
Bayesian Deconvolution of Astronomical Images with Diffusion Models: Quantifying Prior-Driven Features in Reconstructions

Alessio Spagnoletti^{1,2} Alexandre Boucaud¹ Marc Huertas-Company³
Wassim Kabalan¹ Biswajit Biswas¹

¹Université Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France

²Université Paris Cité, CNRS, MAP5, F-75006 Paris, France

³Instituto de Astrofísica de Canarias (IAC), Canarias, Spain

{alessio.spagnoletti, alexandre.boucaud, wassim.kabalan,
biswajit.biswas}@apc.in2p3.fr
mhuertas@iac.es

Abstract

Deconvolution of astronomical images is a key aspect of recovering the intrinsic properties of celestial objects, especially when considering ground-based observations. This paper explores the use of diffusion models (DMs) and the Diffusion Posterior Sampling (DPS) algorithm to solve this inverse problem task. We apply score-based DMs trained on high-resolution cosmological simulations, through a Bayesian setting to compute a posterior distribution given the observations available. By considering the redshift and the pixel scale as parameters of our inverse problem, the tool can be easily adapted to any dataset. We test our model on Hyper Supreme Camera (HSC) data and show that we reach resolutions comparable to those obtained by Hubble Space Telescope (HST) images. Most importantly, we quantify the uncertainty of reconstructions and propose a metric to identify prior-driven features in the reconstructed images, which is key in view of applying these methods for scientific purposes.

1 Introduction

The analysis of astronomical observations is affected by multiple sources of degradation and noise, such as the point spread function (PSF) due to atmospheric interference and optical effects [37]. Although widely used, traditional deconvolution methods often encounter significant limitations, including noise amplification, loss of resolution, and sensitivity to model assumptions [19, 5, 37, 21].

In this paper, we explore and evaluate the application of diffusion models (DMs) [34, 33, 35, 36, 14, 23], for astronomical image deconvolution, through the diffusion posterior sampling (DPS) algorithm [7]. DMs have gained attention in recent years for their ability to model complex data distributions through iterative denoising processes, often outperforming generative adversarial networks in image generation tasks [8]. They provide a flexible and powerful approach to image restoration, either through the use of plug-and-play algorithms [40, 12] or by directly conditioning the diffusion process using a log-likelihood guidance term, $\nabla_x \log p(\mathbf{y}|\mathbf{x}_t)$ [36]. DPS follows this latter approach.

The objective is to leverage realistic simulations to train a score-based diffusion model that can be applied for deconvolution and noise removal on real images. We choose TNG100 simulations from the TNG Illustris project [22, 6] to train our DM and HSC-PDR3 Survey [3] data to test the DPS.

We highlight the pros and cons of our approach, examining factors such as image quality, computational efficiency, robustness to noise, and tendency to hallucinate [18, 42, 29]. While our preliminary

results suggest that our framework holds promise for improving deconvolution outcomes, we also identify several areas where the model struggles, particularly in keeping physical consistency and achieving real-time processing speeds.

2 Model and Datasets

Score-based diffusion model We follow the variance-preserving stochastic differential equation (VP-SDE) approach from [36], focusing on the Denoising Diffusion Probabilistic Model (DDPM) framework [33]. The general idea is to define the forward process \mathcal{F} that perturbs the original data \mathbf{x}_0 into a noisy version of itself \mathbf{x}_t at time t , and a backward process \mathcal{B} that solves a reverse-time stochastic process [4], written as:

$$\mathcal{F} : d\mathbf{x} = -\frac{\beta(t)}{2}\mathbf{x} dt + \sqrt{\beta(t)} d\mathbf{w}, \quad \mathcal{B} : d\mathbf{x} = \left(-\frac{\beta(t)}{2}\mathbf{x} - \beta(t)\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \right) dt + \sqrt{\beta(t)} d\bar{\mathbf{w}}$$

where $\beta(t)$ controls how much variance is added at each step, \mathbf{w} is standard Wiener process and $\bar{\mathbf{w}}$ is a backward Wiener process defined in [4].

When conditioning on an observation \mathbf{y} , the posterior term $p(\mathbf{x}_t|\mathbf{y})$ is introduced to guide the reverse process [36]. Applying Bayes' rule yields

$$d\mathbf{x} = \left(-\frac{\beta(t)}{2}\mathbf{x} - \beta(t) [\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t)] \right) dt + \sqrt{\beta(t)} d\bar{\mathbf{w}}. \quad (1)$$

The DPS algorithm approximates the likelihood $p(\mathbf{y}|\mathbf{x}_t)$ (intractable analytically [36, 7]) with $p(\mathbf{y}|\hat{\mathbf{x}}_0)$, where $\hat{\mathbf{x}}_0$ is an estimate of the recovered data. Thanks to Tweedie's formula [27, 38, 9, 15], using the score of the prior learned during the training process $s_\theta^* \simeq \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ and a cumulative noise schedule $\bar{\alpha}(t)$ that depends on $\beta(t)$ [36], its expression is

$$\hat{\mathbf{x}}_0 \simeq \frac{1}{\sqrt{\bar{\alpha}(t)}} [\mathbf{x}_t + (1 - \bar{\alpha}(t)) s_\theta^*(\mathbf{x}_t, t)].$$

Our publicly available DDPM and DPS implementation¹ is based on [23] and [7], respectively, for the training phase and deployment. Implementation details can be found in Appendix A.

TNG Illustris simulations The selected TNG100 simulated images have been made from snapshots of the evolution of the simulated universe at specific redshifts $z = 0.1527, 0.1693$ and 0.1804 with camera field-of-view face-on (also referred to as v0 in fits files). These snapshots have a fixed physical resolution of 0.4 kpc/pixel. The images are cropped at center to a size of 256 pixels, and we consider only the g, r and i bands. We refer to these images as the *idealized* TNG. This dataset is composed of 17852 images, that is split between a training and test with a 0.9/0.1 ratio. To increase the dynamic range, we normalize the images and use an `asinh` stretch with a scale of 0.01.

To deconvolve real images \mathbf{y} from the HSC survey, we define the operator \mathcal{A} for which we assume that the observation is defined as $\mathbf{y} = \mathcal{A}(\mathbf{x}) + \sigma_{\mathbf{y}}\mathbf{n}$ assuming \mathbf{n} a standard Gaussian noise, and $\sigma_{\mathbf{y}}$ is the noise variance. Our goal is thus to recover \mathbf{x} .

Due to projection effects, the apparent size of galaxies (i.e. the size to which they appear to the camera) at various redshifts is different; Since the model is trained with a pixel scale of fixed physical units per pixel (the idealized TNG data), one must then incorporate in the operator \mathcal{A} a redshift-dependent scaling to match their apparent size. Furthermore, the matching of HSC pixel scale $p = 0.17$ arcsec/pixel must be done before the images are convolved with the HSC PSF Φ to simulate an observation. In Figure 1, we produce a schematic of that process, where the matching operations are grouped under the $f_{z,p}$ function, and $m := m(z, p)$ is the size of the observed image.




Figure 1: Schematic of the process to transform an idealized TNG simulated image into a HSC-like observation through the operator \mathcal{A} .

¹<https://github.com/astrodeepnet/diffusion4astro.git>

HSC - PDR3 Survey data The Hyper Suprime-Cam Subaru Strategic Program (HSC-SSP)² is an astronomical survey conducted using the Hyper Suprime-Cam (HSC) on the Subaru Telescope, a powerful 8.2-meter optical telescope located in Hawaii. The data from this survey, particularly from its third public data release (PDR3) [3], has been widely used in various research fields. We fetch HSC images along with their approximated PSF³ and computed photometric redshift [39, 24, 25].

3 Experiments

The training process for the model required a total of **133 hours** on **32 V100 GPUs**. Following 300k training steps, an analysis revealed that the model exhibited a tendency toward mode collapse [11]. This was identified by projecting the training dataset and the model predictions onto a lower-dimensional manifold using UMAP [20] (refer to Appendix B for corresponding visualizations). To mitigate this issue, the training was extended by an additional 225k steps, accompanied by an increase in the batch size from 64 to 128 and the introduction of a random shift of up to 10 pixels in any direction for the training images. This randomness enabled the model to generate galaxies that were not perfectly centered, aligning more closely with real-world observations where object centers are often uncertain.

To qualify the performance of our model, we computed the Fréchet Inception Distance (FID) [13], a metric that indicates how well the generated samples match the features of the training images. The values computed on 1024 images sampled from our models are shown on Table 1. The slight increase between 450k and 525k may be explained by the random shift, which was not present in the training set, for which all galaxies were perfectly centered. Since our work is the first to use this specific part of the TNG Illustris dataset, we cannot compare these numbers with others, but they may be helpful to set a benchmark future works.

Model training steps	300k	375k	450k	525k
FID value ↓	317.3	235.3	119.4	132.6

Table 1: Performance of our model as a function of training steps (lower is better).

3.1 Results on simulated observations

Since our model has only been trained on simulated images, a first measure of the performance of our algorithm is to deconvolve a simulated TNG image that has previously been convolved with a realistic PSF to simulate an observation. This is possible due to the publicly available TNG dataset [6] which not only provides the simulations we used to train our model but also contains HSC-like images. These are indeed simulated TNG images that have been convolved with HSC PSF, set to the HSC pixel scale, and put at a specific location in the HSC-SSP survey, providing a real background to the simulated objects.

In Figure 2 we show respectively: the final step of the diffusion process $\mathbf{x}_0^{\text{Diff}}$, the re-convolved final step that can be compared to the TNG-HSC observation \mathbf{y} , and the true TNG $\mathbf{x}_0^{\text{Ideal}}$. Both the TNG-HSC and TNG images are stretched with asinh and scale 0.01. We notice how the re-convolved final step image, i.e. $\mathcal{A}(\mathbf{x}_0^{\text{Diff}})$, is close to the observation \mathbf{y} , which is an empirical proof that the reconstruction process did not hallucinate by creating new information. In Figure 3 we show the residual between the reconstructed $\mathbf{x}_0^{\text{Diff}}$ and the $\mathbf{x}_0^{\text{Ideal}}$, from which we can see that the galaxy is well reconstructed since most residuals are due to the artificially added HSC background.




Figure 2: Comparison of DDPM TNG-HSC deconvolved image (leftmost) with *Idealized* TNG counterpart (rightmost). $z = 0.1768$

Figure 3: Residual

²Official website: <https://hsc.mtk.nao.ac.jp/ssp/>

³PSF Picker tool: <https://hsc-release.mtk.nao.ac.jp/psf/pdr3/>

3.2 Comparing with space-based imaging

Thanks to the data publicly available through the PDR3 Survey UltraDeep layer [3], we can crop areas⁴ at known astronomical coordinates to retrieve the *gri* bands, the associated PSFs and photometric redshifts. We chose an object in a location in the sky where we both have images from the COSMOS project [31] and PDR3. We then apply the DPS algorithm with our DDPM model trained on *idealized* TNG images and compare the result with the equivalent Hubble Space Telescope (HST - ACS) images, downsampled at the HSC pixel-scale. The HST is only taken at channel *i*, i.e., the F814W band, to provide a visual comparison.

In Figure 4 and 6, we show the ground-based observation (HSC) y , the space-based reference (HST), and the final step of the diffusion process x_0 . Both the HSC and HST images are stretched with `asinh` and scale 0.01. Each evaluation of the DPS takes ~ 100 seconds on a **single A100 GPU**. Indeed, the time consumption is a negative aspect of the DPS, which doubles the DDPM sampling time due to the extra computation of the log-likelihood gradient term in Equation (1).




Figure 4: Comparison of DDPM HSC deconvolved image (rightmost) with HST counterpart (middle). Ra-dec coords (10h1m49.30s, 2°23'18.60") $z = 0.2168$

Figure 5: Posterior mean (left) and variance-ratio (right) of pixel values over 256 samples from Figure 4.




Figure 6: Comparison of DDPM HSC deconvolved image (rightmost) with HST counterpart (middle). Ra-dec coords (10h 1m 34.77s, 2°4'19.02") $z = 0.1665$

Figure 7: Posterior mean (left) and variance-ratio (right) of pixel values over 256 samples from Figure 6.

To provide a more insightful analysis, we sample 256 images from the posterior distribution and compute the pixel-wise mean and variance (ratio). Figures 5, 7 show that the posterior mean represents a reliable output for the inverse problem since it averages out the stochastic fluctuations. On the right, the variance ratio image is produced by dividing the posterior variance by the prior variance, giving an idea of the areas in which the uncertainty on the generated image is higher, i.e., the prior-dominated areas. The fact that the central part of the variance ratio image is close to 0, means that the model is not hallucinating when reconstructing the galaxy. However the reconstructed outskirts of the galaxy seem to be much more prior-dominated, probably due to the nature of the TNG training dataset.

3.3 Limitations

Despite the good results, the model has limitations when the redshift z gets too high. Indeed, at $z > 0.7$, we get images of size $m(z, p) < 80$. This means that when we try to solve the inverse problem, we are both solving the deconvolution and the upsampling tasks, and an upscaling factor higher than a factor 2 is problematic. We notice on Figure 8 how a single sampled galaxy at $z > 0.8$ is highly prior-dominated. In particular, a high number of features are created (hallucinated) all around the galaxy, and are invisible on the HST counterpart. This is common when the solution to the inverse problem requires injecting too much information from the prior (TNG idealized images).

Considering again 256 posterior samples for this galaxy and computing the mean and variance ratio, demonstrate this prior domination on a large part of the image on Figure 9. The posterior mean is still quite robust with fewer small TNG features than a single sample.

⁴Visual tool: <https://hsc-release.mtk.nao.ac.jp/hscMap-pdr3/app>




Figure 8: Comparison of DDPM HSC deconvolved image (rightmost) with HST counterpart (middle). Ra-dec coords (10h0m2.86s, $2^{\circ}2'13.92''$) $z = 0.8463$.

Figure 9: Posterior mean (left) and variance-ratio (right) of pixel values over 256 samples from Figure 8.

3.4 Quantification of hallucinations

To better qualify the impact of the prior on the deconvolution, we define a metric based on the posterior variance. Building on the results we explored on Figures 5, 7 and 9, we take the pixel-wise posterior variance and prior variance over 256 samples, take their ratio and compute the average value on a centered aperture or radius r (here we chose $r = 30$ pixels) that contains the central part of the galaxy. To verify such metric is positively correlated with a prior-dominated reconstruction, we create a test using simulated TNG images, which we artificially convolve with a Gaussian PSF and add a fixed noise level σ_y measured on HSC images. The main point is to vary the signal-to-noise ratio (SNR) by setting specific magnitudes to the galaxy (it will be lower when the magnitude is higher, i.e. the galaxy fainter). We vary the magnitude from 18 to 23, and report the metric results in Table 2, where we confirm such positive correlation when the noise starts dominating the galaxy signal.

We refer to Appendix C for the visual results at all magnitudes.

Simulated magnitude	18	19	20	21	22	23
SNR	699.3	278.4	110.8	44.1	17.6	7.0
Variance Ratio ($r = 30$)	0.0114	0.0116	0.0168	0.0346	0.0560	0.0685

Table 2: Signal-to-noise and variance ratio as a function of the simulated TNG galaxy magnitude.

4 Conclusion

By exploring diffusion-based methods in the context of astronomical image deconvolution, we aim to provide a foundation for future research and development of signal processing in astronomy. Our findings contribute to the growing body of works on applying advanced machine learning techniques to astrophysical data analysis [16, 26, 1, 32]. In particular, one of our inspirations comes from the work of [2], who first implemented DMs to tackle the astronomical deconvolution problem. Our work builds on it by including the redshift and pixel scale as parameters of the inverse problem to adapt the model to any dataset. We also propose an implementation of a general-purpose algorithm, the DPS, to solve the deconvolution of astronomical observations and make it publicly available, allowing the reproducibility of the results. Finally, we provide a metric to identify prior-driven features in the generated solutions. We hope to take a step toward providing confidence to the scientific community for applying these models to real-case scenarios.

Broader impact

Generative modelling for improving data quality in astrophysics has been investigated quite extensively over the past decade [30, 41, 17, 10]. However, despite promising results, very few works use these techniques for scientific applications because it is difficult to identify whether the model outputs can be trusted. In this work, we try to go a step forward towards deployment by investigating ways to track model hallucinations and prior-driven features. This opens the door to a systematic application of these deconvolution methods on large-scale statistical samples of celestial objects to infer physical properties. We plan to investigate the use of score-based deconvolution techniques to resolve the internal structure of galaxies.

Acknowledgments and Disclosure of Funding

This work was granted access to the HPC resources of IDRIS under the allocation 2024-AD011014557 made by GENCI. This research was supported by the Data Intelligence Institute of Paris (diiP), and IdEx Université de Paris (ANR-18-IDEX-0001). This research has been partly funded by the France 2030 research programme on artificial intelligence, via grant PDE-AI (ANR-23-PEIA-0004). BB acknowledges funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 945304 – COFUND AI4theSciences hosted by PSL University, and Agence Nationale de la Recherche grant ANR-19-CE23-0024 — AstroDeep.

Bibliography

- [1] A. Adam, A. Coogan, N. Malkin, R. Legin, L. Perreault-Levasseur, Y. D. Hezaveh, and Y. Bengio. Posterior samples of source galaxies in strong gravitational lenses with score-based priors. *ArXiv*, abs/2211.03812, 2022.
- [2] A. Adam, C. Stone, C. Bottrell, R. Legin, Y. Hezaveh, and L. Perreault-Levasseur. Echoes in the noise: Posterior samples of faint galaxy surface brightness profiles with score-based likelihoods and priors, 2023.
- [3] H. Aihara, Y. AlSayyad, M. Ando, R. Armstrong, J. F. Bosch, E. Egami, H. Furusawa, J. Furusawa, S. Harasawa, Y. Harikane, B. C. Hsieh, H. Ikeda, K. Ito, I. Iwata, T. Kodama, M. Koike, M. Kokubo, Y. Komiyama, X. Li, Y. Liang, Y.-T. Lin, R. H. Lupton, N. B. Lust, L. A. Macarthur, K. Mawatari, S. Mineo, H. Miyatake, S. Miyazaki, S. More, T. Morishima, H. Murayama, K. Nakajima, F. Nakata, A. J. Nishizawa, M. Oguri, N. Okabe, Y. Okura, Y. Ono, K. Osato, M. Ouchi, Y.-C. Pan, A. A. P. Malag’ on, P. A. Price, S. L. Reed, E. S. Rykoff, T. Shibuya, M. Simunovic, M. A. Strauss, K. Sugimori, Y. Suto, N. Suzuki, M. Takada, Y. Takagi, T. Takata, S. Takita, M. Tanaka, S. Tang, D. S. Taranu, T. Terai, Y. Toba, E. L. Turner, H. Uchiyama, B. Vijarnwannaluk, C. Waters, Y. Yamada, N. Yamamoto, and T. Yamashita. Third data release of the hyper supprime-cam subaru strategic program. 2021.
- [4] B. D. O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12:313–326, 1982.
- [5] M. Bertero, P. Boccacci, and M. Robberto. Inversion method for the restoration of chopped and nodded images. In *Astronomical Telescopes and Instrumentation*, 1998.
- [6] C. Bottrell, H. M. Yesuf, G. Popping, K. C. Omori, S. Tang, X. Ding, A. Pillepich, D. Nelson, L. Eisert, H. Gao, A. D. Goulding, B. S. Kalita, W. Luo, J. E. Greene, J. Shi, and J. D. Silverman. Illustrating in the hsc-ssp: image data release and the major role of mini mergers as drivers of asymmetry and star formation. 2023.
- [7] H. Chung, J. Kim, M. T. McCann, M. L. Klasky, and J. C. Ye. Diffusion posterior sampling for general noisy inverse problems. *ArXiv*, abs/2209.14687, 2022.
- [8] P. Dhariwal and A. Nichol. Diffusion models beat gans on image synthesis. *ArXiv*, abs/2105.05233, 2021.
- [9] B. Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106:1602 – 1614, 2011.
- [10] F. K. Gan, K. Bekki, and A. Hashemizadeh. SeeingGAN: Galactic image deblurring with deep learning for better morphological classification of galaxies. *arXiv e-prints*, page arXiv:2103.09711, Mar. 2021.
- [11] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [12] A. Graikos, N. Malkin, N. Jojic, and D. Samaras. Diffusion models as plug-and-play priors. *ArXiv*, abs/2206.09012, 2022.

- [13] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Neural Information Processing Systems*, 2017.
- [14] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *ArXiv*, abs/2006.11239, 2020.
- [15] K. Kim and J.-C. Ye. Noise2score: Tweedie’s approach to self-supervised image denoising without clean images. In *Neural Information Processing Systems*, 2021.
- [16] F. Lanusse, R. Mandelbaum, S. Ravanbakhsh, C.-L. Li, P. E. Freeman, and B. Póczos. Deep generative models for galaxy image simulations. *Monthly Notices of the Royal Astronomical Society*, 504:5543–5555, 2020.
- [17] L. Lauritsen, H. Dickinson, J. Bromley, S. Serjeant, C.-F. Lim, Z.-K. Gao, and W.-H. Wang. Superresolving Herschel imaging: a proof of concept using Deep Neural Networks. *Monthly Notices of the Royal Astronomical Society*, 507(1):1546–1556, 07 2021.
- [18] C. Liu, H. Shum, and W. T. Freeman. Face hallucination: Theory and practice. *International Journal of Computer Vision*, 75:115–134, 2007.
- [19] L. B. Lucy. Image restorations of high photometric quality. In R. J. Hanisch and R. L. White, editors, *The Restoration of HST Images and Spectra - II*, page 79, January 1994.
- [20] L. McInnes and J. Healy. Umap: Uniform manifold approximation and projection for dimension reduction. *ArXiv*, abs/1802.03426, 2018.
- [21] K. Michalewicz, M. Millon, F. Dux, and F. Courbin. Starred: a two-channel deconvolution method with starlet regularization. *J. Open Source Softw.*, 8:5340, 2023.
- [22] D. Nelson, V. Springel, A. Pillepich, V. Rodriguez-Gomez, P. Torrey, S. Genel, M. Vogelsberger, R. Pakmor, F. Marinacci, R. Weinberger, L. Z. Kelley, M. R. Lovell, B. Diemer, and L. E. Hernquist. The illustriatng simulations: public data release. *Computational Astrophysics and Cosmology*, 6:1–29, 2018.
- [23] A. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. *ArXiv*, abs/2102.09672, 2021.
- [24] A. J. Nishizawa, B. C. Hsieh, M. Tanaka, and T. Takata. Photometric redshifts for the hyper supprime-cam subaru strategic program data release 2. 2020.
- [25] A. J. Nishizawa, B.-C. Hsieh, M. Tanaka, and the HSC Collaboration. Photometric redshifts for the hyper supprime-cam subaru strategic program data release 3. *Publications of the Astronomical Society of Japan*, 74:247, 2022. HSC PDR3 photometric data release.
- [26] M. Ntampaka, C. Avestruz, S. Boada, J. Caldeira, J. Cisewski-Kehe, R. D. Stefano, C. Dvorkin, A. E. Evrard, A. Farahi, D. P. Finkbeiner, S. Genel, A. A. Goodman, A. D. Goulding, S. Ho, A. B. Kosowsky, P. L. Plante, F. Lanusse, M. Lochner, R. Mandelbaum, D. Nagai, J. A. Newman, B. Nord, J. E. G. Peek, A. Peel, B. Póczos, M. M. Rau, A. Siemiginowska, D. J. Sutherland, H. Trac, and B. D. Wandelt. The role of machine learning in the next decade of cosmology. *arXiv: Instrumentation and Methods for Astrophysics*, 2019.
- [27] H. E. Robbins. An empirical bayes approach to statistics. 1956.
- [28] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. *ArXiv*, abs/1505.04597, 2015.
- [29] M. L. Sampson and P. Melchior. Spotting hallucinations in inverse problems with data-driven priors. 2023.
- [30] K. Schawinski, C. Zhang, H. Zhang, L. Fowler, and G. K. Santhanam. Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit. *Monthly Notices of the Royal Astronomical Society: Letters*, 467(1):L110–L114, Jan. 2017.

- [31] N. Scoville, H. Aussel, M. Brusa, P. L. Capak, C. M. Carollo, M. Elvis, M. Giavalisco, L. Guzzo, G. Hasinger, C. D. Impey, J.-P. Kneib, O. LeFèvre, S. J. Lilly, B. Mobasher, A. Renzini, A. Renzini, R. M. Rich, D. B. Sanders, E. Schinnerer, E. Schinnerer, D. Schminovich, P. Shopbell, Y. Taniguchi, and N. de Grasse Tyson. The cosmic evolution survey (cosmos): Overview. *The Astrophysical Journal Supplement Series*, 172:1 – 8, 2006.
- [32] M. J. Smith, J. E. Geach, R. A. Jackson, N. Arora, C. Stone, and S. Courteau. Realistic galaxy image simulation via score-based generative models. *Monthly Notices of the Royal Astronomical Society*, 511(2):1808–1818, 01 2022.
- [33] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. *ArXiv*, abs/2010.02502, 2020.
- [34] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. In *Neural Information Processing Systems*, 2019.
- [35] Y. Song and S. Ermon. Improved techniques for training score-based generative models. *ArXiv*, abs/2006.09011, 2020.
- [36] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations, 2021.
- [37] J.-L. Starck, E. Pantin, and F. Murtagh. Deconvolution in astronomy: A review. *Publications of the Astronomical Society of the Pacific*, 114:1051 – 1069, 2002.
- [38] C. M. Stein. Estimation of the mean of a multivariate normal distribution. *Annals of Statistics*, 9:1135–1151, 1981.
- [39] M. Tanaka, J. Coupon, B. C. Hsieh, S. Mineo, A. J. Nishizawa, J. S. Speagle, H. Furusawa, S. Miyazaki, and H. Murayama. Photometric redshifts for hyper supprime-cam subaru strategic program data release 1. *arXiv: Astrophysics of Galaxies*, 2017.
- [40] S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg. 5-29-2013 plug-and-play priors for model based reconstruction. *IEEE*, 2013.
- [41] A. Vojtekova, M. Lieu, I. Valtchanov, B. Altieri, L. Old, Q. Chen, and F. Hroch. Learning to denoise astronomical images with U-nets. *Monthly Notices of the Royal Astronomical Society*, 503(3):3204–3215, 11 2020.
- [42] N. Wang, D. Tao, X. Gao, X. Li, and J. Li. A comprehensive survey to face hallucination. *International Journal of Computer Vision*, 106:9–30, 2013.


A Training details

The model used to approximate the score function introduced in Section 2 is the UNet [28] model defined in [23]. To train the network we used the following hyperparameters:


- `-image_size 256`
- `-num_channels 192` This flag indicates the base number of channels used in the UNet’s convolutional layers.
- `-num_res_blocks 3` This specifies the number of **residual blocks** to use at each level in the UNet architecture.
- `-diffusion_steps 1000`
- `-noise_schedule cosine` The model uses a **cosine noise schedule**. The cosine schedule has been shown [23] to improve model performance by making the noise added to the images more evenly distributed across the steps (as opposed to the traditional linear schedule).
- `-lr 1e-4` This flag specifies the learning rate.
- `-batch_size 64 or 128` This flag specifies the batch size used during training. We train our model for $300k$ steps using a batch size of 64, and then we improve the results by extending the training to $525k$ steps using a batch size of 128.

B UMAP representation of the data

We use the low-dimensional manifold projection called UMAP [20] to visualize on the same manifold the data distribution from the training set (blue dots) and the one sampled from the generative model (red dots). We use the fact that the distribution of the sampled data is not as homogeneous as the one from the training as a sign that the training is incomplete or collapses into specific modes as shown in Figure (a). After additional training, the behavior improves as seen on Figure (b).



(a) UMAP 3D representation of the training dataset (blue dots), as well as 1000 samples (red dots) generated from the model after $300k$ steps. We can spot a mode collapse at the top right.



(b) UMAP 3D representation of the training dataset (blue dots), as well as 1000 samples (red dots) generated from the model after $525k$ steps. Much more homogeneous.

C Visual quantification hallucinations

The Figures below start from the same idealized TNG image, whose magnitude is artificially changed (second column). The two right-most images are the pixel-wise mean and variance ratio computed over 256 samples from the model applied on the second image. The Figures are corresponding to the values computed in Table 2.




Figure 11: *Left to right*: Idealized - Simulated at **mag 18** - Pixel-wise mean - Variance ratio




Figure 12: *Left to right*: Idealized - Simulated at **mag 19** - Pixel-wise mean - Variance ratio




Figure 13: *Left to right*: Idealized - Simulated at **mag 20** - Pixel-wise mean - Variance ratio




Figure 14: *Left to right*: Idealized - Simulated at **mag 21** - Pixel-wise mean - Variance ratio




Figure 15: *Left to right*: Idealized - Simulated at **mag 22** - Pixel-wise mean - Variance ratio




Figure 16: *Left to right*: Idealized - Simulated at **mag 23** - Pixel-wise mean - Variance ratio