

Antonios Tragoudaras

🏠 antragoudaras.github.io | 📧 antragoudaras | 📄 [Antonios-Tony-Tragoudaras](#) | 🐦 [TragosML](#)

Education

Inria Paris – Astra

Paris, France

PHD IN ARTIFICIAL INTELLIGENCE | *Research focus: Closing the physics gap in Vision Foundation Models* | FUNDED BY

Nov. 2025 – Present

PR[AI]RIE-PSAI

Physics-Grounded Vision Foundation Models Architecting physics-grounded Vision Foundation Models (VFMs) that transition beyond visual mimicry.

University of Amsterdam (UvA)

Amsterdam, The Netherlands

MSC IN ARTIFICIAL INTELLIGENCE

Aug. 2023 – Oct. 2025

Research Focus: Physics-Grounded Learning, Embodied AI, and World Foundation Models.

Academic Standing: GPA 8.35/10 (Cum Laude).

ELLIS MSc Honors Award | Awarded for excellence in AI research under joint supervision across two ELLIS units.

King Abdullah University of Science and Technology (KAUST)

Jeddah, Saudi Arabia

VISITING STUDENT & RESEARCH ASSISTANT

Aug. 2022 – Aug. 2023

Neural Architecture Search, Efficient Deep Learning & TinyML, Voice Activity Detection, and Visual Perception for Autonomous Driving Applications

University of Thessaly

Volos, Greece

BENG IN ELECTRICAL & COMPUTER ENGINEERING (5-YEARS STUDIES; 300ECTS)

Sep. 2016 – Nov. 2021

Thesis Grade: 10/10 (Excellent), GPA: Upper Second-Class Honors.

Undergraduate Research Project: Design Space Exploration of a Sparse MobileNetV2 Using High-Level Synthesis and Sparse Matrix Techniques on FPGAs.

Research Experience

Physics-Grounded Vision Foundation Models (PhD Research)

Inria Paris (Astra)

ADVANCING EXPLICIT PHYSICAL REASONING IN LARGE-SCALE VISUAL MODELS

Nov. 2025 – Present

Advisor: Raoul de Charette, Supervisors: Tuan-Hung Vu, Andrei Bursuc

TL;DR: Moving beyond ‘emergent’ physics via massive data scaling to build *explicitly* causal Vision Foundation Models with an intrinsic understanding of physical dynamics.

Physics-Informed Representation Alignment (PIRA)

University of Amsterdam (UvA) / ISTA

INSTILLING TARGETED, PHYSICAL KNOWLEDGE INTO PRE-TRAINED VIDEO DIFFUSION MODELS

Feb. 2025 – Sep. 2025

[[project page](#)] [[paper](#)] | Supervisors: A. Zadaianchuk, D. Cherniavskii, E. Gavves. Visiting Student @ ISTA in F. Locatello’s group.

TL;DR: Designed **PIRA**, an elegant representation alignment framework that distills physics-proxy signals (like motion consistency) into Video Generative Models, bridging the gap between visual mimicry and dynamics understanding.

- **Motivation:** The key idea is to use the VDM’s latent VAE encoder to process proxy physics-centric signals (like motion, depth, and segmentation masks), for crafting representations that the model can easily learn from. Such structural signals enforce structural and geometrical consistency into the DiT following contemporary advances in representation alignment.
- **Scientific Discovery:** Evaluated on a custom 10,000 synthetic samples (created with Google’s Kubric simulator), demonstrating that differing proxies can teach distinct inductive biases (e.g., depth teaches governing ODEs). The method outperformed explicit-control baselines and standard supervised fine-tuning of the VDM backbone, in OOD generalization and physical adherence.

Physical Reasoning of Video Generative Models (MORPHEUS)

University of Amsterdam (UvA)

QUANTIFYING PHYSICS-REASONING AS AN EMERGENT PROPERTY IN VIDEO GENERATIVE MODELING.

Jul. 2024 – Feb. 2025

[[project page](#)] [[paper](#)]

TL;DR: Introduced **MORPHEUS**, a pioneering benchmark evaluating whether state-of-the-art video models learn underlying Newtonian dynamics or merely interpolate pixel patterns.

- **Technical Pipeline:** Transform pixels to suitable intermediate representations of dynamics, corresponding to videos generated by leading flow-matching/diffusion models conditioned on controlled and clean physics instantiations. Introduced a diverse dataset featuring *Newtonian Dynamics* across diverse settings, including chaotic systems (pendulums), and object-interactions (collisions).
- **Scientific Impact:** Moving beyond flawed trajectory-matching and subjective VQA with strict *physics-informed metrics* (fitting Physics-Informed Neural Networks and measuring physical invariant conservation). Exposed a critical lack of causal physical understanding in modern VGMs.

Neural Architecture Search, Meta-Heuristic Optimization, Transformers for Brain Signal Decoding

King Abdullah University of Science and Technology (KAUST)

GRADUATE STUDENT IN APPLIED AI, SUPERVISED BY POSTDOC FELLOW.

Sep. 2022 – Feb. 2023

Projects/Publications:

- Enhancing DNN models for EEG/ECoG BCI with a Novel Data-Driven Offline Optimization Method.
- Data-Driven Offline Optimization of Deep CNN Models for EEG and ECoG Decoding.

AutoML, Efficient Deep Learning Techniques, and Voice Activity Detection

VSPR – KAUST

VISITING STUDENT RESEARCH INTERN, MENTORED BY POSTDOC FELLOW CHARALAMPOS ANTONIADIS

Feb. 2022 – Jul. 2022

Projects/Publications:

- TinyML for EEG Decoding on Microcontrollers.
- Audio-visual Speaker Diarization: Improved Voice Activity Detection with CNN based Feature Extraction.

Publications

Full list is available at [Google Scholar](#)

- **A. Tragoudaras***, D. Cherniavskii*, C. Zhang*, A. Vozikis, T. Nijdam, D. W. E. Prinzhorn, M. Bodraccka, N. Sebe, **A. Zadaianchuk†**, **E. Gavves†**. “Morpheus: Benchmarking Physical Reasoning of Video Generative Models with Real Physical Experiments” **Accepted in 43rd International Conference on Machine Learning (ICML), 2026.** [[project page](#)] [[paper](#)]
- **A. Tragoudaras**, T. Aslanidis, E. G. Lionis, M. Orozco González, P. Eustratiadis. “Information Leakage of Sentence Embeddings via Generative Embedding Inversion Attacks” in *48th International ACM Conference on Research and Development in Information Retrieval (SIGIR), 2025.* [[paper](#)]
- **A. Tragoudaras**, C. Antoniadis, Y. Massoud. “Enhancing DNN models for EEG/ECoG BCI with a Novel Data-Driven Offline Optimization Method” in *IEEE Access, 2023.* [[paper](#)]
- **A. Tragoudaras**, C. Antoniadis, Y. Massoud. “TinyML for EEG Decoding on Microcontrollers” in *IEEE 56th International Symposium on Circuits and Systems (ISCAS), 2023.* [[paper](#)]
- **A. Tragoudaras**, P. Stoikos, K. Fanaras, A. Tziouvaras, G. Floros, G. Dimitriou, K. Kolomvatsos, G. Stamoulis. “Design Space Exploration of a Sparse MobileNetV2 Using High-Level Synthesis and Sparse Matrix Techniques on FPGAs” in *MDPI Sensors, 2022.* [[paper](#)]
- K. Fanaras, **A. Tragoudaras**, C. Antoniadis, Y. Massoud. “Audio-visual Speaker Diarization: Improved Voice Activity Detection with CNN based Feature Extraction” in *IEEE 65th International Midwest Symposium on Circuits and Systems (MWSCAS), 2022.* [[paper](#)]

Awards

Nov. 2025

– Present [PR\[AI\]RIE-PSAI](#): Awarded a grant from PR[AI]RIE-PSAI for funding my PhD.

Inria Paris (Astra)

[ELLIS MSc Honours](#): Awarded to excellent MSc students who have conducted their thesis

under joint supervision of the UvA ELLIS unit at an ELLIS partner institution outside of the Netherlands.

ELLIS

Teaching Experience

Graduate Teaching Assistant at UvA’s MSc AI Program

UvA, Amsterdam, Netherlands

ASSISTED IN TEACHING GRADUATE-LEVEL (FIRST-YEAR) COURSES BY MAKING SURE STUDENTS UNDERSTOOD THE MATERIAL, ANSWERING THEIR QUESTIONS, CREATING ASSIGNMENTS, GIVING FEEDBACK, AND GRADING EXAMS.

Aug. 2024 – Mar. 2025

Courses:

- Computer Vision 1 (MSc AI) (Aug. 2024 - Oct. 2024)
- Deep Learning 1 (MSc AI) (Oct. 2024 – Dec. 2024)
- Fairness, Accountability, Confidentiality & Transparency in AI (MSc AI) (Jan. 2025 - Feb. 2025)
- Information Retrieval 1 (MSc AI) (Feb. 2025 - Mar. 2025)

Reviewing Services

June. 2026 **NeurIPS 2026 Reviewer**

January. **CVPR 2026 Reviewer**
2026

April. 2026 **ECCV 2026 Reviewer**