Theodoros Kouzelis

PhD Candidate NTUA & Archimedes RU
Email: theodoros.kouzelis@athenarc.gr
Website: https://thkouz.github.io/

Scholar
GitHub
Twitter
LinkedIn

Education

National Technical University of Athens — Athens, Greece
PhD Candidate in Machine Learning — Sep 2024–Present

- Fellow at **Archimedes Research Unit**
- Topic: Deep Generative Models and Representation Learning
- Supervisors: Nikos Komodakis, Spyros Gidaris

National Technical University of Athens — Athens, Greece
MEng in Electrical and Computer Engineering — 2014–2021

Selected Publications

- Coevolving Representations in Joint Image-Feature Diffusion — **ECCV 2026**
  T. Kouzelis, S. Gidaris, N. Komodakis
- Boosting Generative Image Modeling via Joint Image-Feature Synthesis — **(Spotlight) NeurIPS 2025**
  T. Kouzelis, E. Karypidis, I. Kakogeorgiou, S. Gidaris, N. Komodakis
- EQ-VAE: Equivariance Regularized Latent Space for Improved Generative Image Modeling — **ICML 2025**
  T. Kouzelis, I. Kakogeorgiou, S. Gidaris, N. Komodakis

Employment

Runway — New York City, USA
Position: Research Intern supervised by Jonathan Granskog — Oct 2025 – Feb 2026

- **Project:** Improving video autoencoders by incorporating motion priors for more efficient diffusion model training.

Institute of Language and Speech Processing, Athena RC — Athens, Greece
Position: Researcher Assistant — Sep 2022 - Aug 2024

- **Project:** Unsupervised discovery of interpretable directions in the latent space of image diffusion models. Results published at BMVC 2024
- **Project:** Controllable text-to-music diffusion models. Results published at ICASSP 2024
- **Project:** Unsupervised Domain Adaptation for ASR models under low-resource settings. Results published at TASLP
- **Project:** Audio representation learning via language supervision. Results published at DCASE Workshop 2023

Institute of Language and Speech Processing, Athena RC — Athens, Greece
Position: Research Intern — Sep 2021 - Aug 2022

- **Project:** Weakly-supervised speech-to-text alignment of disfluent speech. Results published in a conference paper accepted at Interspeech 2023

Invited Talks

- **NeurIPS in Paris (Oral Presentation)** 2025
  Boosting Generative Image Modeling via Joint Image-Feature Synthesis
- **Valeo AI** 2025
  Self-Supervised Representation Learning for Better Diffusion Models.
- **RunwayML** 2025
  Boosting Generative Image Modeling via Joint Image-Feature Synthesis

Teaching Experience

- **Teaching Assistant** at Department of Informatics, NKUA Spring 2025, 2026
  Lectures on Variational Autoencoders and Diffusion Models for Computer Vision and Deep Learning course.

Reviewerships

ICASSP 2024-25; CVPR 2026 (Outstanding reviewer); ECCV 2026; NeurIPS 2026; ICLR 2027; TMLR

Technical Skills

- **Programming Languages**: Python, C/C++, Haskell
- **Deep learning**: PyTorch, Jax, Wandb, Hydra
- **Other**: LaTeX, TikZ, Bash, SLURM, AWS

Languages

Greek, English, French (Beginner)