

KONSTANTINOS GEORGIU

Machine Learning Engineer

+1(865) 978-9244

https://gkos.dev

gkos.mldev@gmail.com

linkedin.com/in/gkos

United States

github.com/drkostas

SUMMARY

- Ph.D. Machine Learning Engineer with 9+ years building production ML systems for high-impact applications in Large Language Models (LLMs), Computer Vision, and data engineering.
- Delivered ML products with measurable business outcomes, backed by peer-reviewed research and [open-source contributions](#) on GitHub.

EXPERIENCE

- May - Aug 2025** | **Applied Scientist - L5**
Amazon - US
Skills: LLMs · Code Generation · Agentic Frameworks · MCP · Research
- Architected a multi-agent LLM framework for natural language to executable code, achieving 82% functional success on smart home automations and 87.6% on MBPP.
 - Engineered an impossibility detector that flagged over 93% of unfeasible requests (4.5x over baseline), preventing hallucinated code.
 - Surfaced a critical misalignment between heuristic LLM scores (98%) and functional correctness (79% pass), exposing limits of current code assessment.
- 2024 - 2025** | **Co-Founder & ML Engineer**
XPensAI Ltd - United Kingdom
Skills: Computer Vision · Generative AI · SaaS Development · APIs · Cloud Infrastructure · AWS · Azure
- Launched an AI-powered SaaS platform used by over 30 small and medium-sized businesses, reducing manual expense entry by an estimated 65%.
 - Led the development and deployment of core AI algorithms for automated expense tracking, real-time analytics, and receipt processing, improving processing speed by 120% over the baseline implementation.
- 2021 - 2026** | **Machine Learning Researcher**
University of Tennessee - US
Skills: PyTorch · Deep Learning · Computer Vision · LLMs · Finetuning · Research · DDP · Foundation Models · Multi-modal Datasets
- Developed a multi-objective CLIP distillation framework using Soft MoE for per-patch loss weighting (85.3% accuracy, 52.8% mIoU on ImageNet, MEDiC arXiv 2026, ExPLoRe ECCV 2026).
 - Engineered Cross-Scale MAE, a multi-resolution masked autoencoder for remote sensing that bridges scale gaps between satellite imagery resolutions, achieving 5% accuracy improvement over state-of-the-art across 4 benchmarks (NeurIPS 2023), with follow-up fine-tuning strategies reducing training time by 32% (IEEE IGARSS 2024).
 - Authored a comparative analysis of LLM code generation security across 4 major models and 9 tasks ("Occasionally Secure," arXiv 2024, 14 citations), uncovering up to 22 vulnerabilities per model and a 40% security regression in Gemini when using security-focused prompts.
- 2019 - 2021** | **Data Engineer**
Performance Technologies S.A - Greece
Skills: TensorFlow · Data Engineering · Software Engineering · Machine Learning · GCP · Apache Spark · Docker · ETLs
- Led the rapid completion of a critical terabyte-scale data replication project for Greece's leading telecommunications provider, reducing replication time from days to minutes and ensuring real-time views for ETL and analytics.
 - Spearheaded the development of a machine learning model to predict order fulfillment times, which, following a comprehensive analysis of business operations and consultation with clients, resulted in a 34% reduction over previous baseline.
- 2018 - 2019** | **Machine Learning Researcher**
University of Patras - Greece
Skills: Python · Community Detection · Algorithm Design · Machine Learning · Apache Spark · Docker · Graph Models · Research
- Developed a distributed hybrid Girvan-Newman variant on Apache Spark, achieving an 84% runtime reduction on large social graphs without accuracy loss (Algorithms journal, 2019).
- 2017 - 2018** | **Software Engineer**
Global Voices Ltd - UK
Skills: Python · Software Engineering · Operating systems · SQL
- Shipped 20+ features for the company's translation CMS, building Python services and SQL-backed workflows used by 300+ translators across 150+ language pairs.

CONTRACT ROLES

- Jan - May 2026** | **AI Engineer (B2B Contract)**
FleetSmart.ai - Remote
Skills: LLMs · Prompt Engineering · FastAPI · Next.js · GCP · Terraform · PostgreSQL
- Built an AI decision engine for maritime vessel positioning, combining deterministic market filtering with multi-provider LLM reasoning (Gemini, OpenAI, Claude) to generate 5-10 ranked recommendations per vessel with structured output validation via Pydantic schemas.
 - Developed real-time market intelligence pipelines integrating Baltic Exchange rates, bunker futures (FIS Live Fuel Oil), port congestion predictions, and live news sentiment signals for decision context.

EDUCATION

- 2026** | **PhD in Data Science & Engineering**
University of Tennessee
- Fellowship Award recipient at UTK, with doctoral research in Self-Supervised Learning, LLM Security, and Trustworthy AI.

- Coursework included Machine Learning, Neural Networks, Data Mining, Pattern Recognition, and Computer Vision, among many others.

PUBLICATIONS

- ExPLoRe: Loss-Coupled MoE for Masked Image Modeling. - [Under review ECCV 2026](#)
- MEDiC: Multi-objective Exploration of Distillation from CLIP. - [arXiv 2026](#)
- Trustworthy AI for Early Dementia Detection: Robust Feature Masking and Clinical Interpretability. - [CHASE 2025](#)
- Adding a Teaching Assistant: Improving the Quality of Pseudo-labels for Semi-Supervised Object Detection. - [Tec. en Marcha 2025](#)
- Advancing Multi-scale Remote Sensing Analysis through Self-Supervised Learning Fine-tuning Strategies. - [IEEE IGARSS 2024](#)
- Koopman-based Transition Detection in Satellite Imagery. - [IEEE IGARSS 2024](#)
- Occasionally Secure: A Comparative Analysis of Code Generation Assistants. - [arXiv 2024](#)
- Cross-Scale MAE: A Tale of Multi-scale Exploitation in Remote Sensing. - [NeurIPS 2023](#)
- Semantic Segmentation in Aerial Imagery using Multi-level Contrastive Learning with Local Consistency. - [WACV 2023](#)
- A Distributed Hybrid Community Detection Methodology for Social Networks. - [Algorithms 2019](#)