

# Enhancing Galaxy Classification with U-Net Variational Autoencoders. II. JWST High Redshift Galaxy Sample (Postprint)

**Authors:** S. S. Mirzoyan

**Date:** 2026-01-28T11:17:09+00:00

## Abstract

Building on our previous work, we apply a U-Net Variational Autoencoder (VAE) framework to denoise galaxy images from the James Webb Space Telescope and to enhance morphological classification. This study focuses on galaxies observed up to a redshift of ( $z \approx 8$ ), capturing them at early evolutionary stages where their faintness and structural complexity pose significant challenges for traditional classification methods. By mitigating observational noise, our approach enables the more reliable identification of morphological features, particularly in distinguishing between disk and non-disk galaxy types. We evaluate the denoising performance using standard image quality metrics and demonstrate that the enhanced images lead to improved classification accuracy across multiple deep learning models.

Our analysis of a sample of 292 galaxies up to ( $z = 7.69$ ) yields 83 galaxies classified as disk-like by the GCNN model with high confidence, of which approximately 70%-80% are at redshifts greater than 3. These findings suggest that disk-like structures can be prevalent in the early universe. The results highlight the potential of VAE-based denoising as a robust pre-processing step for analyzing high-redshift galaxy populations in ongoing and future astronomical surveys.

## Full Text

### Preamble

Research in Astronomy and Astrophysics, 26:025010 (5pp), 2026 February  
© 2025. National Astronomical Observatories, CAS and IOP Publishing Ltd. All rights, including for text and data mining, AI training, and similar technologies, are reserved. Printed in China.

<https://doi.org/10.1088/1674-4527/ae25c6>  
CSTR: 32081.14.RAA.ae25c6

## Enhancing Galaxy Classification with U-Net Variational Autoencoders. II. JWST High Redshift Galaxy Sample

S. S. Mirzoyan<sup>1,2</sup>

<sup>1</sup>Center for Cosmology and Astrophysics, Alikhanyan National Laboratory, 2 Alikhanyan Brothers str., Yerevan 0036, Armenia; mserg@yerphi.am

<sup>2</sup>Yerevan State University, Yerevan 0025, Armenia

Received 2025 July 30; revised 2025 October 30; accepted 2025 November 25; published 2025 December 24

### Abstract

Building on our previous work, we apply a U-Net Variational Autoencoder (VAE) framework to denoise galaxy images from the James Webb Space Telescope and enhance morphological classification. This study focuses on galaxies observed up to redshift  $z \approx 8$ , capturing them at early evolutionary stages where their faintness and structural complexity pose challenges for traditional classification methods. By mitigating observational noise, our approach enables the identification of morphological features, particularly in distinguishing between disk and non-disk galaxy types. We evaluate the denoising performance using standard image quality metrics and demonstrate that the enhanced images lead to improved classification accuracy across multiple deep learning models. Our analysis of a sample of 292 galaxies up to  $z = 7.69$  shows 83 galaxies classified as disk-like by the GCNN model with high confidence, and of those approximately 70%–80% are at redshifts greater than 3. These findings suggest that disk-like structures can be prevalent in the early universe. The results highlight the potential of VAE-based denoising as a robust pre-processing step for analyzing high-redshift galaxy populations in ongoing astronomical surveys.

**Key words:** methods: data analysis -techniques: image processing -Galaxy: disk -Galaxy: general -galaxies: distances and redshifts

### 1. Introduction

Machine learning (ML) has become deeply integrated into modern astronomical research, offering transformative capabilities for data analysis and interpretation [?, ?, ?, ?, ?, ?, ?, ?, ?]. Deep learning, in particular, has transformed the landscape of galaxy morphology studies, enabling scalable and accurate classification across large astronomical datasets. Recent efforts have demonstrated the efficiency of convolutional neural networks (CNNs) and group-equivariant architectures in capturing complex morphological features [?, ?, ?]. However, the performance of these models is often constrained by the quality of input data, particularly in the presence of noise, artifacts, and low signal-to-noise ratios—conditions that are especially prevalent in high-redshift observations.

The James Webb Space Telescope (JWST) has opened a new window into the early universe, providing unprecedented imaging depth and resolution. Recent catalogs, such as those compiled by [?], have enabled structural analysis of galaxies out to redshifts beyond  $z > 10$ , revealing a diverse population of compact, irregular, and evolving systems. These early-stage galaxies pose unique challenges for automated classification due to their morphological ambiguity and faintness. In addition to morphological studies, JWST data have spurred the development of ML tools for object identification across stellar and extragalactic domains. For instance, [?] introduced SESHAT, an XGBoost-based classifier capable of distinguishing young stellar objects, field stars, and galaxies using synthetic photometry adapted to JWST filters, achieving over 80% recall across classes without spatial priors. Such approaches highlight the growing role of data-driven methods in maximizing JWST's scientific return.

Spectroscopic surveys have also advanced our understanding of galaxy evolution during the reionization era. [?] present the largest spectroscopic catalog of [O III]+H $\beta$  emitters at  $6.75 < z < 9.05$ , constraining the [O III] luminosity function and revealing an accelerated decline in luminosity density between  $z \sim 7$  and  $z \sim 8$ . These findings underscore the importance of wide-area JWST slitless surveys for mapping large-scale structure and refining models of nebular line emission in the early universe.

Beyond the reionization epoch, JWST programs are probing galaxy properties at Cosmic Noon. [?] describe the Blue Jay survey, a deep NIRSpec spectroscopic campaign targeting 153 galaxies at  $1.7 < z < 3.5$ , providing full rest-frame optical coverage and enabling robust measurements of stellar populations and gas content. This dataset offers a benchmark for understanding galaxy assembly during peak star formation. Meanwhile, [?] leverage JWST/PRIMER imaging to reassess the abundance of quiescent galaxies (QGs) at  $z \sim 4 - 8$ . Their analysis reveals that previous estimates were inflated by misclassified Little Red Dots, and that high-mass QGs decline sharply beyond  $z > 4$ , while low-mass QGs maintain nearly constant densities. These results challenge current quenching models and highlight the critical role of mid-infrared data in disentangling galaxy populations at early times.

In this work, we build upon our previously proposed denoising framework based on U-Net Variational Autoencoders (VAEs) [?], which was shown to enhance classification accuracy by suppressing noise while preserving critical structural features. Here, we apply this approach to JWST NIRCam imaging data, focusing on galaxies up to redshift  $z \approx 8$ . Our goal is to evaluate the adaptability of the U-Net VAE model to next-generation datasets and assess its impact on the classification of disk and non-disk galaxy types in the early universe.

By integrating advanced denoising techniques with state-of-the-art classification models, this study contributes to the growing body of work aimed at improving the reliability of morphological analysis in deep-field surveys. The results offer insights into the structural evolution of galaxies at cosmic dawn and demonstrate the utility of generative models in high-redshift astrophysical research.

This paper is structured as follows: Section 2 provides an overview of the dataset employed in our research. In Section 3, we elaborate on the methodologies utilized for data processing and galaxy classification. Section 4 presents the analysis results and their significance. Lastly, Section 5 concludes the paper with a summary of our findings.

## 2. Dataset

This study is based on a subset of the publicly released dataset from the UNCOVER survey [?], which provides ultradeep spectroscopic observations using the JWST NIRSpec instrument in PRISM/CLEAR mode. The full dataset includes 668 galaxies in the A2744 field, spanning a redshift range of  $z \sim 0.3$  to  $z \sim 13$ , as described in [?].

Out of the 668 galaxies, 406 have reliable spectroscopic redshift measurements. We restrict our analysis to this subset, using it to investigate galaxy morphology with the goal of classifying galaxies as either disk-like or non-disk systems. To this end, we performed a visual inspection of each galaxy, selecting only those for which structural features could be discerned despite high background noise or low surface brightness. This selection ensures that morphological classification is based on identifiable features in the imaging data, allowing for a more robust analysis of galaxy structure across cosmic time.

The visual inspection resulted in a final sample of 292 galaxies, with redshifts ranging from  $z = 0.172194$  to  $z = 7.68681$ . These galaxies were selected based on the visibility of morphological structure, enabling a meaningful classification. The redshift distribution of the final sample is summarized in [Figure 1: see original paper].

## 3. Method

This study continues the methodology introduced in [?], where a U-Net VAE is employed as a pre-processing step to denoise galaxy images prior to classification. The goal is to mitigate the impact of astrophysical and instrumental contamination that can obscure galaxy morphology and degrade the performance of CNNs.

We begin by selecting 1400 uncontaminated galaxy images from the EFIGI dataset in [?]. Using the PyAutoLens package, we simulate realistic contamination by randomly inserting 1-3 non-target galaxies or stars per image, with varied light profiles, orientations, intensities, and Sérsic indices. This ensures spatial and morphological diversity, emulating real observational conditions and dataset features [?, ?].

The dataset is divided into 1000 images for training/validation and 400 for testing. To enhance generalization, each training image undergoes 10 contamination simulations, resulting in a 10,000-image augmented dataset.

The images are padded to  $256 \times 256 \times 3$  pixels and rescaled to the  $[0, 1]$  range to match the Galaxy10 DECaLS format and stabilize training. The U-Net VAE architecture consists of an encoder with four convolutional blocks (ReLU activations, max pooling) and a decoder with transposed convolutions for upsampling. The model is trained using a combined Binary Cross-Entropy (BCE) reconstruction loss and Kullback-Leibler (KL) divergence to regularize the latent space. The BCE and KL divergence components are given by:

$$L_{\text{BCE}} = - \sum_{i=1}^n [x_i \log(\hat{x}_i) + (1 - x_i) \log(1 - \hat{x}_i)]$$

$$L_{\text{KL}} = - \frac{1}{2} \sum_{j=1}^d [1 + \log(\sigma_j^2) - \mu_j^2 - \sigma_j^2]$$

where  $\mu$  and  $\sigma$  are the mean and standard deviation vectors of the latent variables respectively, and  $d$  is the dimensionality of the latent space. The total loss function for the VAE can then be expressed as:

$$L_{\text{VAE}} = L_{\text{BCE}} + \beta L_{\text{KL}}$$

where  $\beta$  is a weighting factor to balance the reconstruction loss and the KL divergence. The efficiency and performance quality of this approach are described in detail in the first article of this series [?].

In contrast with the latter, where classification was performed across 10 galaxy types, this study focuses on a binary classification task: distinguishing between disk and non-disk galaxies. This simplification narrows the classification space and aligns with our scientific objective of identifying disk-like galaxies in the early universe. The JWST dataset, which includes galaxies up to redshift 13, provides a unique opportunity to explore the emergence of disk structures at cosmic dawn.

For the binary classification task, we employ Group Convolutional Neural Networks (GCNNs), which have demonstrated strong performance in prior astronomical applications. GCNNs incorporate geometric symmetries, such as rotations and reflections, into the model architecture, enhancing robustness to spatial transformations. Following [?], we use the escnn library to implement GCNNs with D16 symmetry, which includes eight rotations and eight reflections. This design ensures equivariance to a broad range of transformations, making the model particularly well suited for analyzing morphologically diverse, high-redshift galaxies.

#### 4. Analysis

Trained on the Galaxy10 DECaLS dataset, our GCNN model achieved a prediction accuracy of 98.9% on the training set and 98.6% on the test set, supported

by an F1-score of 96.5%. These results confirm the model's strong generalization capability and its effectiveness in identifying disk-like structures in noisy, high-redshift galaxy images.

As this study involves a binary classification task—distinguishing between disk-like and non-disk-like galaxies—the model outputs a two-dimensional probability vector for each input image, representing the likelihood of the galaxy belonging to either class. Our analysis revealed that out of a total of 292 galaxies in the JWST preselected sample, the model classified 83 as disk-like. This outcome is particularly significant given the challenges associated with identifying disk features at high redshifts, where morphological signatures are often faint or distorted due to observational limitations.

While a subset of these classifications may appear visually ambiguous—especially in cases where the galaxy structure is compact or irregular—the model's predictions remain consistent with its learned latent representations and symmetry-aware architecture. This consistency suggests that the GCNN is capable of capturing subtle structural cues that may not be immediately apparent through visual inspection alone.

Figure 2: see original paper shows a scatter plot of spectroscopic redshift versus the predicted probability of disk-like morphology for 83 galaxies identified by the GCNN model. Each point represents a galaxy classified as disk-like, with predicted probabilities exceeding 0.5. The red ellipse marks a region of interest, extending from redshifts greater than 3 and encompassing relatively higher predicted probabilities, indicating a concentration of high-confidence disk-like galaxies at intermediate to high redshifts.

To further investigate the relationship between redshift and the model's confidence in disk-like morphology classification, we computed the mean predicted probability within 11 redshift bins. As shown in Figure 2: see original paper, each point represents the average probability in a bin, with vertical error bars indicating the 95% confidence interval around the mean. This analysis reveals how the GCNN model's predictions vary across cosmic time, highlighting regions of stable or elevated confidence in disk-like structures, particularly at intermediate redshifts.

To facilitate further investigation and validation, a complete list of galaxy IDs identified as disk-like is provided in . This table serves as a reference for astronomers interested in conducting follow-up studies, including spectroscopic analysis, kinematic modeling, or deeper morphological assessments using complementary datasets.

## 5. Conclusions

In this study, we presented a robust framework for enhancing morphological classification of high-redshift galaxies using a U-Net VAE for image denoising, followed by a GCNN for binary classification of disk-like versus non-disk-like

systems. By applying this pipeline to JWST NIRCcam imaging data from the UNCOVER survey, we demonstrated that denoising significantly improves the visibility of structural features, enabling more reliable classification even in the presence of observational noise and low surface brightness.

Our analysis focused on a visually curated sample of 292 galaxies with reliable spectroscopic redshifts, spanning  $z = 0.17$  to  $z = 7.69$ . Out of these, 83 galaxies were classified as disk-like by the GCNN model. Notably, a substantial fraction of these—approximately 70%-80%—are located at redshifts greater than 3 and were predicted with high confidence (probabilities well above 0.7). This finding is particularly compelling, as it suggests that disk-like structures may be more prevalent in the early universe than previously assumed, and that their morphological signatures can be effectively recovered through advanced denoising and symmetry-aware classification techniques.

These results underscore the potential of combining generative models with equivariant deep learning architectures to probe galaxy evolution at cosmic dawn. The ability to identify disk-like galaxies at such early epochs opens new avenues for studying the formation and stability of galactic disks under extreme physical conditions. Future work will extend this approach to larger samples and incorporate kinematic data to further validate the morphological classifications presented here.

## Acknowledgments

We are thankful to the referee for helpful suggestions and comments. We acknowledge the use of EFIGI and Galaxy10 DECaLS datasets, which enabled the development and validation of this research. Our sincere thanks are due to Artur Hakobyan for valuable insights and support in guiding the dataset selection and contributing to key discussions.

## References

- [?] Aguilar-Argüello, G., Fuentes-Pineda, G., Hernández-Toledo, H. M., et al. 2025, MNRAS, 537, 876
- [?] Atemkeng, M. T., Chuma, C., Zaza, S., Nunhokee, C. D., & Smirnov, O. M. 2025, arXiv:2502.17207
- [?] Baillard, A., Bertin, E., De Lapparent, V., et al. 2011, A&A, 532, A74
- [?] Belli, S., Bugiani, L., Park, M., et al. 2025, arXiv:2510.11775
- [?] Cromptoets, B. L., Kirk, H., Gutermuth, R., & Di Francesco, J. 2025, arXiv:2510.07747
- [?] De Paolis, F., Gurzadyan, V. G., Nucita, A. A., et al. 2014, A&A, 565, L3
- [?] Genin, A., Shuntov, M., Brammer, G., et al. 2025, A&A, 699, A343
- [?] Georgiou, C., Chisari, N. E., Bilicki, M., et al. 2025, A&A, 699, A252
- [?] Ghaderi, H., Alipour, N., & Safari, H. 2025, ApJS, 277, 10
- [?] Gurzadyan, V. G. 1999, EL, 46, 114
- [?] Hui, W., Robert Jia, Z., Li, H., & Wang, Z. 2022, JPhCS, 2402, 012009

- [?] Masters, K. 2025, arXiv:2502.09610
- [?] Meyer, R. A., Wang, F., Kakiichi, K., et al. 2025, arXiv:2510.11373
- [?] Mirzoyan, S. 2025, RAA, 25, 095006
- [?] Mirzoyan, S. S., Khachatryan, H., Yegorian, G., & Gurzadyan, V. G. 2019, MNRAS: Letters, 489, L32
- [?] Pandya, S., Patel, P., Blazek, J., et al. 2023, arXiv:2311.01500
- [?] Petrillo, C. E., Tortora, C., Chatterjee, S., et al. 2017, MNRAS, 472, 1129
- [?] Price, S. H., Bezanson, R., Labbe, I., et al. 2025, ApJ, 982, 51
- [?] Qian, Y. 2023, JPhCS, 2580, 012009
- [?] Sanders, J. S., Bahar, Y. E., Bulbul, E., et al. 2025, arXiv:502.02239
- [?] Tian, C., Urry, C. M., Ghosh, A., et al. 2025, ApJ, 981, 5
- [?] Weaver, J. R., Cutler, S. E., Pan, R., et al. 2023, ApJS, 270, 7
- [?] Yang, T., Wang, T., Xu, K., et al. 2025, arXiv:2510.12235

*Note: Figure translations are in progress. See original paper for figures.*

*Source: ChinaXiv – Machine translation. Verify with original.*