Fourth Semester Report

Student Information

Name	Killian Maxwell
Department	Doctoral School of Physics
Email	maxwelki@student.elte.hu
Neptun ID	FZ72X9
Supervisors	Andras Kovacs, PhD; Istvan Csabai, PhD
Thesis Title	Mapping the Cosmic Web: From Simulations to the Deepest Galaxy Surveys

Overview

One overarching goal of modern cosmology is to reconstruct how miniscule density fluctuations seen in the cosmic microwave background (CMB) anisotropies grow to the intricate cosmic web. The Lambda Cold Dark Matter (ΛCDM) cosmological model has shown remarkable explanatory power over a variety of cosmic scales and epochs, and it narrates a reassuring story of a universe currently filled primarily with dark matter and dark energy. However, this explanation is not completely satisfactory. The true nature of the dark components remains a puzzle. Within my PhD project I seek to develop new tools capable of probing the elusive dark energy component. My particular approach is to utilize modern, state-of-the-art simulations and machine learning techniques to prepare a Simulation Based Inference (SBI) model optimized for predictions on the upcoming survey data from the Euclid Telescope. Over the course of the previous four semesters, I have made great strides towards this research goal.

My first semester focused primarily on building my background. I took multiple classes covering both basics and the forefront exploration of machine learning and AI. I also participated in multiple online journal review groups and seminars; and I got involved with the Euclid Consortium by joining their online platform and taking on the task of managing their LinkedIn page. For my second semester, much of these efforts continued. I narrowed down my choice of neural networks to explore in deeper detail, broadened my knowledge of the current research similar to my own, and I participated in more seminars and conferences. By the completion of my third semester, I had a well functioning machine learning model capable of making map-level inferences of cosmological parameters–albeit still with room for improvement.

This last semester (my fourth) brought great strides in accomplishment for my research. I finished up my coursework with the second part of Radio Astronomy, added a couple more textbooks to my background, got more involved in the various teams utilizing SBI within the Euclid Consortium, and presented my current progress at an international seminar at Konkoly. However, the greatest amount of progress gained was in the improvements to my SBI model. The posterior predictions it could infer, and the constraints on those predictions were greatly improved. As a result I was able to get involved in a supplementary project with my team. We are currently on the verge of completing a paper utilizing a combination of my team's exploration and my SBI pipeline. We'd hoped to publish the work in May, however the Euclid Consortium meeting in Barcelona, and my impending complex exam, required postponement of publishing until June.

I will detail all of the aforementioned work and research I have completed over the course of the past two academic years in the sections below.

Details

Coding Classes & Projects

Throughout the course of the previous four semesters I completed online coding classes (this was in addition to the classes I took at ELTE) and additional coding projects. These classes and projects focused on utilizing python libraries specifically designed to implement machine learning (Tensorflow, PyTorch, etc.). They included everything from preparing data in an optimal way for AI to learn patterns, to altering layer structures and hyperparameters to improve the learning process, to constructing proper loss functions that best suit the learning task. These classes and projects provided great insight into my continued work with AI, and were key in guiding me in many of the crucial steps I took.

Literature Review

Throughout the four semesters, I deepened my understanding of mainstream research in machine learning and cosmology by delving into multiple forms of literature. I began with “Cosmology,” by Daniel Baumann, and I regularly returned to it for reference and reminder. Additionally, I utilized the free web-based textbook “Deep Learning” provided by a few collaborating MIT professors: Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Beyond textbooks, I have continued regularly reading papers covering content on machine learning models, exploration of the ISW effect, and various cross-correlations already being made by various teams around the world. I made sure to give specific focus to articles by teams that utilized DeepSphere. My supervisor and I concluded that seeking a way to optimally construct an AI-driven analysis of correlations between density maps and ISW maps was a gap that could be filled in much of the known literature. Thus, a great deal of my literature review has been geared toward that end. This semester particularly, much insight was gleaned from the paper that introduced the Gower Street Simulations (N. Jeffrey, et. al. (arXiv: 2403.02314v1)). Within this paper, a key method for SBI utilizing their simulations was demonstrated, and this became the basis for my model moving forward.

University Coursework

The list of courses I enrolled in over the previous semesters is as follows: New Results in Machine Learning, Data Mining and Machine Learning, Astrostatistics I, Astrostatistics II, (Exo)Planetary Atmospheres II, Compact Stars, Radio Astronomy I, and Radio Astronomy II. Each offered insight, tools, and practices useful to my research. Each of their individual impacts on my progress are detailed in my previous reports, so I will not elaborate on them here.

Seminars and Groups

Since the first semester I have been part of multiple seminars and groups. I joined CosmoVerse–a group of scientists from across Europe that hold regular virtual meetings to explore and discuss current research and new journal articles. I also joined a group at the Konkoly observatory (CSFK) that meets every other week to discuss the application of AI in scientific exploration. I have also continued to participate in weekly meetings with both my supervisors and all the members of our extended team. Additionally, I had multiple opportunities to present my research at seminars and conferences. The most recent of these were Dark Map (a conference I attended and helped facilitate) and the Konkoly Extragalactic Workshop (a conference I had the opportunity to present at). I was also invited to present my work at the WSCLAB GPU Day at Wigner Research Center, but due to a schedule conflict I was unable to do so. However, I hope to present at this conference in the future. Furthermore, I intend to apply for attendance and presentation of my work at the 2027 Euclid Consortium Conference, and one additional conference–hopefully the 2027 ICASS or something similar.

Research Progress

In this section I will detail the culmination of progress made on my research throughout the course of the previous four semesters. The progress already previously reported in semesters one, two, and three will be included along with the most recent accomplishments from the fourth semester. Furthermore, the additional project in which I am involved will be detailed here; it is related heavily to my own research topic, and will serve as a stepping stone into further publications I intend to complete.

One of the first challenges faced when trying to implement Machine Learning within a field is preparing data such that it is optimized for a machine to learn from. Often this means running a number of statistical procedures on the data (such as normalization), converting the data into a form that the machine can take as input (often a vector, but sometimes a map or image), and/or just ensuring there is enough data to learn from. Thanks to the form of the Gower Street simulations, introduced by N. Jeffrey, et. al. (arXiv: 2403.02314v1), the data utilized for machine learning in my exploration was already in the form it needed to be–I merely needed to apply some statistical procedures in preparation, and then find a way to expand the data suite into even more data. The first round of Gower Street simulations came with 791 simulations, each a result of different seed cosmological parameters. Each simulation was finalized as a series of 100 lightcone shells spanning from redshift 0 to redshift 40. These lightcone shells were saved as HEALPix maps, allowing me to utilize the 'healpy' package of Python to read them directly and feed them as input into my neural network model of choice.

I decided to utilize the Graph Convolutional Neural Network (GCNN) developed by Nathanaël Perraudin, et. al. (arXiv: 1810.12186v2) called DeepSphere. The structure of this network allows for convolution to be performed directly on the surface of a sphere (whole or in-part) while maintaining near-complete independence from rotation. Below is an image from their article detailing the structure of DeepSphere. It consists of 3 convolutional layers followed by fully connected (dense) layers. Each convolutional layer implements five operations: convolution, non-linearity, batch normalization, down-sampling, and pooling.

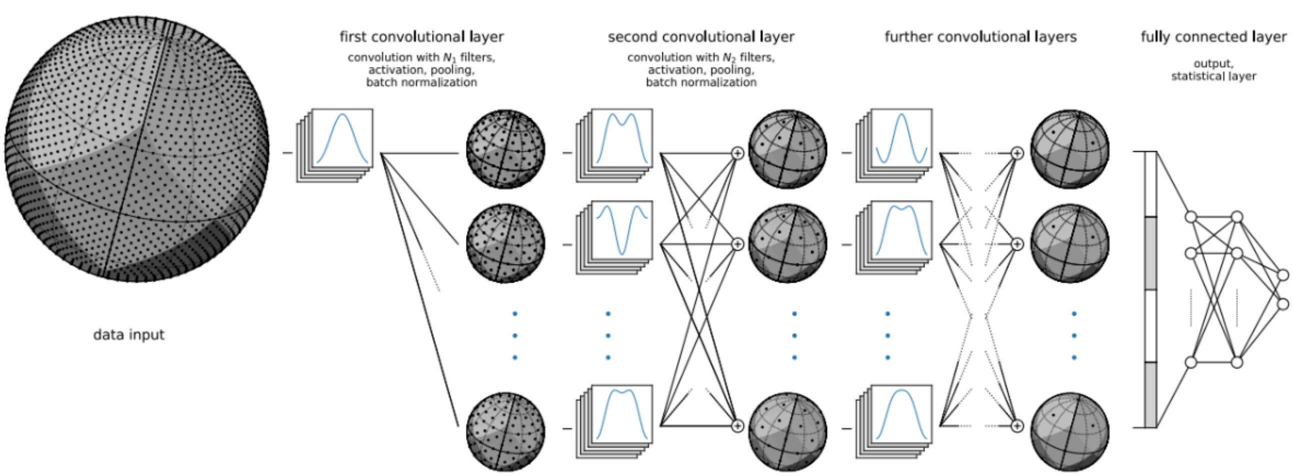

Despite the fact that the maps (redshift shells) provided by the Gower Street suite were already in a usable form, because of computational limitations I needed to downsample them. They came as 2048-sided HEALPix maps, which is over 50 million pixels, and a whopping 200 megabytes per map! So I began the process of downsampling them to 128-sided HEALPix maps. However, before doing so, I made sure to convert them into density contrast maps–they were initially dark matter count maps. By doing this conversion first, the integrity of the data was better maintained during downsampling. Furthermore, utilizing all 100 maps of each simulation–even after downsampling–would have posed further memory issues; so I combined multiple shells into redshift-binned maps. Furthermore, any shells beyond a redshift of 1 were of little impact in my work, which focused on data similar to what will result from the Euclid Telescope Survey. Thus, for each simulation, the 100 initial shells were cut and condensed down to 5 binned shells across redshifts 0 to 1, each a summed map over a delta of 0.2 redshift.

In addition to formatting the data such that it would be usable within my computational memory limitations, I needed more data to work with. Normally, neural networks require thousands to millions of samples to be trained on, so the 791 simulations needed to be expanded in some way. To do this, I sectioned each map within each simulation into 20 equal-area portions, mimicking the geometric structure of an icosahedron. This mapping scheme was chosen due to HEALPix having a similar scheme in how its pixels are laid out on the sphere. Below is an image from the article by K. M. Górski, et. al. (arXiv: astro-ph/0409513v1) detailing the HEALPix partition of the sphere at side numbers 1, 2, 4, and 8; and to its right is an example of 1 of the 20 sections of a single density bin map from a random simulation. At the base schematic, a HEALPix map with 1 side contains 12 pixels, and these pixels align with the vertices of an icosahedron. Furthermore, splitting each map into 20 sections doubled as a method for simulating access to only a portion of the data. This is similar to how telescope surveys often only provide portions of the data contained in the full sky.

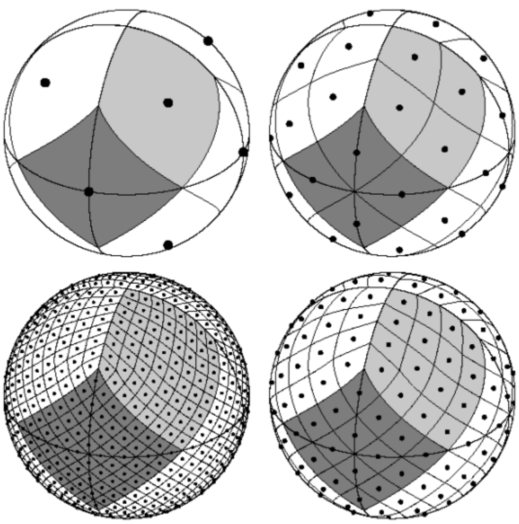

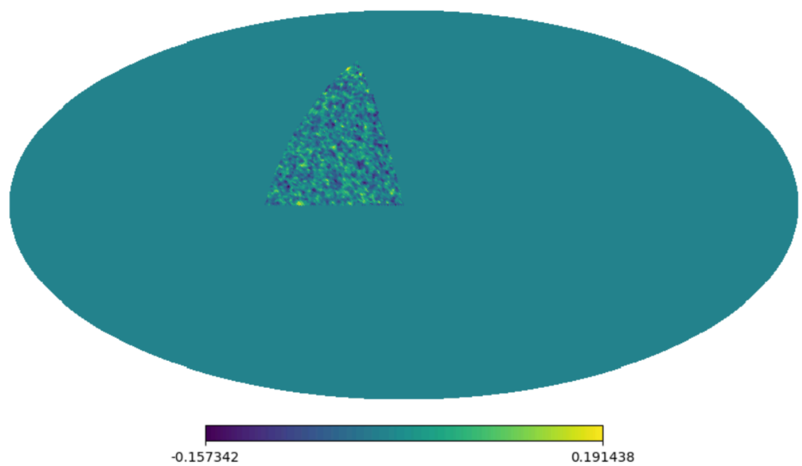

With the data curated into 20 samples of each of the 791 simulations (using the 0.2 redshift binned shells as separate channels of information for each sample), I moved on to training the neural network. Initially, I set up DeepSphere using the basic structure described by Nathanaël Perraudin, et. al. The training rate was set to 0.001, and the network utilized MAE as the loss function. The first training goal was, for each sample, to predict seven cosmological parameters–the same seven utilized in the creation of the Gower Street catalog: Ω_m, σ_8, w, h, n_s, m_ν, and Ω_b. After training for 50 epochs, using 90% of the data for training, 5% for validation, and 5% for testing, for most of the parameters the neural network was only able to predict the mean among all possible true values. However, this was not the case for Ω_m and σ_8; for these two, the network displayed a learning trend towards accurate predictions. In light of this, I narrowed down the training to target just these two parameters, altered the training rate to 0.0001, included a smoothing layer at the beginning, and increased the number of features extracted during the first layers of convolution. Overall, this brought about drastic improvement in predictions of both Ω_m and σ_8. Furthermore, I sought insight from the works of others, hoping to find additional ways to improve DeepSphere's performance. In an article by A. Thomsen, et. al. (arXiv: 2511.04681v1), a slightly different layer structure of DeepSphere was implemented in an application similar to my own. The structure they used is detailed in the table below.

Layer Type	Output Shape	# Parameters
Input	$(N_b, 458\,752, 8)$	0
Smoothing	$(N_b, 458\,752, 8)$	0
Pseudo convolution	$(N_b, 114\,688, 32)$	1 056
Pseudo convolution	$(N_b, 28\,672, 64)$	8 256
Pseudo convolution	$(N_b, 7\,168, 128)$	32 896
Chebyshev convolution	$(N_b, 7\,168, 256)$	163 840
Layer-normalization	$(N_b, 7\,168, 256)$	512
Pseudo convolution	$(N_b, 1\,792, 256)$	262 400
Chebyshev convolution	$(N_b, 1\,792, 256)$	327 680
Layer-normalization	$(N_b, 1\,792, 256)$	512
Pseudo convolution	$(N_b, 448, 256)$	262 400
Residual layer	$(N_b, 448, 256)$	656 896
	$\vdots$	
Residual layer	$(N_b, 448, 256)$	656 896
Flatten	$(N_b, 114\,688)$	0
Layer-normalization	$(N_b, 114\,688)$	229 376
Fully connected	$(N_b, 20)$	2 293 780

I augmented my structure to match theirs, and ran the training with the same settings I'd previously been using, still with the main goal of predicting Ω_m and σ_8. Below are the results of Ω_m predictions (post training) from the two different DeepSphere structures: original DeepSphere structure on the left, and the A. Thomsen, et. al. structure on the right.

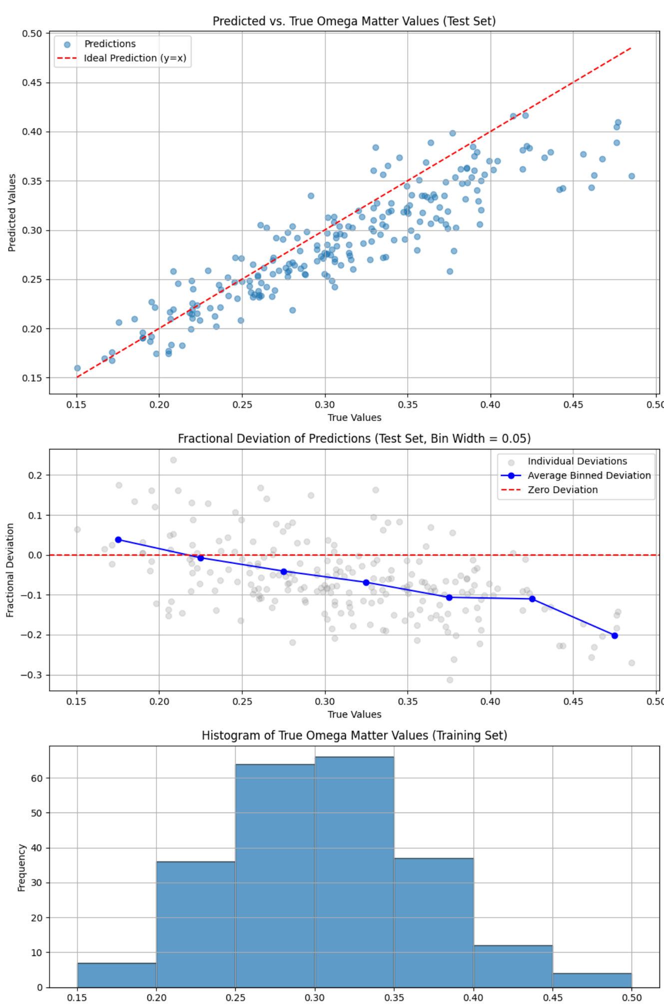

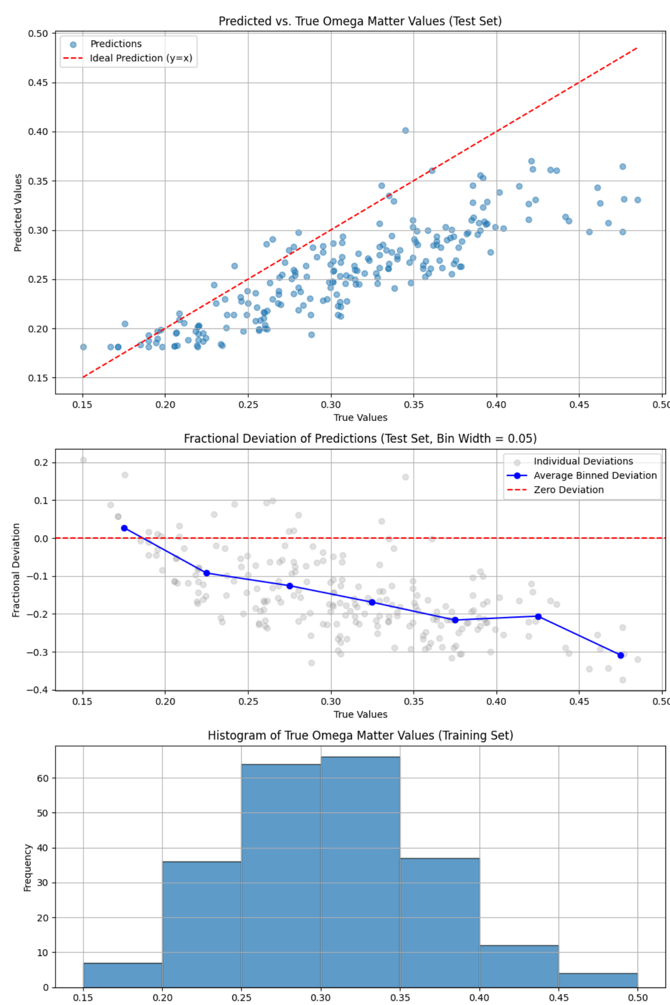

The top panels show a scatter of predictions compared to the true values; the middle panels show the deviation from the true values, with a trendline following the binned averages; and the bottom panels show a histogram of the true values included within the training sets in each case. At this point both DeepSphere structures were demonstrating learning of the trends within the maps. However, the A. Thomsen, et. al. structure demonstrated an increased bias towards lower predictions. There was a noticeable bias in both models at high Ω_m values; however, this can primarily be attributed to the reduced amount of training data available at those values. Nonetheless, because of the overall better performance of the original DeepSphere structure, I continued to use it moving forward.

At this point, the model was demonstrating the capability of inferring cosmological information from the map level. However, in order to properly carry out an analysis of the model's ability to infer the underlying posterior likelihoods for Ω_m and σ_8, both the structure and learning target (loss function) of the model needed to be altered. Within the same paper that introduced the Gower Street Simulations (N. Jeffrey, et. al. (arXiv: 2403.02314v1)) a method for Simulation Based Inference (SBI) utilizing a Masked Autoregressive Flow (MAF) was demonstrated. This became the groundwork for restructuring DeepSphere from the regression model I was using it as, to an inference model.

Additionally, my supervisor and I decided to implement ISW maps as part of the training and testing data for the neural network. This would serve as a novel application within the realm of cosmological map-level SBI. Thanks to the concurrent work of my team in their exploration of the ISW effect, ISW temperature maps had been created for each simulation within the Gower Street suite (Mina Ghodsi Yengejeh, et. al. (arXiv: 2512.24369)). I passed these through my data preparation pipeline to create an appropriate ISW temperature map to go along with the five redshift-binned density contrast maps for each simulation (each remained sectioned into 20 parts of the full sky).

DeepSphere became the first part of the entire model; its role was to derive a conditioning vector from the input maps (redshift-binned density contrast and ISW) which was then passed to the MAF. The MAF would then apply a series of transformations to a generic 2-D Gaussian distribution based on that conditioning vector. Finally, the loss function of the entire model was altered from MAE to Negative Log-Likelihood (NLL). This effectively trained the combined network to infer 2-D joint posterior likelihood contours between Ω_m and σ_8 from redshift-binned density data and ISW temperature information within only a 20th of the sky. These alterations provided a two-fold improvement in my approach. Not only did the model become more accurate in its predictions, but it also allowed for the eventual direct comparison to leading joint contour predictions from existing probes (Planck, DES, KiDS, etc.). Below are plots showing the improved prediction capabilities and contour inferences from the DeepSphere-MAF structure.

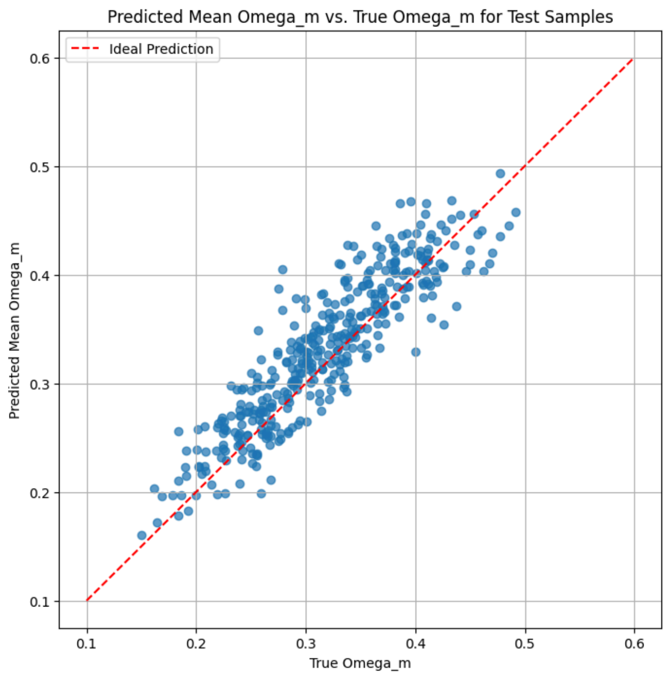

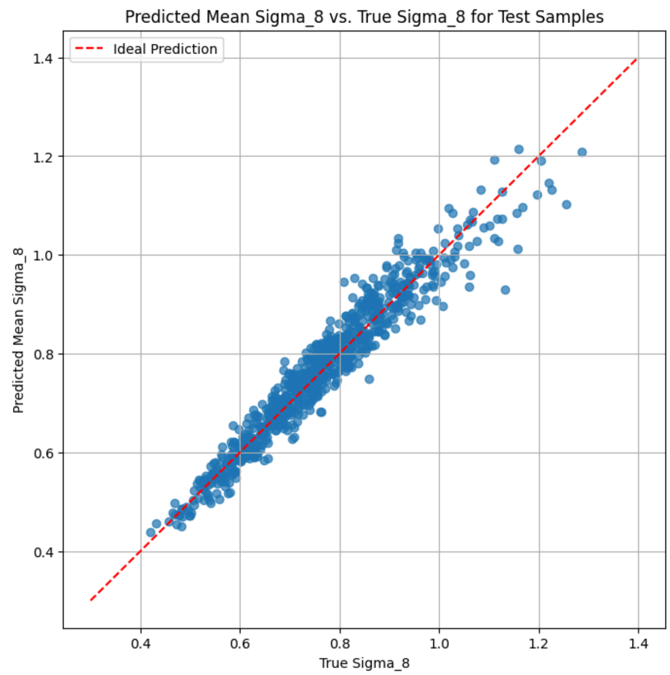

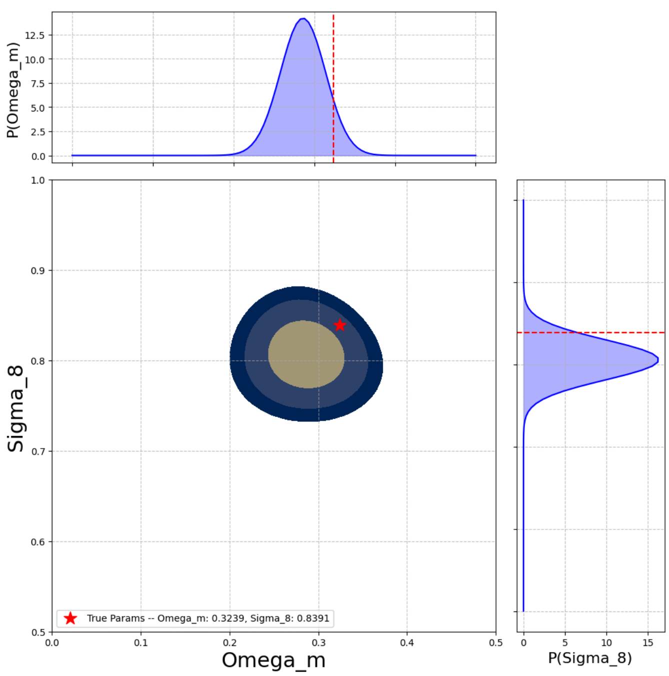

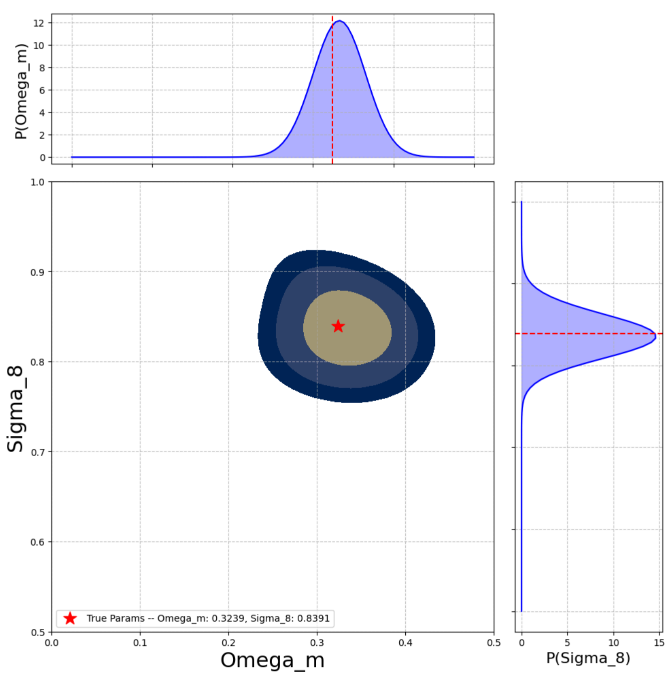

The top left and top right plots show the separate prediction accuracies of the means of Ω_m and σ_8, respectively. The bottom plots show a resulting prediction for the underlying joint posterior between Ω_m and σ_8 of the chosen fiducial cosmology–which was chosen as the simulation within the first Gower Street suite closest to ΛCDM: $\Omega_m = 0.3239$, $\sigma_8 = 0.8391$. The difference between the bottom two is the presence or absence of ISW information. The plot on the bottom left does not include the ISW maps within the data, and the plot on the bottom right does. Even though the inclusion of the ISW information does not currently display an obvious increase in constraints on cosmological parameters, it does indicate an improvement in accuracy of predictions by the model; most notably revealed by the reduction of bias present within the predictions. This is most likely due to the increase in correlated data between density contrast maps and ISW maps that can assist in inferences made by DeepSphere.

As this research into map-level SBI continues to progress, I have simultaneously been working with my team to apply its current level of success in further exploration of the ISW effect–with particular focus on cold spots and voids. By using the ISW maps they previously created, they have generated a suite of stacked void profiles (one for each simulation). These stacked profiles utilize a pipeline that identifies voids within the density contrast maps, locates and takes snapshots of them on the corresponding ISW maps, then stacks those snapshots into a single cold spot image, and finally measures a void profile from that stacked image. Ultimately these profiles come out to just 25 points used to form the profile curve, and these points can be used as the data for training an SBI model. Because they are much simpler than an entire map, the model used to derive the conditioning vector fed to the MAF does not need to be as complex as DeepSphere. Therefore, I utilized a simple Multi-Layer Perceptron (MLP) with three layers (input, hidden, and output) to derive the conditioning vector. Below are some graphs showing the current progress of this endeavor.

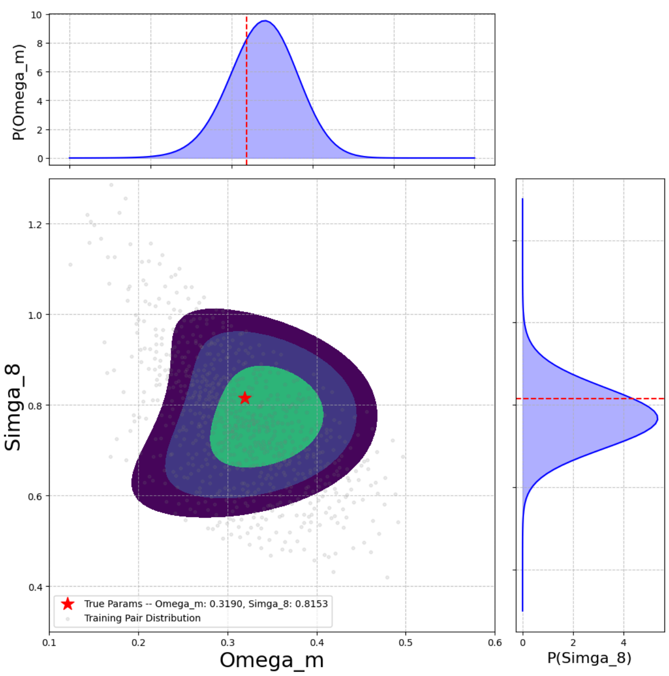

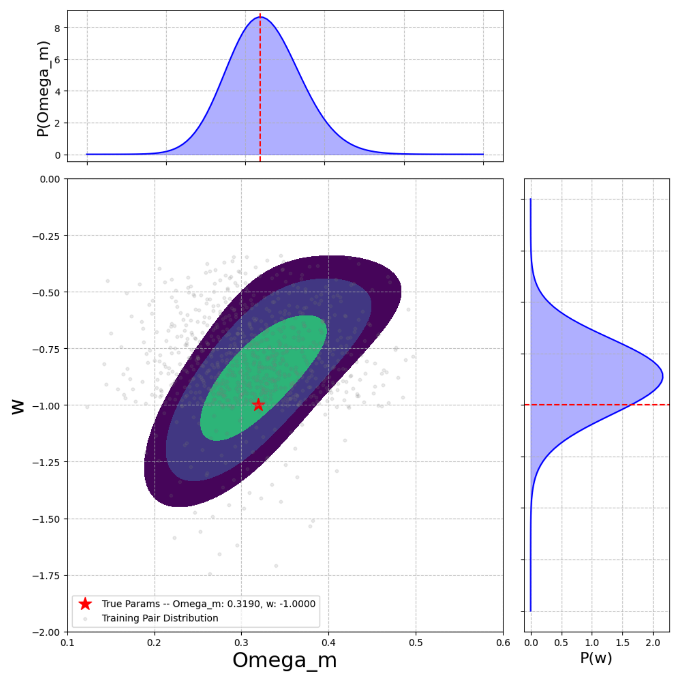

For this project I trained the neural network two separate times. First (as shown on the left), to infer the joint posterior between Ω_m and σ_8, and second (as shown on the right), to infer the joint posterior between Ω_m and w. In both cases the model was able to infer some existing degeneracies between the respective parameters, and, though it exhibited a small bias in each case, the fiducial cosmology was within the 1σ region of the predictions. Furthermore the constraints are within the range of those exhibited by other modern probes of cosmology. All-in-all, the paper for the results from this exploration was intended to be submitted for review by the end of May or early June; however, due to conferences and my impending exam, it needed to be postponed until mid-to-late June. I will be a co-author on this paper, and its utilization of my SBI approach on a reduced level will prove as a stepping stone toward the additional papers involving ISW map-level inference I intend to submit in the Fall and Winter.

Conclusion

Progress on my research has included the continual expansion of my own skills and knowledge, deeper exploration of GCNNs, incorporation of DeepSphere as part of an MAF in SBI, and an implementation of similar methodology in SBI utilizing void profiles. Looking forward, I intend to complete my part in the publication of my team's work in the analysis of void profiles next month, and then utilize that work as support for my continued progress in application of SBI on the map level. My goal is to complete at least two additional publications as the first author within the next year. The first will include a detailing of my methodology and the beneficial effects of incorporation of ISW maps and/or void information within the SBI application; and the second will primarily focus on the incorporation of Euclid-like effects within the data, and an improvement of the model for the eventual use on observational data from the Euclid survey.