Elias Ramzi | AI Research Scientist

Paris, France

☐ +33 6 28 25 40 88 ✉ elias.ramzi@valeo.com ⌂ elias-ramzi

Researcher specialized in Deep Learning and Computer Vision.

Experiences

valeo.ai (Paris, France) - full-time *2024-*
AI research scientist
Research in deep learning and computer vision for autonomous vehicles.

Cnam & Coexya (Paris, France) - 3 year PhD *2021-2024*
PhD student
My PhD in deep learning is focused on computer vision and image retrieval:

- Published four papers in major international machine learning conferences (NeurIPS, ECCV, ICML), three of which as first author.

- Worked in computer vision for image retrieval, on topics such as ranking losses and hierarchical data from both mathematical and experimental aspects.

- Collaborated with an industrial researcher (Google research) to release the first hierarchical landmark image retrieval dataset ↗ as part of a submission to a journal (under-review, TPAMI). It has 1.4m images and three levels of hierarchies: 100k unique landmarks, 78 super-categories and 2 final labels.

- Used the public HPC cluster Jean Zay ↗ to train neural networks in a distributed fashion with up to 4 nodes totaling 16 GPUs.

- Collaborated with two PhD students at Cnam, leading to a TMLR paper on collaborative filtering using graph neural networks, an ICML'23 publication on out-of-distribution detection using energy-based models and an ECCV'24 on prompt learning for VLMs.

- Adapted HAPPIER (ECCV 2022) to trademark logo retrieval, HAPPIER is now used to train models that go in production in Coexya's Acsepto.

Cnam (Paris, France) - 5 month internship *2020*
Research intern
Deep learning for 3D medical image segmentation. I used a learned confidence measure (ConfidNet) to combine segmentation results from different 2D U-Nets.

Sancare (Paris, France) - 6 month internship *2019*
Data engineer and data scientist intern

Balto (Sydney, Australia) - 6 month internship *2018*
Data engineer intern

Education

PhD in Computer Science, Cnam & Coexya (France) *2021-2024*
I investigate deep learning approaches to image retrieval. Specifically, I work on designing appropriate ranking losses to train deep neural networks for image retrieval. I also contributed on out-of-distribution detection using energy-based models, worked on collaborative filtering recommendation using graph neural networks and on VLM adaptation with prompt learning.
Supervisor: Nicolas Thome ↗ (Sorbonne Université, Paris, France).
Advisors: Nicolas Audebert ↗ (IGN, Paris, France); Clément Rambour ↗ (Cnam, Paris, France).
Industrial advisor: Xavier Bitot ↗ (Coexya, Paris, France).

Engineering diploma, CentraleSupélec (France) *2016-2020*
Courses in signal processing (sound, image, speech), machine learning, statistical models (estimators, Bayesian learning), data science (data mining, sparse data representation, NLP).

Technical skills

Open source projects

- Code to train global and local prompts for vision-language models: ⦿ GalLoP
- Code to train out-of-distribution detectors with HEAT (ICML'23): ⦿ HEAT.
- First of its kind hierarchical landmark dataset (TPAMI, under review): ⦿ google-landmark & ⦿ SupRank.
- Code to train models with HAPPIER (ECCV'22) for hierarchical image retrieval: ⦿ HAPPIER.
- Code to train models with ROADMAP (NeurIPS'21) for image retrieval: ⦿ ROADMAP.

Technologies

- My programming language: Python (PyTorch, NumPy, Hydra, Pandas, Scikit-learn).
- My work environment : Linux/macOS + Git + VS Code & GitHub Copilot + Jupyter Notebook.
- Other tools: LaTeX, Shell script, SLURM, Nextflow.

Contributions to open-source

- Pytorch metric learning: ⦿ pytorch-metric-learning.
- Implementation of Mixture of Experts: ⦿ mixture-of-experts.
- Pytorch-lightning: ⦿ lightning.
- Universal Image Embeddings: ⦿ Universal-Image-Embeddings.

Miscellaneous

- French is my mother tongue and I speak English at C1 level.

- I enjoy running and have run (semi-)marathon and trail running races. I also like hiking and climbing.

Publications

Amaia Cardiel, Eloi Zablocki, **Elias Ramzi**, Oriane Siméoni, and Matthieu Cord. "LLM-wrapper: Black-Box Semantic Aware Adaptation of Vision-Language Models for Referring Expression Comprehension." in Proceedings of the 13th International Conference on Learning Representations (ICLR, 2025). online: `https://arxiv.org/abs/2409.11919`

Elias Ramzi, Nicolas Audebert, Clément Rambour, André Araujo, Xavier Bitot and Nicolas Thome. "Optimization of Rank Losses for Image Retrieval." in IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI, 2025). online: `https://arxiv.org/pdf/2309.08250.pdf`

Marc Lafon, **Elias Ramzi**, Clément Rambour, Nicolas Audebert, Nicolas Thome. "GalLoP: Learning Global and Local Prompts for Vision-Language Models." in Proceedings of the 18th European Conference on Computer Vision (ECCV, 2024). online: `https://arxiv.org/abs/2407.01400v1`

Yannis Karmim, **Elias Ramzi**, Raphaël Fournier S'niehotta, and Nicolas Thome. "ITEM: Improving Training and Evaluation of Message-Passing based GNNs for top-k recommendation." in Transactions on Machine Learning Research (TMLR, 2023). online: `https://arxiv.org/abs/2407.07912`

Marc Lafon, **Elias Ramzi**, Clément Rambour, Nicolas Thome. "Hybrid Energy Based Model in the Feature Space for Out-of-Distribution Detection." in Proceedings of the 40th International Conference on Machine Learning (ICML, 2023). online: `https://arxiv.org/abs/2305.16966`

Elias Ramzi, Nicolas Audebert, Nicolas Thome, Clément Rambour, and Xavier Bitot. "Hierarchical Average Precision Training for Pertinent Image Retrieval." in Proceedings of the 17th European Conference on Computer Vision (ECCV, 2022). online: `https://arxiv.org/abs/2207.04873`

Elias Ramzi, Nicolas Thome, Clément Rambour, Nicolas Audebert, and Xavier Bitot. "Robust and Decomposable Average Precision for Image Retrieval." in Proceedings of the 34th Advances in Neural Information Processing Systems (NeurIPS, 2021). online: `https://arxiv.org/abs/2110.01445`