Houssam Eddine Boukhalfa

LLM Reasoning · Deep Learning · Reinforcement Learning

Master's Student, Université Paris-Saclay | b.houssameddine19@gmail.com | Google Scholar | LinkedIn | GitHub | Website

Education

Université Paris-Saclay, co-accredited with Institut Polytechnique de Paris, France Sept. 2026 – Sept. 2027
M2 Data Science (TRIED), expected September 2027

National Higher School of Artificial Intelligence (ENSIA), Algeria Sept. 2021 – June 2026
Engineering Degree and Master 2 in Artificial Intelligence and Data Science (five-year program) | Final year project: 17.5/20
Coursework: Probability & Statistics, Optimization, Machine Learning, Deep Learning, Reinforcement Learning, NLP.

Publications

- O. F. Zouak, **H. E. Boukhalfa**, S. Lakehal, S. Katiyar, S. Nefti-Meziani. *Counterexample Generation via Per-Theorem Symbolic Verifiers: When Imitation Hurts and Reinforcement Repairs.* **EMNLP 2026 (Findings)**, Budapest. [pdf] [code]
- **H. E. Boukhalfa**, S. Mani, L. Said Chikh. *Crafting Culturally Aligned Narratives: Large Language Models for Arabic Children's Story Generation.* **Muslims in ML Workshop, NeurIPS 2025**, San Diego. [pdf] [demo]
  Fine-tuned Gemini 2.0 and Silma 9B on 714 stories; expert-rated cultural alignment 9.4/10 ($\overline{\text{ICC}} = 0.847$); open dataset.

Research and Engineering Experience

End-of-Studies Project – ENSIA & University of Birmingham (remote) July 2025 – June 2026

- With Prof. Samia Nefti-Meziani, Dr. Shiv Ashutosh Katiyar (Birmingham) and Dr. Soumaya Lakehal (ENSIA): counterexample-guided RL with verifiable rewards (**CE-RLVR**) for LLM reasoning, using symbolic verifiers as rewards.
- **Causal world models:** owned the full pipeline (PDDL data preparation, SFT, GRPO-based CE-RLVR training, evaluation) on Qwen3-4B; exact next-state accuracy **89.2% seen-domain / 96.3% out-of-domain** vs. 84.7% / 86.7% for SFT.
- **Mathematical reasoning (EMNLP 2026 Findings):** set up the training pipeline and trained Qwen3-4B with SFT and GRPO-based RLVR; contributed to data preparation. RLVR reversed an SFT-induced collapse in true-theorem recognition (base 0.27, SFT 0.00, **RLVR 0.66**) and outperformed every evaluated 7B open-weights math specialist. [paper] [code] [thesis] [defense]

Research Intern (remote) – LIX, École Polytechnique (IP Paris) Sept. 2025 – June 2026

- Built an end-to-end multimodal pipeline to detect information influence operations on TikTok (video extraction, transcription, translation, text and video-stream analysis). Supervision: Dr. Oana Goga and Gabriel Ben Zenou (Inria Paris-Saclay).
- Compared what video descriptions claim with what the videos actually show, and analysed view and engagement ratios across candidates in the 2024 Romanian presidential election; paper in preparation.

R&D Engineer, AI Team Lead – Nojoom.AI Sept. 2025 – May 2026

- Led the AI team that trained and evaluated the company's Arabic LLMs (pretraining, instruction tuning, GCP deployment).
- **AQL-42B**: MoE LLM depth-extended from Qwen3-30B-A3B ($48 \to 66$ layers, 4.5B active), Arabic tokenizer (28% fewer tokens), 275B-token CPT; Arabic avg. **79.17 vs. 50.48** (base) on the seven-benchmark Open Arabic LLM Leaderboard v2.
- **AQL-Lite**: 5-stage Arabic data pipeline ($885\text{B} \to 215\text{B}$ tokens) and structured PEFT CPT on Gemma 3 12B (RS-LoRA, selective unfreezing, layer-wise LRs), **recovering 94% of full-CPT gains in 52% of the training time**.

Earlier Experience 2024 – 2025

- **Teaching Assistant, Limitless Learning** (2025): practical sessions on deep learning, Transformers, LoRA, RAG.
- **AI/ML Intern, Cosider Agrico** (2024): RAG system for question answering over a proprietary document corpus.
- **AI/ML Intern, Think Touch Solution** (2024): accelerated LLM inference with MPI on an HPC architecture.

Selected Projects

Roman Mosaic Restoration with GANs – Stanford University (supervisor: Dr. Mourad Bouache) 2025

- Designed a GAN with ViT and contextual attention to automatically restore damaged Roman mosaics. [demo] [preprint]

Reinforcement Learning for Autonomous Racing (DQN) 2025

- Full DQN agent: environment modeling, sensor engineering, reward design, and policy robustness analysis. [demo] [code]

Quality Takeoffs – Blueprint Analysis with Computer Vision and AI Agents 2025 – Present

- Detects legend tables, labels fixtures, measures linear elements and generates takeoff reports from blueprints. [website]

Skills

- **AI:** Deep Learning, Reinforcement Learning, LLMs, NLP, Computer Vision, Multimodal AI, Neuro-Symbolic AI.
- **Frameworks:** PyTorch, DeepSpeed, Hugging Face (Transformers, TRL, PEFT), Unsloth, Scikit-learn, Google Cloud.
- **Programming:** Python, C/C++, CUDA | **Languages:** Arabic (native), English (C1), French (C1).

Activities

Co-founder of the ENSIA Tech Community Club; winner of three national AI hackathons; supervised the student teams organizing the Algerian AI Council (Dec. 2024) and the AI for Africa conference series (Dec. 2023).

References

Dr Mourad Bouache
Google ML, Systems & Cloud
Lecturer, Stanford University
Advisor, Nojoom.AI

Prof. Samia Nefti-Meziani OBE
Director, Birmingham Robotics Institute
University of Birmingham

Dr Soumaya Lakehal
Associate Professor
ENSIA, Algeria