

Naoufel AZIZI

Stage de fin d'études — AI / Machine Learning Research · 6 months from March 2027

Paris, France · azizi.naoufel2004@gmail.com · +33 7 71 88 10 89 · github.com/a-naoufel · naoufelazizi.me · linkedin.com/in/naoufel-azizi

Profile

M2 student in Machine Learning, NLP & Data Science at Sorbonne Université. Research experience in **medical image analysis**, with a focus on whether models generalise across acquisition sources rather than on in-distribution headline scores. Comfortable across the full pipeline — dataset construction and auditing, split design, training on HPC, per-stratum evaluation and scientific writing. Also ships production software (Flutter, PHP, Docker), which is unusual alongside the research side.

Education

Master's in Computer Science — MIND (Machine Learning, NLP & Data Science) <i>Sorbonne Université, Paris</i>	Sept 2025 – Jun 2027
Deep learning · computer vision · NLP · image processing · data science · optimisation	
Bachelor's in Computer Science <i>Université Paris-Est Créteil</i>	2024 – 2025
Classe Préparatoire aux Grandes Écoles — Computer Science <i>ESTIN, Béjaïa, Algeria</i>	2022 – 2024

Research Experience

Research Intern — Medical Image Analysis <i>LIP6 (CNRS), Sorbonne Université — team SYEL, Paris</i>	Jun 2026 – Jul 2026
<i>Classifying colorectal polyps from endoscopic images.</i> Supervised by Garance Lucas (PhD candidate) and Prof. Bertrand Granado.	
- Screened 23 public colonoscopy datasets for licence, modality and annotation scheme; established that only 8 carry polyp-class labels and assembled them into a 93,156-image, lesion-annotated pool. - Audited the pool and removed 16,432 files (15% of everything on disk) — two datasets turned out to be duplicate re-releases of sources already included, which would have placed the same physical lesion on both sides of every split and silently inflated every score in the study. - Designed lesion-disjoint, source-balanced 5-fold splits and a leave-one-dataset-out protocol; established lesions rather than images as the unit that sets statistical power (adenocarcinoma: 965 images but only 25 lesions, one patient supplying 65% of them). - Ran four controlled balancing experiments plus LODO — three frozen ImageNet backbones (ResNet-50, ViT-B/16, ConvNeXt-T) × two hyperparameter configurations × five folds — on a SLURM GPU cluster. Principal finding: three-class macro F1 reaches 0.739 in-distribution, but balanced accuracy above chance falls 81% (+0.159 → +0.031) once the test source is held out, and sessile recall collapses from 0.71 F1 to near zero. The models are substantially recognising which dataset an image came from, not the pathology. - Showed fourfold augmentation of minority classes is worth +0.0013 macro F1: the binding constraint is lesion count, not view count. Recommended per-source and LODO reporting as the default output of every run. - Full LaTeX report with a reproducibility appendix (manifest build, dataset generation, verification, sbatch arrays).	

Selected Projects

Knowledge-Graph Extraction from Historical Text with LLMs <i>Personal project — Studium Parisiense corpus</i>	2026
- Entity/relation extraction pipeline over prosopographical records using google/gemma-3-27b-it, with prompt design tuned for schema consistency. - Deployed and served quantised LLMs up to 70B parameters (Hermes-3-Llama-3.1-70B-GPTQ, Qwen2.5-32B-GPTQ) with vLLM on a university A100 HPC cluster, working within storage and network limits.	

- Benchmarked TF-IDF, fastText, BiLSTM+GloVe, SetFit, DistilBERT and DeBERTa under a single evaluation protocol; **91.3% accuracy with SetFit**, analysing the accuracy / complexity / data-efficiency trade-off.

- Embedded vision pipeline (ROI extraction → preprocessing → digit recognition) with a fine-tuned **TensorFlow Lite** model deployed on **ESP32-CAM**, under real memory, latency and illumination constraints.

Engineering Experience

- Built enterprise VPN infrastructure (pfSense, OpenVPN); developed and maintained features for an internal VoIP application on Linux; Docker containerisation and deployment tooling.

- Cross-platform applications in Flutter and React Native, end to end: requirements, implementation, REST integration, deployment and support.

Technical Skills

ML / DL	PyTorch, scikit-learn, Keras, TensorFlow Lite · transfer learning, fine-tuning, experimental design, cross-validation, per-stratum evaluation
Vision & NLP	OpenCV, segmentation, feature extraction, ViT, ConvNeXt, ResNet · Hugging Face Transformers, BERT, SetFit, vLLM, quantised LLM serving
Languages	Python, C, C++, Java, JavaScript, SQL, Bash · NumPy, Pandas
Infrastructure	Linux, Git, Docker, SSH, SLURM / HPC, PostgreSQL, MySQL, MongoDB, LaTeX
Also	Flutter, React Native, React, Node.js, REST APIs

Languages: French (C2) · English (C2) · Arabic (native) | **Interests:** medical AI, distribution shift and evaluation, vision-language models, edge computing, open source