`DiffLense`: A Conditional Diffusion Model for Super-Resolution of Gravitational Lensing Data

Pranath Reddy
University of Florida
Gainesville, FL 32611, USA
kumbam.pranath@ufl.edu

Michael W. Toomey
Center for Theoretical Physics,
Massachusetts Institute of Technology,
Cambridge, MA 02139, USA
mtoomey@mit.edu

Hanna Parul
Department of Physics & Astronomy,
University of Alabama,
Tuscaloosa, AL 35401, USA
hparul@crimson.ua.edu

Sergei Gleyzer
Department of Physics & Astronomy,
University of Alabama,
Tuscaloosa, AL 35401, USA
sgleyzer@ua.edu

Abstract

Gravitational lensing data is frequently collected at low resolution due to instrumental limitations and observing conditions. Machine learning-based super-resolution techniques offer a method to enhance the resolution of these images, enabling more precise measurements of lensing effects and a better understanding of the matter distribution in the lensing system. In this work, we introduce `DiffLense`, a novel super-resolution pipeline based on a conditional diffusion model specifically designed to enhance the resolution of gravitational lensing images obtained from the Hyper Suprime-Cam Subaru Strategic Program (HSC-SSP). The diffusion model, trained to generate Hubble Space Telescope (HST) data, is conditioned on HSC data pre-processed with denoising techniques and thresholding to significantly reduce noise and background interference. This process leads to a more distinct and less overlapping conditional distribution during the model's training phase. We demonstrate that `DiffLense` outperforms existing state-of-the-art single-image super-resolution techniques, particularly in retaining the fine details necessary for astrophysical analyses.

1 Introduction

Gravitational lensing, the bending of light from a distant source by a massive object between a source and the observer, is a powerful tool in astrophysics. Strong gravitational lensing in particular allows us to study the distribution of dark matter on subgalactic scales but also provides a magnified view of background sources which serves as a critical probe of the high redshift Universe. For detailed studies of background sources and the lens itself, high resolution and high quality data is imperative. However, the number of high-resolution gravitational lensing data available is often limited in number, largely due to limitations in the capabilities of the observing instruments and adverse observing conditions. Thus, the generation of high-resolution data is imperative for future detailed studies of galaxies.

Despite these shortcomings, strong gravitational lensing has already shown significant potential in uncovering hints about the nature of dark matter through its substructures, evidenced by analyses of lensed quasars [1, 2, 3], observations from ALMA [4], and extended lensing images [5, 6, 7], among

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

others. Indeed, various studies have explored anticipated signatures from ΛCDM and its extensions to derive information regarding the underlying dark matter distribution, e.g. [8, 9, 10, 11].

Recently, there has been a surge in the use of machine learning to tackle questions in lensing [12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28] and other scientific domains [29, 30]. Machine learning is well suited in this context as the analysis of even a single lens can be quite computationally taxing. Example applications of machine learning in this context include classification [12, 24, 31], regression [32, 20], segmentation analysis [27], domain adaptation [13], and anomaly detection [14]. So far, research has predominantly applied these techniques to simulations, primarily due to the limited availability of strong lensing data. This situation is expected to improve soon with the commissioning of the Vera C. Rubin Observatory and the launch of Euclid [33, 34]. Most previous studies have relied on simulation data as a proxy for the absence of plentiful high quality lenses. One possible work around to this issue is the implementation of super-resolution techniques applied to plentiful, lower quality data.

Super-resolution techniques, particularly those based on machine learning, have shown promise in enhancing the quality of low-resolution astronomical images more generally [35, 36, 37].

Traditional super-resolution techniques typically involve learning a mapping from low-resolution (LR) to high-resolution (HR) images by optimizing a fixed distance function. However, these methods often struggle to capture the intricate details necessary for astrophysical analysis, as the rigid nature of fixed distance optimization can lead to the loss of critical structural features.

In this paper, we introduce `DiffLense`, a novel super-resolution pipeline based on a conditional diffusion model specifically designed to enhance the resolution of gravitational lensing images obtained from the Hyper Suprime-Cam Subaru Strategic Program (HSC-SSP) [38]. Our model leverages high-resolution Hubble Space Telescope (HST) data to condition the diffusion process, enabling the generation of HR images while preserving fine astrophysical details. The domain-specific preprocessing pipeline, which includes noise reduction and background suppression, provides a clearer and more distinct conditional distribution, allowing the model to outperform existing state-of-the-art single-image super-resolution techniques.

In Sec. 2, we provide a comprehensive overview of the dataset utilized in this study. Sec. 3 details the models and methods employed in our analysis. We present the main findings in Sec. 4 and conclude in Sec. 5.

2 Data

We have constructed a dataset containing images of strong galaxy-galaxy gravitational lenses observed with instruments with different resolution. We compiled a list of lens candidates from the literature [39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55] and crossmatched them with archival data. As a low resolution part, we utilized *i*-band images from the third data release of Hyper Suprime-Cam Subaru Strategic Program (HSC-SSP), which has resolution of 0.168"/pix. For high resolution counterparts we searched archival Hubble Space Telescope (HST) data available at MAST [1] and made cutouts from ACS/WFC images in F814W filter with 0.05"/pix resolution. The final dataset contains 173 objects.

3 Methodology

Diffusion models [56] represent a class of generative models that convert noise into structured outputs through a Markov chain-based process inspired by non-equilibrium thermodynamics [57]. In the forward process, Gaussian noise is iteratively added to data, approximating a complex distribution with simpler ones at each noise level. The reverse process, learned by a neural network, progressively removes the noise to yield coherent images. For our model, this process is represented as,

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1-\alpha_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(x_t, t, c) \right) \tag{1}$$

[1] `https://mast.stsci.edu/search/ui/`

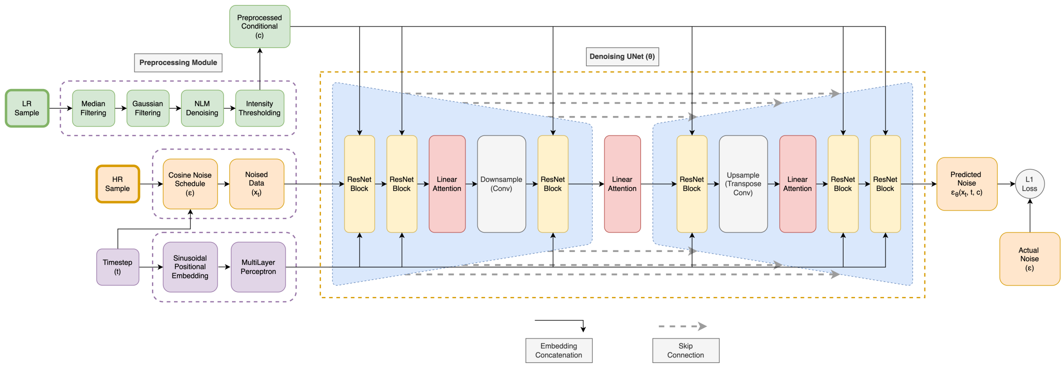

Figure 1: Overview of the `DiffLense` architecture.

where x_t represents the data at time step t, α_t is the noise scale at step t, and $\epsilon_\theta(x_t, t, c)$ is the noise predicted by the network parameterized by θ, conditioned on the low-resolution image c. The model is conditioned on low-resolution images from HSC, while the high-resolution counterparts from HST serve as the ground truth. This conditional approach ensures that the generated high-resolution images retain the astrophysical characteristics of the original low-resolution data.

The denoising of the low-resolution HSC images is crucial for improving the model's performance. The denoising pipeline consists of several steps combined into the following process in the approximate order of their computational complexity:

$$I_{\text{denoised}}(x, y) = \mathcal{T}\left(\mathcal{NLM}\left(G_\sigma\left(\text{Median}\left(I(x, y)\right)\right)\right)\right), \tag{2}$$

where:

- $\text{Median}(I(x, y))$ is the median filter to remove salt-and-pepper noise,
- G_σ is Gaussian smoothing with standard deviation σ,
- $\mathcal{NLM}$ represents Non-Local Means (NLM) denoising, which reduces noise by averaging similar image patches,
- $\mathcal{T}$ is thresholding, setting pixel intensities below a threshold T to zero:

$$\mathcal{T}(I(x, y)) = \begin{cases} I(x, y), & \text{if } I(x, y) \geq T \\ 0, & \text{if } I(x, y) < T \end{cases} \tag{3}$$

The filtering methods in the preprocessing pipeline were chosen to balance noise reduction and detail preservation. Median filtering removes salt-and-pepper noise while preserving edges, Gaussian filtering smooths the image without losing key features, and Non-Local Means (NLM) denoising averages similar patches to maintain structural integrity. This sequential approach, increasing in computational complexity, ensures robustness against various noise types, providing clean conditional inputs for the diffusion model.

In an ablation study included in the Appendix of this paper, we confirmed the importance of our denoising pipeline, which provided the best results when all denoising steps were applied.

The model architecture is based on a U-Net structure [58], and trained on 2880 pairs of low-resolution 64×64 HSC images and high-resolution 128×128 HST images, normalized to the range $[-1, 1]$. We train the model for 2000 epochs using the Adam optimizer [59] with a learning rate of 2×10^{-5} and a batch size of 10. The loss function minimizes the L1 loss between predicted and actual noise, guiding the reverse diffusion process. A cosine noise scheduler [60] is used for smooth noise scaling across 1000 timesteps and the model is trained on two NVIDIA Tesla A100 GPUs using the PyTorch framework [61].

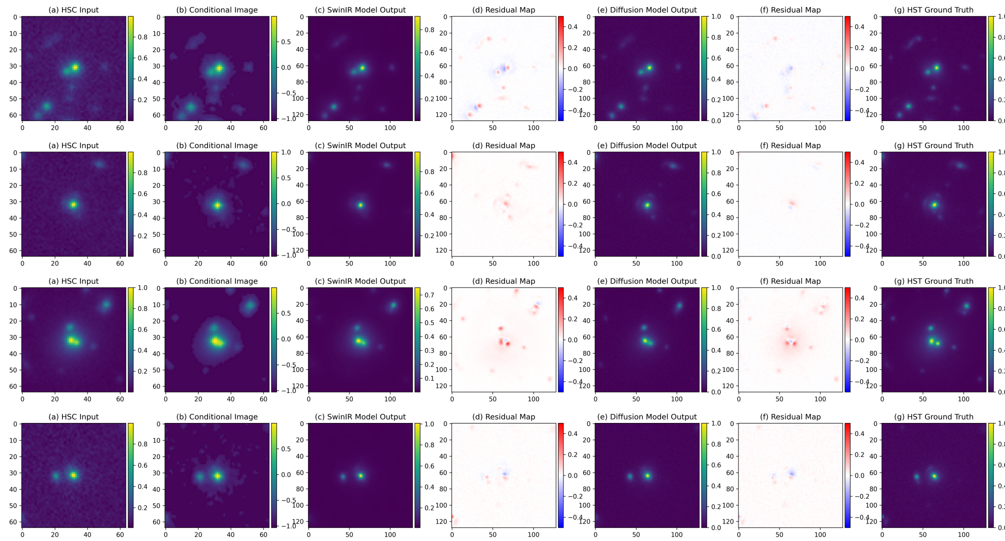

Figure 2: `DiffLense` Output Examples w/ Residual Maps. From left to right: HSC input sample, conditional image, SwinIR output, SwinIR residual map, `DiffLense` output, `DiffLense` residual map, and the ground truth.

4 Results & Discussion

We tested the performance of `DiffLense` on real gravitational lensing images from the Hyper Suprime-Cam Subaru Strategic Program (HSC-SSP) and compared it with state-of-the-art convolution and transformer models such as Residual Dense Network (RDN) [62], Residual Channel Attention Network (RCAN) [63], SwinIR [64], and HAT [65]. We have used Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) for quantitative evaluation.

The quantitative results presented in Table 1 show that `DiffLense` achieved higher PSNR and comparable SSIM values, indicating an improvement in both noise reduction and structural preservation. The reduced SSIM score could potentially be attributed to the misalignment in the intensity range between the super-resolved images and the ground truth. This discrepancy could be due to the generative nature of the diffusion model. Unlike other models that directly calculate loss between predicted and actual images leading to better intensity alignment, the diffusion model clamps the intensities to a fixed range of $[-1, 1]$ and normalizes them, potentially causing some slight misalignment. Additionally, the model's reverse process, which iteratively predicts and subtracts the noise based on the outputs from prior time steps, could lead to the compounding of errors. Inaccuracies in early steps may escalate, increasing the numerical error observed in later outputs. Although the SSIM score does not correlate well with the perceived visual quality, the PSNR score is significantly improved, indicating superior reconstruction quality and reduction of noise.

Table 1: Comparative Performance of Super-Resolution Models

Metric	RDN	RCAN	SwinIR	HAT	**DiffLense**
PSNR	32.94286	33.80036	34.63678	34.01312	**35.06834**
SSIM	0.86701	0.86768	0.86866	0.86565	**0.83937**

Visual inspection of the super-resolved images as shown in 2 confirms that `DiffLense` was more effective in reconstructing fine features of the lensed galaxies, while the other models tended to oversmooth the output or failed to retain the structural details.

5 Conclusion & Future Work

The application of super-resolution techniques to gravitational lensing images represents a potentially significant advancement in astrophysical imaging. In this study we introduced `DiffLense`, a conditional diffusion model that enhances the resolution of gravitational lensing images by conditioning on pre-processed low-resolution inputs. `DiffLense` outperformed state-of-the-art models by preserving fine astrophysical details and reducing noise, making it a powerful tool for lensing analysis. `DiffLense`'s modular preprocessing pipeline and generative capabilities make it a suitable and robust approach for a wide range of applications in astrophysical imaging. For our future work, we intend to explore score-based diffusion models and time-aware learnable denoising pipelines, similar to an approach explored by authors in [66] improving over the fixed denoising pipeline in our current implementation.

Acknowledgements

We acknowledge useful conversations with Stephon Alexander. P.R. was a participant in the Google Summer of Code 2023 program. S.G. was supported in part by U.S. National Science Foundation award No. 2108645. Portions of this work were conducted in MIT's Center for Theoretical Physics and partially supported by the U.S. Department of Energy under grant Contract Number DE-SC0012567. M.W.T is supported by the Simons Foundation (Grant Number 929255).

References

[1] S. Mao and P. Schneider. Evidence for Substructure in lens galaxies? *MNRAS*, 295:587–594, 1998. arXiv.

[2] J.W. Hsueh et al. SHARP - IV. An apparent flux ratio anomaly resolved by the edge-on disc in B0712+472. *MNRAS*, 469(3):3713–3721, 2017. arXiv.

[3] N. Dalal and C.S. Kochanek. Direct Detection of CDM Substructure. *ApJ*, 572:25–33, 2002. arXiv.

[4] Y.D. Hezaveh et al. Detection of Lensing Substructure Using ALMA Observations of the Dusty Galaxy SDP.81. *ApJ*, 823(1):37–56, 2016. arXiv.

[5] S. Vegetti and L.V.E. Koopmans. Bayesian strong gravitational-lens modelling on adaptive grids: objective detection of mass substructure in Galaxies. *MNRAS*, 392(3):945–963, 2009. arXiv.

[6] L.V.E. Koopmans. Gravitational imaging of cold dark matter substructures. *MNRAS*, 363(4):1136–1144, 2005. Oxford Journals.

[7] S. Vegetti and L.V.E. Koopmans. Statistics of mass substructure from strong gravitational lensing: quantifying the mass fraction and mass function. *MNRAS*, 400:1583–1592, 2009. arXiv.

[8] Tansu Daylan, Francis-Yan Cyr-Racine, Ana Diaz Rivero, Cora Dvorkin, and Douglas P. Finkbeiner. Probing the small-scale structure in strongly lensed systems via transdimensional inference. *Astrophys. J.*, 854(2):141, 2018.

[9] S. Vegetti, L. V. E. Koopmans, A. Bolton, T. Treu, and R. Gavazzi. Detection of a dark substructure through gravitational imaging. *MNRAS*, 408(4):1969–1981, November 2010.

[10] Tansu Daylan et al. Probing the Small-scale Structure in Strongly Lensed Systems via Transdimensional Inference. *ApJ*, 854(2):141, February 2018.

[11] S. Vegetti et al. Detection of a dark substructure through gravitational imaging. *MNRAS*, 408(4):1969–1981, November 2010.

[12] Stephon Alexander, Sergei Gleyzer, Evan McDonough, Michael W. Toomey, and Emanuele Usai. Deep Learning the Morphology of Dark Matter Substructure. *Astrophys. J.*, 893:15, 2020.

[13] Stephon Alexander, Sergei Gleyzer, Hanna Parul, Pranath Reddy, Marcos Tidball, and Michael W. Toomey. Domain Adaptation for Simulation-based Dark Matter Searches with Strong Gravitational Lensing. *Astrophys. J.*, 954(1):28, 2023.

[14] Stephon Alexander, Sergei Gleyzer, Hanna Parul, Pranath Reddy, Michael W. Toomey, Emanuele Usai, and Ryker Von Klar. Decoding Dark Matter Substructure without Supervision. *arXiv:2008.12731 [astro-ph, physics:hep-ph]*, August 2020. arXiv: 2008.12731.

[15] Johann Brehmer, Siddharth Mishra-Sharma, Joeri Hermans, Gilles Louppe, and Kyle Cranmer. Mining for dark matter substructure: Inferring subhalo population properties from strong lenses with machine learning. *The Astrophysical Journal*, 886(1):49, nov 2019.

[16] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Science*, 117(48):30055–30062, December 2020.

[17] Adam Coogan, Konstantin Karchev, and Christoph Weniger. Targeted Likelihood-Free Inference of Dark Matter Substructure in Strongly-Lensed Galaxies. *arXiv e-prints*, page arXiv:2010.07032, October 2020.

[18] Noemi Anau Montel, Adam Coogan, Camila Correa, Konstantin Karchev, and Christoph Weniger. Estimating the warm dark matter mass from strong lensing images with truncated marginal neural ratio estimation. , November 2022.

[19] Yashar D. Hezaveh, Laurence Perreault Levasseur, and Philip J. Marshall. Fast automated analysis of strong gravitational lenses with convolutional neural networks. , 548(7669):555–557, August 2017.

[20] Laurence Perreault Levasseur, Yashar D. Hezaveh, and Risa H. Wechsler. Uncertainties in Parameters Estimated with Neural Networks: Application to Strong Gravitational Lensing. , 850(1):L7, November 2017.

[21] S. Schuldt, S. H. Suyu, T. Meinhardt, L. Leal-Taixé, R. Cañameras, S. Taubenberger, and A. Halkola. HOLISMOKES. IV. Efficient mass modeling of strong lenses through deep learning. , 646:A126, February 2021.

[22] Bryan Ostdiek, Ana Diaz Rivero, and Cora Dvorkin. Extracting the Subhalo Mass Function from Strong Lens Images with Image Segmentation. *arXiv e-prints*, page arXiv:2009.06639, September 2020.

[23] Joshua Yao-Yu Lin, Hang Yu, Warren Morningstar, Jian Peng, and Gilbert Holder. Hunting for Dark Matter Subhalos in Strong Gravitational Lensing with Neural Networks. *arXiv e-prints*, page arXiv:2010.12960, October 2020.

[24] Ana Diaz Rivero and Cora Dvorkin. Direct detection of dark matter substructure in strong lens images with convolutional neural networks. , 101(2):023515, January 2020.

[25] Sreedevi Varma, Malcolm Fairbairn, and Julio Figueroa. Dark Matter Subhalos, Strong Lensing and Machine Learning. *arXiv e-prints*, page arXiv:2005.05353, May 2020.

[26] Kyriakos Vattis, Michael W. Toomey, and Savvas M. Koushiappas. Deep learning the astrometric signature of dark matter substructure. *Phys. Rev. D*, 104(12):123541, 2021.

[27] Bryan Ostdiek, Ana Diaz Rivero, and Cora Dvorkin. Image segmentation for analyzing galaxy-galaxy strong lensing systems. *Astron. Astrophys.*, 657:L14, 2022.

[28] Sebastian Wagner-Carena, Jelle Aalbers, Simon Birrer, Ethan O. Nadler, Elise Darragh-Ford, Philip J. Marshall, and Risa H. Wechsler. From Images to Dark Matter: End-to-end Inference of Substructure from Hundreds of Strong Gravitational Lenses. , 942(2):75, January 2023.

[29] Optogpt: A foundation model for inverse design in optical multilayer thin film structures.

[30] Yuxiang Chen, Fengyu Zhang, Zhibo Dang, Xiao He, Chunxiong Luo, Zhengchang Liu, Pu Peng, Yuchen Dai, Yijing Huang, Yu Li, and Zheyu Fang. Chiral detection of biomolecules based on reinforcement learning, 2023.

[31] Sreedevi Varma, Malcolm Fairbairn, and Julio Figueroa. Dark Matter Subhalos, Strong Lensing and Machine Learning. *arXiv e-prints*, page arXiv:2005.05353, May 2020.

[32] Johann Brehmer, Siddharth Mishra-Sharma, Joeri Hermans, Gilles Louppe, and Kyle Cranmer. Mining for Dark Matter Substructure: Inferring subhalo population properties from strong lenses with machine learning. *Astrophys. J.*, 886(1):49, Nov 2019.

[33] Aprajita Verma, Thomas Collett, Graham P. Smith, Strong Lensing Science Collaboration, and the DESC Strong Lensing Science Working Group. Strong Lensing considerations for the LSST observing strategy. *arXiv e-prints*, page arXiv:1902.05141, February 2019.

[34] Masamune Oguri and Philip J. Marshall. Gravitationally lensed quasars and supernovae in future wide-field optical imaging surveys. *MNRAS*, 405(4):2579–2593, July 2010.

[35] Kevin Schawinski, Ce Zhang, Hantian Zhang, Lucas Fowler, and Gokula Krishnan Santhanam. Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit. , 467(1):L110–L114, May 2017.

[36] John Soltis, Michelle Ntampaka, John F. Wu, John ZuHone, August Evrard, Arya Farahi, Matthew Ho, and Daisuke Nagai. A Machine-learning Approach to Enhancing eROSITA Observations. , 940(1):60, November 2022.

[37] Sam F. Sweere, Ivan Valtchanov, Maggie Lieu, Antonia Vojtekova, Eva Verdugo, Maria Santos-Lleo, Florian Pacaud, Alexia Briassouli, and Daniel Cámpora Pérez. Deep Learning-Based Super-Resolution and De-Noising for XMM-Newton Images. In *SciOps 2022: Artificial Intelligence for Science and Operations in Astronomy (SCIOPS)*, page 19, May 2022.

[38] Hiroaki Aihara, Yusra AlSayyad, Makoto Ando, Robert Armstrong, James Bosch, Eiichi Egami, Hisanori Furusawa, Junko Furusawa, Sumiko Harasawa, Yuichi Harikane, Bau-Ching Hsieh, Hiroyuki Ikeda, Kei Ito, Ikuru Iwata, Tadayuki Kodama, Michitaro Koike, Mitsuru Kokubo, Yutaka Komiyama, Xiangchong Li, Yongming Liang, Yen-Ting Lin, Robert H. Lupton, Nate B. Lust, Lauren A. MacArthur, Ken Mawatari, Sogo Mineo, Hironao Miyatake, Satoshi Miyazaki, Surhud More, Takahiro Morishima, Hitoshi Murayama, Kimihiko Nakajima, Fumiaki Nakata, Atsushi J. Nishizawa, Masamune Oguri, Nobuhiro Okabe, Yuki Okura, Yoshiaki Ono, Ken Osato, Masami Ouchi, Yen-Chen Pan, Andrés A. Plazas Malagón, Paul A. Price, Sophie L. Reed, Eli S. Rykoff, Takatoshi Shibuya, Mirko Simunovic, Michael A. Strauss, Kanako Sugimori, Yasushi Suto, Nao Suzuki, Masahiro Takada, Yuhei Takagi, Tadafumi Takata, Satoshi Takita, Masayuki Tanaka, Shenli Tang, Dan S. Taranu, Tsuyoshi Terai, Yoshiki Toba, Edwin L. Turner, Hisakazu Uchiyama, Bovornpratch Vijarnwannaluk, Christopher Z. Waters, Yoshihiko Yamada, Naoaki Yamamoto, and Takuji Yamashita. Third data release of the Hyper Suprime-Cam Subaru Strategic Program. *PASJ*, 74(2):247–272, April 2022.

[39] R. Canameras, S. Schuldt, Y. Shu, S. H. Suyu, S. Taubenberger, T. Meinhardt, L. Leal-Taixe, D. C. Y. Chao, K. T. Inoue, A. T. Jaelani, and A. More. VizieR Online Data Catalog: HSC-SSP lens candidates from neural networks (Canameras+, 2021). *VizieR Online Data Catalog*, pages J/A+A/653/L6, September 2021.

[40] R. Cañameras, S. Schuldt, S. H. Suyu, S. Taubenberger, T. Meinhardt, L. Leal-Taixé, C. Lemon, K. Rojas, and E. Savary. HOLISMOKES. II. Identifying galaxy-scale strong gravitational lenses in Pan-STARRS using convolutional neural networks. , 644:A163, December 2020.

[41] H. T. Diehl, E. J. Buckley-Geer, K. A. Lindgren, B. Nord, H. Gaitsch, S. Gaitsch, H. Lin, S. Allam, T. E. Collett, C. Furlanetto, M. S. S. Gill, A. More, J. Nightingale, C. Odden, A. Pellico, D. L. Tucker, L. N. da Costa, A. Fausti Neto, N. Kuropatkin, M. Soares-Santos, B. Welch, Y. Zhang, J. A. Frieman, F. B. Abdalla, J. Annis, A. Benoit-Lévy, E. Bertin, D. Brooks, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, C. E. Cunha, C. B. D'Andrea, S. Desai, J. P. Dietrich, A. Drlica-Wagner, A. E. Evrard, D. A. Finley, B. Flaugher, J. García-Bellido, D. W. Gerdes, D. A. Goldstein, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, D. J. James, K. Kuehn, S. Kuhlmann, O. Lahav, T. S. Li, M. Lima, M. A. G. Maia, J. L. Marshall, F. Menanteau, R. Miquel, R. C. Nichol, P. Nugent, R. L. C. Ogando, A. A. Plazas, K. Reil, A. K. Romer, M. Sako, E. Sanchez, B. Santiago, V. Scarpine, R. Schindler, M. Schubnell, I. Sevilla-Noarbe, E. Sheldon, M. Smith, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, D. Thomas, A. R. Walker, and DES Collaboration. The DES Bright Arcs Survey: Hundreds of Candidate Strongly Lensed Galaxy Systems from the Dark Energy Survey Science Verification and Year 1 Observations. , 232(1):15, September 2017.

[42] Emily O. Garvin, Sandor Kruk, Claude Cornen, Rachana Bhatawdekar, Raoul Cañameras, and Bruno Merín. Hubble Asteroid Hunter. II. Identifying strong gravitational lenses in HST images with crowdsourcing. , 667:A141, November 2022.

[43] X. Huang, C. Storfer, V. Ravi, A. Pilon, M. Domingo, D. J. Schlegel, S. Bailey, A. Dey, R. R. Gupta, D. Herrera, S. Juneau, M. Landriau, D. Lang, A. Meisner, J. Moustakas, A. D. Myers, E. F. Schlafly, F. Valdes, B. A. Weaver, J. Yang, and C. Yèche. Finding Strong Gravitational Lenses in the DESI DECam Legacy Survey. , 894(1):78, May 2020.

[44] X. Huang, C. Storfer, A. Gu, V. Ravi, A. Pilon, W. Sheu, R. Venguswamy, S. Banka, A. Dey, M. Landriau, D. Lang, A. Meisner, J. Moustakas, A. D. Myers, R. Sajith, E. F. Schlafly, and D. J. Schlegel. Discovering New Strong Gravitational Lenses in the DESI Legacy Imaging Surveys. , 909(1):27, March 2021.

[45] C. Jacobs, T. Collett, K. Glazebrook, E. Buckley-Geer, H. T. Diehl, H. Lin, C. McCarthy, A. K. Qin, C. Odden, M. Caso Escudero, P. Dial, V. J. Yung, S. Gaitsch, A. Pellico, K. A. Lindgren, T. M. C. Abbott, J. Annis, S. Avila, D. Brooks, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, L. N. da Costa, J. De Vicente, P. Fosalba, J. Frieman, J. García-Bellido, E. Gaztanaga, D. A. Goldstein, D. Gruen, R. A. Gruendl, J. Gschwend, D. L. Hollowood, K. Honscheid, B. Hoyle, D. J. James, E. Krause, N. Kuropatkin, O. Lahav, M. Lima, M. A. G. Maia, J. L. Marshall, R. Miquel, A. A. Plazas, A. Roodman, E. Sanchez, V. Scarpine, S. Serrano, I. Sevilla-Noarbe, M. Smith, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, V. Vikram, A. R. Walker, Y. Zhang, and DES Collaboration. An Extended Catalog of Galaxy-Galaxy Strong Gravitational Lenses Discovered in DES Using Convolutional Neural Networks. , 243(1):17, July 2019.

[46] R. Li, N. R. Napolitano, C. Spiniello, C. Tortora, K. Kuijken, L. V. E. Koopmans, P. Schneider, F. Getman, L. Xie, L. Long, W. Shu, G. Vernardos, Z. Huang, G. Covone, A. Dvornik, C. Heymans, H. Hildebrandt, M. Radovich, and A. H. Wright. High-quality Strong Lens Candidates in the Final Kilo-Degree Survey Footprint. , 923(1):16, December 2021.

[47] Milad Pourrahmani, Hooshang Nayyeri, and Asantha Cooray. LensFlow: A Convolutional Neural Network in Search of Strong Gravitational Lenses. , 856(1):68, March 2018.

[48] K. Rojas, E. Savary, B. Clément, M. Maus, F. Courbin, C. Lemon, J. H. H. Chan, G. Vernardos, R. Joseph, R. Cañameras, and A. Galan. Search of strong lens systems in the Dark Energy Survey using convolutional neural networks. , 668:A73, December 2022.

[49] Yiping Shu, Raoul Cañameras, Stefan Schuldt, Sherry H. Suyu, Stefan Taubenberger, Kaiki Taro Inoue, and Anton T. Jaelani. HOLISMOKES. VIII. High-redshift, strong-lens search in the Hyper Suprime-Cam Subaru Strategic Program. , 662:A4, June 2022.

[50] George Stein, Jacqueline Blaum, Peter Harrington, Tomislav Medan, and Zarija Lukić. Mining for Strong Gravitational Lenses with Self-supervised Learning. , 932(2):107, June 2022.

[51] C. Storfer, X. Huang, A. Gu, W. Sheu, S. Banka, A. Dey, A. Jain, J. Kwon, D. Lang, V. Lee, A. Meisner, J. Moustakas, A. D. Myers, S. Tabares-Tarquinio, E. F. Schlafly, and D. J. Schlegel. New Strong Gravitational Lenses from the DESI Legacy Imaging Surveys Data Release 9. *arXiv e-prints*, page arXiv:2206.02764, June 2022.

[52] Leonidas Moustakas. The Master Lens Database and The Orphan Lenses Project. HST Proposal ID 12833. Cycle 20, October 2012.

[53] A. Sonnenfeld, A. Verma, A. More, E. Beaten, C. MacMillan, K. C. Wong, J. H. H. Chan, A. T. Jaelani, C. Lee, M. Oguri, C. E. Rusu, M. Veldthuis, L. Trouille, P. J. Marshall, R. Hutchings, C. Allen, J. O'Donnell, C. Cornen, C. P. Davis, A. McMaster, C. Lintott, and G. Miller. VizieR Online Data Catalog: SuGOHI VI. List up to 2020 (Sonnenfeld+, 2020). *VizieR Online Data Catalog*, pages J/A+A/642/A148, August 2020.

[54] Anton T. Jaelani, Cristian E. Rusu, Issha Kayo, Anupreeta More, Alessandro Sonnenfeld, John D. Silverman, Malte Schramm, Timo Anguita, Naohisa Inada, Daichi Kondo, Paul L. Schechter, Khee-Gan Lee, Masamune Oguri, James H. H. Chan, Kenneth C. Wong, and Kaiki T. Inoue. Survey of Gravitationally Lensed Objects in HSC Imaging (SuGOHI) - VII. Discovery and confirmation of three strongly lensed quasars. , 502(1):1487–1493, March 2021.

[55] Kenneth C. Wong, James H. H. Chan, Dani C. Y. Chao, Anton T. Jaelani, Issha Kayo, Chien-Hsiu Lee, Anupreeta More, and Masamune Oguri. Survey of Gravitationally lensed objects in HSC Imaging (SuGOHI). VIII. New galaxy-scale lenses from the HSC SSP. *PASJ*, 74(5):1209–1219, October 2022.

[56] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. *arXiv e-prints*, page arXiv:2006.11239, June 2020.

[57] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. *arXiv e-prints*, page arXiv:1503.03585, March 2015.

[58] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv e-prints*, page arXiv:1505.04597, May 2015.

[59] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, page arXiv:1412.6980, December 2014.

[60] Alex Nichol and Prafulla Dhariwal. Improved Denoising Diffusion Probabilistic Models. *arXiv e-prints*, page arXiv:2102.09672, February 2021.

[61] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv e-prints*, page arXiv:1912.01703, December 2019.

[62] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2472–2481, 2018.

[63] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image Super-Resolution Using Very Deep Residual Channel Attention Networks. *arXiv e-prints*, page arXiv:1807.02758, July 2018.

[64] Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR: Image Restoration Using Swin Transformer. *arXiv e-prints*, page arXiv:2108.10257, August 2021.

[65] Xiangyu Chen, Xintao Wang, Jiantao Zhou, Yu Qiao, and Chao Dong. Activating More Pixels in Image Super-Resolution Transformer. *arXiv e-prints*, page arXiv:2205.04437, May 2022.

[66] Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin C. K. Chan, and Chen Change Loy. Exploiting Diffusion Prior for Real-World Image Super-Resolution. *arXiv e-prints*, page arXiv:2305.07015, May 2023.

A Ablation Study

In this section, we conduct an ablation study to verify the effectiveness of the denoising pipeline used for the low-resolution HSC samples. The preprocessing pipeline plays an important role in reducing background interference and providing a clearer context for the diffusion model. By systematically isolating individual components of the pipeline, we aim to quantitatively assess the contribution of each denoising step. We first retrain our model without any kind of denoising of the samples, followed by applying a single denoising step of either Gaussian filtering or non-local mean (NLM) denoising, and then compare the results with those obtained using the full denoising pipeline. Our objective is to demonstrate the necessity of denoising the low-resolution samples and the robustness offered by applying a stack of filtering steps. The UNet architecture and other model hyperparameters were kept constant to ensure a fair comparison.

The results of the ablation study are presented in Table 2. Initially, the model was evaluated without any preprocessing, and the results indicate a noticable compromise in structural integrity and an

Table 2: Comparative Performance of Different Denoising Steps

Denoising	L1	SSIM	PSNR
No Denoising	0.0521169	0.77484	32.68117
NLM Denoising	0.0637655	0.79436	32.90603
Gaussian Filtering	0.0426754	0.82149	34.17958
`DiffLense` Pipeline	0.0393266	0.83937	35.06834

elevated noise level in the output images, underscoring the necessity of preprocessing. Introducing NLM denoising as a standalone step yielded a modest improvement in the metrics. However, we noticed that the outputs were inconsistent. While this step produced samples that were close in quality to those generated using the full denoising pipeline, there were also samples with significantly lower quality, which resulted in reduced average quantitative metric results. This inconsistency highlights a lack of robustness when using only the NLM denoising step.

When Gaussian filtering was employed instead of NLM, the model exhibited further improvement in the metrics, underscoring the Gaussian filter's effectiveness in smoothing out noise while preserving critical image features. Lastly, the full denoising pipeline used in our approach achieved the best results. These findings clearly highlight that the full preprocessing pipeline contributes meaningfully to the model's ability to generate super-resolved images and is crucial for optimizing the model's ability to accurately preserve structural details and minimize noise.