

Mehdi Azabou - Graduating Summer 2024

www.mehai.dev · mazabou@gatech.edu · Nationality: Tunisian

Bio

I am a fourth-year Machine Learning Ph.D. candidate at Georgia Tech, advised by Dr. Eva L. Dyer. My main areas of interest are Representation Learning, Data-Centric AI, and Computational Neuroscience. I am actively working on developing methods for self-supervised representation learning for time-series and graphs, and developing new frameworks to build large-scale multimodal foundation models to advance scientific discovery.

Education

- PhD in Machine Learning**, Georgia Institute of Technology Aug 2020 – Present
Advisor: Prof. Eva L. Dyer
- Masters in Computer Science**, Georgia Institute of Technology Aug 2018 – Dec 2020
Specialization: Machine Learning, GPA: 4.0
- Masters in Engineering**, CentraleSupélec Sep 2016 – May 2019
CentraleSupélec is a top French graduate engineering school of Paris-Saclay University.
3rd year Specialization: Interactive Systems and Robotics.
Project: Prediction of Chemical Reaction Outcomes using Graph Neural Networks.
- Classes Préparatoires (CPGE)**, IPEST Sep 2014 – May 2016
Intensive training in mathematics, physics and chemistry to prepare for competitive entrance exams.
-

Research Experience

- Georgia Institute of Technology *Atlanta, GA, United States*
Research Assistant, NerDS Lab Aug 2020 – Present
Working with Dr. Eva L. Dyer on a wide range of problems in artificial and natural intelligence.
- Developed a large transformer model trained on a large collection of high-dimensional multi-variate neuronal timeseries, totalling 27,373 unique channels and around 1 billion tokens. Training datasets exhibited high levels of heterogeneity with various recording technologies, sampling rates and experimental protocols (**NeurIPS 2023**)
 - Developed a state-of-the-art multi-timescale representation learning method that captures latent behavior of quadruped robots and animals by training a model to predict future behavior under an optimal transport objective (**NeurIPS 2023 Spotlight**)
 - Revealed the existence of computational fingerprints of transcriptomic cell types, during neuronal activity across diverse contexts (**Cell Reports, 2023**)
 - Developed Half-Hop, a plug-and-play augmentation for graph learning, now included as a standard augmentation in the PyTorch Geometric library (**ICML 2023**)
 - Co-developed a state-of-the-art self-supervised learning method for large graphs (BGRL) (**ICLR 2022**)
 - Co-developed a spatio-temporal transformer model for multi-variate timeseries data (**NeurIPS 2022**)
 - Co-developed a generative model with a disentangled latent space for neural data (**NeurIPS 2021 Oral**)
- Research Assistant**, Smart Cities Jan 2020 – Jul 2020
- Worked on real-time prediction systems on edge devices used by Georgia Department of Transport (DOT) for the inventorying and assessment of road assets including traffic signs and pavements. Redesigned the evaluation pipeline to better track performance of developed models.

Industry Experience

IBM Research

New York, United States

AI Research Scientist Intern

May 2023 – Aug 2023

- Developed a novel framework for generative time series foundation models. Trained a unified model with multiple public datasets from various sources. Invention disclosure submitted. Manuscript to appear soon.

Parrot Drones

Paris, France

Computer Vision Intern

May 2019 – Nov 2019

- Worked on semantic segmentation tasks for drone aerial imagery. Benchmarked multiple state-of-the-art architectures. Produced software to systematically generate and validate data from simulated off-the-shelf environments, and evaluated domain adaptation methods to address the sim2real gap.

Cleed (Startup)

Paris, France

Deep Learning Intern

Jun 2018 – Sep 2018

- Led efforts to develop a virtual clothing try-on tool. Designed and implemented a generative model that performs garment swapping. Improved model performance by introducing more modalities like dense human pose estimate and clothing segmentation map. Collected data by scraping retail websites.

Publications

(Google Scholar; DBLP)

⁺contributed equally as co-first authors, ^{*}co-senior authors

Papers

2023

- **Azabou, M.**, Arora, V., Ganesh, V., Mao, X., Nachimuthu, S., Mendelson, M., Richards, B.A., Perich, M., Lajoie, G., Dyer, E.L.. *A Unified, Scalable Framework for Neural Population Decoding*, Advances in Neural Information Processing Systems **NeurIPS 2023**.
- **Azabou, M.**, Mendelson, M., Ahad, N., Sorokin, M., Thakoor, S., Urzay, C., Dyer, E.L.: *Relax, it doesn't matter how you get there: A new self-supervised approach for multi-timescale behavior analysis*, Advances in Neural Information Processing Systems **NeurIPS 2023 (Spotlight)** (3% acceptance rate).
- **Azabou, M.**, Ganesh, V., Thakoor, S., Lin, C. H., Sathidevi, L., Liu, R., Valko, M., Veličković, P. and Dyer, E. L. *Half-Hop: A graph upsampling approach for slowing down message passing*. International Conference on Machine Learning **ICML 2023**.
- Schneider, A.⁺, **Azabou, M.**⁺, McDougall-Vigier, L., Parks, D.B., Ensley, S., Bhaskaran-Nair, K., Nowakowski, T.J., Dyer, E.L.^{*}, and Hengen, K.B.^{*} *Transcriptomic cell type structures in vivo neuronal activity across multiple time scales*. **Cell Reports**, Volume 42, Issue 4, April 2023.
- M. Mendelson⁺, **M. Azabou**⁺, S. Jacob, N. Grissom, D.P. Darrow, B. Ebitz, A. Herman, E.L. Dyer. *Learning signatures of decision making from many individuals playing the same game*, 11th IEEE EMBS Conference on Neural Engineering **NER 2023**.
- C. Urzay⁺, N. Ahad⁺, **M. Azabou**, A. Schneider, G. Atmakuri, K.B. Hengen, E.L. Dyer. *Detecting change points in neural population activity with contrastive metric learning*, 11th IEEE EMBS Conference on Neural Engineering **NER 2023**.

2022

- Liu, R., **Azabou, M.**, Dabagia, M., Xiao, J., and Dyer, E. L. *Seeing the forest and the tree: Building representations of both individual and collective dynamics with transformers*. Advances in Neural Information Processing Systems **NeurIPS 2022**.
- Quesada, J., Sathidevi, L., Liu, R., Ahad, N., Jackson, J. M., **Azabou, M.**, Xiao, J., Liding, C., Urzay, C., Gray-Roncal, W., Johnson, E. C. and Dyer, E. L. *MTNeuro: A Benchmark for Evaluating Representations of*

Brain Structure Across Multiple Levels of Abstraction. Advances in Neural Information Processing Systems, Datasets and Benchmarks Track **NeurIPS 2022**.

- Thakoor, S., Tallec, C., Azar, M.G., **Azabou, M.**, Dyer, E.L., Munos, R., Veličković, P., and Valko, M. *Large-Scale Representation Learning on Graphs via Bootstrapping*. In Proceedings of the International Conference on Learning Representations, **ICLR 2022**.
- **Azabou, M.**, Mendelson, M., Sorokin, M., Thakoor, S., Ahad, N., Urzay, C., and Dyer, E. L. *Learning Behavior Representations Through Multi-Timescale Bootstrapping*. Workshop on Multi-Agent Behavior, **CVPR 2022**, selected for Oral.

2021

- Liu, R., **Azabou, M.**, Dabagia, M., Lin, C.H., Gheshlaghi Azar, M., Hengen, K., Valko, M., and Dyer, E. L. *Drop, swap, and generate: A self-supervised approach for generating neural activity*. Advances in Neural Information Processing Systems, **NeurIPS 2021 (Oral)** (1% acceptance rate).
- Lin, C. H., **Azabou, M.**, and Dyer, E. L. *Making transport more robust and interpretable by moving data through a small number of anchor points*. In Proceedings of the International Conference on Machine Learning, **ICML 2021**.
- **Azabou, M.**, Azar, M.G., Liu, R., Lin, C.H., Johnson, E.C., Bhaskaran-Nair, K., Dabagia, M., Avila-Pires, B., Kitchell, L., Hengen, K.B. Gray-Roncal, W., Valko, M., and Dyer, E. L. *Mine your own view: Self-supervised learning through across-sample prediction*. Workshop on Self-supervised Learning: Theory and Practice, **NeurIPS 2021**, selected for Oral.
- **Azabou, M.**⁺, Dabagia, M.⁺, Liu, R.⁺, Lin, C. H., Hengen, K. B., and Dyer, E. L. *Using self-supervision and augmentations to build insights into neural coding*. Workshop on Self-supervised Learning: Theory and Practice, **NeurIPS 2021**.

Abstracts

- **Azabou, M.**, Arora, V., Mineault, P., Ganesh, V., Mao, X., Nachimuthu, S., Mendelson, M., Richards, B.A., Perich, M., Lajoie, G., Dyer, E.L. *Large-scale pretraining on neural data allows for transfer across subjects, tasks and species*. To appear at the Computational and Systems Neuroscience conference (COSYNE), Lisbon, Portugal, 2024.
- Urzay, C.⁺, Ahad, N.⁺, **Azabou, M.**, Schneider, A., Atmakuri, G., Hengen, K.B., and Dyer, E. L. *Detecting change points in neural population activity with contrastive metric learning*. To appear at the Cognitive and Computational Neuroscience (CCN) Annual Meeting, San Francisco, CA, August, 2022. (+ co-first authors)

Awards & Funding

- Outstanding ECE Graduate Research Assistant Award (2024), Georgia Institute of Technology.
 - Institute for Data Science and Engineering (IDEaS) Cyberinfrastructure Proposal (Amount awarded: 600 hours on Nvidia DGX (8x H100) shared system, equivalent of 58,992 US\$ in computing costs); Proposal Title: “Towards a multi-species, multi-region foundation model for neuroscience”.
 - NSF CloudBank Award Request to support scaling my PhD research project (42,000 US\$ in cloud computing credits awarded by the NSF).
 - Travel award to attend CVPR 2022 to present at the “Multi-Agent Behavior Workshop” (Amount awarded: 2,000US\$).
 - First place, GT Convergence Innovation Competition, Climate Solutions category, Fall 2018.
 - French-Tunisian merit-based Scholarship for Grandes Écoles, 2016-2019 – awarded to 40 students per year.
 - Tunisian Excellence Scholarship, 2014-2016 – highest undergraduate scholarship nationally.
-

Talks

- “A Unified, Scalable Framework for Neural Population Decoding” Invited Talk, Trends in NeuroAI reading group, **MedARC, Stability AI**, February 22, 2024.
 - “A Unified, Scalable Framework for Neural Population Decoding” Invited Talk, Neuro+ML Journal Club, **California Institute of Technology**, February 22, 2024.
 - “A Unified, Scalable Framework for Neural Population Decoding” Invited Talk, Neurotheory/Machine learning Journal Club, **Baylor, Rice Univ., Univ. of Houston**, January 26, 2024.
 - “Foundation models for Neuroscience”, co-presented with Blake Richards, **CIFAR Learning in Machines & Brains**, New Orleans, LA, December 7, 2023.
 - “A Unified, Scalable Framework for Neural Population Decoding”, UNIQUE Scientific Retreat, Mont-Tremblant, **Université de Montréal**, October 21, 2023.
 - “A Unified, Scalable Framework for Neural Population Decoding”, NeuroAI Workshop, **Mila**, October 13, 2023.
 - “Mine your own view: Self-supervised learning and evaluation for neural data”, Summer School on Data Science, AI, and Neuroscience, **California Institute of Technology**, Chen Institute for Neuroscience, Pasadena, CA, July 14, 2022.
 - “Learning Behavior Representations Through Multi-Timescale Bootstrapping”, Workshop on Multi-Agent Behavior, New Orleans, LA, **CVPR 2022**.
 - “Mine your own view: Self-supervised learning through across-sample prediction”, Workshop on Self-supervised Learning: Theory and Practice, Virtual, **NeurIPS 2021**.
-

Teaching

- Co-Instructor at the *Caltech/Chen Institute’s Data Science and AI for Neuroscience Summer School*, 2022. Developed and led the representation learning hands-on session.
 - Content Developer and Guest Lecturer for *BMED 6517 Machine Learning in Biosciences at Georgia Tech*, 2021.
 - Developed and led a Python bootcamp session for *DL@MBL: Deep Learning for Microscopy Image Analysis Course at the Marine Biological Laboratory*, MBL, Woods Hole, Massachusetts, Sep 25 - Oct 6, 2021.
 - Teaching Assistant for *CS 4261 Mobile applications and Services at Georgia Tech*, Spring 2019. Guided students through the different stages of app creation: ideation, market research, and front-end and back-end development.
-

Academic Service

- Meeting reporter for CIFAR’s Learning in Machines & Brains program, New Orleans, LA, December 7, 2023.
 - Reviewer for *Neural Information Processing Systems (NeurIPS)* 2021, 2022, 2023.
 - Reviewer for *International Conference on Learning Representations (ICLR)* 2023, 2024.
 - Reviewer for *Computer Vision and Pattern Recognition (CVPR)* 2023, 2024.
 - Reviewer for *International Conference on Machine Learning (ICML)* 2023.
 - Reviewer for *Learning on Graphs Conference (LOG)* 2022, 2023.
 - Reviewer for *Cell Patterns*, 2022.
 - Sub-reviewer for *Neuron*, 2021.
 - Reviewer for *International Conference on Artificial Intelligence and Statistics (AISTATS)* 2021.
-

Mentorship

- Vinam Arora, Master’s in ECE,
Project: A scalable framework for neural population decoding.

2023

- Venkataramana Ganesh, Masters in CS 2022 – 2023
Project: Designing data augmentations for graph representation learning.
- Michael Mendelson, Undergrad in BME 2021 – 2023
Project: Using deep learning to decode signatures of exploration and exploitation in human decision making, Received President's Undergraduate Research Award (PURA) at Georgia Tech.
- Santosh Nachimuthu, Undergrad in BME, 2023
Project: A Scalable Framework for Neural Population Decoding.
- Carolina Urzay, Undergrad in BME 2021 – 2022
Project: Detecting change points in neural population activity with contrastive metric learning.
- Zijiang Wu, Undergrad in CS/Math 2020 – 2021
Project: Revealing aging dynamics in *C. elegans* by modeling behavior dynamics throughout lifespan.

Skills

★★Expert, ★Advanced

- **Programming Languages** Python^{★★}, Java, JavaScript, C++, Matlab, Mathematica.
 - **ML Frameworks** PyTorch^{★★}, PyG^{★★}, raytune[★], TensorFlow[★], scikit-learn[★], OpenCV.
 - **Misc.** Docker[★], Bokeh[★], Flask[★], gcloud, aws, PostgreSQL, Android, ROS, threeJS.
 - **Languages** Fluent English, French and Arabic. Intermediate Spanish. Beginner Mandarin.
-