

Dhia Garbaya

dhia.garbaya@eleves.enpc.fr

dhia680.github.io | github.com/dhia680

Profile

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

Open to research and PhD opportunities in 2026.

Education

École Normale Supérieure (P-S), MVA: MSc [Math, Vision, Learning] 2025 – 2026

- **Coursework:** Convex optimization · Reinforcement learning · Probabilistic graphical models · Bayesian ML · Geometric deep learning · Graphs in ML · LLMs for code and proof · Online algorithms · Robotics...

École des Ponts et Chaussées (IP Paris), MEng [Applied Math, 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

AI scientist Internship || Mistral AI - FR Apr'26 - Oct'26 (upcoming)

- **Team:** Reasoning/RL
- **Supervisor:** Albert Jiang (Tutor: Michal Valko)

Research Visit/Internship || EPFL - Switzerland Feb'25 - Jan'26

- **Teams:** MLO lab & Swiss AI Initiative
- **Supervisor:** Prof. Martin Jaggi
- **Focus:** efficient LLM pre-training (4k GPUs), knowledge distillation, model architecture.
- **Recently:** over-thinking + exploration-exploitation in GRPO.

Research Internship || TII - UAE Jul'24 - Jan'25

- **Teams:** AI theory & Falcon LLM
- **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

Terminator : Learning Optimal Early Stopping in CoT Reasoning · **ICLR'26** Openreview
UT Austin & EPFL - 2026 Soon on arXiv

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

Apertus* : Democratizing Open, Compliant, and Multilingual LLMs arXiv
Swiss AI team - 2025 under review

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

* first-authorship

Coding skills

Languages: Python, C++, Julia, R

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

Academic Projects

Solving PDEs with PINNs, w/ 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 machine translation 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

RL for push recovery of a wheeled robot, w/ Willow team (INRIA) 2025

- PPO, Reward shaping, curriculum, inductive bias...

Molecule captioning 2026

Distinctions, Online certificates

- **French government excellence scholarship:** Among 7 national holders of this scholarship, 2022-2026
- **Valedictorian** in both high school and Prépa
- **Completed certificates:** NLP with Python (Udemy) · Supervised Learning, RL (deeplearning.ai)
- **Research grant:** part of a large project on model compression, funded by the Swiss AI Initiative and led by D. Alistarh, M. Jaggi, T. Hoefler

Languages

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

Volunteering

- **Reviewer** for MenaML'25, ICLR'26.
- **Project Manager** in the Junior Enterprise of Ecole des Ponts (2023 - 2024).