

Dhia Garbaya

dhia.garbaya@eleves.enpc.fr

dhia680.github.io | github.com/dhia680

Profile

Currently focusing on *distillation, reasoning, and pre-training* of language models.

Interested in many question about learning (e.g., experience, imitation, online, self-improvement...).

I am seeking a **5-6 month research internship** (April 2026), ideally with a **PhD** follow-up.

An ideal topic could be centered around *RL, representation learning* and grounded in principled approaches.

Education

ENS Paris-Saclay, MVA: MSc (Math, Vision, Learning) 2025 – 2026

- **Coursework:** Optimal transport · Convex optimization · Reinforcement learning · RMT · Probabilistic graphical models · Deep learning theory · Geometric deep learning · Graphs in ML · Vision · Robotics ...

Ecole Nationale des Ponts et Chaussées (IP Paris), MEng (Applied Math and ML) 2022 – 2026

- **Coursework:** Machine learning (JAX) · Deep learning (torch) · Optimization · Operations research (Julia) · Advanced algorithmics (C+++) · Vision · Stochastic processes · Statistical physics · PDEs · Game theory...

Esprit Prépa, Tunis, Prépa (Math, Physics) 2020 – 2022

- **Coursework:** Algebra/ Topology/ Analysis/ Probability/ Theoretical Physics/ Algorithms...
- **Project:** Optical and thermodynamical optimization of solar cells.

Experience

Research Intern, EPFL – Switzerland Feb 2025 - Now

- **Teams:** **MLO lab & Swiss AI Initiative**
- **Supervisor:** Prof. Martin Jaggi
- **Focus:** efficient pre-training (1k+4k GPUs), knowledge distillation, model architecture.
- **Recently:** over-thinking · exploration in RL (context: reasoning LMs).

Research Intern, TI Institute – UAE Jul 2024 – Jan 2025

- **Teams:** **AI theory & Falcon**
- **Supervisor:** Dr. MEA Seddik
- **Focus:** knowledge distillation, scalable optimization algorithms (1st and 2nd order)
- Large-scale parametrization (MuP), param-free learning, pre-training (5k GPUs)

Publications

FOG Architectures* : Toward Pure FP8 LLM Training at Scale · **NeurIPS'25** Paper v2
EPFL & ETH Zurich - 2025

Apertus* : Democratizing Open, Compliant, and Multilingual LLMs Technical Report
Swiss AI team - 2025

Falcon3 family of Open Models* : Showcasing transfer learning efficiency Official blogpost
Falcon LLM team - 2024 (undisclosed report)

* first-authorship

Coding skills

Languages: Python, C++, Julia, R.

Tools, frameworks: Pytorch, JAX, Megatron-LM, NeMo, HF libraries, GIT, AWS, Slurm.

Academic Projects

Solving PDEs with PINNs, ENPC+Airbus

github/dhia680/pinns-24

- Studied Physics-Informed-Neural-Networks for solving Helmholtz PDEs.

Deep RL for optimizing traffic using autonomous vehicles (exploratory)

github/.../HighwayEnv

- Inspired from CIRCLES project (Berkley, Rutgers, Ecole des Ponts..).
- Studied phantom traffic jam, drivers behaviour. Used existing codebase to train a policy with DQN, PPO.

RNN-based NMT model

github/dhia680/NMT

- Trained an LSTM-based machine translation tiny model and integrated it in a web interface

Operations Research

2023-2024

- Optimizing an offshore wind electrical network for RTE, FR
- Optimizing a car manufacturing chain for Renault, FR

Denoising score matching / RL for wheeled-bipedal push recovery

In progress

Distinctions, Online certificates

- **French government excellence scholarship:** Among 7 national holders of this scholarship, 2022-2026
- **Valedictorian** in both high school and Prépa
- **Completed:** NLP with Python (Udemy) · Supervised Learning, RL (deeplearning.ai)
- **In progress:** Advanced RL in python, DQNs (Udemy) · Diffusion Models (Nvidia)

Languages

- **Arabic** (Native) · **English** (Fluent) · **French:** (Fluent) · **German:** (Intermediate)

Volunteering

- **Reviewer** for MenaML'25.
- **Project Manager** in the Junior Enterprise of Ecole des Ponts (2023 - 2024).