# Andreas Varvarigos

New Haven, CT, USA | andreas.varvarigos@yale.edu | (203)-627-8783 | LinkedIn | Github | Google Scholar

## Research Interests

I am interested in Large Language Models (LLMs) and multimodal foundation models, including time-series models, and in extending and aligning them through post-training reinforcement learning and geometric and graph-structured representations, with applications in network observability and finance.

## Education

**Yale University**, Ph.D. in Electrical and Computer Engineering — Aug 2024 – Present

- **Advisor:** Leandros Tassiulas

**Yale University**, M.Sc. in Electrical and Computer Engineering — Aug 2024 – May 2026

- GPA: **4.0 / 4.0**
- **Courses:** Deep Learning for Graph-Structured Data, Intermediate Machine Learning, Building Distributed Systems, Network Algorithms and Stochastic Optimization, Deep Learning Theory and Applications, Trustworthy Deep Learning etc.

**Imperial College London**, M.Eng. in Electronic and Information Engineering — Sept 2020 – June 2024

- **Grade:** First Class Honors, Dean's List Prize Years 1-3
- **Courses:** Machine Learning, Deep Learning, Optimization, Computer Vision and Pattern Recognition, Signal Processing and Machine Learning for Finance, Mathematics for Machine Learning, Communications Networks, etc.

## Work Experience

**Logic Design Engineer**, NVIDIA — Aug 2023 – Oct 2023

- Implemented scheduling algorithms in a NIC that can deliver total packet order at 100 Gbps, outperforming, throughput-wise, SOTA designs with Lamport clocks
- Used DPDK, a set of libraries/drivers for fast packet processing that allows networking applications to run much faster

**Software Developer**, OpenLightComm Ltd. — Dec 2022 – Sep 2023

- Designed and implemented in Python an access planning tool for optimally deploying the optical front/mid/backhaul of 5G networks, to serve residential and 5G mobile users, different base station types (macro and small cells), wireless environments, and functional split choices

**Research Intern**, Harokopio University of Athens — June 2022 – Aug 2022

- Implemented SOTA intent-based allocation algorithm for cloud computing resources using Q-Learning
- Implemented and evaluated three major edge computing storage systems for IoT, namely MinIO, IPFS, and BigchainDB, by considering transaction rate, number of user requests, response time, and resource utilization

**Research Intern**, National Technical University of Athens — June 2021 – Aug 2021

- Implemented image segmentation model based on attention recurrent residual U-Net for visual road crack inspection

## Selected Publications

- LitBench: A Graph-Centric Large Language Model Benchmarking Tool For Literature Tasks; **A. Varvarigos**, A. Maatouk, J. Zhang, N. Bui, J. Chen, L. Tassiulas, R. Ying, *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2026, Jeju, South Korea
- TelecomTS: Observability Dataset for Multi-Modal Time-Series and Language Analysis; A. Feng*, **A. Varvarigos***, I. Panitsas, D. Fernandez, Y. Guo, J. Wei, J. Chen, A. Maatouk, L. Tassiulas, R. Ying, submitted to ICML 2026
- Intent-based Allocation of Cloud Computing Resources Using Q-Learning; P. Kokkinos, **A. Varvarigos**, D. Konidaris, K. Tserpes, *International Symposium on Algorithmic Aspects of Cloud Computing*, Netherlands, 184–196, Sept. 2023

## Teaching Experience

**Graduate Teaching Assistant, Yale University, New Haven, CT, USA**

- ECE 4551 / 9631 Network Algorithms and Stochastic Optimization — Jan 2026 – May 2026
- ECE 2020 Introduction to Communications and Control — Aug 2025 – Dec 2025

## Honors/Awards

- Bodossaki Foundation Scholarship for Graduate Studies, 2024-2029
- Columbia Global Centers Lab-to-Market; selected as one of 25 teams, 2025-2026

- Dean's List Years 1, 2, 3 awarded by Electrical & Electronic Engineering Dept of Imperial to the students that are within the top 10% and have achieved the equivalent of a 1st for the year, 2020-2023
- Member of National Math Team of Greece for International Math Olympiad and Balkan Math Olympiad, 2020
- Gold (11th grade, 2019; 12th grade, 2020) and Bronze (9th grade, 2017) medals in the Euclid Math Competition; Silver (2020) and Bronze (2019) medals in the Archimedes Math Olympiad (10th–12th grade)

## Selected Projects

**Time-Series Multimodal LLM for Telecom Network Observability Data, 2025**

- Implemented a TS multimodal LLM for telecom data that processes 18 network KPIs to answer diverse QA tasks.
- Supports network condition questions (e.g., jamming presence, congestion, user activity), anomaly-related QAs (e.g., detection, root cause, bounds), and numerical reasoning (e.g., mean, periodicity, trend of KPIs)
- Implemented reinforcement learning via Guided Reasoning PPO (GRPO) with cold-start training to support chain-of-thought reasoning using the VERL framework.

**Hybrid Graph-Language Model for Chemical Property Prediction, 2025** GitHub

- Developed a hybrid model combining Graph Neural Networks (GNNs) and Large Language Models (LLMs) for molecular property prediction by injecting GNN-derived graph embeddings and textual context into the LLM
- Enabled accurate natural language generation of chemical descriptors (e.g., molecular weight, polarity, atom counts) across diverse molecules

**GNN-Policy RL for Delay-Aware Packet Routing, 2025** GitHub

- Designed an RL agent for adaptive packet routing using a GNN to make per-node forwarding decisions
- Simulations on small-world topologies under varied traffic loads demonstrated our framework's effectiveness

**Temporal GNNs for Stock Return Prediction, 2024** GitHub

- Built a framework for stock return prediction on S&P 500 data using temporal Graph Neural Networks (GNNs)
- Constructed multiple stock dependency graphs and tested GCN and GAT architectures on each
- Integrated LSTM modules with GNNs to capture both temporal and relational features
- Benchmarked standard GNNs, LSTMs, and temporal GNNs, showing the latter outperformed the others

**Semantic Video Reconstruction with Diffusion Models (Thesis), 2023 - 2024** Report | GitHub

- Proposed a high-compression video transmission framework, sending segmentation maps instead of full frames
- Designed a diffusion model that reconstructs each frame using its segmentation map and a sequence of previously reconstructed frames (to-be-deployed in receiver side)
- Used Stable Diffusion to condition on prior frames and ControlNet on the current frame's segmentation map
- Enabled efficient reconstruction in bandwidth-constrained scenarios such as drone video streaming

**Music synthesizer, 2023** GitHub

- Implemented, in a team of 4, embedded software for a music synthesizer using an STM32
- Features included different octaves, volume levels, waveforms, simultaneous key presses, a joystick used as pitch bender and volume modifier, record and playback operations, and user-friendly, multi-page display

**Autonomous Pet Caretaker (iFeed), 2023** Marketing | GitHub

- Designed and 3D-printed an automated pet feeder with web-based scheduling, multi-pet profiles, and audio alerting pets; integrated computer vision to verify the correct pet before dispensing; included built-in live stream
- Built on Raspberry Pi Zero with MQTT communication to a custom webserver; full-stack system developed in Vue (front-end) and Python (back-end)

## Skills

**Programming Langs:** Python (incl. PyTorch, TensorFlow, Keras, SciKit-learn), C++/C, MATLAB, Go, SQL, Verilog, Bash

**Technologies:** verl, DeepSpeed, vLLM, Docker, Kubernetes, Linux, FPGA, Arduino

## Languages

- English (fluent), Greek (native), German (good), French (basic)
- F-1 Student Visa, Greek Citizenship, UK Settled Status