

Candidate Overview: I am currently an undergraduate student at Northeastern University studying Mathematics and Physics. I am applying to Harvard's CS doctoral program so that I can pursue a PhD developing methods to make graph neural networks more scalable. This includes a characterization of graph oversmoothing and architectural remedies to the problem. I am fortunate to have had broad exposure through my undergraduate research, including work on approximation error bounds of equivariant functions, property prediction models for atomic data, computational imaging, and cosmology.

My research has resulted in several first author papers. One paper is in submission to Transactions on Machine Learning Research (TMLR) and two others are published in some of the best astrophysics venues, the Astronomical Journal and the Open Journal of Astrophysics. I also have a co-first author paper that was awarded an Oral Spotlight as part of the top 4/300+ submissions at the NeurIPS Machine Learning and the Physical Sciences workshop, which I wrote with one other undergraduate student and no advisory supervision. Additionally, I have over 14 total publications, including co-authorship, highlighted by a paper accepted and to appear in Nature Astronomy. I am the first student in my university's Physics department to be nominated for the computing research association outstanding undergraduate research award. Similarly, I am the first Northeastern student to have received national recognition for research, service, and academic performance through the leadership scholarship presented by the Society of Physics Students National Chapter.

Research Interests: Based on my research interests and preparation for graduate research, **I am most interested in working with Professor Melanie Weber.** I would like to spend my PhD researching over smoothing in graph neural networks. This research thread is both timely and impactful: Over-smoothing is one of the defining limitations of graph neural networks, as it hinders both scalability and expressivity. While recent works have proposed various graph rewiring methods to alleviate this issue, causes of over-smoothing beyond network depth are still largely mysterious. I am particularly interested in developing architectures that use soft priors that discourage over-smoothing. This is in contrast to approaches such as unitary convolution, which strictly preserve the smoothness of a graph even when *some* amount of smoothing is desirable. Relatedly, I am also interested in the approximation and generalization error of smoothness preserving functions, and understand to what extent unitary convolution networks can be contextualized over-constraining.

In the long term, my goal is to become a professor and lead my own research lab. For the reasons mentioned above, I think that the CS doctoral program at Harvard is uniquely suited to prepare me as a researcher and academic.

Research Background: The primary thread of my undergraduate research has studied the generalization and approximation error of equivariant functions. In particular, I was concerned with how well equivariant functions could minimize calibration error, which is loosely defined as the discrepancy between a model's accuracy and purported confidence. The main contribution of this work was a theorem that placed calibration error bounds under various assumptions of symmetry breaking between the model and the data, expressed through the language of correct, incorrect, and extrinsic equivariance. I also showed experimentally how symmetry mismatch can cause a detriment to model calibration in both regression and classification settings. This led to a first author publication currently in submission to Transactions on Machine Learning Research (TMLR). My contributions were especially significant

considering that I proposed the main idea, what to prove and how, how to substantiate the results experimentally, and wrote up the final paper myself.

Eager to demonstrate the applicability and impact of my theory, I joined the Harvard and Smithsonian Center for Astrophysics as a visiting scientist with the AstroAI group, where I was advised by Cecilia Garraffo and worked closely with professors Bill Freeman and Sara Seager. During this time, I showed that my previous work was applicable to the biomarkers retrieval problem. The biomarkers retrieval problem aims to identify signs of life in exoplanet atmospheres in the form of Carbon based macromolecules. Specifically, chemical spectra and uncertainty information are used with the Hamiltonian Monte Carlo (HMC) algorithm to estimate relative chemical abundances. I contributed to this problem by developing an E(3)-invariant machine learning model that predicted these chemical properties *and* embedded uncertainty estimates into the prediction using a technique called evidential regression. This allows us to look for molecules whose spectra may not be measured on Earth and expensive to simulate numerically. I assessed the reliability of the uncertainties in terms of the calibration error bounds derived in my previous work. Once I deemed that the uncertainties were well calibrated, I used them to build the covariance information necessary to perform HMC. A paper is currently in preparation.

A parallel thread of my undergraduate research was focused on developing algorithms for large scale astronomical surveys in order to perform weak gravitational lensing analysis. Specifically, I developed a method for modeling optical blurring artifacts intrinsic to the James Webb Space Telescope for use by the COSMOS-Web astronomical survey collaboration. This led to two first author publications in the *Astronomical Journal* and the *Open Journal of Astrophysics*. I am proud of this work not only because I conceived of the idea itself but also because I have already seen an immediate impact on the field: The suite of benchmarks I developed for my paper were used in subsequent analysis; this includes the discovery of >100 rare strong lens systems, the highest resolution map of dark matter ever measured, and the public data release of the COSMOS-Web survey.

My most recent research has examined over-smoothing in graph neural networks in the context of partial differential equation solving and aims to elucidate both theoretically and empirically the impact of unitarity on learning. It is this thread that I want to continue into my PhD, and a paper is in preparation for ICML.