

Amar Ali-bey

Montréal, Canada

amar.alibey@gmail.com · LinkedIn · GitHub · Google Scholar

EXPERIENCE

Machine Learning Applied Research Scientist · *CoactiveAI, Montreal, QC* Feb 2025 – Present
Work on search, ranking, and video understanding systems within the ML research team at a Series-B startup building multimodal search for enterprise media.

- **Algorithmic Innovation.** Invented a vector-arithmetic method for semantic search with negation, enabling zero-latency attribute exclusion on standard vector DBs without re-indexing. This unlocked a key primitive for agentic search.
- **Inference Optimization.** Cut post-retrieval inference cost by 100× by replacing an LLM-in-the-loop bias filter with a lightweight custom cross-encoder.
- **Entity Recognition Pipeline.** Built a multi-stage entity recognition pipeline spanning detection, alignment, embedding, and indexing; achieved 60× speedup over baseline and delivered a production-ready implementation.
- **Hybrid Search & Calibration.** Unified multimodal search over tens of millions of records by calibrating Sentence-Transformer and CLIP score distributions for reliable single-list fusion in production.
- **Video Understanding.** Replaced shot-level summarization with hierarchical scene-level VLM/LLM understanding, reducing compute by 10×, and built an LLM-as-a-judge benchmark to validate quality gains on movie data.

Machine Learning Consultant (*Part-time, during Ph.D.*) 2018 – 2023
Built reproducible data curation, training, and inference pipelines for object detection, semantic segmentation, and cross-modal retrieval. Delivered PoCs for industrial safety, agriculture, and autonomous navigation.

SELECTED FIRST-AUTHOR PUBLICATIONS & OPEN SOURCE

BoQ: A Place is Worth a Bag of Learnable Queries *CVPR 2024*
A Transformer-based architecture that learns a set of global queries to aggregate features for visual retrieval. Still SOTA in visual place recognition two years after publication. [90+ citations]

MixVPR: Feature Mixing for Visual Place Recognition *WACV 2023*
Lightweight feature-mixing global descriptor. Deployed in production by multiple companies for its extreme speed and efficiency. [340+ citations]

GSV-Cities: Toward Appropriate Supervised Visual Place Recognition *Neurocomputing 2022*
Large-scale training dataset for place recognition. 5,000+ downloads; de facto standard for training supervised visual localization models. [200+ citations]

Global Proxy-based Hard Mining for Visual Place Recognition *BMVC 2022*
Optimized contrastive model training using high-efficiency hard mining and tiny batch-sizes. [15+ citations]

Open Source Projects

- **OpenVPRLab:** End-to-end VPR training and benchmarking framework used widely by the community.
- **NanoCLIP:** Efficient text-to-image retrieval model (DINOv2 + Sentence-Transformers) via custom finetuning.

800+ total GitHub stars across all repositories, widely adopted in academic research and industry.

EDUCATION & ACADEMIC ACTIVITIES

Ph.D. in Computer Vision and Machine Learning · *Laval University* 2017 – 2024
Thesis: *Deep Representation Learning for Visual Place Recognition.* · TA/Lecturer for 5 ML & Robotics courses, Mentored M.Sc. students; Reviewer for CVPR, ECCV, NeurIPS, TPAMI, IROS, ICRA.

M.S. Program in Computer Science · *Laval University, Québec, Canada* 2015 – 2017
Transitioned to Ph.D.: Completed masters-level coursework and transferred directly into the doctoral program.

B.S. in Computer Engineering · *ESI, Algiers, Algeria* 2014
Valedictorian (1st / 200+). Ranked 4th in nationwide competition for PhD fellowship.

TECHNICAL DEPTH

Core ML: Model Training & Finetuning (DDP/LoRA), Contrastive Learning, Search & Retrieval, Video Understanding, LLM & VLM.

Frameworks & Tools: PyTorch, PyTorch Lightning, HuggingFace, W&B, ONNX/TensorRT, Quantization, vLLM.

Infrastructure: Vector DBs (FAISS, HNSW, BM25), Anyscale/Ray, AWS (EC2, S3), Docker, Linux.

Programming: Python (Expert), C/C++, SQL · **Spoken Languages:** English (Fluent), French (Native).