Dhia Garbaya

FRA - CHE | dhia.garbaya@eleves.enpc.fr | in/dhia-garbaya

dhia680.github.io | github.com/dhia680

Profile

Currently focusing on transfer learning, pre-training language models, optimization, and reasoning.
With a broad interest in Deep Learning, particularly in RL.

Education

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

- **Coursework:** Optimal transport · Reinforcement learning · Theory of deep learning · Convex optimization · Stopping times and online algorithms · Random Matrix Theory · Robotics...

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

- **Coursework:** Machine Learning (JAX) · Deep Learning (torch) · Convex Optimization · Operations Research (Julia) · Advanced Algorithmics (C++) · Stochastic Processes · Statistical Physics · Stats in high dimensions · PDEs · Vision · Game theory...

Esprit Prépa, Tunis, Prepa for Grandes Ecoles Sept 2020 – July 2022

- **Coursework:** General, Linear Algebra/ Topology/ Analysis/ Theoretical Physics/ Eng sciences
- **Project:** Optical and thermodynamical optimization of solar cells.
- **Grade:** Valedictorian

Experience

Visiting Student | Research Intern, EPFL, Switzerland Feb 2025 - Now

- **Lab**: MLO
- **Supervisor**: Prof. Martin Jaggi
- Core member of **Swiss AI Initiative** LLM team (training on ALPS cluster).
- Focus: pre-training, transfer learning, low-precision training, model architecture.
- Member of Apertus LLMs core team.

Research Intern, Technology Innovation Institute – Abu Dhabi Jul 2024 – Jan 2025

- Focus: Knowledge distillation, Optimization algorithms at scale (1st and 2nd order)
- Knowledge distillation, Large-scale transformer parametrization, Param-free learning
- Core contributor to **Falcon3** family of models.

Club Project Manager, Junior Entreprise – Ile-de-France 2023 – 2024

- Responsible for technical and AI related studies

Publications

FOG Architectures: Towards Pure FP8 LLM Training at Scale · **NeurIPS'25** Preprint - V1

Apertus: Democratizing Open, Compliant, and Multilingual LLMs Technical Report

Falcon3 family of Open Models: Showcasing transfer learning efficiency Undisclosed Report
Blogpost

Coding skills

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

Tools, frameworks: GIT, AWS, cluster computing, linux, Pytorch, JAX, Megatron-LM, NeMo, HF libraries

Academic Projects

Solving PDEs with PINNs, ENPC+Airbus github/dhia680/pinns-24

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

RL for optimizing traffic using autonomous vehicles 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
- Tools: Pytorch, Sumo simulator.

RNN-based NMT model github/dhia680/NMT

- Trained an LSTM-based machine translation toy model (20M) and integrated it in a web interface
- Tools: Tensorflow, Pandas.

Operations Research 2023-2024

- Optimizing an offshore wind electrical network for RTE, FR
- Optimizing a car manifacturing chain for Renault, FR
- Tools: Python, C++, Julia, Gurobi (MIP), S.Annealing

Distinctions, Online certificates

- **French government excellence scolarship**: Among 7 national holders of this scolarship, 2022-2026.
- **NLP** with Python, Udemy.
- **Supervised Learning, RL**, deeplearning.ai.
- **In progress**:
  - Accelerated computing with cuda (python and C++), Nvidia.
  - Advanced RL in python, DQNs, Udemy.
  - Diffusion Models, Nvidia.

Languages

- **Arabic**: Native
- **French**: Fluent
- **English** Fluent
- **German**: Intermediate