs and expandable to regulated sectors.

- Contributed to the **AI safety layer**, including risk detection logic, automated monitoring signals, and alignment-oriented evaluation for LLM behaviour.

Tech (across projects): Python, SQL, (some Java/C++) | PyTorch, TensorFlow, scikit-learn | Transformers/LLMs (BERT-style encoders, LLaMA/GPT-style integration) | RAG + Context Engineering | AI Safety (Red-Teaming, Monitoring, Eval Harnesses, Policy Alignment) | Multimodal ML (audio/text/video) | RL (DQN, Multi-Agent RL) | Graph Neural Networks | MLOps/CI-CD | AWS/Azure | Docker/Kubernetes | ROS (vSLAM work)

WORK EXPERIENCE

Lead AI Researcher / Engineer, Fountech AI **Sept 2024- Present**

- Led end-to-end delivery of multiple greenfield AI projects, building production systems from scratch and leading cross-functional teams of **31** engineers/researchers.
- Developed and deployed **LLM/Transformer** solutions across **generative AI**, **multimodal foundation models**, and **agentic/AGI-style** workflows.
- Ran **red-teaming and adversarial evaluation** campaigns to uncover jailbreaks, policy bypasses, and failure modes using automated test frameworks and LLM-based grading.
- Improved **model alignment** via **preference tuning**, prompt optimization, and **context engineering** to meet safety and reliability requirements.
- Built **data pipelines** for training/benchmarking: dataset querying, curation, labeling strategies, and quality gates to produce high-signal evaluation corpora.
- Implemented evaluation harnesses and reporting for safety, robustness, and quality metrics to support release readiness and governance.

Tech: **Python, SQL, Java** | LLMs/Transformers | RAG | Prompting | Safety/Alignment

Machine Learning Engineer - CTO, Briteyellow **Feb 2021-2024**

- Led development of production **ML + LLM** systems from prototype to cloud deployment, avoiding additional hardware spend and saving **$500K+** while improving business outcomes.
  Built and evaluated **NLP/LLM** solutions (incl. **LLaMA, GPT-3**) using structured prompting, context engineering, and model optimization to drive measurable efficiency and cost reductions.
- Owned **AI safety & reliability** workstreams: red-teaming/adversarial testing, policy-aligned prompting, and automated evaluation to reduce unsafe or off-policy outputs.

- Developed scalable real-time and batch inference pipelines on **AWS (SageMaker, Lambda, S3)** and **Azure (Azure ML, AKS)** with monitoring and performance/latency tuning.
  Partnered with DevOps to implement **CI/CD for ML**, improving reproducibility and cutting deployment time by **30%**.
- Designed and shipped **personalized recommender systems** (collaborative filtering + deep learning), increasing engagement **20%** and sales **25%**.
- Improved model quality through feature/data refinement and tuning, achieving **+20%** accuracy and stronger decision-making signals for stakeholders.
- Developed a novel **vSLAM** framework for indoor mapping (train stations) using **ROS, C++, Python**; published **4 papers** in top venues and contributed to **5 patents**.
- Communicated complex ML/LLM and safety trade-offs to non-technical stakeholders, supporting roadmap and risk decisions.

Tech: Python, SQL, Java | PyTorch, TensorFlow, scikit-learn | Transformers, LLMs (LLaMA, GPT-3) | Prompting & Context Engineering | AI Safety (Red-Teaming, Alignment, Eval Harnesses) | MLOps/CI-CD | AWS (SageMaker, Lambda, S3) | Azure (Azure ML, AKS) | Docker/Kubernetes | ROS, C++

Machine Learning Engineer NLP, ImpacTech **Sept 2017- Aug 2018**

- Applied **NLP + sentiment analysis** to extract actionable insights from customer interactions on the Impact AI platform.
- Built an end-to-end pipeline for call analytics: **speech-to-text**, text cleaning/normalization, tokenization/lemmatization, and feature engineering.
- Developed and deployed sentiment classification models to detect **positive/negative/neutral** intent and track sentiment trends across sales and support calls.
- Analyzed call transcripts to surface **keywords/phrases** correlated with successful sales outcomes and higher customer satisfaction.
- Integrated NLP models into the production platform for **near real-time insights**, partnering with data scientists and software engineers on deployment and monitoring.
- Improved sentiment model performance by **+8% accuracy** through enhanced feature engineering and model iteration.
- Identified language patterns linked to customer frustration; informed agent training updates that reduced repeat support issues by **15%**.
- Delivered a **keyword-spotting / compliance flagging** model to detect potential policy violations and enable proactive review by the compliance team (AI safety/compliance).

Tech: Python | NLP (spaCy, NLTK) | ML (scikit-learn) | Speech-to-Text (ASR pipelines) | Text Classification | Feature Engineering | Real-time Inference | Monitoring & Evaluation | Data Processing (Pandas/SQL)

RELEVANT SKILLS

Programming Languages

- Python, C++, ROS, R, MATLAB, Java
- Proficient in computational libraries, OOP, TDD, and multi-processing

Machine Learning

- Supervised and unsupervised learning algorithms.
- Bayesian Machine Learning.
- Deep learning frameworks like PyTorch, TensorFlow, Keras
- Hands-on experience with DL libraries: ONNX, OpenVino
- Scikit-learn for traditional machine learning algorithms.
- XGBoost and LightGBM for gradient boosting.
- Data Forecasting and Recommender Systems.

Natural Language Processing (NLP):

- NLTK (Natural Language Toolkit) and spaCy for NLP tasks.
- Transformers (Hugging Face Transformers) for working with Large Language Models.
- LLMs (GPT-3/4, BERT, LLaMA and Bard).
- RAG Framework.
- Gensim for topic modelling.

Data Science

- Data cleaning and pre-processing.
- Exploratory data analysis (EDA).
- Data visualization using tools like Matplotlib, Seaborn, Plotly and Tableu.
- NumPy, Pandas, SciPy and Matplotlib.
- Statistical analysis and hypothesis testing.
- Data mining and pattern recognition.

Computer Vision

- SLAM, vSLAM, feature extraction, image registration, and camera calibration.
- Visual odometry algorithms such as feature-based (FAST, SIFT) and direct methods (Direct Sparse Odometry).
- Sensor fusion techniques (cameras, LiDAR, IMU).
- Convolutional Neural Networks (CNNs) for image analysis.
- Object detection and image segmentation techniques.
- Experience with CV libraries: OpenCV, Pillow, torchvision

Statistics

- Descriptive statistics: Mean, median, mode, standard deviation, variance, percentiles, and quartiles.
- Inferential statistics: Hypothesis testing, p-values, confidence intervals, and effect sizes.
- Probability theory: Probability distributions (e.g., normal, binomial, Poisson, exponential, and uniform), Bayes' theorem, and conditional probability.
- Random samples and sampling distributions, Parameter estimation and Bayesian inference.

Reinforcement Learning

- OpenAI Gym.
- Q-Learning, Deep Q-Network (DQN), Policy Gradient Methods, SARSA, Monte Carlo Methods, Temporal Difference (TD) Learning, Asynchronous Advantage Actor-Critic (A3C).

Tools and Libraries

- Jupyter Notebook or JupyterLab for interactive data analysis.
- Version control systems like Git and Docker for containerization.
- Database systems like SQL for data retrieval.
- Cloud platforms for data storage and processing, such as AWS, Azure, or Google Cloud.

EDUCATION

PhD Computer Vision and Machine Learning, University of Bedfordshire **2021-2024**

MSc Data Science and Machine Learning, University of Bath **2019-2020**

BA (Hons) Degree Computer Science, Kingston University **2015-2019**

A Levels: Mathematics, Computer Science **2010-2012**

Agios Ioannis Lyceum: Further Maths, Physics and Computer Science **2010-2013**

RELEVANT SKILLS & COURSES

- Microsoft Office (ECDL completed 2010); daily use of Microsoft Word and LaTeX (thesis and publications), PowerPoint (conference presentations), and Excel (data analysis).
- Completed courses in quantitative finance and algorithmic trading.
- Strong background in machine learning algorithms and quantitative research methodologies.
- Current driving licence.

INTERESTS

- Volunteering – peer reviewer at MDPI journal, I am responsible for reviewing and evaluating research papers in the fields of machine learning, robotics, data science, and artificial intelligence.
- Gym – Daily gym goer.

REFEREES

- Upon Request.