LensPINN: Physics Informed Neural Network for Learning Dark Matter Morphology in Lensing

Ashutosh Ojha
Indian Institute of Technology Dhanbad
Dhanbad, Jharkhand 826004, India
ojhaaashutosh1005@gmail.com

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

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

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

Abstract

We present LensPINN, a Physics-Informed Neural Network architecture for studying dark matter through strong gravitational lensing images. Our approach integrates the gravitational lensing equation into the model, combining the capabilities of Vision Transformer (ViT) and CNN frameworks. The architecture incorporates the lensing equation to perform lensing inversion for the reconstruction of the source galaxy. The deflection angle due to gravitational lensing is learned by the Vision Transformer Encoder, and the information from the source image is then passed through the architecture to enhance model learning, leading to improved performance in tasks related to strong gravitational lensing and dark matter localization. In this paper, we focus on a classification task that distinguishes between simulations of different dark matter models. We compare the performance of our model, LensPINN, with previous state-of-the-art models and other leading architectures. We propose two versions of the LensPINN model: LensPINN_small, which is highly efficient, having only half the number of parameters while performing on par with all other models, and LensPINN_large, which has the same number of parameters as existing models but surpasses all of them across various metrics.

1 Introduction

Despite over a century of research, our understanding of the fundamental nature of dark matter remains as elusive as it was at the time of its discovery. Numerous experimental efforts have sought to detect leading dark matter candidates, yet none have yielded definitive results to date [1, 2, 3, 4, 5, 6]. While much of the search has focused on potential interactions between dark matter and the Standard Model, another promising approach is to explore its gravitational effects. A particularly active area of research involves using strong galaxy-galaxy lensing data to study and constrain dark matter properties. The sensitivity of extended lensing arcs to small perturbations in the lensing potential, caused by substructure in the main halo, offers a unique opportunity for probing dark matter. Studies have demonstrated that differences in the distribution and morphology of these substructures can help distinguish between different dark matter models [7, 8, 9, 10, 11, 12, 13]. However, the task is far from simple. The complexity of the problem, especially when dealing with sub-dominant lensing signals, makes extracting meaningful information particularly challenging.

Machine Learning and the Physical Sciences Workshop, NeurIPS 2024.

The inherent challenges of working with lensing data have led to extensive exploration of machine learning algorithms in this field. Currently, only a limited number of high-quality lensing datasets are available for training purposes. Fortunately, upcoming initiatives like the Vera Rubin Observatory (VRO) [14] and Euclid [15] are expected to generate thousands of images, thereby increasing the demand for more advanced analytical methods. In the meantime, simulations have become essential for investigating the integration of machine learning with strong lensing studies. Numerous research efforts have demonstrated that machine learning holds significant potential for extracting dark matter information from lensing data, as evidenced by studies such as [16, 17, 18, 19, 20]. Specifically, these investigations have consistently found that convolutional neural networks (CNNs) are particularly well-suited for strong lensing applications. This suitability is unsurprising because CNNs possess two key attributes: 1) they exploit correlations within image data, and 2) they exhibit translation invariance. A major advantage of CNNs is their ability to automatically filter out irrelevant or redundant "degrees of freedom" due to their inherent translational invariance. Moreover, incorporating known redundancies from a dataset directly into the network architecture has been shown to enhance model performance, as demonstrated by the use of equivariant neural networks for domain adaptation in simulated strong lensing data [21].In addition the Vision Transformers (ViT)[22] utilizes the self-attention mechanism to capture long-range dependencies between pixels, making them excellent at recognizing global patterns across the entire image.

Along similar lines, we have seen the emergence of integrating physics into neural network architectures, leading to the development of Physics-Informed Neural Networks (PINNs). PINNs incorporate essential physical information into the network, improving its learning capabilities. In the context of dark matter morphology, PINNs have been applied, such as in the Lensformer [23]. In this work, we propose a Physics-Informed Neural Network called `LensPINN`, which leverages the strengths of both the self-attention mechanism of Vision Transformers (ViTs) and CNNs. The gravitational lensing equation [24] is directly embedded into its architecture. To demonstrate `LensPINN`'s superior performance, we benchmarked its ability to classify mock Hubble observations of galaxy-galaxy strong lensing across different realizations of dark matter.

2 Dataset

To evaluate our proposed architecture, we simulate galaxy-galaxy strong lensing data resembling observations from an HST-like survey. These datasets are generated using the publicly available code `lenstronomy`[25], producing 64×64 pixel images, each in a single channel. The background galaxies, which are lensed in our simulations, are modeled using a Sersic light profile. We create three distinct classes of dark matter models.

The first model represents the standard cold dark matter (CDM) scenario, where the main dark matter halo is modeled with a spherical isothermal profile and includes subhalos drawn from a standard subhalo mass distribution (refer to [16] for details). The second model simulates the lensing signature of very light axion dark matter, with a particle mass of approximately 10^{-23} eV. In this case, substructure formation is significantly suppressed, and the primary observable feature arises from topological defects in the halo, manifesting as vortices [16]. The third model is a dark matter scenario devoid of any substructure. While not realistic, as observational data clearly rule it out, it serves as a useful comparison against the other substructure classes in simulations.

In total, we generated 3,000 simulations per class (axion, cold dark matter, no substructure), with an 80:20 train-test split applied across the three dark matter models.

3 Preprocessing

We have applied a physics-informed preprocessing technique to the image:

$$\text{preprocessed image} = \left| \tanh \left(\nabla_x \nabla_y \left(\log \left(\frac{I_{\max}}{I} \right) \right)^2 \right) \right|$$

The $\frac{I_{\max}}{I}$ is used for contrast enhancement highlighting faint structures. The logarithm function is applied to compress and filter the dynamic range of intensities, allowing subtle variations to be more noticeable. Squaring is then used to ensure all the values are positive, amplifying intensity differences.

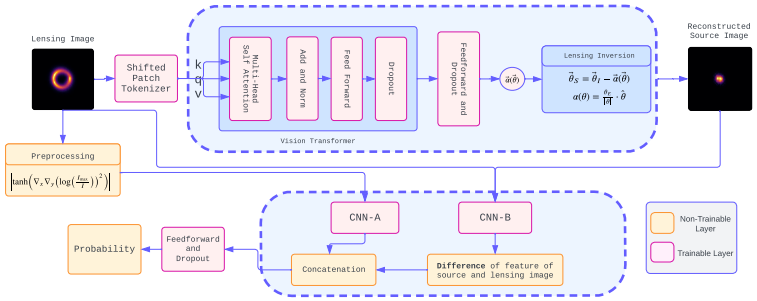

Figure 1: `LensPINN` Architecture.

Double differentiation is performed to identify sharp changes in intensity(edge detection). Finally, normalization is done using the $\tanh$ function, improving model convergence and stability.

4 `LensPINN` Architecture

The proposed architecture follows an encoder-decoder framework. The encoder uses a Vision Transformer architecture to encode the angular deflection and applies the gravitational lensing equation for lensing inversion, thereby reconstructing the source image. The decoder, which uses Convolutional Neural Networks (CNN), takes the reconstructed source image, the original image, and the preprocessed image as inputs. By providing the decoder with these three inputs, we aim to enhance the model's ability to make informed decisions.

4.1 Physics Informed Encoder

The encoder section of the model uses the Vision Transformer (ViT)[22] architecture. It utilizes the gravitational lensing equation for reconstructing the source galaxy image (refer to [24] for details).In the Inverse Lens Layer (see Fig. 1), we apply the lens equation in its dimensionless form: $\vec{\theta}_S = \vec{\theta}_I - \vec{\alpha}(\vec{\theta}_I)$[26]. Here, $\vec{\theta}_S = (x_s, y_s)$ denotes the dimensionless source position in the source plane, corresponding to the source galaxy, while $\vec{\theta}_I = (x_i, y_i)$ represents the dimensionless image position in the image plane, corresponding to the observed image.The term $\vec{\alpha}(x_i, y_i) = (\alpha_x(x_i, y_i), \alpha_y(x_i, y_i))$ represents the deflection angle. The deflection angle, $\vec{\alpha}(\vec{\theta})$, depends on the mass distribution of the lens and is expressed as the gradient of the gravitational potential produced by the lens: $\vec{\alpha}(\vec{\theta}_I) = \vec{\nabla}\psi(\vec{\theta}_I)$ where $\psi(\vec{\theta}_I)$ represents the gravitational potential, which includes contributions from both the lensing galaxy and potential dark matter. For a singular isothermal sphere (SIS) model, the deflection angle

$$\alpha(\theta) = \theta_E \frac{\theta}{|\theta|} [24]$$

is directly proportional to the Einstein radius θ_E and is directed radially towards the lens. The lens equation then becomes

$$\beta = \theta - \theta_E \frac{\theta}{|\theta|}$$

where β is the source position and θ is the observed position[24](after lensing effect). The process begins by encoding the Einstein Radius using the ViT, which is then utilized to determine the source coordinates via the lensing equation. If the source coordinates fall outside the grid of the source plane, the grid is expanded to accommodate the magnification effects of strong gravitational lensing. When multiple points on the image plane map to the same point on the source plane—commonly observed in strong gravitational lensing (formation of multiple images occurs due to the strong gravitational field of the lensing galaxy)—pixel values are averaged to manage this overlap effectively. At the end, after establishing the relationship between the source and observed image coordinates, we fill in the pixel values on the source plane, resulting in the reconstructed source image.

Table 1: Performance of Different Models (Parameters in millions; Best in **bold**, 2nd best underlined)

Model Name	Parameters(M)	Accuracy	Micro F1 Score	ROC AUC Score		
				No Subs.	CDM	Axion
ResNet18	11.17	0.818	0.817	0.97	0.85	0.95
ViT	13.72	0.863	0.864	0.99	0.66	0.96
CaiT	13.76	0.878	0.871	0.99	0.64	0.96
ViTSD	13.73	0.867	0.868	1.00	0.66	0.97
Lensformer	15.7	0.957	0.959	1.00	0.99	0.99
LensPINN_small	**7.17**	0.956	0.957	1.00	0.99	0.99
LensPINN_large	14.56	**0.996**	**0.996**	**1.00**	**1.00**	**1.00**

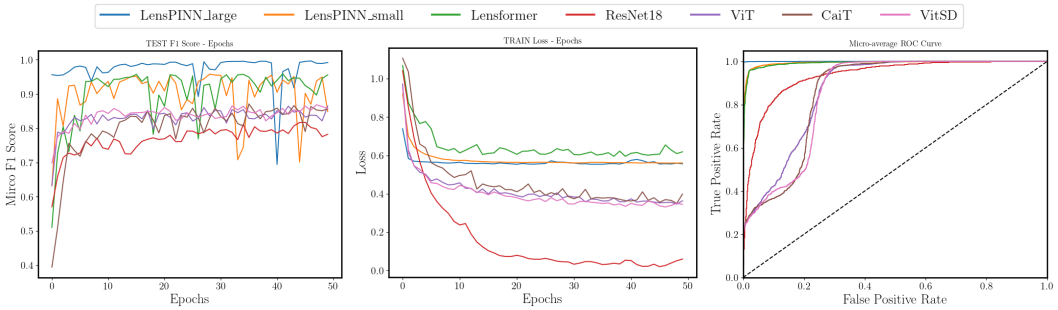

Figure 2: Model Performance Metrics. The left plot shows Micro F1 score vs epochs, the middle plot shows training loss vs epochs and the right panel displays the Micro-Averaged ROC-AUC curves for various models.

4.2 Decoder

The decoder section of the model consists of two Convolutional Neural Networks (CNNs), which excel at local feature extraction. The goal is for the model to learn from the differences between the source image and the observed lensed image, helping it identify the factors causing these differences, specifically the lensing galaxy and its dark matter substructure. The first CNN extracts features from the source and lensed images, while the second CNN performs feature extraction on the preprocessed image. For `LensPINN_small`, we used MobileNetv3[27] as the CNN, and for `LensPINN_large`, we used EfficientNetB0[28]. We then obtain three feature vectors: one for the source image, one for the lensed image, and one for the preprocessed image. The difference between the features of the source and lensed images is computed and concatenated with the features of the preprocessed image. This combined feature vector is then passed through fully connected layers for classification.

5 Experiments

We conducted a comprehensive study comparing various models, including Convolutional Neural Networks (CNNs) and Transformer models such as the Vision Transformer (ViT) [22], CaiT [29], ViTSD [30], ResNet18 [31], and Lensformer(a PINN) [23]. All models had the same number of parameters (around 14 million). They were fine-tuned and trained for 50 epochs, allowing the loss curve to converge and to avoid overfitting. We used the AdamW optimizer with a learning rate of 1e-4. The evaluation metrics were the micro F1 and ROC AUC scores. Both `LensPINN_small` and `LensPINN_large` outperformed the others, with `LensPINN_small` achieving a 0.957 F1 score and `LensPINN_large` scoring 0.996. `LensPINN_small` performed on par with the others but had half the number of parameters. Both `LensPINN_small` and `LensPINN_large` also converged significantly faster (see Fig.3) than any other model. For a detailed comparison of results, see Table 1. These results clearly demonstrate that integrating physics into the model enhances its learning ability, leading to quicker convergence (see Fig. 3) and superior performance.

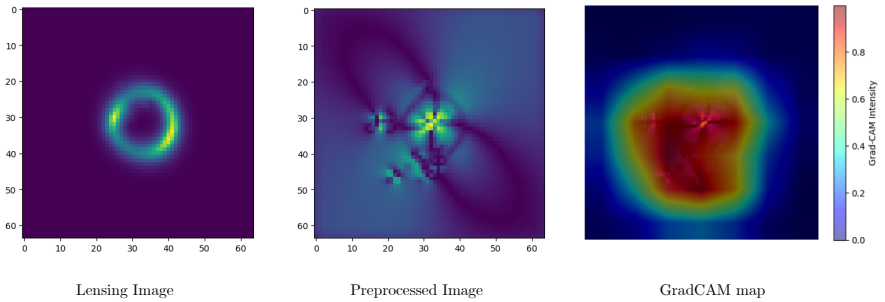

Figure 3: GradCAM visualization of a CNN model trained on lensing images, with the final image showing the GradCAM overlay on the preprocessed image.

6 Ablation Study

Model	Training Data	Micro F1-Score
MobileNet	Raw Lensing Image	0.750
MobileNet	Physics-Informed Preprocessing	0.970
LensPINN_small	Physics-Informed Preprocessing + Raw Lensing Image	0.952
EfficientNet	Raw Lensing Image	0.850
EfficientNet	Physics-Informed Preprocessing	0.980
LensPINN_large	Physics-Informed Preprocessing + Raw Lensing Image	0.997

Table 2: Comparison of Model Performance over preprocessed vs not preprocessed dataset

6.1 Importance of Preprocessing

The preprocessed images highlight regions that assist the model in making better decisions. This is evident from the fact that when training a CNN solely on the original images, the Grad-CAM maps reveal that the areas the model focuses on are aligned with those emphasized in the preprocessed images. This suggests that the preprocessing highlights key regions that the model would likely focus on even without any physics-based information. Therefore, the preprocessing significantly aids the model in making more informed decisions.

6.2 Importance of PINN

The table 2 shows that the CNN trained on preprocessed images performs better than those trained solely on the min-max scaled raw lensing images, and Lens_PINN large outperforms all of them. This can be attributed to the fact that the decoder part of LensPINN receives context from both the lensing image and the preprocessed image, as well as the source image. However, MobileNet outperforms `LensPINN_small`, which may be due to the fact that in the decoder part, a single CNN (MobileNet) is tasked with extracting feature maps from both the source and lensing images using the same CNN module. This setup may be suboptimal, as MobileNet's simpler architecture struggles to adapt to the diverse features of both the source and lensing images simultaneously. When MobileNet focuses solely on the preprocessed image, a more uniform and simplified representation, it performs better. On the other hand, in `LensPINN_large`, which uses EfficientNet for feature extraction, which handles feature representations more effectively than MobileNet, which is why LensPINN_large outperforms `LensPINN_small` and all other models.

7 Result and Discussion

The key takeaway from our experiments is that the `LensPINN` model not only **converges rapidly** but also demonstrates **superior performance**, despite having **half the number of parameters** compared to other models (see Table. 1). This improvement can be attributed to the incorporation of physics information into the model, which enhances its decision-making abilities. The `LensPINN`'s ability to learn from the difference between the source and observed images is crucial to its success.

Acknowledgements

A.O. was a participant in the Google Summer of Code 2024 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. acknowledges financial support from the Simons Foundation (Grant Number 929255).

References

[1] A. K. Drukier, Katherine Freese, and D. N. Spergel. Detecting Cold Dark Matter Candidates. *Phys. Rev.*, D33:3495–3508, 1986.

[2] Mark W. Goodman and Edward Witten. Detectability of Certain Dark Matter Candidates. *Phys. Rev.*, D31:3059, 1985. [,325(1984)].

[3] D. S. Akerib et al. Results from a search for dark matter in the complete LUX exposure. *Phys. Rev. Lett.*, 118(2):021303, 2017.

[4] Xiangyi Cui et al. Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment. *Phys. Rev. Lett.*, 119(18):181302, 2017.

[5] E. Aprile et al. Dark Matter Search Results from a One Ton-Year Exposure of XENON1T. *Phys. Rev. Lett.*, 121(11):111302, 2018.

[6] Morad Aaboud et al. Constraints on mediator-based dark matter and scalar dark energy models using $\sqrt{s} = 13$ TeV pp collision data collected by the ATLAS detector. *JHEP*, 05:142, 2019.

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

[8] 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.

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

[10] 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.

[11] 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.

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

[13] 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.

[14] 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, 2019.

[15] Masamune Oguri and Philip J. Marshall. Gravitationally lensed quasars and supernovae in future wide-field optical imaging surveys. *Monthly Notices of the Royal Astronomical Society*, pages no–no, apr 2010.

[16] Stephon Alexander, Sergei Gleyzer, Evan McDonough, Michael W. Toomey, and Emanuele Usai. Deep learning the morphology of dark matter substructure. *The Astrophysical Journal*, 893(1):15, apr 2020.

[17] Siddharth Mishra-Sharma and Ge Yang. Strong Lensing Source Reconstruction Using Continuous Neural Fields. In *39th International Conference on Machine Learning Conference*, 6 2022.

[18] Gemma Zhang, Siddharth Mishra-Sharma, and Cora Dvorkin. Inferring subhalo effective density slopes from strong lensing observations with neural likelihood-ratio estimation. *Monthly Notices of the Royal Astronomical Society*, 517(3):4317–4326, 10 2022.

[19] Thuruthipilly, Hareesh, Zadrozny, Adam, Pollo, Agnieszka, and Biesiada, Marek. Finding strong gravitational lenses through self-attention - study based on the bologna lens challenge. *A&A*, 664:A4, 2022.

[20] Ana Diaz Rivero and Cora Dvorkin. Direct Detection of Dark Matter Substructure in Strong Lens Images with Convolutional Neural Networks. *Phys. Rev. D*, 101(2):023515, 2020.

[21] 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.

[22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.

[23] Lucas J. Velôso, Sergei Gleyzer, and Michael W. Toomey. Lensformer: A physics-informed vision transformer for gravitational lensing, 2023. Presented at NeurIPS ML4PS 2023 Workshop.

[24] Joachim Wambsganss Peter Schneider, Christopher Kochanek. Gravitational lensing: Strong, weak, and micro, 2006.

[25] Simon Birrer, Adam Amara, and Alexandre Refregier. Gravitational lens modeling with basis sets. *The Astrophysical Journal*, 813(2):102, nov 2015.

[26] Ramesh Narayan and Matthias Bartelmann. Lectures on gravitational lensing, 1997.

[27] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for mobilenetv3, 2019.

[28] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks, 2019.

[29] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers, 2021.

[30] Seung Hoon Lee, Seunghyun Lee, and Byung Cheol Song. Vision transformer for small-size datasets, 2021.

[31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.