Mohammad Fahes

mfahes.github.io ✉ mohamad-fahes@outlook.com

Education

PhD | ***Vision-Language Models*** 2022 – 2026
Inria Paris, France

Master Research MVA | ***Mathematics, Vision, Learning*** 2020 – 2021
ENS Paris-Saclay Gif-sur-Yvette, France

Master's Degree (Diplôme d'ingénieur) | ***Geostatistics and Applied Probabilities*** 2019 – 2021
Mines Paris-PSL Paris, France

Master's Degree (Diplôme d'ingénieur) | ***Mechanical Engineering*** 2015 – 2021
Lebanese University – Faculty of Engineering Beirut, Lebanon

Professional Experience

PhD at Inria | ***Vision-Language Models*** 10/2022 – 09/2026
Efficient adaptation of vision-language models.
Advisors: Tuan-Hung Vu, Andrei Bursuc, Raoul de Charette, Patrick Pérez.

Student Researcher at Google DeepMind | ***Large Language Models*** 08/2025 – 01/2026
Training large language models with soft tokens.
Advisors: Anurag Arnab, Cordelia Schmid.

Research Intern at Inria | ***Vision-Language Models*** 04/2022 – 09/2022
Zero-shot domain adaptation with vision-language models.
Advisors: Tuan-Hung Vu, Andrei Bursuc, Raoul de Charette, Patrick Pérez.

Research Intern at Télécom Paris | ***Machine Learning and Signal Processing*** 04/2021 – 09/2021
Deep unrolling for sparse matrix factorization with application to blind source separation in astrophysics. Advisors: Christophe Kervazo, Jérôme Bobin.

Publications and Preprints

M. Fahes, C. Schmid, A. Arnab. Soft Teacher Forcing. *Under review, 2026*.

M. Fahes, T. Vu, A. Bursuc, P. Pérez, R. de Charette. Domain Adaptation with a Single Vision-Language Embedding. *International Journal of Computer Vision (IJCV), 2026*.

M. Fahes, T. Vu, A. Bursuc, P. Pérez, R. de Charette. CLIP's Visual Embedding Projector is a Few-shot Cornucopia. *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 2026*.

Y. Benigmim, M. Fahes, T. Vu, A. Bursuc, R. de Charette. FLOSS: Free Lunch in Open-vocabulary Semantic Segmentation. *IEEE/CVF International Conference on Computer Vision (ICCV), 2025*.

M. Fahes, T. Vu, A. Bursuc, P. Pérez, R. de Charette. A Simple Recipe for Language-guided Domain Generalized Segmentation. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024*.

M. Fahes, T. Vu, A. Bursuc, P. Pérez, R. de Charette. PØDA: Prompt-driven Zero-shot Domain Adaptation. *IEEE/CVF International Conference on Computer Vision (ICCV), 2023*.

M. Fahes, C. Kervazo, J. Bobin, F. Tupin. Unrolling PALM for Sparse Semi-blind Source Separation. *International Conference on Learning Representations (ICLR), 2022*.

Academic Services

Conference Reviewer: CVPR 2024–2025, ECCV 2024, ICLR 2025, ICCV 2025, NeurIPS 2025–2026, ICLR 2027.
Journal Reviewer: IEEE Transactions on Multimedia.

Awards

Outstanding/Top Reviewer: ECCV 2024, CVPR 2025, ICCV 2025, NeurIPS 2026.

Skills

Programming Languages: Python
Libraries & Frameworks: PyTorch, JAX, NumPy
Languages: Arabic (Native), French (Fluent), English (Fluent)