Vasileios Lioutas

Applied Machine Learning Research Scientist · Vancouver, Canada
contact@vlioutas.com · vlioutas.com · LinkedIn · GitHub · Google Scholar +1 (613) 581-8083

Summary

Research Scientist in AI with 8+ years of experience translating cutting-edge machine learning research into production systems used by industry leaders. Published at ICML, ICLR, NeurIPS, and ACL with expertise in deep generative models, probabilistic inference and reinforcement learning. Proven ability to bridge rigorous research and real-world deployment by shipping models that power commercial products at scale. Driven by a passion for solving hard technical problems at the intersection of research and real-world impact, I'm excited to push the boundaries of what AI systems can do next.

Core Competencies

Machine Learning (ML) · Artificial Intelligence (AI) · Generative AI · Research & Innovation · Quantitative Analysis
Software Engineering · Distributed Computing · Scalable ML Systems · End-to-End Solutions Development
Technical Leadership · Strategic Insight · Global Interdisciplinary Collaboration · Mentoring

Experience

Research Scientist · Inverted AI · Vancouver, BC *Sep 2020 – Present*

- Led the design, development, and maintenance of machine learning models for realistic driving simulation powering the full commercial product catalog (DRIVE, INITIALIZE, SCENARIO), adopted by AV industry leaders including NVIDIA and the Toyota Research Institute.
- Designed and trained early-stage efficient driving planner models for on-board testing. Served as a key contributor to ITRA Verify, a composable tool suite for validating AV planners against safety and behavioral requirements.
- Improved the performance of deployed simulation products across diverse client needs through proprietary and published innovations, including inference-time safety planning methods that enforce hard constraints (e.g., no infractions) and strengthen safety compliance.
- Trained and deployed a broad range of models for internal use, including diffusion models, VAEs, transformers, and RL-based policies.
- Supervised co-op students and interns, and strengthened the internal codebase through code reviews, testing, and engineering best practices.
- Co-authored papers published at major venues and contributed to a wide range of additional internal research and development efforts.

ML Research Associate · Huawei Noah's Ark Lab · Montreal, QC *May 2019 – Aug 2020*

- Designed compressed neural embedding models achieving 4-7× size reduction, meeting production requirements for edge deployment.
- Developed methods for compressing neural machine translation systems and large language models, contributing to Huawei's production NLP stack.
- Implemented many research methods from the literature.
- Published 3 papers at ACL and EMNLP; filed 1 patent (granted US20230222353A1, 2023).

Machine Learning Engineer · MEDIAFORCE.ca · Ottawa, ON *Feb 2018 – Aug 2018*

- Built two real-time recommendation systems (k-NN and deep learning), owning the full pipeline from data collection and model training to API serving.

Technical Skills

ML: Diffusion Models/Flow Matching · Variational Autoencoders · Autoregressive Models · Planning & Model Predictive Control · Reinforcement Learning · Transformers · World Models · Knowledge Distillation · LLM Fine-tuning (SFT · RLHF)
Engineering: PyTorch · CUDA · Distributed Training · HuggingFace · AWS/GCP · Docker · Weights & Biases
Languages: Python (Expert) · C++ · SQL

Education

Ph.D., Computer Science · University of British Columbia · Vancouver, BC *2020 – 2025*
Thesis: *Towards Realistic Controllable Driving Simulators* · Advisor: Dr. Frank Wood

M.Sc., Computer Science · Carleton University · Ottawa, ON *2018 – 2020*
Thesis: *Sequence Modeling with Linear Complexity* · GPA 12.0/12.0

B.Sc., Computer Science · Aristotle University of Thessaloniki · Greece *2011 – 2016*

Key Projects

TorchDriveSim · Open-Source 2D Driving Simulation (43 stars on GitHub) *2023*

- Fully differentiable execution supporting rendering ego-rotated and ego-centric bird's-eye views, various kinematic models, agent types, traffic controls, and goals.
- Adopted by the research community working in autonomous driving.
- Utilized by TorchDriveEnv, a lightweight reinforcement learning benchmark for autonomous driving.

TaLK Convolutions · Open-Source Novel Neural Architecture (29 stars on GitHub) *2020*

- Developed a sequence modeling architecture with linear time complexity in the number of tokens.
- Implemented custom CUDA kernels for PyTorch to support parallelization.
- Produced state-of-the-art results in various natural language tasks.

Selected Publications (full list: Google Scholar · 320+ citations · h-index 9)

- **V. Lioutas**, A. Scibior, M. Niedoba, B. Zwartsenberg, and F. Wood, "Control-ITRA: Controlling the Behavior of a Driving Model," *arXiv: 2501.12408*, 2025.
- J. Lavington, K. Zhang, **V. Lioutas**, M. Niedoba, Y. Liu, D. Green, S. Naderiparizi, X. Liang, S. Dabiri, A. Scibior, B. Zwartsenberg, and F. Wood, "TorchDriveEnv: A Reinforcement Learning Benchmark for Autonomous Driving with Reactive, Realistic, and Diverse Non-Playable Characters," *arXiv:2405.04491*, 2024.
- **V. Lioutas**, J. Lavington, J. Sefas, M. Niedoba, Y. Liu, B. Zwartsenberg, S. Dabiri, F. Wood, and A. Scibior, "Critic Sequential Monte Carlo," *ICLR*, 2023.
- M. Niedoba, J. Lavington, Y. Liu, **V. Lioutas**, J. Sefas, X. Liang, D. Green, S. Dabiri, B. Zwartsenberg, A. Scibior, and F. Wood, "A Diffusion-Model of Joint Interactive Navigation," *NeurIPS*, 2023.
- **V. Lioutas**, A. Scibior, and F. Wood, "TITRATED: Learned Human Driving Behavior without Infractions via Amortized Inference," *TMLR* and *ML4AD workshop at NeurIPS*, 2022.
- A. Scibior*, **V. Lioutas***, D. Reda, P. Bateni, and F. Wood, "Imagining The Road Ahead: Multi-Agent Trajectory Prediction via Differentiable Simulation," *ITSC* and *ADP3 workshop at CVPR (Best Paper Award)*, 2021.
- A. Rashid*, **V. Lioutas***, and M. Rezagholizadeh, "MATE-KD: Masked Adversarial Text, a companion to Knowledge Distillation," *ACL*, 2021.
- A. Rashid, **V. Lioutas**, M. Rezagholizadeh, and A. Ghaddar, "Towards Zero-Shot Knowledge Distillation for Natural Language Processing," *EMNLP (Oral Presentation)*, 2021.
- **V. Lioutas** and Y. Guo, "Time-aware Large Kernel Convolutions," *ICML*, 2020.
- **V. Lioutas**, N. Passalis, and A. Tefas, "Explicit ensemble attention learning for improving visual question answering," *Pattern Recognition Letters*, 2018.

Selected Patents

- **V. Lioutas**, A. Rashid, and M. Rezagholizadeh, "Method and System for Training a Neural Network Model Using Adversarial Learning and Knowledge Distillation," *US Patent App. US20230222353A1*, March 2023.

Awards & Recognition

- NSERC CGS-D Doctoral Scholarship ($35,000/yr) 2021–2024
- Best Paper Award, Autonomous Driving Workshop @ CVPR 2021
- Graduate Support Initiative (GSI) Award, UBC 2020–2024
- Senate Medal for Outstanding Academic Achievement, Carleton University 2020