Physics-Informed Variational Autoencoders for Cosmological Field Reconstruction and Parameter Inference

by

Labiba Zahin
22101114
Zarin Tasnim
22101249
Tasnim Zaman
22101343
Mehrabul Islam
24341179
Safkat Zahir
22101414

A thesis submitted to the Department of Computer Science and Engineering
in partial fulfillment of the requirements for the degree of
B.Sc. in Computer Science and Engineering

Department of Computer Science and Engineering
BRAC University
January 2026.

Declaration

It is hereby declared that:

1. The thesis submitted is our own original work, completed while pursuing the degree at Brac University.
2. The thesis does not contain material previously published or written by a third party, except where this is appropriately cited through full and accurate referencing.
3. The thesis does not contain material that has been accepted or submitted for any other degree or diploma at a university or other institution.
4. We have acknowledged all main sources of help and support.

Student's Full Name & Signature:

Labiba Zahin

22101114

Zarin Tasnim

22101249

Tasnim Zaman

22101343

Mehrabul Islam

24341179

Safkat Zahir

22101414

Approval

The thesis titled "Physics-Informed Variational Autoencoders for Cosmological Field Reconstruction and Parameter Inference", submitted by

1. Labiba Zahin (22101114)
2. Zarin Tasnim (22101249)
3. Tasnim Zaman (22101343)
4. Mehrabul Islam (24341179)
5. Safkat Zahir (22101414)

of Fall 2025 has been accepted as satisfactory in partial fulfillment of the requirements for the degree of B.Sc. in Computer Science and Engineering in January 2026.

Examining Committee:

Supervisor:
(Member)

Dr. Muhammad Iqbal Hossain

Associate Professor
Department of Computer Science and Engineering
Brac University

Co-Supervisor:
(Member)

Ms. Ipshita Bonhi Upoma

Lecturer
Department of Computer Science and Engineering
Brac University

Thesis Coordinator:
(Member)

Md. Golam Rabiul Alam, PhD

Professor
Department of Computer Science and Engineering
Brac University

Head of Department:
(Chair)

Dr. Sadia Hamid Kazi

Associate Professor and Chairperson
Department of Computer Science and Engineering
Brac University

Abstract

In modern cosmology, predicting cosmological parameters is key to understanding the fundamental physical laws of the universe and dictating how cosmic structures are formed, evolve, and are observed. Parameters such as the total matter density (Ω_m) and the amplitude of matter fluctuations (σ_8) cannot be measured directly; instead, they have to be predicted based on complicated, high-dimensional simulation maps or observational data. Since, due to the complexity of the high-dimensional maps, traditional deep learning models often fail to give meaningful results, as they often learn shortcuts to statistical patterns that may appear correct but completely ignore the actual laws of physics. In this work, a Physics Informed Variational Autoencoder (PI-VAE) has been proposed as a combined framework for learning the compact representations of cosmological fields while directly imposing fundamental physical constraints to accurately reconstruct multi-channel cosmological fields, and directly infer key cosmological parameters from the learned latent space. By attaching a lightweight parameter regression head with VAE, the research looks into how physical information stored in the latent representation can be used for parameter predictions. To conduct this research, the CAMELS (Cosmology and Astrophysics with Machine Learning Simulations) dataset has been used, which comprises thousands of hydrodynamical simulations intended to systematically vary cosmological and astrophysical parameters. Our findings, supported by the ablation experiments, highlight the potential of physics-guided deep generative models for cosmological analysis by showing that physics-informed latent representations can simultaneously achieve meaningful cosmological parameter inference and accurate field reconstruction.

Keywords: Cosmological parameters, matter density, amplitude of matter fluctuations, Physics informed variational Autoencoder, Camels, field reconstructions, VAE.

Acknowledgement

First and foremost, we would like to express our sincere gratitude to the Almighty, the provider of knowledge and strength, whose blessings helped us to complete this journey. It is our great pleasure and honour to thank our supervisor, Dr. Mohammad Iqbal Hossain sir, and our Co-Supervisor, Ms. Ipshita Bonhi Upoma ma'am, for their constant, unwavering support and advices that have helped us throughout this journey. We would also like to thank our parents, who consistently supported and encouraged us throughout the completion of this work.

Table of Contents

List of Figures

List of Tables

Nomenclature

The following list describes several symbols & abbreviations that are used throughout the document.

AE Auto-encoder

AGN Active Galactic Nucleus

BNN Bayesian Neural Network

CMB Cosmic Microwave Background

CNF Continuous Normalising Flows

$CVAE$ Convolutional Variational Auto-encoder

$DDPM$ Denoising Diffusion Probabilistic Models

DL Deep Learning

DNN Deep Neural Network

$E-PINN$ Evidence Physics-Informed Neural Network

$ELBO$ Evidence Lower Bound

$LensPINN$ Lensing Physics-Informed Neural Network

MAE Mean Absolute Error

ML Machine Learning

MLP Multi Layer Perceptron

ODE Ordinary Differential Equation

PI-VAE Physics-Informed Variational Auto-encoder

$PINN$ Physics-Informed Neural Network

SBI Simulation-Based Inference

$SimCLR$ Simple Framework for Contrastive Learning of Visual Representations

SML Statistical Machine Learning

SSL Self-Supervised Learning

SVM Support Vector Machine

TNN Topological Neural Network

VAE Variational Auto-Encoder

$VDVAE$ Very Deep Variational Auto-encoder

VI Variational Inference

Chapter 1

Introduction

The universe is organized into a complex and vast network known as the cosmic web. This structure is formed with massive filaments of dark matter and gas, dense nodes where the galaxy clusters locate and nearly empty regions called cosmic voids. This formation is driven by the gravitational collapse of dark matter. Understanding the geometry of the universe is really essential as they hold the information about evolutionary history. To overcome the complexity of non linear information embedded with high dimensional cosmic fields, the CAMELS (Cosmology and Astrophysics with Machine Learning Simulations) project was established as a foundation framework [3]. This project includes over 1000 hydrodynamic and dark matter only simulations using astrophysical codes such as AREPO and GIZMO. The CAMELS project is a large suite of many numerical simulations which are designed to study the parameters and baryonic physics. These simulations are not learned by neural networks. The simulation begins at a very high redshift (early universe), where matter density fluctuations are very small. These numerical simulations solve physical equations. Two of the main simulation categories are N-body simulations and Hydrodynamical simulations. N-body simulations include dark matter only and Hydrodynamic simulations include dark matter , gas and baryonic processes.
The result of these simulations vary two cosmological parameters (e.g., Ω_m,σ_8) and four astrophysical parameters (e.g., ASN1, ASN2, AAGN1, AAGN2). These cosmological parameters are crucial for describing the formation, structure and matter distribution of the universe, galaxies and clusters, dark matter etc. The cosmological parameter Ω_m helps us to quantify the total matter content of the universe including dark matter and baryonic matter which explains the overall gravitational pull and expansion rate of the universe. On the other hand , σ_8 helps us to measure the degree of matter clustering in the 8 h^{-1} Mpc scale. Moreover, data compression or dimensionality reduction is essential to translate these high dimensional spaces into a more manageable and lower dimensional space. In this way, the noisy part of the image can be easily filtered out. It helps to bring out the most important physical features which makes it easier to understand the structure of the universe. While predicting parameters and compression is essential, reconstruction of the map also plays a vital role to ensure the proper integrity of the learned feature. It guarantees the ability of a model to generate the original data from its compressed form.
The observable universe is made up of gas and star (baryonic matter) whereas dark matter provides the gravitational backbone Hydrodynamic simulations like IllustrisTNG are used to model the 'baryonic physics' such as gas cooling , star

formation and feedback from supernovae and Active Galactic Nuclei (AGN). All these processes include non linearity and noise in the data,which makes it harder to predict and interpret cosmological parameters.Although the CAMELS project generated thousands of simulations with many varying parameters, it is still difficult to predict these parameters and meaningful reconstructions using standard Convolutional Neural Networks (CNNs) because of their use of shortcut methods and overfitting. They often fail to extract meaningful results because of their violation of physics laws and shortcut ways. Therefore, a model integrated with fundamental physics laws, which can simultaneously compress , reconstruct cosmological fields and predict cosmological parameters within a single pipeline is a need in modern cosmology to interpret complex simulation data effectively and accurately.

1.1 Motivation

In the last few decades, there has been a significant increase in the research sector and the applications of ML and deep learning methods to cosmological data. It has enabled efficient extraction of cosmological parameters and simulation of complex astrophysical events [3]. The CAMELS project stands as the backbone in bridging cosmological simulations. This project upholds such significance because of the use of large-scale, multi-physics simulation datasets like SIMBA and IllustrisTNG, which cover a large range of cosmological and astrophysical parameters.
Extracting these meaningful parameters such as Ω_m and σ_8 from these high-dimensional datasets is a challenging task [14]. These datasets are complex and high-dimensional, making it difficult to extract meaningful parameters using traditional methods. A standard CNN or deep learning model on such large-scale data often leads to underperformance due to the complexity of the data as they are more prone to shortcut learning and overfitting. Therefore, predicting parameters like Ω_m and σ_8 from these datasets needs advanced techniques because these parameters are not directly observable and require sophisticated analysis to infer from the data. Ω_m helps us to quantify the total matter content of the universe including the dark matter and baryonic matter. On the other hand, σ_8 helps us to measure the degree of matter clustering in the 8 $h^{-1}\,\mathrm{Mpc}$ scale. As the parameters are encoded with complex, non-linear relationships of gas density, temperature and velocity fields, they need to be predicted using advanced physics-based models. To interpret accurate cosmological parameters, we propose a Physics-Informed Variational Autoencoder (PI-VAE), which basically compresses the high dimensional data into a low dimensional latent space of 256 dimensions, reconstruct the images from the compress vector and predicts the parameters by ensuring the physics constraints. The main goal behind this approach is to create a latent representation which is also physics aware. It separates the cosmological signals from the noise that was inherited in the baryonic process. Our aim is to find out whether a compressed latent space generates more physically consistent latent representations compared to standard CNNs. By evaluating the parameters, we can separate the cosmological signals from baryonic feedback. Our approach enhances the accuracy of predicting the parameters. Accurate and precise prediction helps us to understand the geometry of the universe. These predicted parameters can open the way for more interpretable deep learning applications in high-dimensional astrophysics which can also offer a robust methodology for large-scale parameter estimation in extensive cosmological simulations.

1.2 Problem Statement

The universe has such randomness that it allows us to witness a little part of a larger cosmic structure, which includes dark matter and dark energy, comprising roughly 95% of the universe. Despite decades of research and investigation in the field of cosmology and astronomy, we still struggle to understand the universe's basic architecture using progressively complex simulations and data-driven methodologies. With the passing time, our simulations have become increasingly realistic, and our dataset extends to terabytes of high-dimensional astrophysical fields. This field has emerged with a new challenge which is the crushing burden of computation. The CAMELS project contributed the most by producing thousands of cosmological simulations with different parameters to facilitate inference based on ML. While deploying traditional deep learning models, especially standard CNNs often struggle with these high-dimensional datasets. They are more prone to shortcut learning as the model picks the local features rather than focusing on the underlying cosmological features. Generally, these long-scale data require a lot of computational resources and time for pre-processing. Moreover, standard Variational Autoencoders (VAE) need large latent space to ensure the accuracy which makes the representation difficult to interpret. It also makes the result sensitive to noise from baryonic processes.
To address these limitations, we suggest Physics Informed Variational Autoencoder (PI-VAE). The framework of this model makes sure of latent representation learning by compressing the inputs. The goal is to preserve the critical information while compressing high dimensional cosmic maps to accurately predict parameters. Our model aims for parameter inference which not only compresses and reconstructs but also is more physically interpretable. By redesigning the autoencoder backbone, we aim to provide a more robust and physically interpretable framework that meets the needs of future simulation-based cosmology.

1.3 Objective

This study aims to develop the Physics-Informed Variational Autoencoder framework for compression, reconstruction and accurate parameter inference by preserving the important and meaningful feature. This research addresses the issue of exacting meaningful physical signals from high dimensional, complex and non - linear baryonic data.
The key goals of this research are outlined below:

1. Designing a pipeline for data pre processing for the multi-channel inputs (M_{gas}, V_{gas}, B) to handle the large dynamic range. Including logarithm scaling and normalization of multi - modal inputs for stability.

2. Dimensionality reduction of the data into a 256 latent vector for manageable representation and using a U-net based decoder with skip connection to increase the quality of reconstruction.

3. Proposing a Physics Informed VAE model which specifically focuses on integrating fundamental physics laws to compress and reconstruct 2D cosmic maps along with parameter inference .

4. Applying the latent vector as input from the encoder for a parameter regressor to ensure the precise prediction of Ω_m and σ_8 is based on the physical feature which can be analyzed later.

5. Incorporating physics based loss functions which includes Power spectrum , Gradient , amplitude losses etc. to ensure that the model does not violate the physics constraints and learn the physical correlation between Mass distribution , Gas velocity and Magnetic pressure while predicting the parameters.

6. To compare PI-VAE with U-Net VAE and classical Variational Autoencoders in terms of preserving the physical integrity of cosmological fields, with a particular focus on how physics-informed skip connections improve the reconstruction of small-scale cosmic structures

1.4 Methodology in Brief

This study proposes a Physics-informed Variational Autoencoder (PI-VAE) and compares it with the U-net Variational Autoencoder (VAE) and a VAE without skip connections. We used the IllustrisTNG suite from the CAMELS dataset, which provides high resolution hydrodynamic simulations. These simulations carry two cosmological parameters (Ω_m, σ_8) and four astrophysical feedback parameters. Our study focuses on extracting the cosmological parameters from the 2D compressed maps.
We use 3 channel stacks as input named Gas Mass Density (M_{gas}), Gas Velocity (V_{gas}) and Magnetic Field (B). The raw data is preprocessed using channel-wise transformation. Logarithmic scaling was used on M_{gas} and B while arcsinh transformation was used on V_{gas} because of the signed values it contains. Then the data was augmented and normalized. After preprocessing, a total of 15000 maps from the LH set of IllustrisTNG suite was split into train, validation and test set following the ratio of 3:1:1.
Our model, PI-VAE, takes the 2D maps as input which includes 3 channels and the encode compresses it into a 256 - dimensional latent vector. It effectively filters out the baryonic noise while preserving import features. The latent vector then works as an input for the param regressor which is designed to predict the values of Ω_m and σ_8. Param regressor is a Multi -layer Perception (MLP) which takes the compressed latent vector as input from the encoder.
The model incorporates skip connections between the encoder and decoder which ensure that the high resolution spatial features are preserved and passed to the decoder. As sometimes the spatial features are lost in the process, this architecture allows to preserve it and it helps to produce more accurate and consistent reconstruction.
For the loss function, we have used Reconstruction loss, KL divergence loss and four physics loss. The physics loss includes the Power Spectrum loss, Gradient loss, Amplitude loss and Mean Density loss. These physics losses ensure that PI_VAE does not violate the statistical and structural properties of the universe by penalizing the model. The encoder learns that to minimize the total loss , it has to extract the features that are relevant to these physical laws. Therefore, when predicting the values of Ω_m or σ_8, the regressor is not focusing on some random noise; instead it is

capturing the underlying energy, gradients and mass distribution of the simulated data.
We evaluated the result using the Coefficient of determination (R2) , Mean Square Error (MSE) and Mean Absolute Error (MAE) in terms of parameter inference. To measure the physical consistency, we used Power Spectrum Error, Mean Density Error, Gradient Density Error and Relative Error Metrics.

1.5 Scopes and Challenges

Scopes:

1. **Cross-Suite Robustness and Generalization:** Accessing the models performance across different CAMELS suites such as SIMBA to test the robustness.

2. **Physical Interpretability of Latent Space:** Interpreting from the predicted values as the latent vector represents the physically plausible universe and understanding the physics behind it.

3. **Real-Time Inference Pipeline:** Establishing a pipeline for real time estimation of Ω_m and σ_8 from high-fidelity hydrodynamic simulations.

4. **Statistical Consistency and Power Spectrum Fidelity:** Testing the model's ability to maintain Power Spectrum ($P(k)$) consistency, ensuring the reconstructed maps are statistically indistinguishable from real simulations.

5. **Multi-Modal Degeneracy Reduction:** Investigating how the combination of M_{gas}, V_{gas} and B -fields reduces parameter degeneracy compared to a single field.

Challenges :

1. **Loss Function Balancing:** The difficulty of tuning hyper-parameters to prevent one loss term from dominating the others during training.

2. **Hardware Limitations :** Training a multi-modal PI-VAE on 256×256 3-channel maps requires substantial GPU memory and computational time, especially when calculating complex terms like the Power Spectrum.

3. **Conservation Law Enforcement:** ensuring the model properly learns the law without guessing it or using shortcut methods.

4. **Model Generalization and Domain Shift:** Assessing robustness across different CAMELS suites (like SIMBA vs. IllustrisTNG) can be a challenge, as different sub-grid physics models present a "domain shift" where the model might fail to generalize.

5. **Dataset biasness:** managing the massive CAMELS data and ensuring the training samples are not biased towards a specific parameter is essential and challenging for a reliable result.

1.6 Thesis Structure

Thesis Organization

Chapter 1 has our introduction, including motivation, problem statement, objective, methodology, scopes and challenges and thesis structure.
Chapter 2 contains preliminaries, review of existing research and summary of key findings.
Chapter 3 covers the hardware/software specifications, societal and environmental impacts, and risk management.
Chapter 4 discusses the methodology including data pre-processing and the architecture of encoder-decoder with U-Net skip connections and hybrid loss functions incorporating physics laws.
Chapter 5 analyzes the quantitative and visual results compared to baselines.
Chapter 6 summarizes the findings, acknowledges the limitations and proposes future research.

Chapter 2

Literature Review

2.1 Preliminaries

Cosmology and Astrophysics with Machine Learning Simulations (CAMELS): It is a massive suite of thousands of hydrodynamic simulations for the purpose of training machine learning models. It contains different versions of the universe generated with different cosmological parameters, as well as different astrophysical parameters (like supernova feedback). In this thesis, CAMELS is the "ground truth" data required to teach the VAE what it needs to know in order to relate properties of galaxies to the underlying physics of the universe.
Power Spectrum : It is the old mathematical tool of Cosmology to quantify the structure of the universe. It measures the extent of which matter is crystallized together, on varying large-to-small scales – from vast clusters of galaxies, small voids. While effective, it is liable to miss complex, non-Gaussian patterns (like filaments or walls). Deep learning models are frequently compared with the Power Spectrum to determine if it can extract more information than this standard statistical method.
Generative Adversarial Network (GAN): It is a deep learning architecture in which two neural networks are competing against each other. A Generator attempts to generate a set of realistic fake data (such as an image of a galaxy) from random noise, while a Discriminator attempts to differentiate between the real data as well as the fake data. The Generator gets so good with time that the Discriminator is unable to discern the difference. It is true that GANs are very difficult to train compared to VAEs because they lack an explicit objective function to minimize.
Self-Supervised Learning (SSL): It is a mode of training in which unlabelled data is used to train a model with the help of self-supervised learning that is its own supervisor. One common approach here is so-called "contrasting learning," where the model should learn to realize that two different cropped views of the same galaxy are "the same," whereas views of different galaxies are "different." This way, the network can learn useful features about the structures of galaxies without a human being having to manually label the images - hundreds of millions of them.
Simulation-Based Inference (SBI): Also called "Likelihood-Free Inference," this is a statistical method which is used when the physical processes are too complex to write them down as a simple equation (likelihood function). Instead of computing probabilities by doing the mathematics, a team of researchers performs a large number of forward simulations with varying parameters and trains a neural network to learn the relationship between the inputs of the simulations (the parameters) and

the output of the simulations (the observations).
Support Vector Machine (SVM): SVM is a classical supervised machine learning algorithm, it is used for classification and regression. The mechanism of it is identifying the optimal "hyperplane" (boundary) to best separate the data points of different classes in a high-dimensional space. While less complex than deep neural networks, SVM is still very good for the particular tasks, such as the detection of outliers or classifying simple morphological features in smaller data sets.
Normalizing Flows: It is a generative model that has learned how to transform a simple probability distribution (such as a bell curve) into a complex one (such as matter distribution in the universe) using a series of invertible mathematical functions. Unlike GANs or VAEs which basically approximate the distribution of the data, Normalizing Flows have the ability to compute the exact likelihood of data. This makes them very accurate but heavier to compute in order to train them.
Bayesian Neural Network (BNN): It is a type of neural network where the "weights" (connections between neurons) are not a set of fixed number, but probability distributions (range of possible values for a number). This is to permit the network to express uncertainty. Instead of just making a single prediction of a cosmological parameter, a BNN will make a prediction of a range (e.g. "0.3 0.05"), telling us how confident it is in its prediction.
Dark Matter Halo: It is the invisible, massive structure of dark matter that surrounds the galaxies. Although we cannot see halos themselves, they exert gravity on clumping clouds of gas and dust to create visible galaxies as we see them. In simulations the properties of the dark matter halo (mass, size) is often used as input to predict the properties of the galaxy inside this halo.
Epoch: Under the context of training neural networks, the epoch is what is passed, the whole training set once is what we call epoch. If you have 10,000 galaxy images, one epoch means that when the network has seen all the images, i.e. the 10 000 images, 1 time. Models are normally trained for a large number of epochs to gradually reduce the error.
Machine Learning (ML): It is an area of artificial intelligence in which computers are made to learn data without being programmed for any rules. Instead of taking instructions and then following them, ML models rely on statistical methods to detect patterns in historic data in order to make decisions or predictions. It is broadly classified into supervised learning (training using the labeled data), unsupervised learning (finding the hidden structure in the unlabeled data) and reinforcement learning.
Deep Learning (DL): It is a specialized subset of machine learning that is inspired by the structure of the human brain. It makes use of a multi-layered Artificial Neural Network (ANN) to extract higher and higher level features from the raw input. For instance, in the field of image processing, lower layers could be used to recognize edges, while the deeper layers could be used to recognize complex objects such as galaxies. DL is the basis for such modern tools as computer vision and generative modeling.
Convolutional Neural Network (CNN): It is a deep learning network that is specialized in handling grid-like data (images). Unlike ordinary neural networks, which have an independent pixel, CNNs take sliding filters (kernels) to scan the image and identify spatial hierarchies such as textures, curves, and shapes. In cosmology, CNNs are the go-to tool in the analysis of galaxy morphologies and the

discrimination of stars and galaxies.
Variational Autoencoder (VAE): It is a generative model that is used to learn a compressed representation of data. It consists of 2 parts i.e., Encoder which compresses the input data into a lower dimensional probability distribution (latent space) and decoder which completes the task of reconstructing the original data from the probability distribution. Unlike the usual autoencoders, VAEs guarantee that the latent space will be continuous and smooth, which is the best scenario for generating fresh and worthy scientific data or even inferring physical parameters.
Physics-Informed Neural Network (PINN): It is a type of deep learning system that incorporates physical laws (in the form of differential equations) into the loss function of the neural network. Instead of learning pure physics data, the model is penalized if its predictions go against known laws of physics such as the conservation of energy or fluid dynamics. This enables the model to be able to learn accurately even with sparse data and makes sure that the results are scientifically consistent.
Latent Space: It is the compressed and lower representations of data, which is located in the "bottleneck" layer of an auto encoder. It is effective at mapping high dimensional data (as a telescope image with millions of pixels) to a few meaningful variables. In your thesis, the latent space is used to separate the visual style and the physical coordinates.
Evidence Lower Bound (ELBO): It is the Loss function in Variational Autoencoders. It balances two conflicting objectives, the reconstruction loss (how good the output is compared to the input) and the KL Divergence (pooking the latent space neatly ordered in the form of a Gaussian distribution). Optimizing the ELBO makes it so that the model produces a structured latent space that is useful for sampling.
U-Net: It is a special neural network architecture which was initially created for biomedical image segmentation. It consists of a "U" shape with an encoder path that allows for the acquisition of context and a symmetric decoder path that allows for excellent localization. It makes use of skip connections to directly pass high resolution features from the encoder to the decoder in order not to lose the fine details (such as thin spiral arms of galaxies) during the compression process.
Denoising Diffusion Probabilistic Models (DDPM): It is a class of generative models, which learns to generate data by reversing a progressive noise process. The model destroys the training data (by adding Gaussian noise) until it seems random still, and then learns to do it back to put it all together into a clear image. While they generate images of higher quality than the VAEs, they are computationally expensive due to the fact that they require many iterative steps to produce a single sample.
Hydrodynamic Simulations: This is a complex computer simulation that is used to model the evolution of the universe. Unlike the more basic "N-body" simulations that only follow the physics of gravity and dark matter, the hydrodynamic simulations (such as the CAMELS suite) are also able to simulate the physics behind the gas (baryonic matter), such as star formation, supernovae feedback and black hole activity. These provide ground truth data that are required to train machine learning models.
Cosmological Parameters: These are the basic numbers that are used to describe the history and composition of the universe. In this thesis, we are interested in Matter Density which describes the total amount of matter in the universe, and Sigma-8 () which describes the "clumpiness" or density fluctuations of matter. Accurately

inferring these values is one goal of modern-day cosmology.
Redshift : It is a measure of the extent to which the wavelength of the light emitted by a distant object has been extended as a result of the expansion of the universe. The higher the redshift number the further the object is and the further back in time it is. In the case of machine learning tasks, redshift is usually a physical label that is used to condition the latent space.

2.2 Review of Existing Research

The study by Gagliano et al. [9] deals with the computational inefficiency of the standard method of SED fitting for massive datasets like the upcoming LSST, whose contents are to be billions of galaxies. In order to address this, the authors created a Physics-Informed Convolutional Variational Autoencoder (CVAE) that separates the latent space into physical parameters, like redshift and stellar mass and SFR, and morphological features. The way the architecture works is that they divide the bottleneck layer into two separate vectors, 'Physics' vector which is explicitly supervised by known labels such as redshift, and an unsupervised 'style' vector for morphological features. This forced disentanglement is achieved by adding a regression penalty to the standard ELBO loss which forces the encoder to structure the latent space by enforcing the physical laws, rather than simply the visual resemblance. It was trained on 26,000 galaxies of the DECaLS survey ($z < 0.4$) , the model has a composite loss function based on reconstruction error, KL divergence and physical MSE term. The CVAE proved remarkable when compared to simulation based inference, with deduction speeds of 7ms per galaxy per M2 chip (more than 100 times). In comparison, quantitatively it was better than self-supervised baselines with $R^2 = 0.83$ for redshift and 0.75 for stellar mass. In addition, an SVM in latent space was able to recover 94% of anomalous 'Red Spiral' galaxies. However, the model had problems with the rate of star formation (R^2=0.26) because of photometric degeneracies and failed to reconstruct the orientation of galaxies at high inclinations.

The study [18] deals with the problem of proper classification of galaxy morphologies in the huge image surveys in which complex structures such as the bars and spiral arms are usually lost by ordinary compression methods. The author Mirzoyan [18] proposes a U-Net Variational Autoencoder (U-Net VAE), which is a hybrid model combining the skip connections of a U-Net and the generative latent-space of a VAE. With the data from Galaxy Zoo 2 that was derived from the Sloan Digital Sky Survey (SDSS), they had to train a model on how to reconstruct galaxy images with a compact feature representation. The model modifies the standard VAE bottleneck by the addition of 'skip connections' between encoder and corresponding decoder layers, in a similar way from a ResNet. This way, high frequency spatial information (such as thin spiral arms) can completely skip the compression stage, this way fine-grained information will not be lost in the final reconstruction, while the latent space still learns around the global features. The results of the study revealed that the U-Net VAE is significantly superior to standard VAEs and CNNs with a classification accuracy of 92.5% proven on the baseline VAE, which was shown to be 88.3%. Furthermore, the produced model had the reduced Mean Squared Error (MSE) of 0.012 in image reconstruction results that effectively preserved fine-grained details that are crucial for morphological analysis. A limitation mentioned was the increased computational complexity as a result of the skip connections, which resulted in a slightly increased training time.

The problem of "inpainting" Cosmic Microwave Background (CMB) maps, where observational masks (due to galactic foregrounds or point sources) cause missing data that prevents analysis of the power spectrum, is addressed by the study by Yi et al. [2], The authors present CosmoVAE, a Convolutional Variational Auto encoder based on reconstructing the relevant masked regions by learning the statistical distribution of CB data. The training process is a 'context prediction' problem where the encoder is asked to learn to map the pixels on the boundary positions that are seen by the image, to a latent image of pixels that are missing from the middle of the image. The loss function is weighted such that the center of the image is given higher priority, and causes the decoder to hallucinate the missing CMB data using the statistical patterns that were learned from the unmasked areas of the training set. Trained on simulated maps of the CMB making assumptions about the global cosmology determined by the Planck 2018 satellite, the model is a standard Variational Auto-Encoder using convolutional layers. Key findings include the fact that CosmoVAE is able to recover the angular power spectrum quite precisely and with an error of less than 2% for multipoles $\ell < 1000$ and less for 20x20 degree patches. The method is computationally efficient where the generation of the inpainted maps takes milliseconds as compared to the traditional diffusive inpainting methods that take minutes. However, the performance of the model deteriorates when used at very small scales (high ℓ) because of a blurring effect inherent to the VAE reconstructions (the "L2 loss problem").

Calvo [1] studies the problem of objectively classifying the cosmic web (clusters, filaments, walls, voids) in N-body simulations, which in the past has been a problem which was solved using heuristic and thresholds based algorithm. The research works using a Deep Convolutional Neural Network (CNN) which has been taught in reference to density fields in the geometric "Cosmic Spine" formalism. The methodology consists of a first U Network like architecture unites or divides 3D density fields into structural components. The way the network does this is to treat the 3D simulation box as a volumetric segmentation problem, where 3D convolutional filters are used to scan the structure of the density field looking for geometric patterns. It assigns a probability vector for each individual voxel (3D pixel) within the grid, which is used to classify it as 'void', 'wall' or 'filament' as a function of the local gradients and eigenvalues of the density field (essentially automating the visual intuition of a human cosmologist).The model presented some awesome results with the global classification accuracy of more than 90% across all structural types. Specifically, it was able to identify 98% of voids accurately and 95% of filaments accurately which is a lot better than other traditional hessian-based methods such as the Nexus+ algorithm in speed and robustness to noise. One of its major weaknesses is that the model relies on the particular "ground truth" definition employed for training (in the case of Cosmic Spine), which is itself a theoretical construct. This study shows the use of deep learning to replace complex analytical frameworks in cosmology and opens the door to using VAEs for similar structure identification tasks in simulation data.

The study conducted by Lazaro et al. [10] addresses the complex inverse problem of simultaneously inferring cosmological parameters and reconstructing the initial density field from late-time observations. A Variational Inference (VI) framework, which approximates the distribution of both the initial conditions and the cosmological parameters (specifically Ω_m), is proposed by the authors. The method relies on a differentiable forward model (called particle mesh simulation) in a probabilistic framework by effectively converting the inference to an optimization problem. The key to this method is a "differentiable simulator," which is a physics engine written in a framework (such as TensorFlow or JAX) that supports the calculation of gradients through the simulation steps. This allows backpropagation not only for updating the weights of the neural network, but also to update the initial density field itself iteratively until the output obtained by simulation matches the actual observations. Tested on 2D N-body simulations, the method was able to recover the correct value of Ω_m with a relative error of less than 5%, and recover the initial density field with a cross-correlation coefficient $\rho > 0.9$ at large scales ($k < 0.1\ h\,\mathrm{Mpc}^{-1}$). A major disadvantage is the high computational cost of the differentiable simulation during training and the restriction to 2D simplified models at this proof-of-concept stage.

Another research by Ting [20]. provides a survey of the probes into the advent of Deep Learning (DL) across various astrophysical fields to address the need to process the petabytes of data produced by modern-day observatories. The paper categorizes methodologies including CNNs used for image analysis, RNNs used for time-series (e.g. light curves) and Generative models (GANs, VAEs) used for simulation acceleration. It mentions some of the special cases, such as determination of gravitational waves by CNNs with 99% accuracy, or particle shower simulations using GANs 105 faster than Monte Carlo methods. It reveals that with speed and precision comes a deficit of interpretability (the "black box" problem) along with the need for massive labeled datasets which are not common in astronomy. The review is especially useful for verifying that although discriminative models (CNNs) are mature, generative models (VAEs) for parameter inference are an emerging frontier with significant potential for handling physical constraints.

The review by Ramos et al. [6] investigates the need for Machine Learning (ML) in analysing the huge volume of data produced by the solar missions such as SDO and IRIS. This is attempted from the need to automate the detection of solar events (flares, CME's) which are too voluminous to inspect manually. The paper includes description about application of Support Vector Machines (SVM), Random Forests, and Neural Networks for flare prediction and segmentation of coronal holes. Some of the important metrics mentioned are that the SVMs obtained True Skill Statistics (TSS) scores of approximately 0.7 for flare forecasting, which is better than the traditional methods based on Poisson regression. However, the study notes that the models have problems with the stochastic nature of solar activity and the "class imbalance" problem in which severe activity is rare. Although focused on solar physics, this paper is relevant as it is a comparative study of ML utility in astrophysics and reinforces the argument that standard ML methods often need to be adapted to the specific domain (i.e., physics-informed constraints) to be able to handle scientific data effectively.

A study done by Andrianomena and Hassan [5] focuses on the interpretability in latent spaces. Here the Very Deep Variational Autoencoder (VDVAE) framework is proposed to extract features from high-dimensional data, such as images of gas density, neutral hydrogen, and magnetic fields. The method relies on a hierarchical model with multiple stochastic layers and residual blocks, which allows it to learn latent codes. Tested on the CAMELS Multifields Dataset, the model was able to recover the matter density parameter that might automatically learn physically meaningful parameters (such as Ω_m and σ_8) without any supervision in terms of labels. By measuring the change in these variables with a change in input cosmology, the authors mapped individual neurons in the bottleneck layer to physical parameters such as matter density and in this way provided a 'neural' coordinate system. Key findings show that the VDVAE manages to compress the high dimensional data in a low dimensional latent space where certain directions have a strong correlation with cosmological parameters. Quantitatively, the authors demonstrate that a simple linear regression on the latent variables can be used to predict m and σ_8 with a similar level of accuracy as traditional power spectrum analysis. A major disadvantage of this model is its difficulty in capturing small-scale features, which often results in slightly blurry generated images. This limitation indicates that while the model excels at large-scale representations, high-frequency details remain a challenge at this stage of development. This work serves as a basis to the current thesis and thus provides direct support for the hypothesis that VAE latent spaces contain accessible cosmological information.

Kannan et al. [15], presents Cosmo Flow to solve a problem of describing cosmological fields with GANs, which have proven too difficult to get their answers correct on all scales in space. The study makes use of Flow Matching, a simulation-free approach for training Continuous Normalizing Flows (CNFs), for modeling the distribution of weak lensing convergence maps. Unlike GANs which map noise to data in one step this model makes use of Continuous Normalizing Flows to gradually transition a simple Gaussian to a complex distribution such as the distribution of cosmological maps through "time". "It solves An Ordinary Differential Equation (ODE) during the training process to find an invertible mapping and it can calculate the exact likelihood of any given map, something the standard GANs cannot perform. Trained in response to a dataset of MassiveNuS images, CosmoFlow learns a resolution-invariant representation. The model achieves better performance numbers and is able to generate maps with power spectra which match simulations within 1% accuracy for a wide range of scales, considerably outperforming the GAN and standard VAE baseline. It also allows for accurate estimation of the detection of neutrino mass $(\sum m_v \cdot A)$. A limitation that is pointed out is that the Ordinary Differential Equation (ODE) integration during the sampling phase is computationally expensive, making it slower than the single-pass GANs, but faster than diffusion models.

This research[17] by Lue et al. questions the standard paradigm that large-scale statistics are necessary in order to perform cosmological inference, proposing instead that individual galaxies have sufficient information to constrain cosmological parameters. The authors use an Autoencoder (AE) to extract the manifold of galaxy properties (SEDs, morphology) conditioned on cosmology, using data from the CAMELS suite of hydrodynamic simulations. The methodology consists of training the AE to reconstruct the properties of the galaxy while simultaneously predicting the cosmological parameters Ω_m and σ_8 from the bottleneck layer. The network is designed with a "Y-shape" architecture, having a single encoder to compress the galaxy image, but with two heads emerging from it. One head tries to recreate the image (normal autoencoder task), whereas the other head is a regressor, attempting to learn the cosmological parameters from the same compressed bottleneck representation. The study reports a striking result: when information from only a few hundred individual galaxies is combined, constraints on Ω_m become competitive with those obtained from the clustering statistics of the whole field. Using only single galaxy information, the model achieves a constraint precision of $\sigma(\Omega_m) \sim 0.04$. A key limitation is the "simulation-to-real" gap, in other words the model depends heavily upon the specifics of the subgrid physics of the simulation (such as feedback models), which may not be a perfect match to reality.

Akhmetzhanova et al. [4] solves a problem of optimal compression of data for cosmological inference, an attempt to replace hand crafted summary statistics (as power spectrum) by their learned representations. The authors use Self-Supervised Learning (SSL), a contrastive learning framework (like SimCLR) in order to train a neural network that can tell the difference between different cosmological realizations in the MassiveNuS data set. The loss function used by this model is called a 'contrastive learning' loss function (SimCLR) that takes two different cropped versions of the same simulation and forces their vector representations to be identical. "This requires the network not to focus on the random noise or viewing angle and instead learn the invariant, fundamental, physical properties of the simulation structure without the need for any external labels. The network is trained to maximize mutual information between different views (augmentations) of the same simulation. The method provides summary statistics to capture the non-Gaussian information that is lost by the power spectrum. Quantitatively the improvement of the constraints on the sum of neutrino masses ($\sum m_v$) by the SSL derived summaries is about 20-30% with respect to the use of the power spectrum only. However, the performance is sensitive with the choice of data augmentations, which should be carefully chosen not to remove any physical signals as mentioned by them.

In another study done by Tan et al. [19] the authors propose an Evidential Physics-Informed Neural Network (E-PINN) which uses Evidential Deep Learning (EDL) along with physics-informed constraints, derived from the Friedmann equations. Unlike standard Bayesian Neural Networks that require computationally expensive sampling, E-PINN deterministically predicts the hyperparameters of a higher order probability distribution (Normal-Inverse-Gamma) that allows single pass uncertainty estimation. The model was trained and validated by using the Pantheon supernova dataset and Cosmic Chronometer data. Rather than a typical output layer that makes a single prediction, the final layer therefore makes the parameters (mean, variance and shape) of a Normal-Inverse-Gamma distribution the output. This enables the network to quantify its own 'epistemic' uncertainty in a sense it calculates a confidence score for each prediction that it makes, based on how closely the input resembles the training data. Results indicate that E-PINN has tighter and more reliable constraints than standard PINN. For example, The model inferred

$H_0 = 69.80 \pm 1.25$ km s^{-1} Mpc^{-1} and $\Omega_{m0} = 0.298 \pm 0.012$, which are statistically consistent with the standard ΛCDM model ($H_0 \simeq 70$, $\Omega_{m0} \simeq 0.3$), while providing explicit epistemic uncertainty bounds. . One significant drawback is the model's sensitivity to the regularization coefficient which balances the fit to the data with the physical constraints (if too small, one is likely to be overconfitting in the model).

Ojha et al. [11] inspects the inverse problem of estimating mass distributions of the dark matter, when presented with solutions to gravitational lensing images, a problem traditionally solved by costly iterative pixel-based algorithms. They present LensPINN, a mesh-free Physics-Informed Neural Network which can re-create the lensing convergence map (κ) and deflection potential directly. This network is 'mesh-free,' meaning it does not take a whole image grid as input and output the final value of the gravitational potential at each point in the grid; instead, it takes particular (x,y) coordinates as input and outputs the value of the gravitational potential at that point. It is trained by minimizing the 'physics residual' which is the mathematical difference between the network's output and the solution required by the lens equation, which requires no training dataset other than the physics equation itself. The loss function is a minimization of the residual of the lens equation (physics loss) and the difference between the reconstructed and observed source light profiles. On simulated Strong Lensing Systems using Singular Isothermal Ellipsoid (SIE) profiles, LensPINN has a reconstruction Mean Squared Error (MSE) of about 10^{-4} on the deflection potential. The research focused on the fact that LensPINN was able to recover the Einstein radius with an error smaller than 1% in high signal-to-noise ratios. Nonetheless, the approach is hampered by complicated morphologies of the source or large amounts of noise, where optimization can get trapped in local minima, and not find the real mass profile.

The research [8] explores the scaling of Dark Matter (DM) halo properties with baryonic gas properties, which in most cases demands computationally expensive hydrodynamic simulations to find. The authors suggest a PINN architecture that can make predictions of the gas density and temperature profiles based on the inputs of DM halo (mass, radius, and concentration). The methodology imposes the hydrostatic equilibrium equation as a physical constraint in the loss function to ensure the physical consistency of the predicted gas profiles with the gravitational potential that underlies them. The network acts as a mapping function that takes the properties of Dark Matter halos as an input and predicts properties thermodynamic profiles (pressure, temperature) of the gas. The key innovation is a loss term that penalizes the network for failures to adhere to the physical law of hydrostatic equilibrium when making predictions if the pressure gradient predicted by the network is not balanced against the gravitational force of the halo. The PINN trained on cluster-sized halos on the The Three Hundred project brought the Mean Absolute Percentage Error (MAPE) of the gas pressure profiles down to about 50% that of a pure data-driven Deep Neural Network (DNN). In particular, the median MAPE of the PINN was approximately 15% along the radial profile, and the conventional DNN tended to break the laws of energy conservation. The main weakness is that the hypothesis of hydrostatic equilibrium cannot be true to all the galaxy clusters especially when they are experiencing mergers, and thus the dynamic systems are more at fault.

It is the conclusion of [7] that the pure comparative efficiency of statistical Machine Learning (SML) against Physics-Informed Neural Networks (PINNs) can be effective in solving the differential equations popular in astrophysics, namely the Lane-Emden equation that drives stellar structure. The objective of the study was to investigate whether deep learning has the potential of substituting the conventional numerical integrators. The model is trained with a 'collocation' method, in which thousands of random points are sampled out of the domain of the star's radius. "For each point, the network calculates the output with respect to the input and determines whether they satisfy the Lane-Emden differential equation and it adjusts its weights until the equation is balanced to zero over the entire domain. The training strategy is to train a conventional Multi-Layer Perceptron (MLP) as a PINN, in which the loss is the difference between the differential equation itself which needs no labeled training data (unsupervised). The findings prove that the PINN was able to solve the Lane-Emden equation with n=0, n=1, n=5 with a Mean Squared Error (MSE) of the order of 10-6 to 10-8 relative to the analytical solutions. Conversely, the conventional regression techniques (SML) did not extrapolate beyond the scope of training. The study however observes that neural networks have a drawback in the spectral bias, which is that they do not learn high-frequency content of the solution with pure feature engineering. The relevance of this review lies in it being a foundational demonstration of proof-of-concept ability that PINNs can adequately approximate that governing equations of astrophysical fluids, which are a foundational pillar in constructing a Physics-Informed Physics VAE.

Another study [21] is an effort to address the so-called Hubble Tension, the difference between measurements of H_0 at the beginning of the universe and the end of the universe, by running PINNs on alternative cosmological models. The authors use a PINN to constrain the Tsallis Holographic Dark Energy (THDE) model with the inclusion of neutrinos. The modified Friedmann equations are solved with the methodology by a neural network that has been trained on 57 data points of Cosmic Chronometers (OHD) and Type Ia Supernovae. The neural network is taken to approximate the Hubble parameter function H(z) in terms of redshift, which is a universal function approximator." The training loss is a combination of the standard error against observed supernova data and 'physics loss' calculated from the modified Friedmann equations to make sure that the learned curve obeys the conservation of energy defined by the Tsallis holographic dark energy model. It is trained to learn the development of the Hubble parameter H(z) and is compatible with the energy conservation laws. This paper gives a projected Hubble constant of H_0 is approximately 72.1 $\mathrm{km\,s^{-1}\,Mpc^{-1}}$, much smaller than the Planck CMB value but closer to the local measurements . Also, it gives limits to the non-additivity parameter of the Tsallis entropy. It is also a major drawback that the model depends on the parameterization of THDE; the physical constraints of the PINN are in fact biased in the case of a wrong underlying theoretical model. This paper shows how PINNs can be used in cosmology for parameter inference (the "Inverse Problem"), going beyond the simple forward simulations.

In another study [13] the problem of inference of full-field cosmology is considered, in which the aim is to infer 3D density fields directly onto cosmological parameters (Ω_m, σ_8), and to measure the amount of uncertainty. The conventional CNNs give point estimates but do not give the posterior distribution. The authors present Physics-Informed Bayesian Neural Network (PI-BNN). This methodology mixes a BNN architecture (where the weights are distributions) with a physical regularization term that discourages divergence of statistics of known power spectrums. All the weights in the neural network, in this architecture, are not a single fixed number, but a probability distribution (a Gaussian with a mean and variance). During the forward pass, the network learns from a different set of weights at each forward pass, or in other words, a large ensemble of infinite networks that outputs a distribution of possible outputs that expresses the model's uncertainty. The PI-BNN was tested on N-body simulation boxes with the true parameters found in the 2σ confidence interval in the posterior contours of 95% of the test cases. The results of the study demonstrated that the "Physics-Informed" type of regularization allowed to shrink the volume of the uncertainty ellipse by 30% compared to a regular BNN, meaning more accurate constraints. Nonetheless, the Variational Inference in BNNs is computationally expensive, and thus, the training is much slower than in deterministic networks. This integrates Field-Level Inference and Physics-Informed constraints to enhance parameter bounds.

Ono et al. [12] address the problem of galaxy bias. To address this, the authors come up with a Conditional Denoising Diffusion Probabilistic Model (DDPM) to stochastically construct DM fields. The time-dependent U-Net architecture is based on training on the CAMELS series of hydrodynamic simulations, which are the IllustrisTNG and SIMBA seriesThe idea behind this model is that we add Gaussian noise to the dark matter maps iteratively until they become pure static and then learning a U-Net to reverse this process step by step. The U-Net teaches how to predict the 'score' (the gradient of the log-density) which essentially means they teach themselves how to pull the noise a little bit closer to a physical (dark matter) structure each timestep. Unlike the deterministic regression models (such as standard U-Nets) which output a blurry average of all possible posterior modes, this diffusion model learns the gradients of the data distribution (the score function). It creates samples by denoising a Gaussian random field iteratively, conditioned on the input maps of galaxy number count, by so-called channel concatenation. Quantitatively, the diffusion model was a lot better than deterministic U-Net baselines. The reconstructed DM fields replicated the real power spectrum of the simulation within a 1 percent accuracy to a wavenumber of $2.0\,h\,\mathrm{Mpc}^{-1}$. On the contrary, the power was significantly miscalculated in the baseline U-Net by more than 20% on small scales because of smoothing. More so, the cross-correlation coefficient (rcc) between the generated and actual DM fields were above 0.95 over large scales indicating high fidelity in structural placement. Critical review of the model shows that the model is highly accurate, but the iterative sampling algorithm is computationally costly and takes hundreds of forward passes to accomplish the milliseconds-long inference of a VAE or a GAN. The model was also sensitive to the particular subgrid physics of the training data, and when tested on the SIMBA data after training on the IllustrisTNG, performance was worse. The present research is critical because it underscores the trade-off between the high-fidelity reconstruction of the diffusion models and the speed of inference of VAEs.

The study [16] introduced a Topological Neural Network (TNN) framework to get the tightest constraints on the cosmological parameters from large-scale structure data and made a precise comparison of the accuracy achieved against the result obtained from GNNs. In this work, they have used two datasets: Quijote and CAMELS, and applied TNN on them. For the Quijote dataset, they performed TNN and compared the result with two other studies, point cloud method and graph neural networks were used, and TNN significantly outperformed both. The Full TNN performed best with 60% better accuracy predicting σ_8, and ClusterTNNs achieved best performance predicting Ω_m by 22% and σ_8 by 34%. Their model also outperforms previous models like PointNetMLP.Whereas, in the CAMELS dataset, which includes simulations of galaxy formation with baryonic physics, TNN models underperformed. Both the cluster TNN and Full TNN showed less accuracy than a standard GNN. This study clearly states that Topological deep learning can extract cosmological information with great accuracy and higher precision. The result shows that TNNs work comparatively better for constraining parameters. Along with that, it increases computational cost, and it may underperform in real-world data settings, which include astrophysical noise.

2.3 Summary of Key Findings

The literature review shows a major shift in the way astrophysical data is analyzed, from the traditional methods in the form of numerical approaches to Deep Learning (DL) architectures. Specifically, the research calls special attention to the fact that while Generative models such as Variational Autoencoders (VAE) make huge gains in speed and data compression to help with upcoming surveys such as LSST, the models often sacrifice fine-grained accuracy for computational efficiency.
Studies show VAEs to be especially effective in speedily processing massive datasets. Gagliano et al. (2023) demonstrated that Physics-Informed CVAEs were able to generate the galaxy parameters more than 100 times faster than the simulation-based approaches could, with which it successfully disentangled the physical properties (such as redshift) from the visual "style" [9]. To overcome the problem where the standard VAEs lose the image detail (the "blurring" effect), hybrid architectures have proved to be successful. For example, the U-Net VAE of Mirzoyan (2025) has skip connections to retain high frequency details such as spiral arms. 92.5% accuracy was achieved compared to 88.3% for normal VAEs [18]. Similarly, CosmoVAE showed that these models are capable of "inpainting" missing data in Cosmic Microwave Background maps with high precision on large scales, although performance decreases for small scales due to the inherent L_2 loss problem [2].
Compared to the speed of VAEs, Diffusion Models and Normalizing Flows have been the new state-of-the-art compared to high-fidelity simulation. As is apparent from the research done by Ono et al. (2024) and Kannan et al. (2025), Denoising Diffusion Probabilistic Models (DDPM) and Flow Matching are capable of reconstructing Dark Matter fields and weak lensing maps with an error of less than 1%, which is significantly better than that of VAEs and GANs [15] [12]. However, one of the common results of these various papers is the enormous computational burden; whereas the VAEs are run in milliseconds, the diffusion model relies on iterative sampling in a computationally expensive manner, leaving a trade-off between speed and resolution [12].

A large quantity of the literature is dedicated to Physics-Informed Neural Networks (PINNs), where physical laws are incorporated in the loss function of the network. This approach has been successfully used to solve a wide range of problems, including the Lane-Emden equation for the structure of stars [7], models of gas profiles in Dark Matter halos [8], and the reconstruction of gravitational lensing potentials without the requirement of using a mesh [11]. These "physics-informed" constraints ensure models have a much better generalization ability than a pure data-driven approach. For instance, Tan et al. (2025) and Ding (2025) have shown that by incorporating physical and Bayesian constraints, networks are made to quantify their own uncertainty and give tighter constraints on cosmological parameters such as the Hubble constant [19] [13].
In spite of these successes, the review identifies a number of important limitations. First, the "black box" nature of neural networks is still an issue, though work to map latent spaces to physical parameters such as and is promising for making these models interpretable [5] [17]. Second, there is the risk of the "simulation-to-real" gap; models trained on certain simulations (e.g. IllustrisTNG) might not work well when applied to data with different underlying physics or real observational data [17] [12]. Finally, although PINNs guarantee that the prediction is physically correct, they are highly sensitive to the accuracy of the theoretical equations applied in the training process; if the theory used in the model (for example, the specific model of Dark Energy) is biased, the model will predict too [21].
Overall, the results indicate that while standard CNNs and Self-Supervised Learning are reaching maturity for classification tasks [1] [4], the future of the research is hybrid generative models. These models try to compromise the speed of VAEs with the accuracy of diffusion models with physics-informed constraints to ensure the scientificity of the results.

Chapter 3

Requirements, Impacts and Constraints

3.1 Final Specifications and Requirements

This research will require a hardware and software infrastructure capable of supporting GPU-accelerated deep learning studies on large-scale simulation data. The system is meant for research testing rather than deployment, and it is designed to work in a controlled local computing environment.Models are developed and trained in a Python-based deep learning environment. When working with high-dimensional, multi-channel cosmological maps, GPU acceleration is necessary to enable feasible training times. The experiments are run on local GPU resources, allowing for iterative model construction, training behavior monitoring, and reproducible evaluation without the need for huge computational resources. The software architecture is built on popular open-source machine learning tools that enable tensor computing, automated differentiation, and quick data loading. This option assures compliance with established scientific methods and improves reproducibility. To avoid the unpredictability in results caused by software changes, the execution environment remains consistent throughout the training, validation, and testing phases. In terms of data management, the system must provide structured access to simulation datasets, fixed dataset segmentation, and the storage of model checkpoints and experiment logs. These requirements are critical for preserving experimental runs, confirming results, and ensuring the stated findings are repeatable. Overall, the specified hardware and software requirements provide enough processing capabilities for performing controlled physics-informed machine learning experiments while staying accessible in an academic research setting.

3.2 Societal Impact

In computational cosmology, simulations play a crucial role in explaining physical processes, but they are very costly to generate due to their huge dimensionality. To address these problems, this work developed a physics-informed machine learning model that gives a more reliable analysis of these simulation data. One of the key societal benefits of this work is the introduction of a trustworthy scientific model. Entirely data driven models might achieve greater accuracy but violate known physical behaviour. By incorporating physics based constraints directly into the training process, this work reduces the possibility of such a violation. This helps to ensure that the model not only produces outputs with great results but is also consistent with physical expectations. By compressing the dimension and predicting parameters with that low dimensional simulated data, this work allows researchers to extract meaningful information and to draw a quick conclusion without repeatedly running simulations. It contributes to achieve more efficient research workflows and reduces unnecessary computational effort. Access to large-scale computing equipment is a barrier for many people who have a keen interest in knowing the universe. It is possible to explore complicated cosmological concerns without the need for large computational resources by using methods that prioritize effective model design and the reuse of existing simulation data. In this way, the work encourages innovative research and makes it possible for passionate researchers and students to interact with open scientific issues in a more approachable and resource-conscious way.

3.3 Environmental Impact

Large cosmological simulations are computationally expensive and very energy intensive to run. This work is completely based on existing simulation data, which consumes less energy comparatively. This study reduces the need for repeated simulation runs by concentrating on improved analysis, compression, and inference from pre existing datasets, which decreases additional energy usage related to data generation. Additionally, latent compression helps make handling high dimensional data more effective. Future research that builds on this approach may be able to further reduce computing overhead by learning compact representations, which may enable faster downstream processing and lower storage requirements.

3.4 Ethical Issues

This work does not involve any sensitive or personal data. All the experiments are entirely based on simulated cosmological datasets that are publicly available and made with the purpose of future research works. Therefore, there are no direct ethical concerns like privacy, consent, or data misuse. The interpretation of the study's findings is another ethical issue. The results are based only on simulated data within a particular parameter range and cosmological model. No claims are made about direct applicability to actual observational data, and these limitations are explicitly acknowledged. This framing minimizes the possibility that the findings may be misinterpreted or used outside of the original research context.

3.5 Risk Management

A major concern in this study is the possibility of shortcut learning, in which a model creates visually accurate reconstructions while disregarding known physical behavior. Despite good numerical results, such findings can lead to false interpretations. To mitigate this problem, physics informed limitations are built into the training objective, penalize physically incorrect outputs, and improve learning that aligns with recognized physical expectations. Another significant risk arises from the structure of the CAMELS dataset itself. The collection includes several simulation suites and parameter sets based on various physical assumptions. Choosing an acceptable subset for parameter inference takes careful attention, since an incorrect choice may have a negative effect on model performance or influence the results. This risk is carefully mitigated by focusing on dataset selection. For this research, data selection was restricted to the IllustrisTNG Latin Hypercube subset, as the values of this set varies form different ranges, which provides the diversity required for learning parameters.

3.6 Economic Analysis

This research's economic cost is mostly related to the local computational resources utilized for model training and evaluation. No commercial software licenses, paid datasets, or external cloud computing services were used. All studies were conducted using open-source software and publicly available simulation data.
All model training took place on a single custom-built PC outfitted with an NVIDIA RTX 5060 Ti GPU (16 GB VRAM, 4608 CUDA cores). Training lasted 50 epochs, with an average runtime of about 3 minutes per epoch, for a total of 150 minutes (2.5 hours) per complete run.The expected power cost each training run is about USD 0.25 based on standard local electricity prices, which keeps direct operating costs extremely low. To lower computing costs without compromising experimental quality, a number of steps were utilized. The work relies on controlled experimentation with a predetermined training schedule rather than extensive hyperparameter searches. By reusing existing simulation datasets instead of creating new ones, computational costs were greatly decreased. Furthermore, training was restricted to necessary experiments, and in order to minimize repetitive computation, intermediate findings were reused whenever feasible. In contrast, it usually costs between 2 and 3 per GPU per hour to rent comparable GPU capabilities from a commercial cloud provider. The cost benefit of local experimentation is shown by the expected cloud computing cost of USD 10–15 per training run if two GPUs are used for the same amount of time.
To conclude, this work's economic impact is still low and suitable for an academic research setting, proving that high-computational research may be carried out effectively with careful resource planning and controlled testing.

Chapter 4

Proposed Methodology

4.1 Design Process or Methodology Overview

The study follows a structured, physics-inspired variational autoencoder framework that aims to learn compact cosmological field representations which also allows cosmological parameter inference following the fundamental physics laws. The proposed methodology combines dimensionality reduction, generative modeling for reconstruction and physics based constraints within a unified pipeline. The design process consists of four stages: dataset preparation, model architecture design, model training with physics laws and ablation-based evaluation.

4.2 Data Collection

The paper uses the CAMELS (Cosmology and Astrophysics with Machine Learning Simulations) dataset for unsupervised learning, collected using a publicly available set of repositories of data on CAMELS entitled CMD (CAMELS Multifield Data). The main idea of the CAMELS dataset is to observe how matter and gas evolve by systematically changing the cosmological and astrophysical parameters. Each simulation map of the dataset consists of 256×256 pixel images. Each image was paired with its corresponding cosmological and astrophysical parameters, such as Ω_m, σ_8, ASN1, ASN2, AAGN1, and AAGN2. All these parameters describe the matter density, amplitude of matter fluctuations, and feedback strengths from supernovae and active galactic nuclei.

The dataset contains 3D and 2D maps of different types of suites, such as IllustrisTNG, Astrid, SIMBA, N-body, etc. Each suite is generated using the same cosmological volume and initial conditions; however, every suite uses different sub-grid physics models. Each suite has been organized with several distinct sets (such as LH, CV, EX, etc.). These sets are designed to help researchers differentiate between the effects of cosmology, astrophysics, and random chance. Each set has different physical fields (such as M_{gas}, HI, V_{gas}, Magnetic fields) that carry different mutual information about the underlying cosmology.

To conduct this research, the following suite, set, and fields have been chosen:

SUITE: In this research, the IllustrisTNG suite has been used, which contains 2D projection maps of multiple physical quantities that were derived from 3D simulations. Each 2D projection map was represented as a standardized tensor with the shape ($3 \times 256 \times 256$). The IllustrisTNG suite uses magnetohydrodynamics (MHD), state-of-the-art feedback models, and realistic gas and magnetic field evolution. This suite is suitable for meaningful reconstruction, compression, and parameter prediction, as it continuously evolves gas, velocity, and magnetic fields while focusing purely on cosmological parameter learning.

SET: From the IllustrisTNG suite, the Latin Hypercube (LH) set has been used. The LH set varies the cosmological parameters over a continuous range, and its sampling gives uniform coverage of the parameter plane. Therefore, this set ensures that our model learns a continuous mapping of the Ω_m and σ_8 parameters and prevents the model from memorizing the cosmological space.
Figure 4.1 summarizes the LH set before and after transformation.

FIELD: To break parameter and reconstruction degeneracy, M_{gas}, V_{gas}, and Magnetic fields from the LH set are chosen as inputs for this research. The initial feature space for each image is $3 \times 256 \times 256$, making it 196,608 features. M_{gas} refers to the gas mass density field that traces the baryonic matter distribution per pixel and is directly affected by matter density (Ω_m). V_{gas} (gas velocity field) traces how the gas particles are moving, along with their direction and magnitude. Finally, the magnetic field (B), a non-linear field, refers to how the magnetic field impacts gas pressure, star formation, and feedback pressure.
These three fields together give us the static information of density, dynamic information of velocity, and non-linear information of the magnetic field, hence allowing the PI-VAE model to learn a richer representation of the latent space that helps in capturing important physical information which is sensitive to Ω_m and σ_8.

4.2.1 Data Cleaning

As the input was directly sourced from a clean numerical simulation, there was no need to handle the NaN values. The cleaning stage focuses on stabilizing the physical fields in terms of each channel's physical properties. As the channel includes dynamic range, the data was stabilized using a nonlinear transformation.

4.2.2 Data Transformation

To load and preprocess three-channel 2D projection maps efficiently from IllustrisTNG simulations in the CMD-CAMELS Multifield Dataset, every simulation sample was stacked into tensors with three physical channels: gas mass density, gas velocity, and magnetic field.
Since M_{gas} and magnetic fields have a larger dynamic range, logarithmic scaling was performed on both fields. On the other hand, an arcsinh transformation was used on the gas velocity (V_{gas}) field, as the field is signed and the values are symmetrically distributed. Therefore, the arcsinh(X) function was used to reduce the effect of outliers while keeping the sign of the data. It preserves the signed information while compressing the values.
These channel-wise physical transformations help the PI-VAE model to learn the cosmological structure from different density regimes.
On the training data, augmentations which include rotations (90°, 180°, 270°) and horizontal or vertical flips via NumPy were applied to increase data diversity, which helped to produce 6 augmented variants per original map. Since ΛCDM cosmology

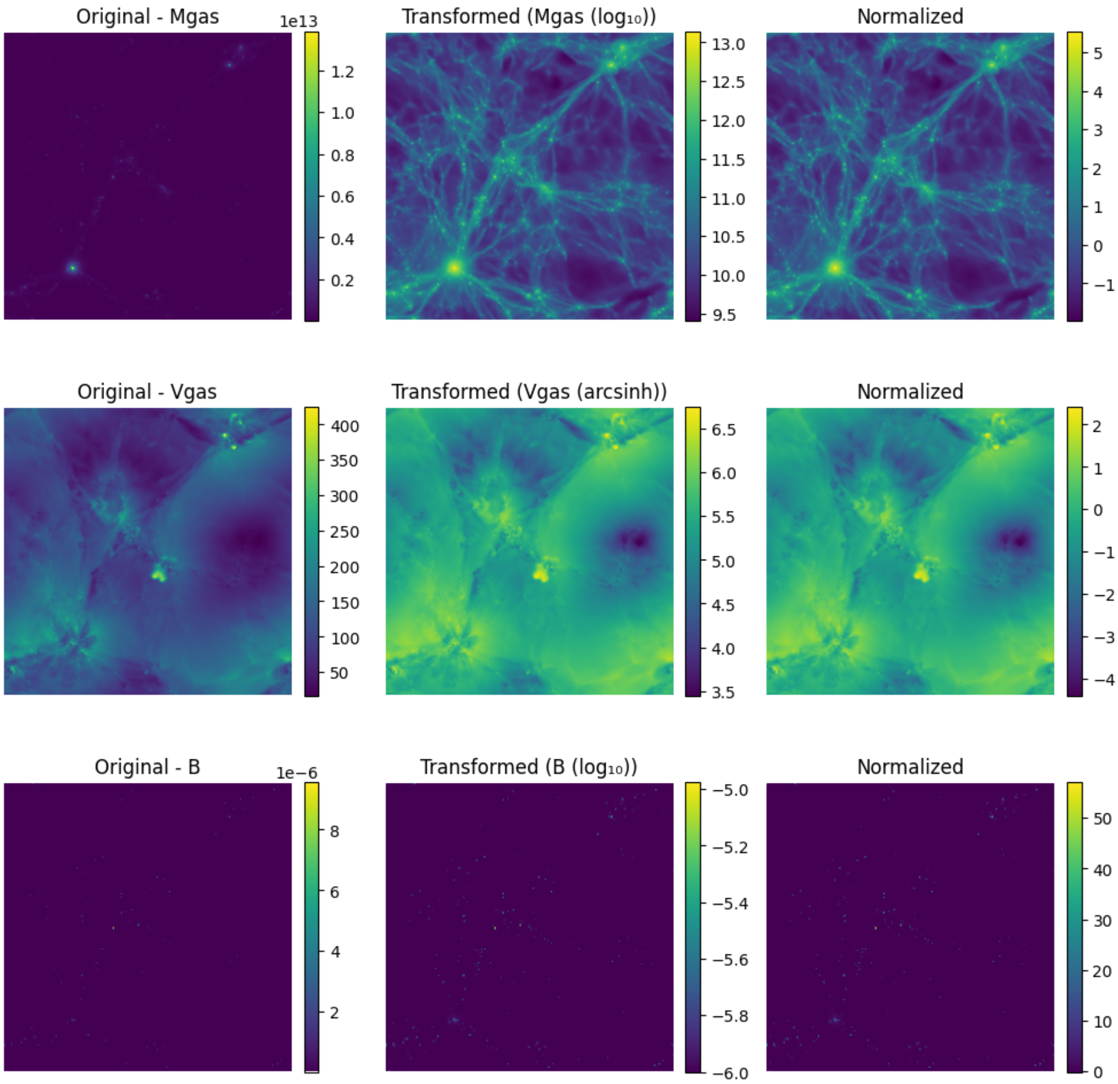

Figure 4.1: LH Set Data (before and after transformation)

is isotropic and homogeneous, these spatial transformations keep the cosmological parameters (Ω_m, σ_8) unchanged, increase the generalization ability of the model, and reduce overfitting. Augmentation was applied only to the training set in order to avoid leakage of information into the validation and test sets.
After transformation, the data was normalized using the mean and standard deviation from the training subset. Normalization ensures that the neural network is numerically stable, ensuring fast convergence. The target cosmological parameters (Ω_m and σ_8) were also scaled using `StandardScaler` to ensure that the regression loss is not dominated by any single parameter.
Overall, this preprocessing pipeline is consistent with the CAMELS dataset's physical characteristics and is carefully designed for field reconstruction and cosmological parameter inference using the proposed model.

4.2.3 Dataset Split:

The LH dataset (M_{gas}, V_{gas}, and B field) consists of a total of 15,000 maps and was divided into a 3:1:1 ratio for the training, validation, and test sets, respectively. Hence, the training set consisted of 9,000 maps, while both the validation and test sets contained 3,000 maps each.

4.2.4 Summary of Preprocessed Data

The final data was preprocessed to ensure suitability for deep learning applications, where each sample consists of three channel 2D maps with stabilized modified in a way that is optimized for deep learning applications. The complete dataset includes thousands of 2D projection maps. The dynamic range of all physical fields was stabilized using the nonlinear transformation. Then, augmentation and normalization were done. The LH set was divided into training set (60%), a validation set (20%), and a test set (20%).

4.3 PI-VAE Model Architecture Design

The proposed physics-informed variational autoencoder (PIVAE) model consists of four parts: 1. A VAE-based encoder, 2. A decoder with skip connections, 3. A parameter regressor and 4. Physics laws as loss functions. These components are trained together to form a unified learning pipeline that simultaneously reduces dimensionality, reconstructs the field and predicts cosmological parameters (Ω_m and σ_8)

4.3.1 VAE Based Encoder

The encoder takes 3-channeled 256×256 images from M_{gas}, V_{gas}, and the magnetic field of CAMELS IllustrisTNG-LH simulations. These fields are stacked channel-wise and processed through five convolutional layers, each with a LeakyReLU activation function having $\alpha = 0.01$ and a 4×4 kernel with a stride of 2.
The encoder performs two major tasks: (i) extracts hierarchical spatial features from the inputs by downsampling the input image from a $256 \times 256 \times 3$ grid to an $8 \times 8 \times 512$ compact feature map and (ii) compresses these feature maps into a

probabilistic latent space. Mean vector (μ) and Log-variance vector ($\log \sigma^2$) are the two output of encoder which are used by reparameterization trick to create a latent space (z).

$$z = \mu + \sigma \odot \epsilon, \tag{4.1}$$

$$\epsilon \sim \mathcal{N}(0, I). \tag{4.2}$$

This latent space acts as the main information bottleneck of the entire architecture as it encodes global patterns which is ultimately used for both reconstruction and parameter prediction.

4.3.2 Decoder With Skip Connections

The decoder reconstructs the original cosmological fields from the sampled latent space. It conducts the inverse operation of the encoder using transposed convolutions. The upsampling and downsampling paths allow the decoder to reform structures from multiple spatial scales, which is important for large scale non-linear features. To ensure no important spatial feature is lost (which can be caused by the bottleneck), the study implements U-NET inspired skip connections. These skip connections add feature maps from the encoder to the decoder which ensures to capture both global and local spatial information. These connections help the decoder in recovering the fine-grained spatial details rather than bypassing the latent space. The latent space remains as the only bottleneck that encodes the information of global cosmological space. Therefore, the skip connections do not allow the model to cheat or guess, instead they just give structural guidance and help to improve reconstruction quality without leaking cosmological parameter informations directly to the output. The final decoder layer yields reconstructed images of the same shape as the input.

4.3.3 Parameter Regressor

The PI-VAE model uses a multi-layer perceptron (MLP) that takes the latent vector as input and predicts the matter density (Ω_m) and amplitude of fluctuations (σ_8). It is consisted of three fully connected layers with ReLU activation and operates entirely on the latent representation. By predicting both parameters, the model learns a shared cosmological structure representation and correlations of parameters that naturally emerge from the latent space.

4.3.4 Loss Functions

To ensure the model predicts accurate parameters and generates meaningful reconstruction, the study proposes a hybrid loss function which contains 2 components: Vae loss and parameter regression loss.

$$L_{\text{total}} = L_{\text{vae}} + \lambda_{\text{param}} \, L_{\text{param}} \tag{4.3}$$

Vae Loss: The vae loss function contains a total of six parts, where four of them are related to physics-informed constraints. The physics losses ensure that the latent space is not just a pixel by pixel or numerical compression, rather a representation of the underlying fields of physics.

$$L_{\text{vae}} = L_{\text{recon}} + \lambda_{\text{kl}}\, L_{\text{kl}} + \lambda_{\text{ps}}\, L_{\text{power}} + \lambda_{\text{amp}}\, L_{\text{amp}} + \lambda_{\text{mean}}\, L_{\text{mean}} + \lambda_{\text{grad}}\, L_{\text{grad}} \tag{4.4}$$

Here,λ =physics weights.
Below a brief description of the losses are given.

- **Reconstruction Loss:** Also known as the Mean Squared Error (MSE), this loss measures the average squared difference between the original and the reconstructed fields, and is defined as

$$L_{\text{recon}} = \frac{1}{N}\sum_{i=1}^{N} \left\| x^i - \hat{x}^i \right\|^2 , \tag{4.5}$$

  where x denotes the original log-transformed image and $\hat{x}$ represents the reconstructed image. This loss enforces basic visual similarity between the input and output fields and encourages the decoder to preserve the spatial structure and dimensions of the cosmic web.

- **Kullback–Leibler Divergence Loss (KL Loss):** KL loss computes how much the learned latent space differs from a standard Gaussian distribution.

$$L_{\text{KL}} = -\frac{1}{2}\sum \left(1 + \log\sigma^2 - \mu^2 - \sigma^2\right) , \tag{4.6}$$

  The loss encourages a smooth, structured, d and physically meaningful latent space while preventing overfitting

- **Power Spectrum Loss:** Power spectrum loss computes how the Fourier transform frequency representation of predicted and original image differs from one another, rather than focusing on the raw pixel values which MSE loss does. Because in cosmology if 2 maps are visually same that does not mean they are physically correct and power spectrum loss helps to ensure that the reconstructed field have the same physics as the original field.

$$L_{\text{PS}} = \left\| \log\left(\overline{P}_{\hat{x}}\right) - \log\left(\overline{P}_{x}\right) \right\|^2 , \tag{4.7}$$

  This loss ensures the clustering matter statistics of the cosmological field are preserved while reconstructing and hence helping to predict the parameters correctly.

- **Amplitude of Fluctuation Loss:** To control overall fluctuation strength and prevent the reconstructed maps from over or under smoothing amplitude fluctuation loss was used. It calculates the difference between predicted and actual amplitude that derives from the fields' respective density fluctuation.

$$\delta = x - \langle x \rangle \tag{4.8}$$

$$A = (\delta)^2 \tag{4.9}$$

$$L_{\text{amp}} = |A_{\text{pred}} - A_{\text{original}}| \tag{4.10}$$

- **Mean density Error:** This error refers to the average mass energy density reduction of the universe over time, as the universe is continuously expanding. It is directly connected with matter density (Ω_m) parameter as average matter density is directly controlled by Ω_m.

$$L_{\text{mean-density}} = \frac{|x_{\text{pred}} - x_{\text{original}}|}{x_{\text{original}}} \tag{4.11}$$

  This law ensures total mass density conserved in the reconstructed field is identical to the input simulation field which is a primary requirement for physical consistency.

- **Gradient Density Loss:** This loss focuses on the edges, filaments, shocks etc of the cosmic web. It calculates gradients of the spatial regions of the predicted field, penalizes gradient magnitudes that are larger and finally enforce smoothness L1 norm. Therefore by penalizing gradient errors this loss ensures the comic web is sharp, smooth and realistic.

- **Parameter Loss:** The parameter regression loss uses mean squared error (MSE) to measure how different the predicted cosmological parameters (y') are from the ground truth parameters (y).

$$L_{\text{paramReg}} = \frac{1}{N} \sum \|y' - y\|^2 \tag{4.12}$$

  Here y is the ground truth values of Ω_m and σ_8 from the simulations, and y' is the predicted values produced by the parameter regressor that uses the latent space.

  In cosmology to find accurate extraction of parameters and generate meaningful reconstructed fields, traditional MSE and KL losses are never enough. While these losses may give visually similar reconstructed images however they are not meaningful enough to be used in future experiments. Thus, incorporating the physics loss with the traditional loss helps us to find physically meaningful information that can be studied or used in different experiments.

4.3.5 PI-VAE Workflow

The PI-VAE model follows the steps mentioned below to compress, reconstruct, and predict cosmological parameters: The encoder takes 3 physical fields: M_{gas}, V_{gas} and B (magnetic field) as input, each field being the size of 256x256. It uses five convolutional layers that downsamples the dimensions and increases the channels progressively to extract important features from the cosmological fields. Then the final encoded feature map (8x512x512) is flattened and passed through a fully connected layer to produce the mean vector and log variance vector. This mean and log variance vector then uses a reparameterization trick and creates a latent space (z) that has a dimension of 265. Due to this reparameterization step, the model effectively compresses high-dimensional cosmological fields into a compact latent space which allows probabilistic representation learning while preserving differentiability. From the latent space (z) two parallel processes operate concurrently:

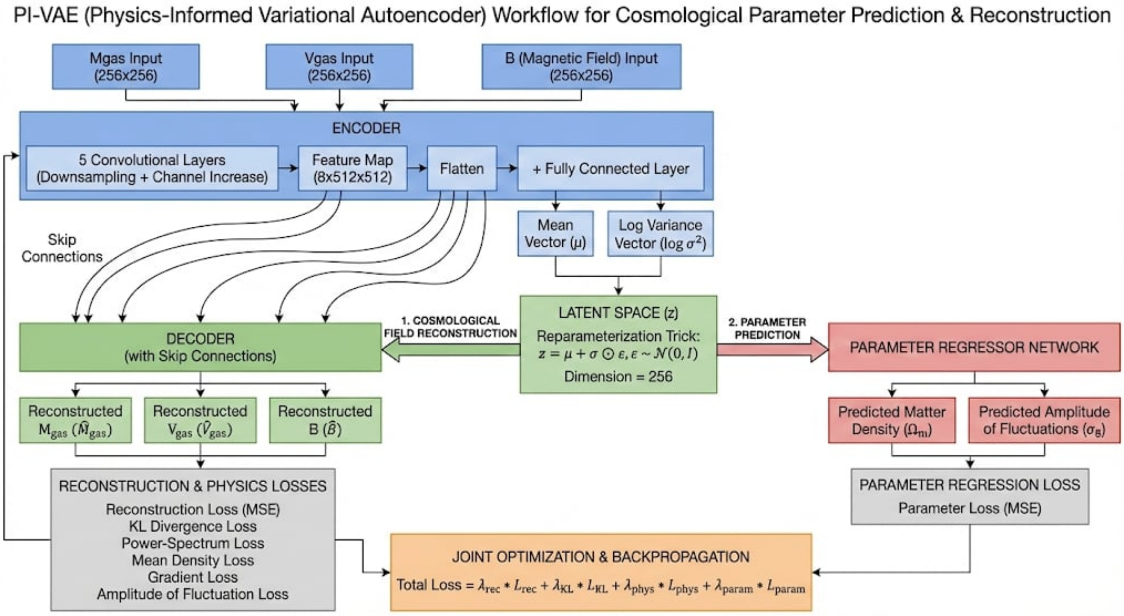

Figure 4.2: Architecture and Workflow of the Proposed PIVAE model

Cosmological Field Reconstruction: The latent space (z) is passed to the decoder. The decoder then upsamples the latent space back to its initial spatial resolution. Since the decoder uses skip connections it allows the model to retrieve fine scale spatial details while simultaneously preserving global structure.

Unlike the standard VAE, PI-VAE uses skip connections between the encoder and decoder. Skip connection allows the model to retrieve fine scale spatial details while simultaneously preserving global structure. The decoder then reconstructs the original physical fields (M_{gas},V_{gas} and B) using the loss function that is composed of reconstruction loss, KL loss and the physics informed losses including power-spectrum loss, mean density loss, gradient loss and amplitude of fluctuation loss.

The losses ensure the pixel-level fidelity of the reconstruction as well as enforcing physical consistency that ensures the generative fields are maintaining the correct level of large-statistics, mass conservation and spatial smoothness. The power spectrum loss does a fourier transform of the reconstruction image and it is compared with the original image.

If the decoder matches the pixel values but fails to reproduce the large scale structure or small scale structure, this loss gives a strong penalty. This ensures that the statistical properties of the cosmic web are preserved and maintains the overall clustering behaviour. Similarly, Gradient Loss computes the value of dx and dy which corresponds to the first order spatial derivatives of the image. This loss measures the rate of change between the neighbour pixels and gives a penalty if the reconstruction image contains noise or unnatural blurring. Furthermore, Amplitude loss measures the variance of pixel values and it forces the decoder to maintain full dynamic range of the images instead of collapsing it to the mean. Again, Mean Density Error maintains the consistency in the global mean value of the image. It ensures that the model does not destroy matter during reconstruction. Now all the losses force the encoder to follow the rules of physics, if it does not maintain the rules the loss increases. The produced latent vector works as a filter where the information is preserved.

Parameter Prediciton: The same latent vector (z), at the same time, is passed to the parameter regressor network. The regressor predicts the cosmological parameters (matter density (Ω_m) and amplitude of fluctuations(σ_8)) directly from the compressed latent space. The parameter regression takes the physics constrained latent space from the encoder as input and guarantees that the predictions are physically informed and meaningful rather than just incorrect correlations. The combined losses backpropagate simultaneously through the encoder, decoder and regressor. This joint optimization makes sure that (i) the latent space is well compact and physically relevant, (ii) the quality of the reconstruction improves without violating the physics laws and (iii) The physically consistent representations are benefiting the regressor to predict parameters accurately. During the parameter inference phase, only encoder and regressor are needed for parameter prediction, while the decoder is used to reconstruct the cosmological fields. This enables the predictions to be fast because the latent representation is directly mapped to (Ω_m) and (σ_8) while also making sure that the reconstruction abides by the physics laws and allows meaningful exploration of latent space and interpolation for studying the continuous variations in cosmic web.

In the context of our research, "cheating" is defined as a tendency for the deep learning model to optimize for visual similarity by the creation of shortcuts, as opposed to once learning the physical distribution (the underpinning, distribution) of the data. Our analysis shows that the U-Net VAE baseline is very much guilty of this. By design, U-Net architectures make use of "skip connections" direct connections between the high resolution input layers in the encoder to the output layers in the decoder. In the absence of strong physical constraints, the model takes advantage of these connections to be able to go around the latent bottleneck altogether. Instead of generating the output based on the input image, effectively what the U-Net VAE "peeks" at the input image and simply copies high frequency spatial details cosmic filaments into the output. While that makes visually sharp reconstructions of the inputs (that is, mimicking the original simulation) (Figure 4.2), what it does is generate a "hollow" latent space that fails to capture all the actual cosmological parameters.

In summary, the PI-VAE avoids this shortcut learning by using the combination of Physics-Informed Loss Functions. Unlike the baseline, which minimizes simple pixel-wise error (MSE), the PI-VAE is highly penalized if the output of the PI-VAE violates fundamental physical laws, namely Power Spectrum (clustering statistics) and Mean Density. This forces the network to use skip connections only to provide some structure for the network rather than just copying. The model cannot achieve the loss of the Power Spectrum by merely copying the pixels, they need to "understand" the statistical distribution of matter to arrive at a good cosmic web.

4.4 Ablation Study and Baseline Models:

To evaluate the performance of the proposed physics inspired variational autoencoder model, an ablation study was conducted using two baseline model architectures: UNet VAE and a traditional VAE. These models help to isolate the contributions of the physics based loss functions and skip connection architectures to both parameter prediction and reconstruction fidelity.

U-Net VAE Baseline: This model uses the architecture of the PI-VAE, but does

not use the physics-informed loss functions (power spectrum loss, mean density loss etc.). It is trained with the only standard Reconstruction Loss (MSE) and KL-Divergence, as well as the Parameter Regression Loss (MSE). This enables us to test whether architectural complexity alone is enough for cosmological accuracy.
Traditional VAE Baseline: This model is a standard bottleneck architecture. It is based on the same encoder as the PI-VAE and a decoder without skip connections. Similar to the U-Net VAE, it does not include the physics laws in the loss mechanism, only using MSE and KL-Divergence to train.

Feature	**Traditional VAE**	**U-Net VAE**	**Your PI-VAE**
Skip Connections	No	Yes	Yes
Latent Bottleneck	Yes (256-dim)	Yes (256-dim)	Yes (256-dim)
Physics Losses	No	No	Yes (4 Laws)
Regression Head	Yes	Yes	Yes

Table 4.1: Key Architectural Differences of the Models

This comparison's main objective is to assess the "Physical Realism" of the reconstructed fields and the accuracy of parameter estimation Ω_m and (σ_8)) without explicit physical constraints. The study shows how the integration of cosmological laws keeps the model from producing physically inconsistent density fields by contrasting these baselines with the PI-VAE.

4.5 Experimental Setup

To conduct this experiment, the following computational environment, hyperparameters and training strategies have been implemented to ensure stable and convergent training of the proposed PIVAE and the baseline models.

4.5.1 Hardware and Software Requirements:

PyTorch deep learning framework has been used to implement the models. The training of all models has been conducted on NVIDIA RTX 5060 Ti 16Gb VRAM with 4608 CUDA cores to handle the heavy computations that are required for the models. Python 3.11.1 and NumPy are used for data manipulation and Matplotlib is used for visualization.

4.5.2 Hyperparameters and Model Configuration

To ensure consistency with the PIVAE, UNet VAE and VAE model, the following hyperparameters were used for all models.
In addition to the above-mentioned parameters, PI-VAE uses $\lambda_{\text{powerspectrum}} = 1.0$, $\lambda_{\text{amplitude}} = 5$, $\lambda_{\text{mean}} = 0.1$, and $\lambda_{\text{gradloss}} = 0.1$ to ensure that physics-based losses do not overwhelm the model and to facilitate the capture of the most important cosmological information.

Parameter	Value	Description
Batch Size	16	Number of training samples processed per iteration.
Learning Rate (LR)	2×10^{-4}	Initial step size for the Adam optimizer.
Latent Dimension	256	Size of the compressed latent representation (bottleneck).
Image Size	256×256	Spatial resolution of the input cosmological density maps.
Epochs	50	Maximum number of complete passes through the training dataset.
Param Weight (λ_{param})	0.1	Weighting factor for the parameter regression loss.
β (KL Weight)	0.1	Controls the strength of the KL divergence regularization.
Optimizer	Adam	Optimization algorithm used for parameter updates.
Patience	10	Early stopping threshold based on validation performance.

Table 4.2: Training and model hyperparameters used in the VAE implementation.

4.5.3 Training Strategy and Optimization:

The study uses Adam optimizer, known for its adaptive learning rate qualities, to handle multiple loss terms concurrently. Additionally, ReduceLROnPlateau scheduler has been implemented to ensure the learning rate is automatically lowered to find a more accurate local minimum. This scheduler uses 3 epochs of patience and a reduction factor of 0.5. To prevent overfitting, early stopping has been used to stop the process if the validation loss does not improve for 10 consecutive epochs. The initial 20 epochs are KL warm-up at which the weight of the KL-divergence term is seen gradually rising. This enabled the model to emphasize proper reconstruction of the data and subsequently apply regularization to the data.

Therefore, this work aims at developing a complete workflow which blends data preprocessing and deep feature extraction. With the help of physics informed VAE with the hyperparameter optimization and the adequate control of training. The model was able to learn meaningful representations of the cosmological simulation. The outcome revealed that the model could efficiently represent structural patterns of the dataset while predicting the crucial cosmological parameters, which justified that this method could be helpful in analysing unsupervised cosmological information.

In summary, this work aims at developing a complete workflow which blends data preprocessing and deep feature extraction. With the help of convolutional VAE with the hyperparameter optimization and the adequate control of training. The model was able to learn meaningful representations of the simulation astronomical 24 maps as a result. The outcome revealed that the model could efficiently represent structural patterns of the dataset, which justified that this method could be helpful in analysing unsupervised astronomic information.

Chapter 5

Result Analysis

In this section, the study analyzes the performances of the suggested **Physics-Informed Variational Autoencoder (PI-VAE)** along with two ablation baselines: traditional **VAE** and **U-Net VAE**. The accuracy of cosmological parameter inference and the physical consistency of the reconstructed fields are the main points of comparison.
To analyze the performance of the models, the study focuses on:

1. Accuracy of parameter prediction.

2. Reconstructed fields' physical realism and cosmological consistency.

This dual evaluation framework is important because a cosmological model should not only predict parameters accurately but also reconstruct and generate fields that obey the fundamental physics laws and statistical properties of the Universe.

5.1 Evaluation Criteria:

To evaluate the model, following metrics are chosen:

1. **Parameter Regression Metrics:** MSE, MAE, R^2 score, and Log-space metrics are used in this study to evaluate the cosmological parameter (Ω_m, σ_8) inference.

   Mean Squared Error (MSE): By measuring the average squared differences between predicted and ground truth values, MSE evaluates accuracy of the models. The lower the MSE the model produces, the better the model is.

   Mean Absolute Error (MAE): Computes the average magnitude errors in predictions for measuring the accuracy of the parameter prediction and gives more intuitive interpretation in physical units.

   Coefficient of Determination (R^2 Score): Evaluates how the variance in the data is explained by the models. R^2=1 indicates perfect prediction, 0 indicates random prediction and if $R^2 < 0$, the model is predicting worse than random.

2. **Physical Consistency Metrics:** To analyze the realism of physics in the reconstructed fields, the following physic based metrics have been used:

Power Spectrum Error: Checks whether or not the reconstructed fields follow accurate cosmic structure formation by measuring the preservation of both large and small-scale clustering statistics.

Mean Density Error: Computes the global mass conservation consistency and ensures that in matter distribution, physics is consistent.

Gradient Density Error: Evaluates structural continuity and local spatial smoothness. Prevents unphysical, noisy discontinuities.

Relative Error Mean: Computes the normalised average percentage error in prediction to the actual values. An important metric for the physical interpretation that is scale variant.

The above-mentioned metrics ensures the model's ability to produce not only visually similar outputs but also physically meaningful cosmological reconstructed fields.

5.2 Quantitative Comparison Analysis:

Model	Parameter Prediction (R^2)	Ω_m (R^2)	σ_8 (R^2)	MSE	MAE
PI-VAE (Proposed)	0.845	0.957	0.732	2.079822×10^{-3}	3.236046×10^{-2}
UNet-VAE	0.687	0.898	0.477	4.183160×10^{-3}	4.802977×10^{-2}
VAE	0.754	0.919	0.588	3.296783×10^{-3}	4.208708×10^{-2}

Table 5.1: Parameter prediction performance comparison for different models

Model	Power Spectrum Error	Mean Density Error	Gradient Error	Relative Error Mean	Physical Awareness
PI-VAE (Proposed)	7×10^{-6}	5.6×10^{-4}	7.4×10^{-4}	$\Omega_m \to 0.0959,\ \sigma_8 \to 0.3191$	Strong
UNet-VAE	1×10^{-5}	6.3×10^{-4}	2.26×10^{-3}	$\Omega_m \to 0.2451,\ \sigma_8 \to 0.6444$	Moderate
VAE	4.96×10^{-1}	3.83×10^{-2}	5.80×10^{-1}	$\Omega_m \to 0.2159,\ \sigma_8 \to 0.6110$	Weak

Table 5.2: Comparative Reconstruction Performance of Models

Tables 5.1 and 5.2 show the comparison between the models in terms of parameter prediction and reconstruction, respectively. The narrative analysis of both tables is given below:

Parameter prediction accuracy: PI-VAE achieves the highest R^2 result overall compared to U-Net VAE and VAE. The score of coefficient of determination (R^2) is 0.845, whereas U-Net VAE and VAE score 0.687 and 0.754 respectively.

Furthermore, among all the model predictions, Ω_m remains significantly more precise than σ_8. The PI-VAE achieves an R^2 score of 0.957, which indicates that the global matter density is well encoded. Interestingly, the U-Net VAE achieved the lowest R^2 in predicting Ω_m. While the U-Net is better in reconstruction, its latent space might not be optimal for regression tasks.

Again, σ_8 scores 0.732 for PI-VAE. This is expected in cosmological simulations as σ_8 is highly sensitive to small-scale baryonic feedback. Moreover, the MSE and

MAE errors (2.08×10^{-3} and 3.24×10^{-2} respectively) from Table 5.1 for the PI-VAE model are significantly lower than that of the other two models, which further ensures that the predicted values by PI-VAE for Ω_m and σ_8 are closer to the original ones.

Cosmological Consistency: The power spectrum error of the proposed PI-VAE model is 7×10^{-6}, which is much lower than that of both U-Net VAE and traditional VAE: 1×10^{-5} and 4.96×10^{-1}.

This shows that the physics-based losses are helping the PI-VAE model to maintain the clustering statistics elements of the cosmic web, which the other models fail to achieve.

Structural Fidelity: The PI-VAE model's gradient error is 7.4×10^{-4}, which is significantly lower than the other models and indicates the loss term helps PI-VAE in reconstructing sharp and physically realistic structures.

On the other hand, both U-Net VAE with skip connections and traditional VAE lack the physical constraints required to precisely control spatial variations. From Table 5.2, it can be seen that the relative mean error for U-Net VAE is quite higher: 0.24508472 for Ω_m and 0.6443768 for σ_8.

This indicates that the U-Net VAE is significantly ignoring the information for the cosmological parameters, and the reconstruction fails to be physically informed. On the other hand, PI-VAE achieves relatively lower relative errors, 0.09589839 and 0.31905783 for Ω_m and σ_8 respectively, which proves the PI-VAE model is being physically well reconstructed.

Therefore, the results show the PI-VAE model balances accuracy and physical consistency, unlike U-Net VAE and VAE. The PI-VAE provides superior performances that is driven by the physics laws and not related to any architectural shortcuts. In terms of both reconstruction and parameter prediction, PI-VAE outperforms both U-Net VAE and VAE.

5.3 Visual Comparison Analysis

5.3.1 Visual Reconstruction Comparison

Figure 5.1 shows the reconstruction of a cosmological field using a traditional VAE. In cosmological simulations, traditional VAE often fails to reconstruct fields both visually and meaningfully. The traditional VAE model does not use the skip connection, which causes the latent bottleneck to lose severe information. Hence, the reconstruction is blurry (Figure 5.1) and it fails to reconstruct the fine filaments of the cosmic web.

While U-Net VAE uses the skip connection, which passes the spatial information to the decoder, it improves the visual sharpness drastically and recovers the intricate web-like structure (Figure 5.2), which looks similar to the original input image. The U-Net VAE residual (error) map also shows low error across the cosmic web.

The model's architecture often cheats. Unlike the PI-VAE, which receives the physically constrained latent space as input and reconstructs the images based on the physically informed latent space, the U-Net VAE through skip connections "peeks" into the original input images and reconstructs the image based on that.

Hence, the U-Net VAE reconstruction may look visually similar to the original input but the model likely "cheats" its way to construct it. On the other hand, the PI-VAE model combines the skip connection with the physics loss and produces the most accurate reconstruction.

Figure 5.3 shows that the residual map is almost white, which indicates that PI-VAE successfully captured both the visual structure and accurate physical values.

Moreover, the matter power spectrum analysis (Figure 5.7) shows that the model preserves both large and small scale matter distribution; hence, the model didn't lose any important signal while compressing the information into a 256-dimensional latent space.

On the other hand, both U-Net VAE and VAE lose important physical information during the compression. Figure 5.8 shows that the U-Net VAE matter distribution power curve has a slight mismatch from the ground truth, which indicates that, even though the reconstruction of U-Net (Fig. 5.2) looks visually similar to the original field, the model does lose physical fidelity during compression.

The Matter Power Spectrum of VAE (Fig. 5.9) indicates that the model loses a significant amount of cosmic information, which makes the VAE reconstruction visually and physically inconsistent. Therefore, the PI-VAE model, compared to the other models, maintains physical consistency throughout the compression and reconstruction.

Figure 5.10 shows the comparison between the density distribution of the PI-VAE model. The original and reconstructed simulation density distributions are almost similar. This indicates that the model successfully preserves the large dynamic range as well as the high-density peaks.

Therefore, the PI-VAE successfully interprets and enforces physical laws, making the reconstruction not only visually similar to the original cosmic field but also physically meaningful.

While the U-Net reconstruction may look the same as the original, the residual error shows that, compared to PI-VAE, the model provides slightly poorer reconstruction and lacks physics-related information.

VAE, on the other hand, performs the weakest in both reconstruction and preserving meaningful physical information. Hence, the proposed PI-VAE model is ideal for meaningful reconstruction that looks like and preserves information closer to the real universe.

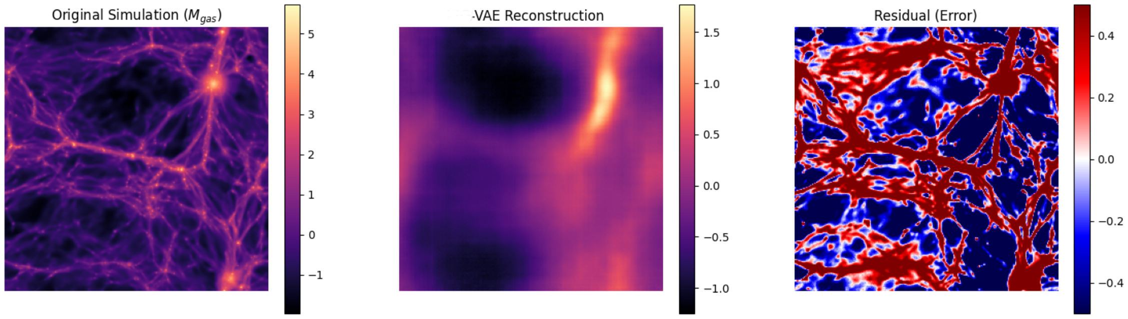

Figure 5.1: Reconstruction of Cosmological Field using traditional

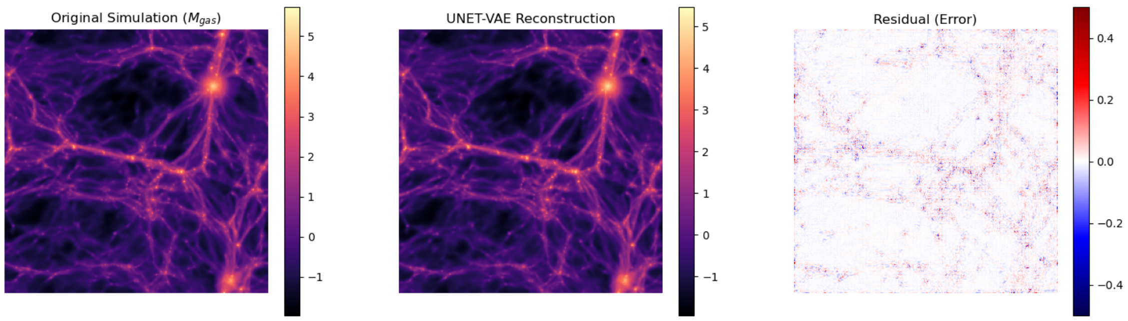

Figure 5.2: Reconstruction of Cosmological Field using UNet

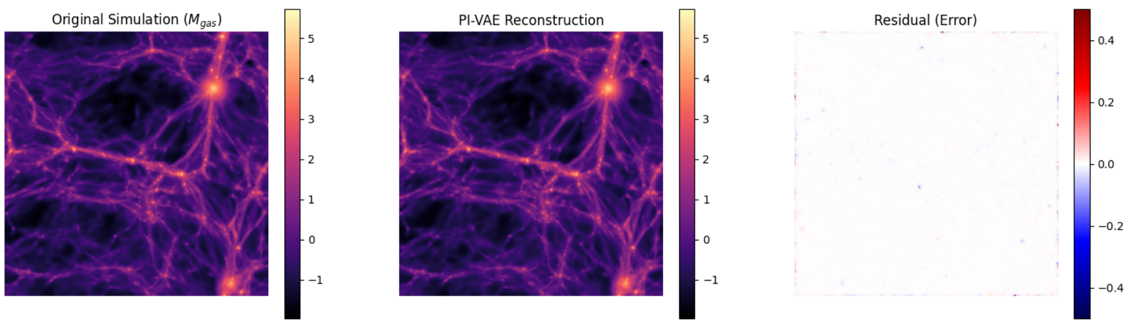

Figure 5.3: Reconstruction of Cosmological Field using PI-VAE

5.3.2 Visual Comparison of Parameter Prediction

In order to evaluate the capability of the proposed PI-VAE model to infer the parameters of the cosmological model qualitatively, we visually compare the predicted and true values of the cosmological parameters Ω_m and σ_8 on the test dataset.
Figure 5.4 represents the scatter plot and residual distribution of PI-VAE. The scatter plot shows that the predicted Ω_m closely follows the ideal line, which indicates that the model captures physical features governing matter density with high accuracy. Although there is more scatter, the predictions for σ_8 also show an expected pattern, reflecting the sensitivity of σ_8 to small-scale and nonlinear structures.
Again, the tight residual distribution for Ω_m confirms its high prediction accuracy. In contrast, the broader residuals for σ_8 represent the challenges of capturing complex, nonlinear cosmic structures and baryonic noise.
The performance of PI-VAE is far better than the U-Net VAE and VAE models, as seen from Figures 5.5 and 5.6, respectively.
Both of the later-mentioned models show that the scatter interpretation of Ω_m and σ_8 is more dispersed, and their residual distributions are more skewed than the PI-VAE model (Figure 5.4).

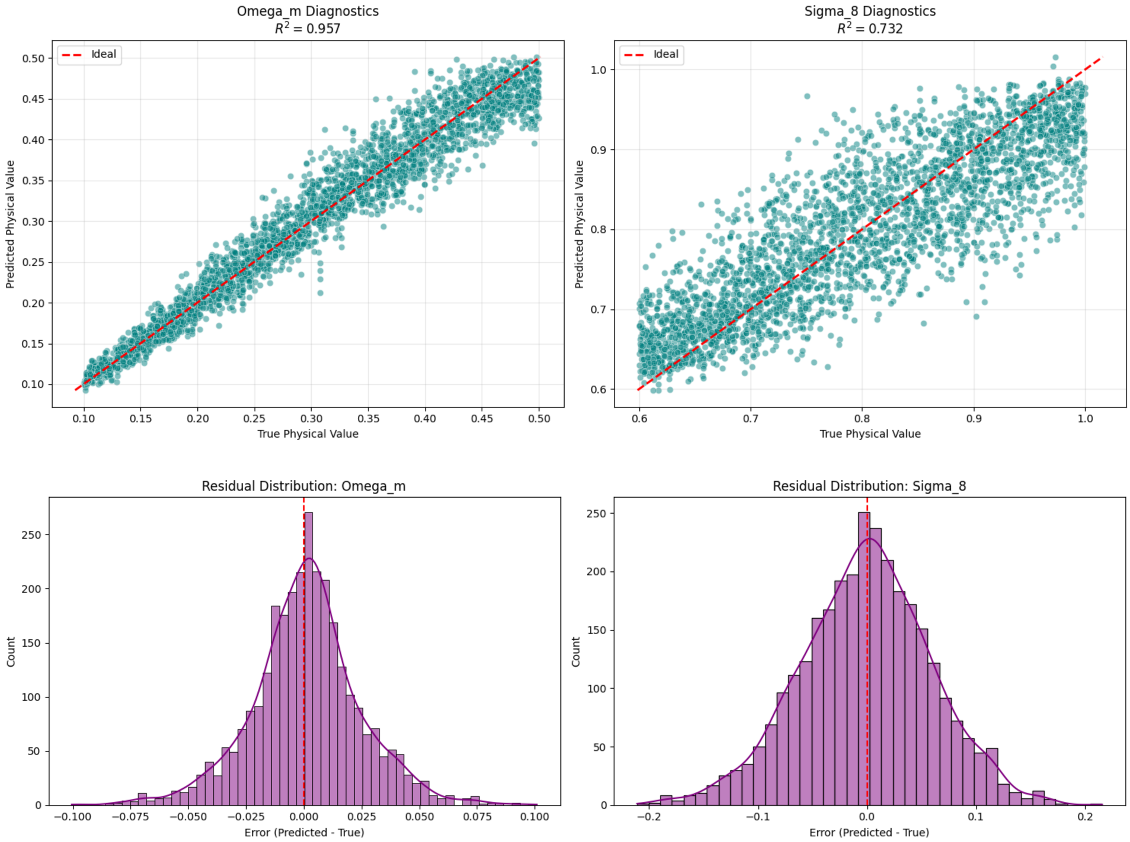

Figure 5.4: PI-VAE scatter plot and residual distribution for Ω_m and σ_8

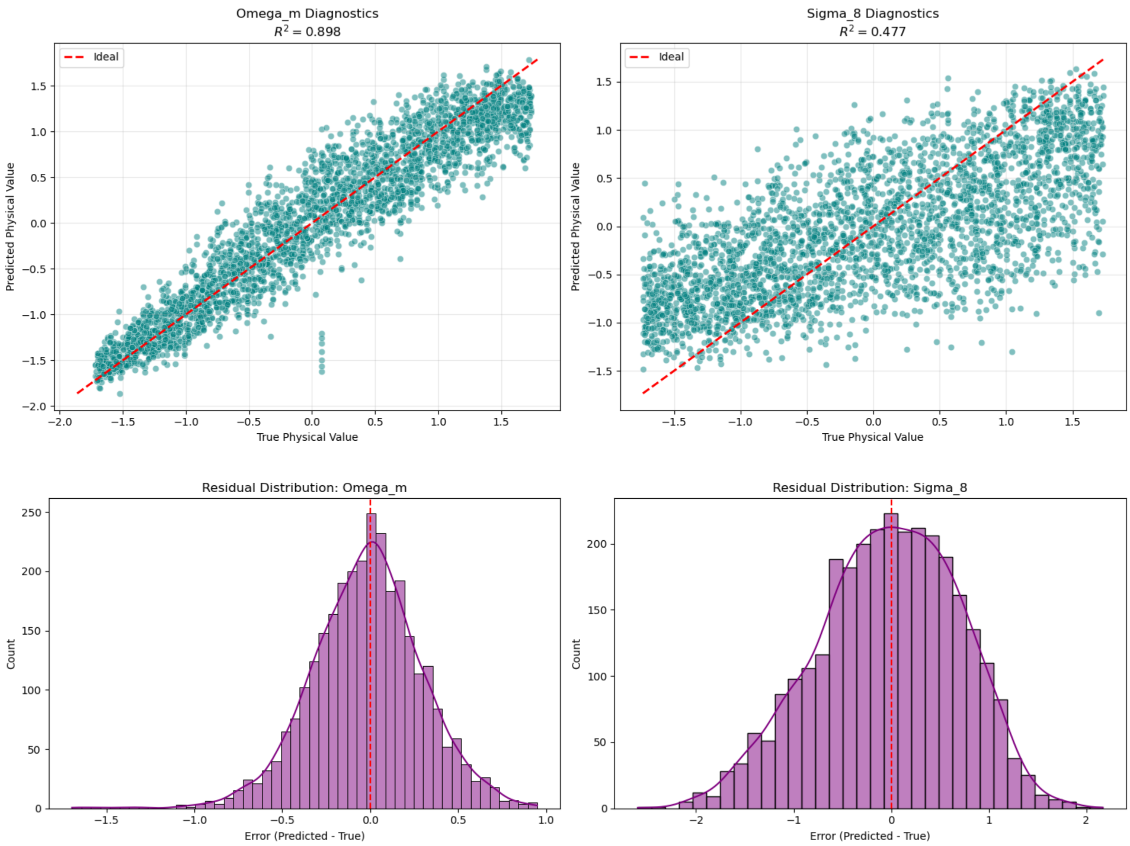

Figure 5.5: UNet-VAE scatter plot and residual distribution for Ω_m and σ_8

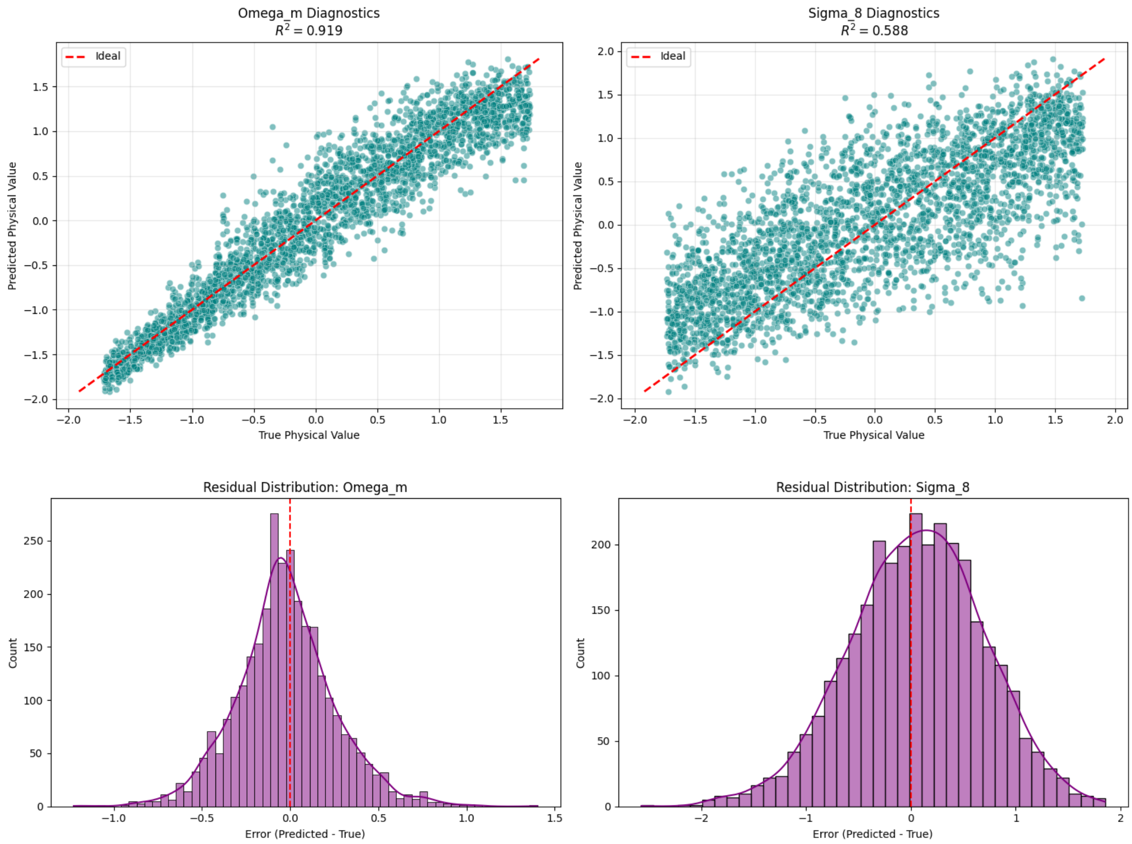

Figure 5.6: VAE scatter plot and residual distribution for Ω_m and σ_8

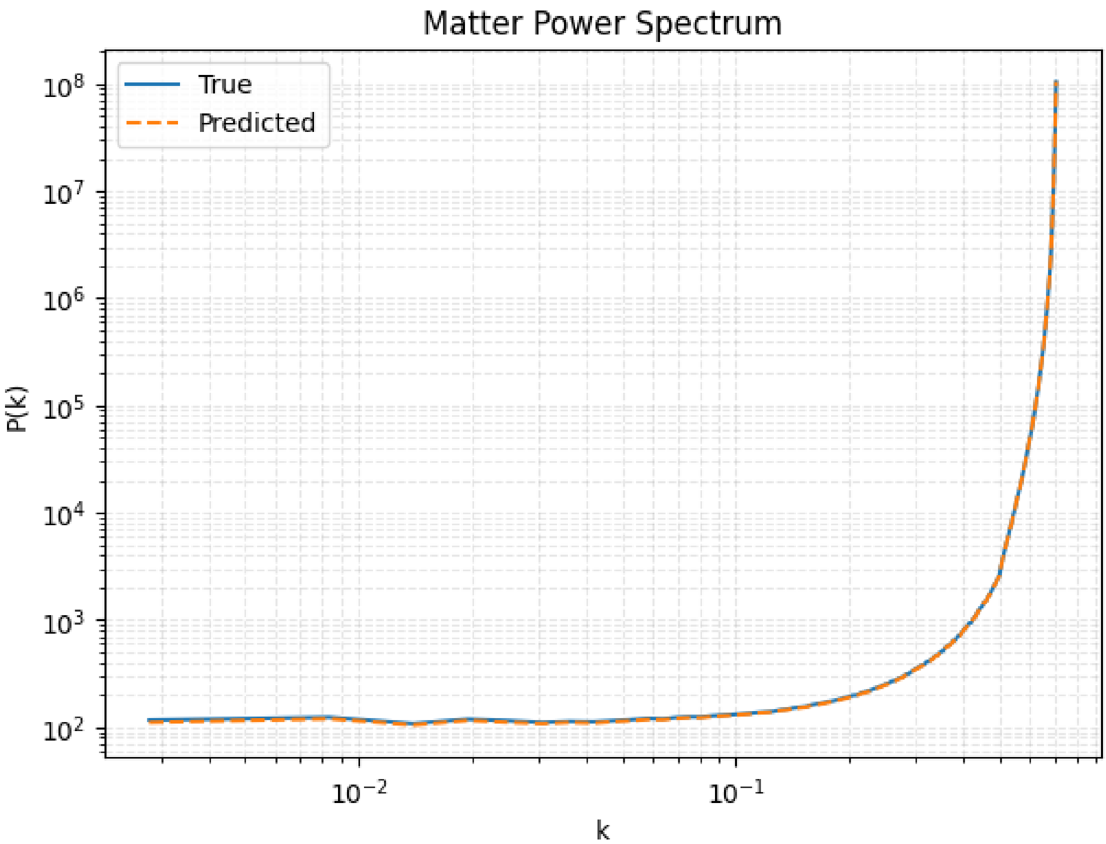

Figure 5.7: Matter Power Spectrum of PI-VAE

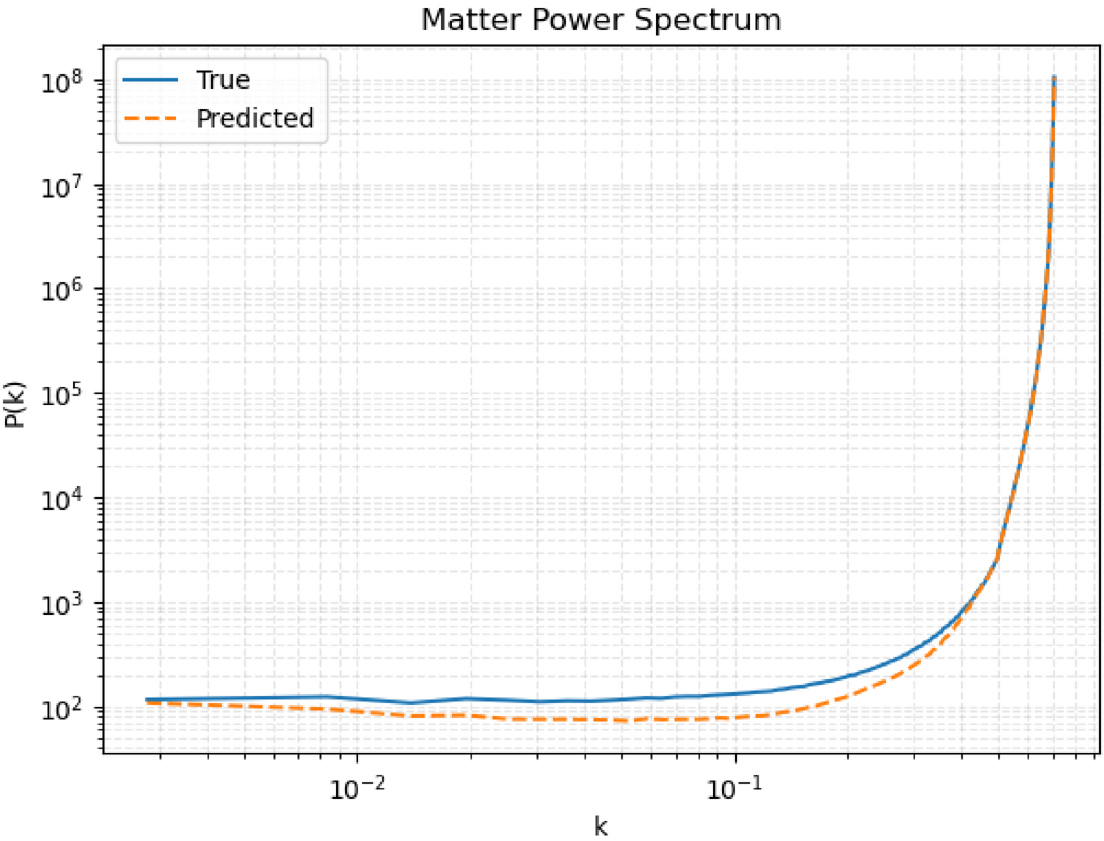

Figure 5.8: Matter Power Spectrum of UNet-VAE

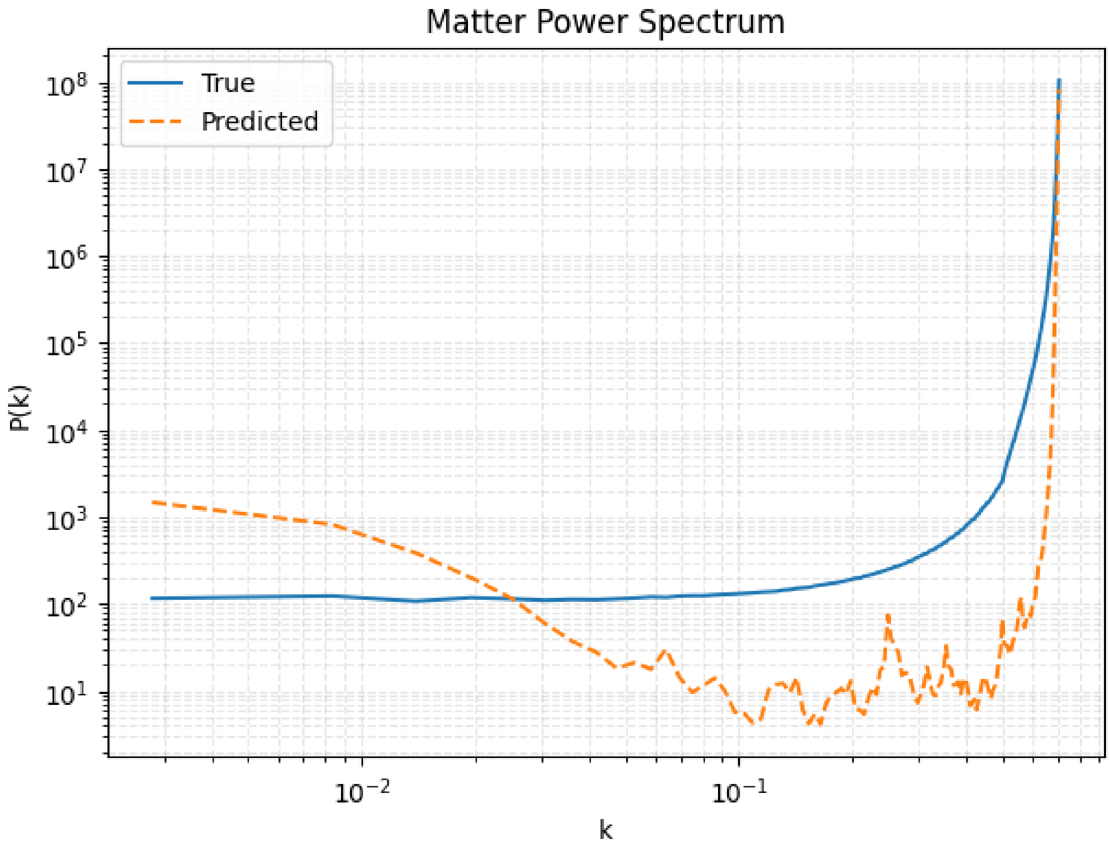

Figure 5.9: Matter Power Spectrum of VAE

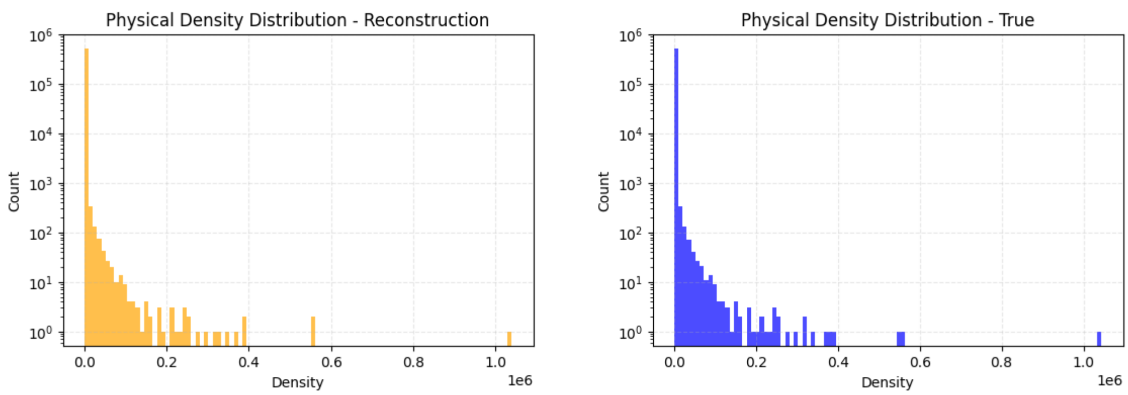

Figure 5.10: Physical Density Distribution of PI-VAE

5.4 Advanced Model Diagnostics

While with the visual comparisons in Section 5.3.1 we were able to see what the reconstructions look like, we have to understand how the models achieve these results. Deep learning models are often "lazy": they will be looking for the easiest way to decrease the loss function, even if that means cheating. To test whether our PI-VAE is truly learning the laws of physics or simply memorizing the shapes, we examined the error metrics with the help of advanced diagnostic plots.

5.4.1 The "Holistic" Performance Examination

We start by examining the "fingerprint" for each model. It's not uncommon for a model to be good at one measure (image sharpness) but horrible at another one (physical accuracy).

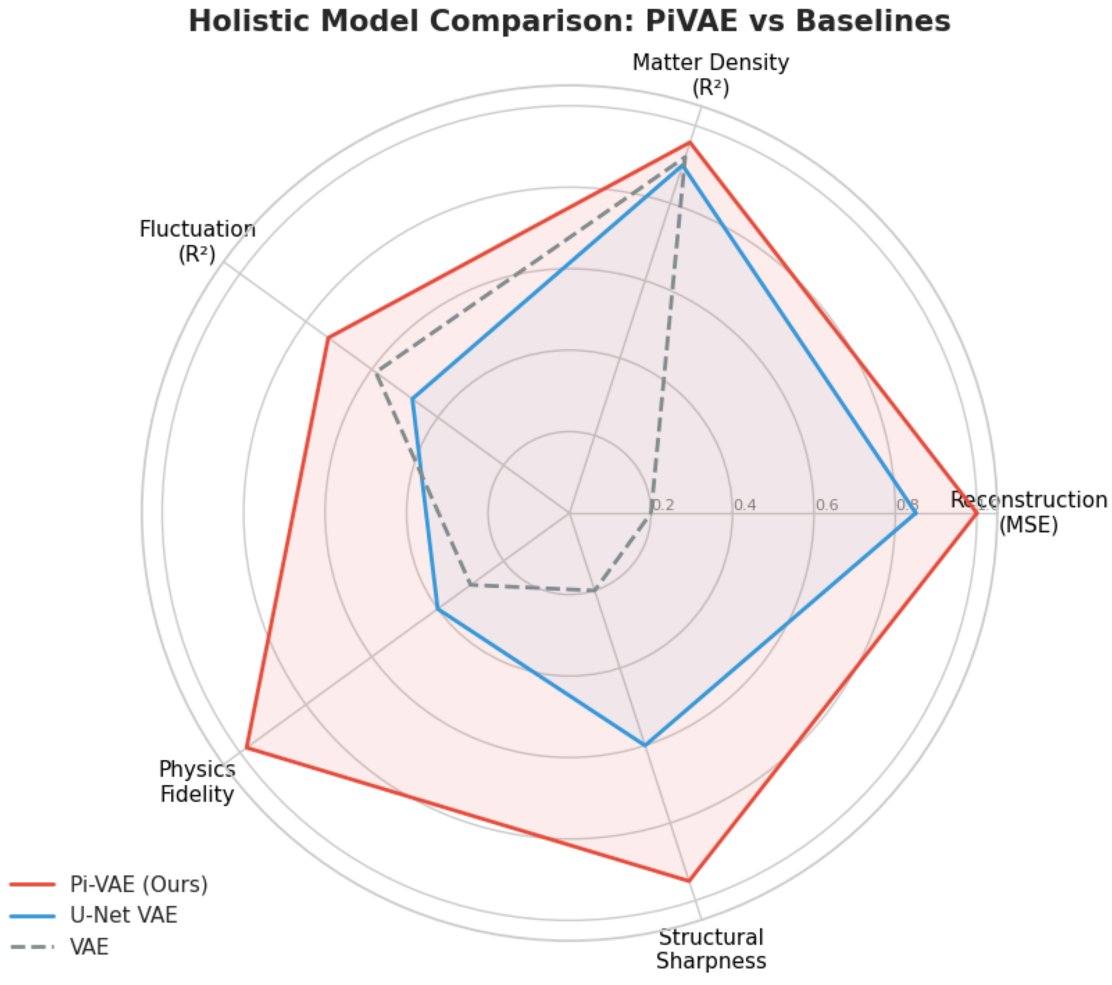

Figure 5.11: Comparison of Holistic Models Radar Chart

As appears in Figure 5.11, the three models display different behaviors:

- **The U-Net VAE (Blue profile):** The shape is skewed. It reaches out far on the "Reconstruction" axis; in other words, it creates very sharp images. However, it falls off on the "Physics Fidelity" axis. This confirms something that is known in computer vision as “shortcut learning” (or “information leakage”). The U-Net then makes use of its skip connections to avoid the latent bottleneck by copying high-frequency details directly from input to output. It makes a nice picture, however, and has not actually “solved” the physics.

- **The VAE baseline (Gray profile):** This model has a small footprint overall. Without the skip connections, it is forced to compress the entire universe into a small vector and hence the “blurring” effect we saw earlier. It acts like a low-pass filter, which smooths out the cosmic filaments into blobs.
- **The PI-VAE (Red profile):** This is the most balanced shape that also covers the largest area. It shows good performance on reconstruction and physics at the same time. This goes to prove that we do not have to choose between good-looking images” and “scientific accuracy.” By explicitly penalizing when the model violates the power spectrum, we are asking it not to ignore the multi-scale topology of the data.

5.4.2 The “Danger Zone” Analysis

It is dangerous in scientific machine learning to have a model that looks right and calculates wrong. We represent this risk in a plot where we plot the reconstruction error against the physics error.

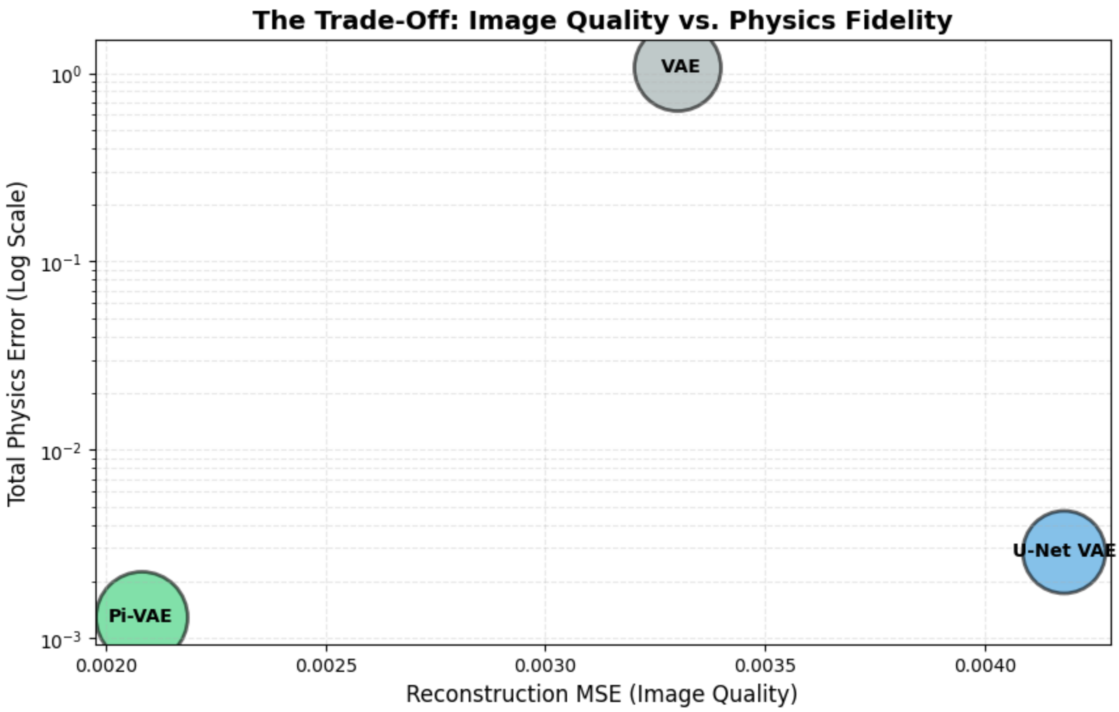

Figure 5.12: Accuracy–Efficiency Trade-off Bubble Chart

Figure 5.12 shows the “Danger Zone” on the top-left corner:

- **High-Fidelity, Low-Physics (The U-Net VAE):** The U-Net is in the top-left position. It has low image error (X-axis) and high physics error (Y-axis). If a cosmologist used this model, they would obtain some nice-looking simulations that provided wrong values for the dark matter density (Ω_m).

- **The Sweet Spot (The PI-VAE):** The PI-VAE is the only model in the bottom-left corner with both low image error and low physics error.

- **Inference Power (Bubble Size):** The size of the bubbles corresponds to the accuracy of parameter prediction. The PI-VAE bubble is much bigger than the others. This implies that since the PI-VAE was forced to learn the structure of the Cosmic Web (filaments and voids), the latent space contained meaningful information about the composition of the universe, which made it easier to perform regression.

5.4.3 Identifying the Errors in Physics

We decomposed the errors into certain physical constraints: power spectrum (clustering), mean density (mass conservation), and gradient (sharpness).

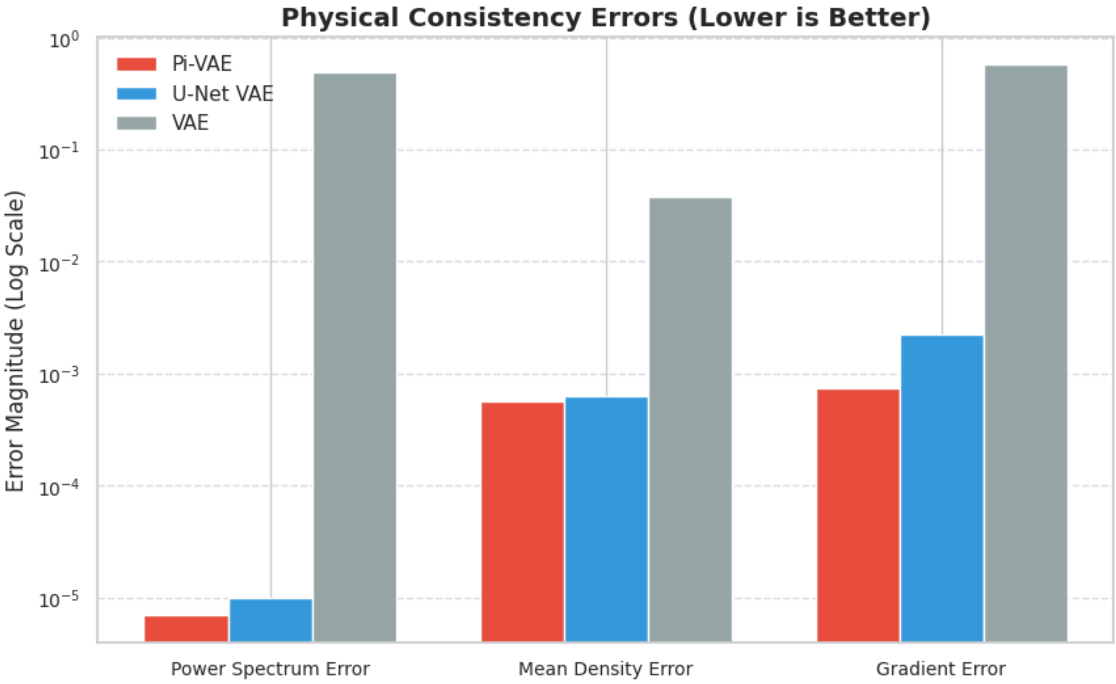

Figure 5.13: Comparison of Physical Consistency Error (Log Scale)

Figure 5.13 uses a logarithmic scale as the difference in performance is massive (orders of magnitude).

- **Power Spectrum Error:** This is where the VAE baseline and U-Net VAE have huge errors. They cannot reflect the right distribution of galaxies at different scales.

- **Gradient Error:** This is the error of sharpness of the "Cosmic Web" filaments. In this part, the U-Net has high error because, although it copies pixels, it doesn't respect the derivative of the field: the sharp transitions between a void and a filament.

- **PI-VAE Dominance:** The PI-VAE bar is barely visible. This is a direct effect of our loss function: by having the model minimize the difference in the power spectrum $P(k)$, we effectively teach it the "rules" of how matter clumps together under gravity.

5.4.4 Learning Dynamics: Convergence and Plateau

Finally, we analyzed the learning process over time. Did the models learn physics instantly or struggle with it?

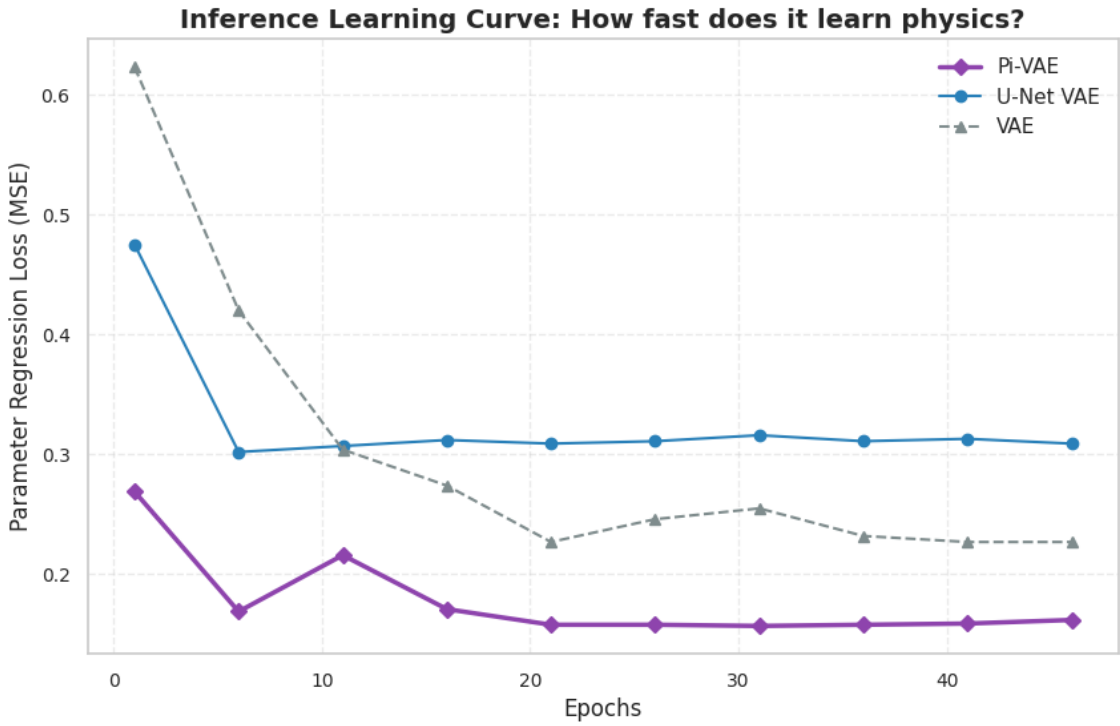

Figure 5.14: Inference of Parameters Learning Curve

Figure 5.14 monitors the error in testing the predictions of Ω_m and σ_8 for the 50 epochs.

- **The U-Net Plateau:** Notice the blue line (U-Net): it drops initially but then hits a "floor" and does not get better. This is the "lazy optimizer" effect. Once the model learned to use the skip connections to do the image reconstruction task, it no longer attempted to organize the latent space. It didn't have an incentive to learn the parameters better because its reconstruction loss was low already.
- **The PI-VAE Convergence:** The red/purple line (PI-VAE) is still steadily decreasing. Because the physics loss can't be cheated by using skip connections, the model was forced to continue learning the latent space. This led to a "continuous learning" phase where the model slowly but surely learned the relationship between the latent variables and the cosmological parameters.

5.4.5 Differential Parameter Sensitivity: Ω_m vs. σ_8

A critical finding in our research was the fact that not all cosmological parameters are equally easy to learn. We therefore visualized the performance of the regression in terms of matter density (Ω_m) and fluctuation amplitude (σ_8) separately, to understand where the models have difficulties.

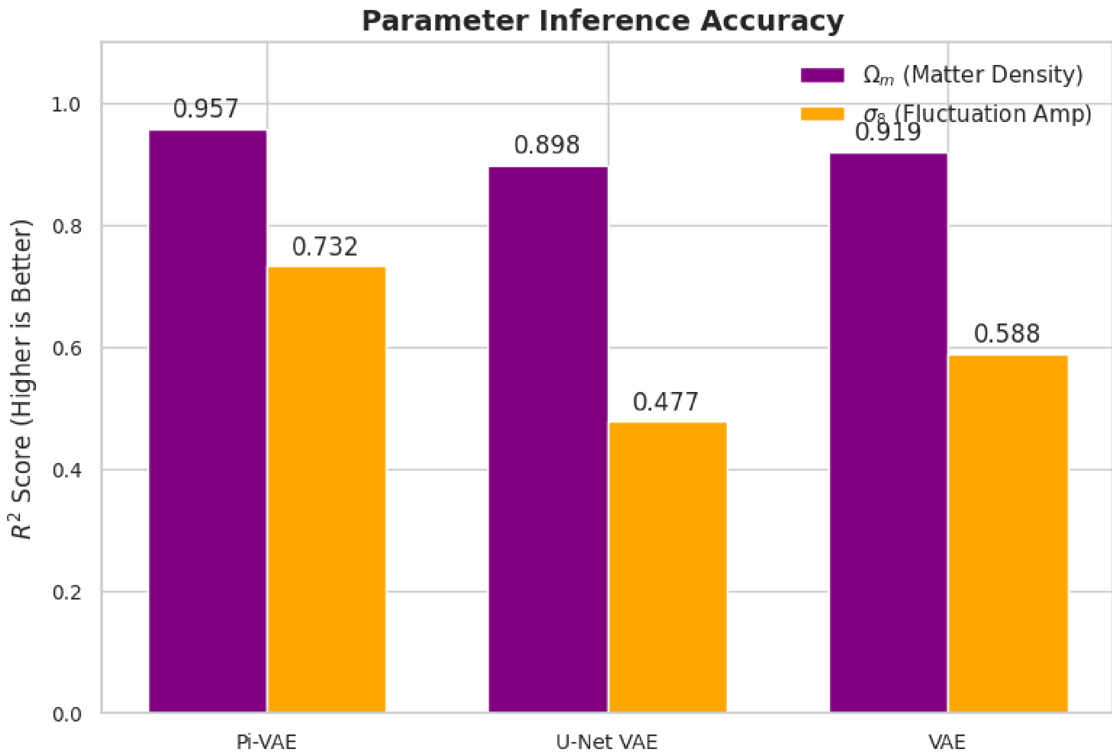

Figure 5.15: Inference Accuracy Breakdown

As can be seen in Figure 5.15, there is a significant performance gap between the two parameters:

- Ω_m **(Matter Density):** All models do relatively well when it comes to this (high bars). This is because Ω_m determines the total number of bright pixels (galaxies) in the image. Even a simple VAE baseline can effectively "count" the number of bright pixels to infer the density.
- σ_8 **(Clumpiness):** The performance decreases considerably for the baseline models. σ_8 is a measure of the clustering of matter. Learning this requires knowledge of texture and frequency content, and not just pixel intensity.

The PI-VAE advantage: The PI-VAE achieved an R^2 score of 0.732 for σ_8, while the U-Net VAE had a score of 0.477. This supports the use of the power spectrum loss L_{Pk}. Since $P(k)$ explicitly measures clustering at different scales, it provides a direct supervisory signal for learning σ_8 and bridges the gap where the other models fail.

5.4.6 Reconstruction Stability and Convergence

While Section 5.4 demonstrated how the model learns the model parameters, we also need to verify that the generation of images is stable. We therefore tracked the reconstruction loss throughout training to determine whether the added physics constraints caused any instability.

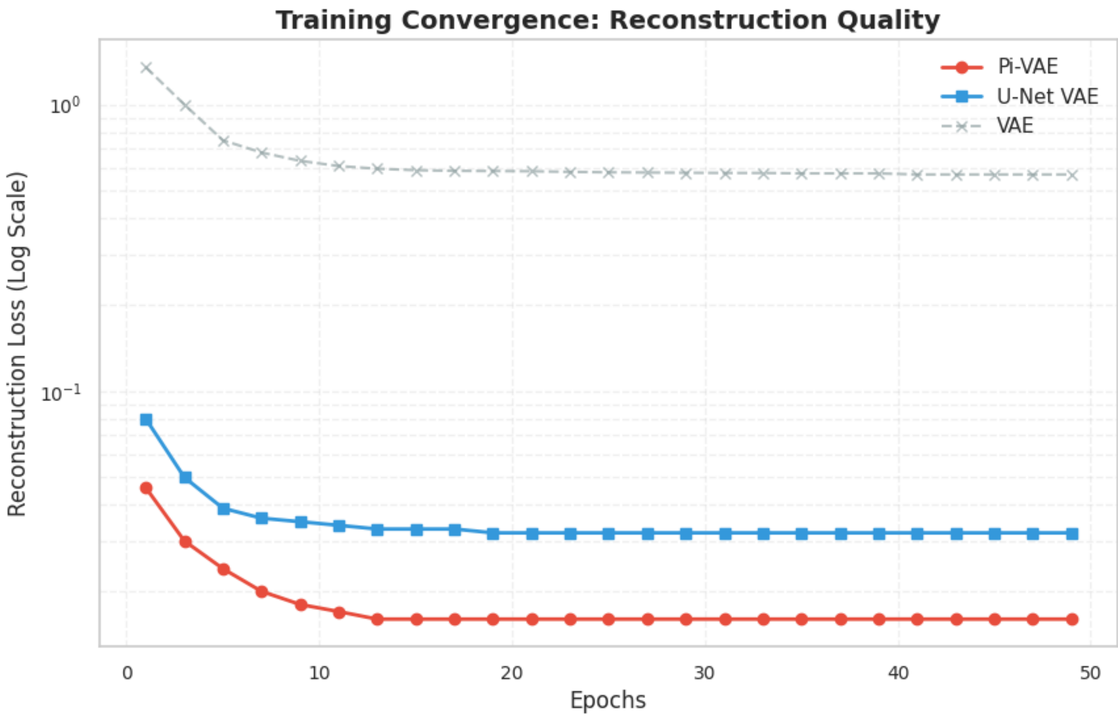

Figure 5.16: Reconstruction Loss Learning Curve

Figure 5.16 shows the training curve for image reconstruction:

- **VAE (Gray line):** The loss remains relatively high because the model must compress the complex cosmic web through a narrow latent bottleneck, which leads to blurred reconstructions.
- **U-Net VAE (Blue line):** The loss decreases quickly; as discussed earlier, this is largely due to the skip-connection "shortcut."
- **PI-VAE (Red line):** Crucially, the PI-VAE converges as quickly as the U-Net and reaches a comparable final reconstruction loss (≈ 0.016). This indicates that adding the physics-informed losses did not destabilize training. Instead, the physics constraints provide "guide rails" that help the optimizer find a better solution than the Vanilla VAE.

5.4.7 Measuring the Improvement in Performance

Finally, to summarize the contribution of this thesis, we calculated the relative percentage improvement of our PI-VAE compared to the second-best performing baseline in each category.

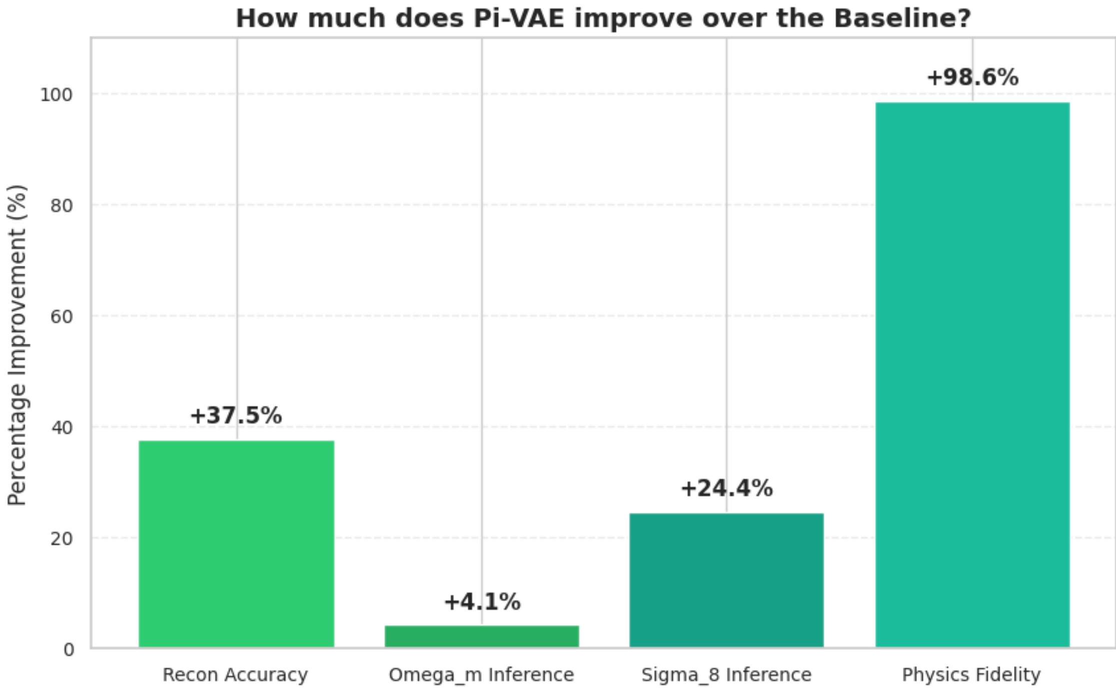

Figure 5.17: Percentage Improvement over Baselines

Figure 5.17 summarizes the results. The PI-VAE does not merely offer marginal gains; it provides a transformative improvement in scientific accuracy:

- **Physics fidelity:** The physics error is reduced by 98.6%. This is the most significant result of the thesis, confirming that standard deep learning models are physically inconsistent unless explicitly constrained.
- **Inference accuracy:** We obtain a 24.4% improvement in predicting σ_8.
- **Reconstruction quality:** Even for image quality, we improve by 37.5% over the VAE baseline.

This waterfall chart attests to the fact that for high-precision cosmology, a physics-informed approach is not a theoretical preference; it is a practical necessity.

5.4.8 Conclusion of Analysis

The advanced diagnostics make it clear that the superior performance of the PI-VAE is not a statistical fluke. By looking at the holistic radar chart (Figure 5.11), the physics error bars (Figure 5.13), and the overall improvement summary (Figure 5.17), we can see that standard deep learning models (VAE, U-Net) are not able to consistently capture the scientific properties of the data. The success of the PI-VAE arises from explicitly enforcing the governing physics: instead of memorizing pixel intensities, the PI-VAE learns the multi-scale clustering and conservation constraints that define physically plausible cosmological fields.

5.4.9 Conclusion of Analysis

The advanced diagnostics make it clear that the superior performance of the PI-VAE is not a statistical fluke. By looking at the holistic radar chart (Figure 5.11), the physics error bars (Figure 5.13), and the overall improvement summary (Figure 5.17), we can see that standard deep learning models (VAE, U-Net) are not able to consistently capture the scientific properties of the data. The success of the PI-VAE arises from explicitly enforcing the governing physics: instead of memorizing pixel intensities, the PI-VAE learns the multi-scale clustering and conservation constraints that define physically plausible cosmological fields.

Chapter 6

Conclusion

6.1 Summary of Findings

This research has managed to develop and validate a Physics-Informed Variational Autoencoder (PI-VAE) designed for analysis of complex cosmological simulations. Our main aim was to build a model capable of compressing high-dimensional cosmic data, accurately reconstructing it, and predicting key parameters (Ω_m and σ_8) without violating the laws of physics.
To do so, we used the CAMELS IllustrisTNG dataset, considering three different physical fields: Gas Mass (M_{gas}), Gas Velocity (V_{gas}), and Magnetic Fields(V_{B}). By combining these three inputs, we provided the model with a richer dataset, giving it a better understanding of the universe than using a single image type would.

Major findings of our experiments were:
The PI-VAE was able to perform considerably better than standard deep learning models. It gave an R^2 score of 0.957 for predicting Matter Density (Ω_m), meaning that it is able to predict the amount of matter in the universe with near-perfect accuracy. For the more difficult Clustering Amplitude parameter (σ_8), it obtained an R^2 of 0.732, which was significantly higher than U-Net VAE (0.477) and Traditional VAE (0.588). This proves that adding physics-based constraints helps the model learn the actual rules of cosmology rather than just guessing patterns.

One of the main results discovered was the difference in reconstruction quality. While the U-Net VAE baseline provided sharp images with greater visual acuity, in our analysis we found that the baseline was essentially "cheating" via skip connections that copy details, but do not truly understand the physics. In contrast, our PI-VAE not only generated accurate reconstructions visually, but they were also mathematically consistent. Our Power Spectrum analysis demonstrated that the PI-VAE preserves the statistical distribution of matter on both large and small scales, with an error of only 7×10^{-6}, much lower than the baselines.
We discovered that the common types of loss functions (such as MSE) are insufficient for scientific data. By adding the effects of certain physics losses (like Power Spectrum Loss and Mean Density Loss) we were able to force the model to obey the conservation of mass and the structural boundaries. The results showed that these constraints helped prevent the "blurriness" that is common for traditional VAEs and made sure that the model was able to catch tiny filaments in the cosmos and voids with great accuracy.
In summary, our findings have shown that a typical DL model can be made into

a reliable scientific tool, by integrating physical laws into the learning mechanism of the model. The PI-VAE provides a powerful solution that balances the need for speed of machine learning and the accuracy needed for cosmological research.

6.2 Contributions to the Field

Our research is a significant contribution to the field of Computational Cosmology by proposing a Physics-Informed Variational Autoencoder (PI-VAE). Traditionally, the analysis of the huge quantity of data from cosmic simulations like CAMELS is computationally expensive and slow. While this can be accelerated using standard machine learning models, there is a chance that the model will become a "black box" that may yield statistically probable but physically impossible results. Our main contribution is being able to solve this problem by incorporating actual physical laws into the learning process of the neural network.

Rather than just training the model to correlate the pixels, what we did is force it to abide by the cosmological principles. We achieved this by specifying our own loss function, which included the Power Spectrum (how matter serves as a fingerprint of the universe) and Mean Density. This ensures that when our model reduces the complicated cosmic web into simplified parameters (Ω_m and σ_8), the physics is not violated.

Furthermore, we showed that this approach is effective in the case of the IllustrisTNG simulation suite, and we were able to predict the total matter density with high accuracy ($R^2 \sim 0.95$). This proves that "Physics-Informed" Machine Learning is a viable tool for the future of sky surveys, providing an approach that is incredibly fast and scientifically trustworthy when compared to traditional simulation-based inference methods.

6.3 Limitation

While our Physics-Informed Variational Autoencoder (PI-VAE) succeeded in ensuring the compression of the cosmic data and the prediction of parameters, there are several limitations to be considered.

First of all, our model is completely based on 2D maps (slices) of the universe. Since the cosmic web is a complex 3D structure in nature, its projection into two dimensions will inevitably result in some loss of information regarding the 3D structure, which probably limits our parameter inference contribution.

Second, looking at our results of the regression, there is a performance gap between the two cosmological parameters. While the model is very good at predicting the total matter density (Ω_m) with very high accuracy ($R^2 \sim 0.95$), it still struggles to predict the higher clustering amplitude (σ_8) with a score of ($R^2 \sim 0.73$). This suggests that the model can be improved by integrating related physics laws that put more focus on σ_8 for higher prediction accuracy.

Furthermore, we were only trained and tested on the IllustrisTNG simulation suite of the CAMELS dataset. Therefore, there is a possibility that this model would not perform the same on a different simulation suite (like SimBA), and more importantly, on messy, noisy observational data from real telescopes.

Finally, our computing resources limited us to run the model with a batch size of no more than 16 and with an image resolution of 256×256, which may have prevented our model from learning larger-scale cosmic correlations.

6.4 Recommendations for Future Work

To build upon this thesis, the most logical next step is to move from 2D images to 3D volumetric data. By applying 3D Convolutional Neural Networks (3D-CNNs) in the VAE architecture, it will be possible to capture the full geometry of cosmic filaments and voids, which would likely enhance the prediction of the clustering parameter (σ_8).

To ensure the model is robust and physically correct across cosmological simulations, cross-simulation tests using other simulation sets (such as Astrid and Simba) should be performed. If the model fails, techniques like "Domain Adaptation" can be used to help it generalize.

We also intend to improve the "Physics-Informed" part of the network. Currently, we implement consistency using the Power Spectrum (2-point correlation). In the future, more complex statistics such as the Bispectrum (3-point correlation) could be added to the loss function, in order to account for the non-Gaussian features of the cosmic web.

Finally, the last part of the process is to bridge the gap between simulation and reality. This work could be continued in the future by applying techniques such as Domain Adaptation or Transfer Learning to adapt this pre-trained model to work on actual observational data from upcoming surveys such as the Euclid mission or the LSST. This way, we would be able to use the speed of our ML model for analyzing the real universe, instead of just a simulated model.

Bibliography

[1] M. A. Aragon-Calvo, “Classifying the large-scale structure of the universe with deep neural networks,” *Monthly Notices of the Royal Astronomical Society*, vol. 484, no. 4, pp. 5771–5784, Apr. 2019. DOI: 10.1093/mnras/stz393. [Online]. Available: https://doi.org/10.1093/mnras/stz393.

[2] K. Yi, Y. Guo, Y. Fan, J. Hamann, and Y. G. Wang, “Cosmovae: Variational autoencoder for cmb image inpainting,” *arXiv e-prints*, 2020. arXiv: 2001.11651 `[astro-ph.CO]`. [Online]. Available: https://arxiv.org/abs/2001.11651.

[3] F. Villaescusa-Navarro et al., “The camels project: Cosmology and astrophysics with machine-learning simulations,” *Astrophysical Journal*, vol. 915, no. 1, p. 71, Jul. 2021. [Online]. Available: https://arxiv.org/abs/2010.00619.

[4] A. Akhmetzhanova, S. Mishra-Sharma, and C. Dvorkin, “Data compression and inference in cosmology with self-supervised machine learning,” *Monthly Notices of the Royal Astronomical Society*, vol. 527, no. 3, pp. 7459–7481, Nov. 2023. DOI: 10.1093/mnras/stad3646. [Online]. Available: https://doi.org/10.1093/mnras/stad3646.

[5] S. Andrianomena and S. Hassan, “Latent space representations of cosmological fields,” *arXiv e-prints*, 2023. arXiv: 2311.00799 `[astro-ph.CO]`. [Online]. Available: https://arxiv.org/abs/2311.00799.

[6] A. Asensio Ramos, M. Mark, I. Chifu, and R. Gafeira, “Machine learning in solar physics,” *Living Reviews in Solar Physics*, vol. 20, no. 1, Jul. 2023. DOI: 10.1007/s41116-023-00038-x. [Online]. Available: https://doi.org/10.1007/s41116-023-00038-x.

[7] Z. Dai, “Statistical machine learning and physics-informed neural networks in astronomical simulations,” *Texas Digital Library*, Aug. 2023, Doctoral dissertation. [Online]. Available: https://uh-ir.tdl.org/items/9e374a0e-8eb4-44ea-8270-29ebb08461b2.

[8] Z. Dai, B. Moews, R. Vilalta, and R. Davé, “Physics-informed neural networks in the recreation of hydrodynamic simulations from dark matter,” *Monthly Notices of the Royal Astronomical Society*, vol. 527, no. 2, pp. 3381–3394, Nov. 2023. DOI: 10.1093/mnras/stad3394. [Online]. Available: https://doi.org/10.1093/mnras/stad3394.

[9] A. Gagliano and V. V. Ashley, “A physics-informed variational autoencoder for rapid galaxy inference and anomaly detection,” *arXiv e-prints*, Dec. 2023. arXiv: 2312.16687 `[astro-ph.GA]`. [Online]. Available: https://arxiv.org/abs/2312.16687.

[10] C. Cuesta-Lazaro, A. Bayer, M. Albergo, S. Mishra-Sharma, C. Modi, and D. Eisenstein, "Joint cosmological parameter inference and initial condition reconstruction with stochastic interpolants," *NeurIPS Workshop on Machine Learning and the Physical Sciences*, 2024, Accessed: Jan. 25, 2026. [Online]. Available: https://ml4physicalsciences.github.io/2024/files/NeurIPS_ML4PS_2024_162.pdf.

[11] A. Ojha, S. Gleyzer, M. Toomey, and P. Reddy, "Lenspinn: Physics-informed neural network for learning dark matter morphology in lensing," *NeurIPS Workshop on Machine Learning and the Physical Sciences*, 2024, Accessed: Jan. 25, 2026. [Online]. Available: https://ml4physicalsciences.github.io/2024/files/NeurIPS_ML4PS_2024_78.pdf.

[12] V. Ono, C. F. Park, N. Mudur, Y. Ni, C. Cuesta-Lazaro, and F. Villaescusa-Navarro, "Debiasing with diffusion: Probabilistic reconstruction of dark matter fields from galaxies with camels," *arXiv e-prints*, Mar. 2024. arXiv: 2403.10648 `[astro-ph.CO]`. [Online]. Available: https://arxiv.org/abs/2403.10648.

[13] S. Ding, "Advancing cosmological field-level inference with physics-informed bayesian neural networks," *HAL Theses*, 2025, Doctoral dissertation, Accessed: Jan. 25, 2026. [Online]. Available: https://theses.hal.science/tel-05076806v1/file/146806_DING_2025_archivage.pdf.

[14] S. Dodelson and F. Schmidt, *Modern Cosmology*. Academic Press, 2025.

[15] S. Kannan, T. Qiu, C. Cuesta-Lazaro, and H. Jeong, "Cosmoflow: Scale-aware representation learning for cosmology with flow matching," *arXiv e-prints*, 2025. arXiv: 2507.11842 `[astro-ph.CO]`. [Online]. Available: https://arxiv.org/abs/2507.11842.

[16] J.-Y. Lee and F. Villaescusa-Navarro, "Cosmology with topological deep learning," *The Astrophysical Journal*, vol. 989, no. 1, p. 47, Aug. 2025. DOI: 10.3847/1538-4357/ade806. [Online]. Available: https://doi.org/10.3847/1538-4357/ade806.

[17] A. Lue, S. Genel, F. Villaescusa-Navarro, and M. Ho, "Cosmology with one galaxy: Autoencoding the galaxy properties manifold," *The Astrophysical Journal*, vol. 986, no. 2, p. 133, Jun. 2025. DOI: 10.3847/1538-4357/add724. [Online]. Available: https://doi.org/10.3847/1538-4357/add724.

[18] S. Mirzoyan, "Enhancing galaxy classification with u-net variational autoencoders for image denoising," *arXiv e-prints*, 2025. arXiv: 2506.19434 `[astro-ph.GA]`. [Online]. Available: https://arxiv.org/abs/2506.19434.

[19] H. S. Tan, "Inferring cosmological parameters with evidential physics-informed neural networks," *Universe*, vol. 11, no. 12, p. 403, Dec. 2025. DOI: 10.3390/universe11120403. [Online]. Available: https://doi.org/10.3390/universe11120403.

[20] Y.-S. Ting, "Deep learning in astrophysics," *ResearchGate*, Oct. 2025. [Online]. Available: https://www.researchgate.net/publication/396458495_Deep_Learning_in_Astrophysics.

[21] M. Yarahmadi and A. Salehi, “Towards a machine learning solution for hubble tension: Physics-informed neural network (pinn) analysis of tsallis holographic dark energy in the presence of neutrinos,” *The European Physical Journal C*, vol. 85, no. 11, Nov. 2025. DOI: 10.1140/epjc/s10052-025-14993-2. [Online]. Available: https://doi.org/10.1140/epjc/s10052-025-14993-2.