AFIF BOUDAOUD

afifboudaoud@gmail.com | LinkedIn | GitHub

RESEARCH EXPERIENCE

ETH Zurich
March 2024 – Present
PhD Student
Zurich, Switzerland
Optimizing High-Performance Automatic Differentiation for Machine Learning and Scientific Computing

- Built ADELIA: the first AD-enabled INLA implementation, replacing finite-difference (FD) gradients with a structure-exploiting multi-GPU backward pass. Achieved 4.2–$7.9\times$ per-gradient speedups and 5–$8\times$ energy savings over DALIA (SOTA), and enabled reliable convergence on large models where FD stalls.
- Built DaCe AD: an efficient automatic differentiation framework that is, on average, over $92\times$ faster than JAX on gradient computation for scientific computing patterns from the NPBench suite.
- Implemented code optimizations to accelerate the ICON weather model, achieving a speedup of up to $1.51\times$ vs the optimized OpenACC version. This work won the 2025 ACM Gordon Bell Prize for Climate Modelling.

New York University Abu Dhabi
October 2022 – February 2024
Research Assistant
Abu Dhabi, UAE
LOOPer: A Learned Automatic Code Optimizer For Polyhedral Compilers

- Improved the existing deep learning cost model accuracy by 5% through feature engineering.
- Expanded the space of possible optimizations through a mathematical representation of code transformations.
- Generated LOOPerSet: a 28 million datapoints open-source dataset for program speed-up prediction.
- Accelerated data generation speed by $5\times$ through pruning and other data engineering techniques.
- Obtained a median speedup of 1.4× vs. Pluto, a SOTA polyhedral auto-scheduler, on the Polybench suite.

Mila – Quebec Artificial Intelligence Institute
June 2020 – September 2020
Research Intern
Montréal, Canada

- Implemented a memory and computation-time efficient alternative to the convolutional layer used in CNNs.
- Adapted the existing algorithm for inference-on-edge.

PUBLICATIONS

- **ADELIA:** Automatic Differentiation for Efficient Laplace Inference Approximations (under review) at: SC 2026.
- **DaCe AD:** Unifying High-Performance Automatic Differentiation for Machine Learning and Scientific Computing in: CLUSTER 2025.
- **LOOPer:** A Learned Automatic Code Optimizer For Polyhedral Compilers in: PACT 2025.
- **PerfDojo:** Automated ML Library Generation for Heterogeneous Architectures in: SC 2025.
- **LOOPerSet:** A Large-Scale Dataset for Data-Driven Polyhedral Compiler Optimization.

EDUCATION

École Nationale Supérieure d'Informatique Alger (ESI)
Sep 2017 – Jul 2022
Master's Degree in Computer Science
Algiers, Algeria