

Dhia Naouali

dhia12naouali@gmail.com | [Homepage](#) | [GitHub](#) | [LinkedIn](#) | [Up-to-date CV](#)

Summary

Final-year SWE student, interested in **representation learning** and **vision**: modeling (SSL, diffusion, ...) and understanding emergent structure in learned representations, with broader interest in **world models** and RL.

Looking for a **research thesis internship**, starting Feb 2027

Publications & Posters

Cross-Time Directional Selection in Diffusion Sampling

[Dhia Naouali](#) | AI for Stochastic Dynamics Workshop, **NeurIPS 2026**

[paper](#)

VLCP: Vision Language Control Policy Closed-Loop Code Replanning for Robot Manipulation

[Dhia Naouali](#), Minghan Wu, Claudia Wong, Abhinav Puthran, Omar G. Younis | EMR Workshop, **ECCV 2026**

[paper](#)

EEML'26 Summer School: presented *Perceptually Controversial Latents* as a poster

[poster](#)

ICLR'26 Representational Alignment Workshop: challenge report accepted as a poster

[poster](#)

NeurIPS'25 EEG Foundation Challenge: ranked 13/1,183 (top 1.1%) on the combined tasks

Experience

Research Intern | Kempner Institute, Harvard University, Prof. Qianqian Wang's group

Remote, USA | **Oct 2026 – Present**

- Working on visual representation learning and latent geometry for diffusion models.

Research Intern | Algoverse

Remote, USA | **Jan 2026 – Jul 2026**

- Developed closed-loop manipulation planning that recovers from mid-episode failures, **10×** open-loop success.

ML Instructor | GoMyCode

Remote, MENA | **Jun 2025 – Dec 2025**

- Taught ML to university students and career-shifting professionals, maintained **90+ NPS** across **50+** sessions

ML Research Intern | LabTIM

Hybrid, Tunisia | **Apr 2025 – Jun 2025**

- Developed local post-hoc explainability for MGMT methylation prediction from pathological scans.

ML Engineer | Omdena

Remote, USA | **Mar 2025 – Jun 2025**

- Trained a plastic debris detection model on collected satellite imagery, now deployed in Italian waters.

ML Engineering Intern | WEDTECT

Remote, Tunisia | **Jan 2025 – Mar 2025**

- Compressed an on-device model with distillation, curriculum learning, and pruning for real-time inference.

Service, Mentoring & Open Source

Reviewer, AI for Stochastic Dynamics Workshop, NeurIPS'26

Mentor, [RLS Research Lab](#): advising undergraduates in latent diffusion, world models, and vision models

Problem Setter & Mentor, GOAT Competition: designed an efficiency-aware monocular depth estimation challenge

Open Source, Kornia ([ADA tutorial](#)); PyTorch [Docathon'25](#) (Silver Medal; explicit imports, RST docs migration)

Projects

Visual Legibility from Noise

[GitHub](#), **Jul 2026**

- Probes **semantic emergence** along the diffusion trajectory of a latent diffusion model on SigLIP2 latents, decoded to both images and language, and compares denoising across the two modalities
- Guides generation with text prompts, training-free, on a class-conditional ImageNet-1k latent diffusion model

Perceptually Controversial Latents

[GitHub](#), **Mar 2026**

- Builds counterexamples to the **Platonic Representation Hypothesis**: optimizes the initial noise and prompts of a frozen **FLUX**, used as a semantic prior, to minimize **CKA** between frozen vision encoders (**DINO**, **JEPA**) on the generated images
- Finds a pattern shared across the optimized images, then trains a contrastive head to separate them from natural images and retrieves the nearest natural images in its space, which show the same pattern

Dinov3-JAX

[GitHub](#), Oct 2025

- Reimplements DINOv3 in JAX/Flax with **FSDP** for accelerator-agnostic distributed pretraining
- Cuts training iteration time **8–12%** over the original PyTorch implementation via JIT-compiled gradient accumulation

Tickling Vision Models

[GitHub](#), Aug 2025

- Extracts **50+** interpretable features from InceptionV1 via **sparse autoencoders** and activation maximization
- Shows causal control over predictions: directional interventions shift class logits, reversibly, via donor-image patching and targeted ablation

CUDA-Accelerated CNN from Scratch (without PyTorch / JAX)

[GitHub](#), Oct 2024

- Implements CNN training from scratch with custom CUDA kernels, profiling-driven to **18×** over the CPU baseline, with **Grad-CAM** to inspect its predictions
- Engineers for **numerical stability** (log-sum-exp, overflow and NaN guards) and **leak-free** GPU memory (fixed, reused activation buffers across CuPy and kernel ops)

Blogs

- Modeling the World | *survey of world model architectures, latent dynamics, and test-time planning* [link](#)
- A Decade of Residuals: History & Effects on Modern ML | *traces residual connections from ResNets to modern models* [link](#)
- A Gentle Introduction to Distributed Training | *data, tensor, and pipeline parallelism with practical intuitions* [link](#)
- Transformers as a Flock of Tokens | *collective token dynamics as an alternative mental model for attention* [link](#)
- A Representation Space Interpretation of Neural Networks | *geometric view of how networks transform data across layers* [link](#)

Technical Skills

Core ML Competencies: visual representation learning, self-supervised learning, world models, diffusion models, RL

Tools & Technologies: Python, **PyTorch**, **JAX/Flax**, NumPy, Scikit-learn, OpenCV, Git, Docker, WandB, Hydra, tmux

Education

Engineering Degree, Software Engineering

Sep 2022 – Jul 2027

Higher Institute of Mathematics & Informatics of Monastir

- **Relevant Coursework:** ML, Distributed Systems, Graph Theory, Advanced Probability, Algebra & Calculus 3, DSA
- **Achievements:** Ranked **5th** in the pre-engineering cycle