Lotfi Abdelkrim Mecharbat

Machine Learning PhD Researcher • Vision Foundation Models, Meta-Learning, Efficient Adaptation, AutoML

Abu Dhabi, UAE • lotfi.mecharbat@mbzuai.ac.ae • +971 56 318 8546 • meclotfi.io • meclotfi • mecharbat-lotfi

Research Focus

- My research studies how pretrained vision foundation models organize knowledge internally, and how much of it can be reused for a new task with limited data and compute. I build methods that locate task-relevant structure inside pretrained weights (channels, layers, adapter placements) and test how well it transfers, using controlled comparisons across tasks, pretraining families, and model scales.
- First-author work at MICCAI (2025, 2026), IJCNN and ESWEEK, with two papers under review at ICLR and AAAI 2027; complemented by industrial deployment experience (ALSTOM) spanning INT8 quantization, TensorRT, and on-device profiling.

Core areas: Vision foundation models, PEFT, neural architecture search, structured channel pruning, meta-learning, medical image analysis, hardware-aware optimization.

Education

PhD, Machine Learning — Sep 2024 — Present
Mohamed bin Zayed University of AI (MBZUAI), Abu Dhabi, UAE
Advisors: Prof. Mohammad Yaqub, Prof. Martin Takáč

- Thesis area: Automatic specialization of vision foundation models into compact task-experts under compute budgets.
- Year-1 coursework: TinyML & LLMs (A+), Continuous Optimization (A+), Advanced ML (A), Causal ML (A).

Engineering Degree & M.Sc. in Computer Science — Sep 2017 — Sep 2022
Higher National School of Computer Science (ESI), Algiers, Algeria
High Honors

- M.Sc. thesis: *Neural Architecture Search of Efficient Vision Transformers for Object Recognition on Edge Devices*.

Experience

Doctoral Researcher — Sep 2024 — Present
Department of Machine Learning, MBZUAI, Abu Dhabi, UAE

- **Channel-Remix** *[Under Review, ICLR 2027]*: Explore whether the knowledge a task needs is already present in a vision foundation model's pretrained channels, so it only has to be selected and recombined rather than relearned. Retained channels absorb sparse, learned combinations of pruned ones; with only these coefficients fitted to task data, the student matches LoRA distillation with **3–4× less data** across 16 tasks and 7 pretraining families. Since all students share one pool of pretrained directions, they can be compared directly to study which knowledge is shared across tasks and which is task-specific.
- **CarveVFM** *[Under Review, AAAI 2027]*: Developed a budget-guided search that measures how much each layer of a pretrained vision foundation model matters for a given task and selects task-specific widths accordingly. From only 512 calibration images it produces a compact dense model in minutes; the resulting 14M-parameter model outperforms natively pretrained ResNet-50, EfficientNet-B4, and ConvNeXt-Pico across 12 tasks.
- **MetaVFM** *[MICCAI 2026]*: Trained a meta-learner across 90+ datasets to predict adapter placement, rank, and initialization for new medical imaging tasks, eliminating the need for per-dataset hyperparameter search. Across 17 evaluation datasets, the predicted configurations achieve the highest average accuracy among evaluated PEFT methods and converge **2–3× faster** than PEFT baselines under a fixed epoch budget.
- **MedNNS** *[MICCAI 2025]*: Developed a shared embedding space for medical datasets and supernet sub-networks, allowing a new dataset to retrieve a matching architecture and initialization without running a new NAS search.
- *Systems & Infrastructure:* Built distributed training pipelines on NVIDIA and AMD ROCm clusters; multi-GPU profiling; mixed-precision (FP16/BF16) scaling for large ViT architectures.

Research Assistant — Oct 2023 — Oct 2024
Center for Quantum and Topological Systems (CQTS), NYU Abu Dhabi, UAE

- **MoENAS** *[Under Review]*: Formulated hardware-aware NAS with sparse Mixture-of-Experts routing for edge devices, optimizing accuracy-latency Pareto frontiers while enforcing algorithmic fairness constraints.
- **CV-SPINNS** *[Preprint]*: Designed complex-valued physics-informed neural networks for quantum spin simulations, embedding unitary constraints directly into the network architecture to avoid numerical drift.

Research Engineer — Oct 2022 — Sep 2023
ALSTOM & LAMIH UMR CNRS 8201, Valenciennes, France

- **DFuseNAS** *[IJCNN 2024]*: Developed continuous neural architecture search driven by diffusion processes, accelerating search convergence over evolutionary and reinforcement-learning baselines.
- **MARVIN** *[SACI 2025]*: Built real-time railway anomaly detection combining self-supervised representations with synthetic defect rendering; deployed via ONNX Runtime and NVIDIA TensorRT on edge hardware.

Research Assistant Sep 2021 — Sep 2022
Embedded Systems Group, LAMIH Laboratory, Valenciennes, France

- **HyT-NAS** *[ESWEEK 2022]*: Hardware-aware search over hybrid CNN–Transformer backbones under sub-watt budgets.
- Performed on-device profiling and INT8 post-training quantization on NVIDIA Jetson (TX2, Nano) and Raspberry Pi.

Publications

Peer-reviewed Conferences

[C5] **L. A. Mecharbat**, I. Almakky, E. Abdelrahman, H. Lashen, A. Almsouti, M. Takáč, M. Yaqub. *MetaVFM: Meta-Learned Efficient Adaptation of Vision Foundation Models for Medical Imaging*. **MICCAI 2026**.

[C4] **L. A. Mecharbat**, S. Niar, H. Ouarnoughi. *MARVIN: Monitoring Anomalies via Railway Video Insight and Neural Networks*. **SACI 2025**.

[C3] **L. A. Mecharbat**, I. Almakky, M. Takáč, M. Yaqub. *MedNNS: Supernet-based Medical Task-Adaptive Neural Network Search*. **MICCAI 2025**.

[C2] **L. A. Mecharbat**, H. Benmeziane, H. Ouarnoughi, K. ElMaghraoui, S. Niar. *DFuseNAS: A Novel Approach to Neural Architecture Search Based on Diffusion Processes*. **IJCNN 2024**.

[C1] **L. A. Mecharbat**, H. Benmeziane, H. Ouarnoughi, S. Niar. *HyT-NAS: Hybrid Transformers Neural Architecture Search for Edge Devices*. **ESWEEK 2022**.

Under Review and Preprints

[P3] **L. A. Mecharbat**, A. Elsayad, I. Almakky, M. Takáč, M. Yaqub. *Channel-Remix: Efficient Task-Specific Knowledge Transfer via Pretrained Channel Remixing*. *Under review at ICLR 2027*.

[P2] **L. A. Mecharbat**, A. Elsayad, I. Almakky, M. Takáč, M. Yaqub. *CarveVFM: Carving Specialists Out of Generalists, Automatic Task-Specific Compression of Vision Foundation Models*. *Under review at AAAI 2027*.

[P1] **L. A. Mecharbat**, A. Marchisio, M. Shafique, M. Ghassemi, T. Elhanai. *MoENAS: Mixture-of-Experts NAS for Accurate, Fair, and Robust Edge DNNs*. *Under review.*

Preprints, project code, and demos: meclotfi.io/publications

Technical Background

Methods: Parameter-efficient fine-tuning (LoRA, adapters), neural architecture search (supernets, hardware-aware), structured pruning, channel remixing, knowledge distillation, meta-learning, multi-objective optimization.

Compression & Systems: Post-training quantization (INT8, FP4, AWQ), mixed-precision (FP16, BF16), weight folding, CUDA kernel development, Triton, NVIDIA TensorRT, ONNX Runtime, multi-GPU distributed training (NVIDIA & AMD ROCm).

Languages & Tools: Python, C/C++, Bash, SQL; PyTorch, Hugging Face, OpenCV; Docker, Slurm, Git, Weights & Biases.

Edge Platforms: NVIDIA Jetson (AGX Orin, TX2, Nano), Raspberry Pi.

Awards & Additional Information

- **2nd Place Worldwide, AutoML 2022 Global Competition:** Bayesian optimization for Transformer configurations.
- **High Honors,** M.Sc. in Computer Science & Software Engineering, ESI Algiers (Class of 2022).
- **Languages:** Arabic (native), English (fluent, academic), French (TCF C1).
- **Outside Research:** Competitive football & applied CV (custom player/ball tracking on amateur match footage), trekking, history.