Alexandre Araujo

aaraujo001@gmail.com • alexandrearaujo.com • Github • Scholar

Summary

I am an accomplished machine learning researcher with over 6 years of experience in designing, training, and evaluating large-scale neural networks, particularly in developing stable, scalable, and robust neural network architectures. I have a proven track record of high-impact research with over 470 citations and an h-index of 10.

Key Accomplishments:

- *Published 18 research papers at top ML conferences, e.g., NeurIPS, ICML, ICLR, UAI, AAAI, etc.*
- *Published multiple papers on AI Safety and Robustness: certified accuracy and adversarial robustness.*
- *Supervised and mentored 3 graduate students and 3 undergraduate interns*
- *Trained large-scale networks (> 1B parameters) on a Slurm cluster on hundreds of GPUs*
- *Designed a large-scale dataset of 150M images for distillation of DINOv2 architecture*
- *Designed a new, stable neural network that allows scaling depth to 1000 layers*
- *Fine-tuned LLM for incorporating multi-modality from image encoders.*

Areas of Expertise: Large-scale neural network design and training, distributed computing for machine learning, adversarial machine learning, self-supervised learning and knowledge distillation.

Education

PhD, Computer Science, PSL Research University, Paris, France 2021
MS, Business Administration, SKEMA Business School, Lille, France 2016
BS, Mathematics, University of Versailles, Versailles, France 2011

Research & Industry Experience

New York University **New York, NY, US**
Postdoctoral Researcher (Advisors: S. Garg, F. Khorrami) *January 2023 – June 2024*

- Designed scalable, stable neural network architectures and introduced a novel metric to improve transformer stability, advancing machine learning and AI safety.

INRIA / École Normale Supérieure **Paris, France**
Postdoctoral Researcher (Advisors: J. Ponce, J. Mairal) *October 2021 – December 2022*

- Conducted research on Focus Stacking from Handheld Raw Image Bursts. Designed a large-scale computer vision dataset to improve recent advancements on Focus Stacking with supervised learning.

PSL Research University **Paris, France**
Ph.D. Candidate (Advisors: Y. Chevaleyre, B. Negrevergne and J. Atif) *September 2017 – June 2021*

- Thesis: Building Compact and Robust Deep Neural Networks with Toeplitz Matrices
- PhD in Deep Learning, specializing in the development of compact and robust neural networks.

Wavestone **Paris, France**
Data Scientist *September 2015 – August 2017*

- Predicted mortgage application acceptance for a mortgage broker using machine learning algorithms. Utilized 5 years of historical data and deployed the model into production.
- Predicted customer churn for an energy company using machine learning algorithms. Constructed a 1 billion row dataset from 3 years of historic data gathered with Hadoop.
- Predicted train breakdowns for a European Railway Company using machine learning algorithms. Analyzed 20 years of historic data to develop predictive models.

Amazon **Luxembourg**
Data Engineer Intern *December 2014 – May 2015*

- Developed and optimized complex SQL queries in Amazon Redshift to generate comprehensive transportation and financial statistics, providing valuable insights for business decision-making.
- Designed and implemented efficient data pipelines to automate the flow of information into Business Intelligence (BI) dashboards, significantly reducing manual data processing and improving reporting efficiency.
- Created robust automated data pipelines to power real-time dashboards, enabling instant visibility into critical transportation and financial metrics across the organization.

Publications

+470 citations, H-index 10, 5 ICLR, 3 NeurIPS, 3 ICML, 1 UAI, 1 AAAI

Conference Papers

Fine-grained Local Sensitivity Analysis of Standard Dot-product Self-Attention
A. Havens, ***A. Araujo***, *H. Zhang, B. Hu* – **ICML 2024**

LipSim: A Provably Robust Perceptual Similarity Metric
S. Ghazanfari, ***A. Araujo***, *P. Krishnamurthy, F. Khorrami, S. Garg* – **ICLR 2024**

The Lipschitz-Variance-Margin Tradeoff for Enhanced Randomized Smoothing
B. Delattre, ***A. Araujo***, *Q. Barthélemy, A. Allauzen* – **ICLR 2024**

Novel Quadratic Constraints for Extending LipSDP beyond Slope-Restricted Activations
P. Pauli, A. Havens, ***A. Araujo***, *S. Garg, F. Khorrami, F. Allgöwer, B. Hu* – **ICLR 2024**

On the Scalability and Memory Efficiency of SDP for Lipschitz Constant Estimation of Neural Networks
Z. Wang, A. Havens, ***A. Araujo***, *Y. Zheng, B. Hu, Y. Chen, S. Jha* – **ICLR 2024**

Exploiting Connections between Lipschitz Structures for Certifiably Robust DEQ models
*A. Havens**, ***A. Araujo****, *S. Garg, F. Khorrami, B. Hu* – **NeurIPS 2023**

Diffusion-Based Adversarial Sample Generation for Improved Stealthiness and Controllability
H. Xue, ***A. Araujo***, *B. Hu, Y. Chen* – **NeurIPS 2023**

Towards Better Certified Segmentation via Diffusion Models
O. Laousy. ***A. Araujo***, *G. Chassagnon, M. Revel, S. Garg, F. Khorrami, M. Vakalopoulou* – **UAI 2023**

Efficient Bound of Lipschitz Constant for Convolutional Layers by Gram Iteration
B. Delattre, Q. Barthélemy, ***A. Araujo***, *A. Allauzen* – **ICML 2023**

A Unified Algebraic Perspective on Lipschitz Neural Networks
A. Araujo*, *A. Havens**, *B. Delattre, A. Allauzen, B. Hu* – **ICLR – Spotlight 2023**

A Dynamical System Perspective for Lipschitz Neural Networks
*L. Meunier**, *B. Delattre**, ***A. Araujo****, *A. Allauzen* – **ICML – Oral 2022**

On Lipschitz Regularization of Convolutional Layers using Toeplitz Matrix Theory
A. Araujo, *B. Negrevergne, Y. Chevaleyre, J. Atif* – **AAAI 2020**

Understanding and Training Deep Diagonal Circulant Neural Networks
A. Araujo, *B. Negrevergne, Y. Chevaleyre, J. Atif* – **ECAI 2020**

Theoretical Evidence for Adversarial Robustness through Randomization
R. Pinot, L. Meunier, ***A. Araujo***, *H. Kashima, F. Yger, C. Gouy-Pailler, J. Atif* – **NeurIPS 2019**

Workshop Papers

Stronger Universal and Transfer Attacks by Suppressing Refusals
D. Huang, A. Shah, ***A. Araujo***, *D. Wagner, C. Sitawarin* – **NeurIPS – Workshop 2024**

R-LPIPS: An Adversarially Robust Perceptual Similarity Metric
S. Ghazanfari, S. Garg, P. Krishnamurthy, F. Khorrami, ***A. Araujo*** – **ICML – Workshop 2023**

Advocating for Multiple Defense Strategies against Adversarial Examples
A. Araujo, *L. Meunier, R. Pinot, and B. Negrevergne* – **ECML – Workshop 2020**

Compact Deep Learning Models for Video Classification using Circulant Matrices
A. Araujo, *B. Negrevergne, Y. Chevaleyre, J. Atif* – **ECCV – Workshops 2018**

Preprints

EMMA: Efficient Visual Alignment in Multi-Modal LLMs
S. Ghazanfari, ***A. Araujo***, *P. Krishnamurthy, S. Garg, F. Khorrami* – **Preprint 2024**

PAL: Proxy-Guided Black-Box Attack on Large Language Models
C. Sitawarin, N. Mu, D. Wagner, ***A. Araujo*** – **Preprint 2024**

Towards Real-World Focus Stacking with Deep Learning
A. Araujo, *J. Ponce, J. Mairal* – **Preprint 2023**

Activities and Services

Applied Projects

Full-stack Search Engine with Retrieval-Augmented Generation (RAG) Approach **Ongoing**

Tech stack: Next.js, Django, Faiss, Boost Beast, Sentence Transformer, PostgreSQL

- Developing a full-stack search engine with a Next.js frontend and a Django backend utilizing Sentence Transformers for query embedding and cross-encoder for result re-ranking.
- Implementing a language model (LM) for query autocomplete
- Designing and deploying a REST API using Boost.Beast to serve search results from an index built with Faiss (Facebook AI Similarity Search) for high-performance, vector-based similarity searches.
- Building a web crawling pipeline with Apache Nutch to gather and embed website data using Sentence Transformers, improving search relevance and accuracy.

Teaching

New York University, New York, NY, US

Graduate Course: Adversarial Machine Learning 2023

PSL Research University, Paris, France

Executive Master: Adversarial Machine Learning 2020, 2021
Master IASD: Data Mining & Machine Learning 2019
Master ID: Data Mining & Machine Learning 2019

École Polytechnique, Paris, France

Data Science & Machine Learning 2016, 2017, 2018, 2019, 2020

Reviewer

International Conference on Learning Representations (ICLR)	2024, 2025
Neural Information Processing Systems (NeurIPS)	2023, 2024
International Conference on Machine Learning (ICML)	2023, 2024
European Conference on Computer Vision (ECCV)	2024
Computer Vision and Pattern Recognition Conference (CVPR)	2023
International Conference on Computer Vision (ICCV)	2023
Artificial Intelligence and Statistics (AISTATS)	2022, 2023
Association for the Advancement of Artificial Intelligence (AAAI)	2022, 2023

Invited Talks

University of Illinois Urbana-Champaign	October 2023
NYU – Center for Data Science	April 2022
INRIA / École Normale Supérieure de Paris	July 2021
École Normale Supérieure de Lyon	July 2021
INSIS – French National Center for Scientific Research	January 2021
PFIA – French AI conference	June 2019, 2020, 2021
International Cybersecurity Forum	January 2020
Limits of AI – BPI Conference	June 2019

Technical Skills

Programming Languages : Python, C++, SQL
HPC Job Schedulers : Slurm, IBM Spectrum LSF
Deep Learning Frameworks : TensorFlow, PyTorch
ML Libraries: XGBoost, LightGBM, Scikit-Learn
Data Science Framework : OpenCV, SciPy, NumPy, Pandas