
Deconvolution of Astronomical Images with Deep Neural Networks

Hong Wang

Computational Science Initiative,
Brookhaven National Laboratory,
Upton, NY 11973-5000, USA
hwang6@bnl.gov

Sreevarsha Sreejith

Physics Department,
Brookhaven National Laboratory,
Upton, NY 11973-5000, USA
ssreejith@bnl.gov

Yuewei Lin

Computational Science Initiative,
Brookhaven National Laboratory,
Upton, NY 11973-5000, USA
ywlin@bnl.gov

Nesar Ramachandra

Computational Science Division,
Argonne National Laboratory,
Lemont, IL 60439, USA
nramachandra@anl.gov

Anže Slosar

Physics Department,
Brookhaven National Laboratory,
Upton, NY 11973-5000, USA
anze@bnl.gov

Shinjae Yoo

Computational Science Initiative,
Brookhaven National Laboratory,
Upton, NY 11973-5000, USA
sjyoo@bnl.gov

Abstract

Optical astronomical images are strongly affected by the point spread function (PSF) of the optical system and the atmosphere (seeing) which blurs the observed image. The amount of blurring depends on both the observed band, and more crucially, on the atmospheric conditions during observation. A typical astronomical image will therefore have a unique PSF that is non-circular and different in different bands. Observations of known stars give us an estimation of this PSF. Any serious candidate for production analysis of astronomical images must take the known PSF into account during image analysis. So far the majority of applications of neural networks (NN) to astronomical image analysis have ignored this problem by assuming a fixed PSF in training and validation. We present a neural network architecture based on Deep Wiener Deconvolution Network (DWDN) that takes the PSF into account when performing deconvolution, a possible approach of leveraging PSF information in neural networks. We study the performance of this algorithm under realistic observational conditions. We employ two regularization schemes and study custom loss functions that are optimized for quantities of interest to astronomers. We show that our algorithm can successfully recover unbiased image properties such as colors, ellipticities and orientations for sufficiently high signal-to-noise. This study represents a comprehensive application of AI in astronomy, where the experimental design, model construction, optimization criteria, error estimation and metrics of benchmarks are all meticulously tailored to the domain problem.

1 Introduction

The advent of large telescopes and big data sets are bringing about a transformable era for astronomical survey science. New datasets obtained from upcoming ground-based and space-based observing

facilities will contain data from large volumes of the observable sky at unprecedented depths and cadences [11, 1, 7, 14, 6]. While the resolution of space-based telescopes is usually diffraction limited, the large apertures of ground based survey telescopes are limited by *seeing*. Seeing is a distortion of the image caused by the perturbation of an optical wavefront as it passes through the turbulent atmosphere. This results in a finite point-spread function (PSF). The ground based surveys often do not employ adaptive optics due to the large field of view, resulting in an arcsecond sized point spread function.

The two main features of the PSF is that it is a variable and that it is known in each exposure. The PSF is a variable because the atmosphere above the telescope fluctuates. It varies in size, shape and orientation. It is typically asymmetric and different in different bands. On the other hand, it is also known in each exposure, because invariably every exposure contains stars that are excellent point sources and therefore can directly be used to probe the PSF.

Recovery of an object’s shape, position, flux and other intrinsic parameters in the presence of noise and varying PSF is goal of astronomical image analysis. Machine learning techniques using convolutional operations have demonstrated promising avenues for speeding up the astronomical processing. Majority of these grapple with the galaxy deblending problem such as [15, 3, 2, 18]. However, most of the current neural-network based galaxy deblenders have a number of simplifying assumptions, including a constant PSF. This problem is usually “swept under the rug” by assuming that several training sets can be provided, one for each PSF. However, this is likely impractical for reasons we will cover in the following sections. It is also clear that simply ignoring this problem will lead to biased and sub-optimally inferred parameters for individual objects. Therefore the proper treatment of PSF is one of the issues that need to be addressed before neural network based approaches can be lifted from the realm of toy problems and benchmarking datasets to real observations from telescopes. Other issues include blending, masking and modeling various detector artifacts. In this work we specifically address the issue of leveraging PSF information and build a network that deconvolves a noisy, convolved input image to produce the true image of the galaxy at the resolution supposed by the PSF.

2 Description of the problem

We use a linear model for the observed astronomical image, a convolution with a point spread function followed by addition of observational noise:

$$I_i = I_{gt,i} \otimes \text{PSF}_i + N_i, \tag{1}$$

where PSF_i is the PSF for the channel i , N_i is the noise, assumed to be homogeneous and normally distributed and index gt refers to ground truth. It is important to note that PSF is band dependent and that the amplitude of noise is also band-dependent. This model is an approximation, because the brighter-fatter effects, amplifier non-linearities, and other detector phenomena make the transformation weakly non-linear. The shape of PSF also varies across the focal plane. However, in this work we assumed that the region of interest is always small enough that the the PSF can be assumed constant.

One of the classical regularization techniques of deconvolution is the Wiener Deconvolution formulated in Fourier space as (dropping the band subscripts for clarity):

$$\text{WD} = \frac{P_s(k)}{P_s(k) + P_n(k)} \text{PSF}^{-1}(k), \tag{2}$$

where $P_s(k)$ and $P_n(k)$ are spectral densities of signal and noise respectively. A neural network extension of the above formulation is the Deep Wiener Deconvolution Network [5].

3 Band-wise Deep Wiener Deconvolution Network

Works using neural networks for deblurring have emerged recently [9, 17], among which the Deep Wiener Deconvolution Network (DWDN, [5]) provides a brand new solution for the non-blind image deblurring problems. Instead of deconvolution in the image space like the existing methods do, DWDN applies Wiener Deconvolution explicitly in the feature space. As a simple but effective integration of classical Wiener Deconvolution and deep learning, DWDN achieves outstanding performance with fewer artifacts. DWDN contains two components, the feature-based Wiener

deconvolution module and a multi-scale feature refinement module, where the former carries out the deconvolution process and the latter restore high-quality images using features from the previous module.

A restriction that limits the direct use of DWDN on astronomical images is that typically, for galaxy images, different bands have different PSFs. However, in DWDN, the feature maps extracted from the input have already mixed up and no longer have distinguishable bands as the image does. Thus, we propose an enhanced version of DWDN that can work on band-dependent PSFs. Figure 1 shows the architecture of our model. In order to apply feature-based deconvolution on different input bands separately, we parallel three deconvolution modules. This number varies based on the number of bands in the input image and PSF. Each band of the input image with the corresponding PSF band will go through one of those modules to complete the feature-based deconvolution. Eq. 3 expresses the deconvolution for l -th feature channel in i -th deconvolution module

$$h_i^l = \mathcal{F}^{-1} \left(\frac{\overline{\text{PSF}_i(k)}}{\overline{\text{PSF}_i(k)\text{PSF}_i(k)} + \frac{s_{i,l}^n}{s_{i,l}^x}} \right) \quad (3)$$

$$H_i = \mathcal{C}[h_i^1, \dots, h_i^L] \quad (4)$$

where h_i^l is the deconvolved feature of the l^{th} channel, and L is the total number of feature channels. \mathcal{F} refers to discrete Fourier transformation and $\overline{\text{PSF}_i(k)}$ is the complex conjugate of $\text{PSF}_i(k)$. $s_{i,l}^n$ and $s_{i,l}^x$ are the variance and standard deviation of this channel for noise and input image separately which can be estimated by blurred features and mean-filtered features. After concatenation as Eq. 4, we get output from i -th deconvolution module H_i , that is, deconvolved features for i -th band of the input image and PSF. H_i will then be passed to the second module \mathcal{R} as in Eq. 5 with C denoting the number of deconvolution modules.

$$\hat{I} = \mathcal{R}(H_1, \dots, H_C) \quad (5)$$

We keep the refining module \mathcal{R} the same except the number of channels of its input feature is C times as original DWDN because of the concatenation. In this way, different bands of the input image will execute their own deep Wiener deconvolution and the deconvolved features are utilized to restore the clean images. More details of the architecture can be found in [5].

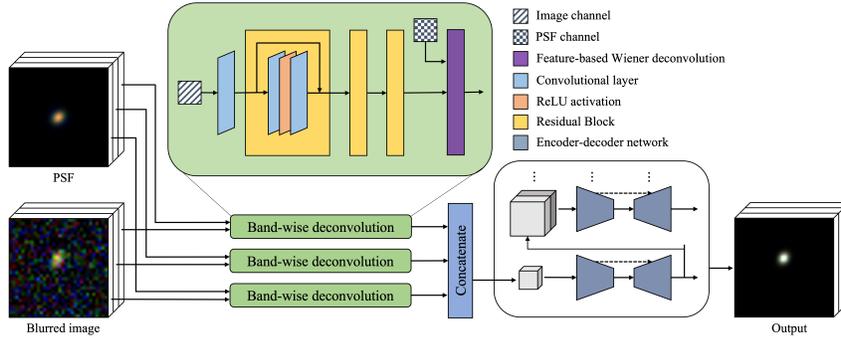


Figure 1: The framework of the model used in this paper. The three band-wise deconvolution modules share the same architecture. The feature refinement module contains convolution-based encoder-decoder network for multiple levels. The kernel size is 5×5 for all convolutional layers.

4 Custom loss function

In majority of image analysis problems, the networks are trained with one of the standard loss-functions that encode some intuitive notion of similarity between the true and reconstructed image. Typical example includes ℓ_1 , ℓ_2 loss functions and the cross-entropy functions.

However, in astronomical image analysis we are interested in particular properties of the output image. In fact the first few image moments described in the next section naturally describe the astronomical

quantities of interest: object brightness, its position and its shape. It is therefore natural to include them into the loss function minimized during training – this way we ask neural network to focus explicitly on the quantities of interest. We will describe this non-standard training procedure next.

4.1 Quantities of interest

Let I_{ij} be flux value in pixel (i, j) . We define raw moments as

$$M_{pq}(I) = \sum_{x=1}^W \sum_{y=1}^H x^p y^q I_{x,y} \quad (6)$$

The following quantities are of interest:

M_{00}	Total flux	
$\langle x \rangle = \frac{M_{10}}{M_{00}}$	Astrometric position x	
$\langle y \rangle = \frac{M_{01}}{M_{00}}$	Astrometric position y	
$\mu_{20} = \frac{M_{20}}{M_{00}} - \langle x \rangle^2$	Second central moment x^2	
$\mu_{11} = \frac{M_{11}}{M_{00}} - \langle x \rangle \langle y \rangle$	Second central moment xy	(7)
$\mu_{02} = \frac{M_{02}}{M_{00}} - \langle y \rangle^2$	Second central moment y^2	
$e_1 = \frac{\mu_{20} - \mu_{02}}{\mu_{20} + \mu_{02}}$	+ Ellipticity component	
$e_2 = \frac{2\mu_{11}}{\mu_{20} + \mu_{02}}$	× Ellipticity component	

These quantities can be trivially calculated from a noise-less image in a way that enable their use in loss function.

4.2 Training procedure and loss function

The model is trained in an end-to-end manner with two phases. In the first phase, the model is pre-trained to acquire the ability to deblur and produce reasonable prediction as output, while in the second phase, the model is fine-tuned with a more complex loss including quantities listed in Section 4.1 for more accurate description of the galaxy.

The loss function for the first phase is formulated in Eq. 10. It contains ℓ_1 -norm between the output and the ground truth as shown in Eq. 8 and the difference of the total flux as Eq. 9. In Eq. 8, f is the network, I and I_{gt} refer to the input and true image. λ_0 controls the trade-off between the two parts.

$$\mathcal{L}_{\ell_1} = \frac{1}{WHC} \sum_{c=1}^C \sum_{x=1}^W \sum_{y=1}^H |f(I)_{x,y,c} - (I_{gt})_{x,y,c}| \quad (8)$$

$$\mathcal{L}_{M_{00}} = \frac{1}{C} \sum_{c=1}^C |M_{00}(f(I)) - M_{00}(I_{gt})| \quad (9)$$

$$\mathcal{L}_1(I, I_{gt}) = \mathcal{L}_{\ell_1} + \lambda_0 \mathcal{L}_{M_{00}} \quad (10)$$

In the second phase, in addition to \mathcal{L}_{ℓ_1} and $\mathcal{L}_{M_{00}}$, the astrometric position $\langle x \rangle$ and $\langle y \rangle$, the second-order central moments μ_{11} , μ_{20} , μ_{02} are utilized to fine-tune the model. The quantities are defined in Section 4.1 and the loss for these quantities has the same form as Eq. 9. Thus, the total loss for the second phase can be written as Eq. 11. λ 's are the weights for different terms and manually adjusted so that those terms are subdominant, but not negligible compared to the dominant ℓ_1 loss.

$$\mathcal{L}_2(I, I_{gt}) = \mathcal{L}_{\ell_1} + \lambda_0 \mathcal{L}_{M_{00}} + \lambda_1 \mathcal{L}_{\langle x \rangle} + \lambda_2 \mathcal{L}_{\langle y \rangle} + \lambda_3 \mathcal{L}_{\mu_{11}} + \lambda_4 \mathcal{L}_{\mu_{20}} + \lambda_5 \mathcal{L}_{\mu_{02}} \quad (11)$$

5 Training and testing data

The dataset used in this work was generated using WeakLensingDeblending package¹ developed in order to test the shear measurement of weak lensing by LSST Data Management. From the galaxy catalog class FixedGalaxyCatalog, composite exponential galaxy images of single galaxy

¹<https://github.com/LSSTDESC/WeakLensingDeblending/>

	Accurate PSF		Average PSF		Accurate PSF (ℓ_1 loss)		Gaussian regularization	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
PSNR	46.408	45.951	45.172	44.991	46.314	45.477	46.273	45.724
1-SSIM	0.011	0.006	0.012	0.007	0.011	0.007	0.011	0.006

Table 1: PSNR(dB) and SSIM

are simulated in *g*, *r* and *i* bands with single exposures per band and Gaussian PSFs that vary across the different bands (using industry standard `galSim`, [16]). For each galaxy, a noisy image, a truth image and a PSF image are generated as a stack of the three aforementioned bands. Since our objective is to ascertain the influence of incorporating PSF in neural networks for denoising purposes of astronomical images, we have set shear to be zero. Moreover, we incorporated a signal-to-noise (SNR) cut of > 10 .

The 100,000 images thus generated for each of the three categories (truth, noisy and PSF) have dimensions of $35 \times 35 \times 3$. The entire dataset is divided into training set with 90,000 images and test sets with 10,000 images. Each image is sized $32 \times 32 \times 3$ and normalized to $[0, 1]$ by dividing its overall maximum pixel value so that colors are preserved. During testing, in order to get the accurate performance of the model, the prediction from the network are rescaled back by exactly same factor which allows a direct comparison with the raw ground truth.

6 Results

6.1 Implementation and comparison models

For experiments, we adopt Adam optimizer [8] and the batch size is 128. During the first phase, the model is trained for 80 epochs with learning rate as 10^{-4} . In the second phase, we fine-tune the model for 10 epochs and the learning rate is set to 10^{-5} . The weights for moment-related loss are as follows: $\lambda_{0,1,2} = 0.01$, $\lambda_3 = 0.002$ and $\lambda_{4,5} = 0.001$.

In order to evaluate the performance of our model, we train 4 different models:

- The fiducial model as explained in the text;
- The *Average PSF* model, where the training and testing datasets remain the same, but instead of fitting actual per-image PSF to the network, we freeze the PSF to its average shape. This allows us to explore the information gain coming from knowledge of the actual PSF;
- The *ℓ_1 loss only* model, where training is not fine-tuned with custom loss function;
- The *Gaussian regularization* model where Wiener deconvolution formula is replaced with a different regularization as described in Appendix.

6.2 Evaluation

In most tests we split the sample into three categories by the value of M_{00} as a proxy for signal-to-noise. Figure 2 visualizes the performance of the fiducial model on the testing set. We select one example from samples with low, medium and high M_{00} separately. From the results, our deblurring model can recover the morphological and color information for images under different conditions, even for images with low fluxes. It can also be observed that the model performs better on images with higher fluxes. We will analyse in detail in the following sections.

We first apply some standard metrics such as the peak signal-to-noise ratio (PSNR) and the structural similarity index (SSIM) [19] to evaluate the quality of images recovered from our model compared with the ground truth images. The results are listed in Table 1, including the mean and median among the testing set. The best results are from the model with accurate PSF and moment loss. Specifically, by comparison between the first and third columns in the table, we can conclude that those moment-related loss terms help improve the performance, reflected by a higher PSNR and SSIM value, but the improvement is modest. The results from models with accurate PSF and average PSF also demonstrate that the quality of the prediction will get degraded if the accurate PSF is not provided for each image again with surprisingly modest improvements.

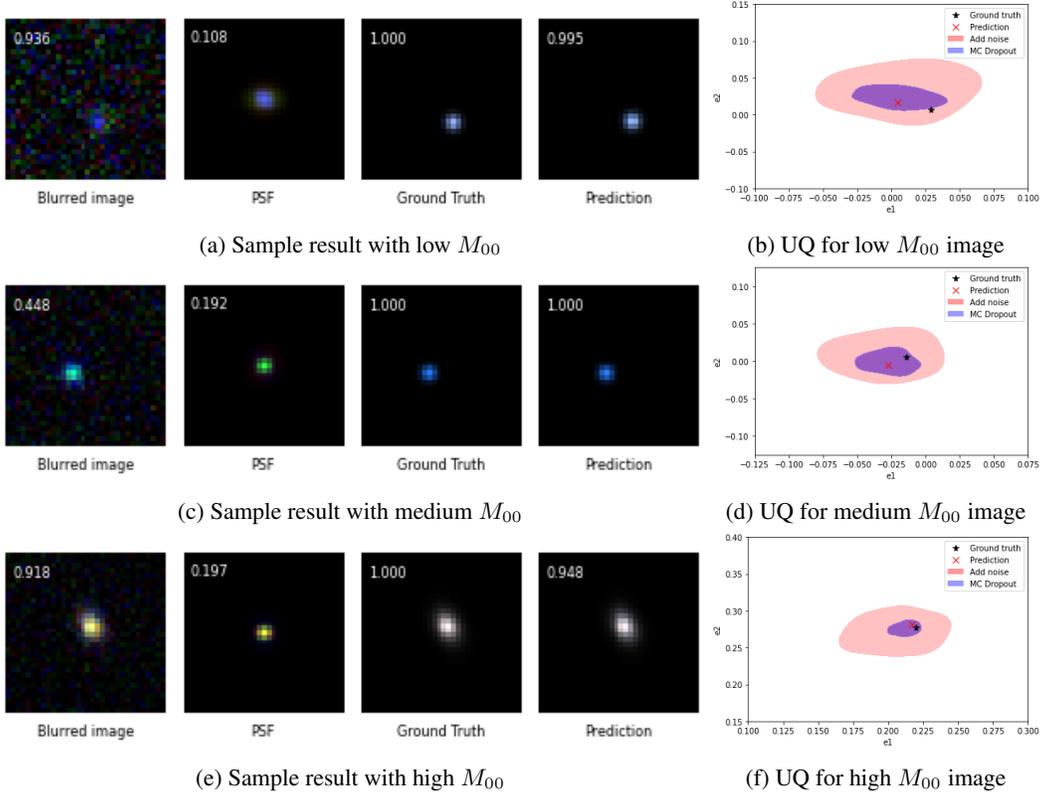


Figure 2: Visualization of typical results. Value of M_{00} is a proxy for object signal-to-noise.

6.3 Recovery of astronomical quantities

Table 2 in Appendix presents the difference between output from the model and true image for astronomical quantities of interest. These values are averaged over the three available bands. For our fiducial network, the relevant scatter plots can also be found in Figure 3. We find that across the noise levels, the networks correctly recover the quantities of interest. The results nicely scatter around $x = y$ line with no obvious bias for the total flux, first moments (astrometric position) and second moments. Ellipticity suffers from small bias, where the network tends to preferably make objects less round. This effect is observed for both e_1 and e_2 . To understand this effect better we have split objects by signal-to-noise and replotted the histograms for e_1 in the Figure 4a. Here we see that high signal-to-noise objects have their shapes recovered without bias, while there is an emergence of "multiplicative" bias for low signal-to-noise where measured ellipticities are systematically low.

Next we want to study how much information we are getting by using a per-band known PSF. Figure 4b plots the histograms for model with average PSF. We see essentially the same structure, but the scatter around the mean is somewhat increased. Similar observations for e_2 can be found in Figure 5 in the Appendix. This is confirmed in the Table 2, where rms errors when using average PSF are systematically larger by about 5-10% compared to those of fiducial network for ellipticity and second moments. We also find that the error on the astrometries are not affected by knowing or not the true PSF – this makes sense because our PSFs are Gaussian with zero mean and therefore using a wrong PSF to deconvolve the image is not going to affect its central position. Somewhat more surprising is that the error on the object flux also seems to be unchanged between the two cases. More analysis on the recovery of shapes and its biases can be found in Appendix.

6.4 Uncertainty quantification (UQ)

Like majority of the deep learning frameworks, our model is also a point-prediction method, i.e., the predictions in the network are not associated with error estimates. Large number of uncertainty quantification approaches are currently applied in the context of AI applications in physics [4, 13, 10],

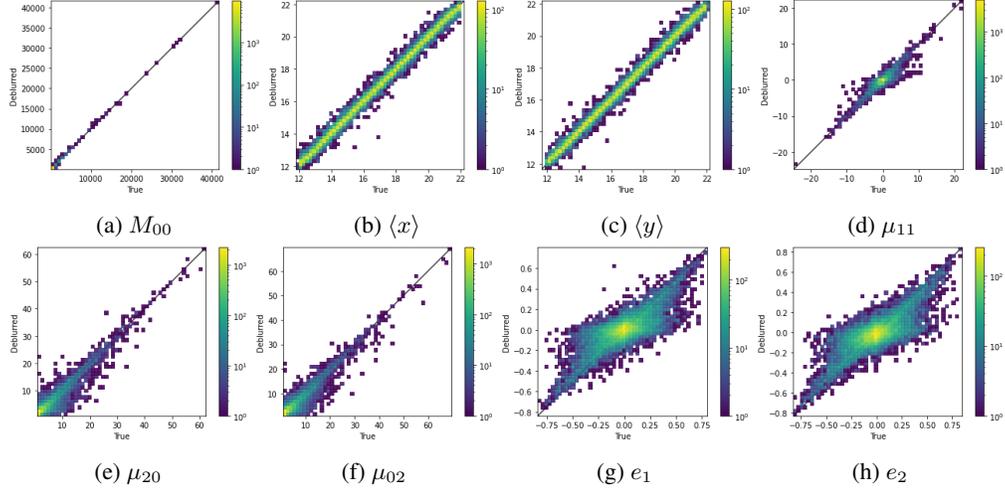


Figure 3: 2-d histograms for moment recovery from model trained with accurate PSF. Different panels are for different quantities as indicated. The x axis was calculated on the true image and the y axis on the recovered image. The density of points is represented with color. Note that the color scaling is logarithmic to bring out the outliers.

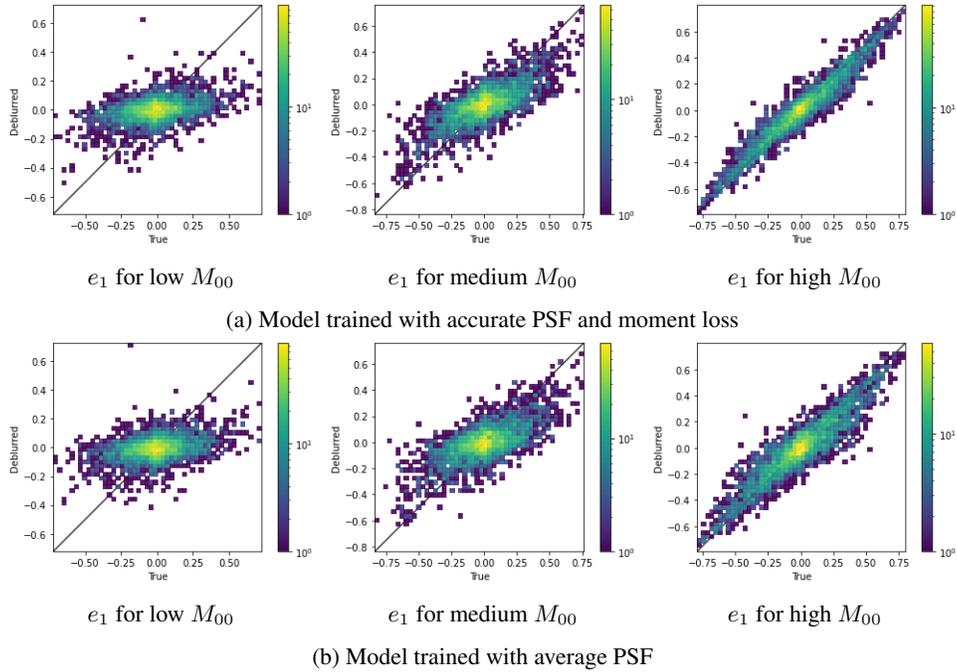


Figure 4: 2D distributions of e_1 recovery split by signal-to-noise.

some of which may be computationally expensive in our problem, where the individual model training requires 25.5 hours on one Nvidia V100 GPU with 32GB memory.

In this paper, we perform a two-fold treatment of uncertainty. The first one aims at estimating the model uncertainties, where we quantify the fluctuations in the predicted images due to imperfect model training. Second, we quantify the uncertainty propagated due to the noisy input data during testing. Both these treatments are approximations aimed at obtaining a rudimentary understanding of the confidence intervals of our methods, within reasonable computational budgets. In applications involving real observations, several systematic effects may have to be considered in a greater detail.

Model uncertainty using Monte Carlo methods: For model uncertainties, we choose the UQ technique of Monte Carlo (MC) Dropout, that provides approximate errors with little computational overhead and minimal changes to the architecture. The implementation of MC Dropout is: First, we train a model without any UQ considerations to obtain a fiducial model. Next, we consider an ensemble of N_{tot} models with the same architecture, only differing from each other by a fraction of trained neurons re-initialized to a random value. We obtain this ensemble of networks and the fraction of randomized (or “dropped-out”) weights is dictated by the dropout rate d . Since the Deep Neural Networks are over-parametrized, a small number of randomized weights will only result in small fluctuations in the output images. Each image in the testing set will be evaluated by those N_{tot} models and obtain N_{tot} outputs. The mean and variance of these outputs will be considered as the uncertainty quantified prediction from the ensemble.

Statistical uncertainty due to noisy image inputs: In addition to model uncertainty which in physicist language corresponds to systematic error, any observation also has noise or random uncertainty. Since neural network results are point estimates, one would ideally process the same image with multiple noise realizations to estimate the noise in any derived quantities. Since this is impossible, we instead generate N_{tot} images with additional noise drawn from the same distribution as the real noise and run those through the network. These form an ensemble from which a noise estimate on the derived quantities can be formed. Since these images now have double the noise variance, the result is strictly speaking an overestimate of measurement uncertainty.

In the experiment, $N_{tot} = 200$ for both approaches and dropout rate $d = 0.2$. Results of these two techniques are presented in the right panel of Figure 2. We see that the truth is always within the estimated uncertainty region, whose size is likely overestimated. We also see that the size of the statistical uncertainty decreases with increasing SNR, but interestingly the size of the systematic uncertainty also somewhat shrinks with an increasing signal to noise. In other words, as the SNR increase, the network seems to internally produce more consistent results.

7 Discussion & Conclusions

Despite numerous applications of neural networks to the astronomical image analysis, the problem of using the information from the known, band dependent and relatively large PSF has not been addressed directly. In this work, we propose a neural network based deconvolution algorithm to address this crucial missing link. It derives from a network that deploys Wiener Deconvolution in the feature space. We have two notable improvements aimed at the properties of astronomical tasks. Firstly, instead of a simplifying assumption of constant PSF, we employ a band-wise deconvolution module. This enables the network to work with each band separately, rendering is compatible with the PDFs of the imaging systems in astronomy. Secondly, we engage quantities representing particular properties of astronomical images into our custom loss function. The model can have better investigation on those properties by directly training on those quantities. Extensive experiments and analysis demonstrate that the model can acquire more better recovery of shape than comparison models, especially for images with higher SNR. This work is the start of incorporating band-separated PSFs information in astronomical image processing. More broadly, this study motivates the need of robust, domain-focussed architectures in AI-for-Astronomy problems.

Acknowledgements

We thank Erin Sheldon for useful discussions. This material is partly based upon the work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and Office of High Energy Physics, Scientific Discovery through Advanced Computing (SciDAC) program on "Accelerating HEP Science: Inference and Machine Learning at Extreme Scales". NR’s work at Argonne National Laboratory was supported by the U.S. Department of Energy, Office of High Energy Physics. Argonne, a U.S. Department of Energy Office of Science Laboratory, is operated by UChicago Argonne LLC under contract no. DE-AC02-06CH11357.

References

- [1] T. M. C. Abbott, F. B. Abdalla, A. Alarcon, J. Aleksić, S. Allam, and et al. Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing. , 98(4):043526, Aug. 2018.

- [2] B. Arcelin, C. Doux, E. Aubourg, C. Roucelle, and LSST Dark Energy Science Collaboration. Deblending galaxies with variational autoencoders: A joint multiband, multi-instrument approach. , 500(1):531–547, January 2021.
- [3] A. Boucaud, C. Heneka, E. E. Ishida, N. Sedaghat, R. S. de Souza, B. Moews, H. Dole, M. Castellano, E. Merlin, V. Roscani, et al. Photometry of high-redshift blended galaxies using deep learning. *Monthly Notices of the Royal Astronomical Society*, 491(2):2481–2495, 2020.
- [4] T. Y. Chen, B. Dey, A. Ghosh, M. Kagan, B. Nord, and N. Ramachandra. Interpretable Uncertainty Quantification in AI for HEP. *arXiv e-prints*, page arXiv:2208.03284, Aug. 2022.
- [5] J. Dong, S. Roth, and B. Schiele. Deep wiener deconvolution: Wiener meets deep learning for image deblurring. *Advances in Neural Information Processing Systems*, 33:1048–1059, 2020.
- [6] J. Green, P. Schechter, C. Baltay, R. Bean, D. Bennett, and et al. Wide-Field InfraRed Survey Telescope (WFIRST) Final Report. *arXiv e-prints*, page arXiv:1208.4012, Aug. 2012.
- [7] C. Hikage, M. Oguri, T. Hamana, S. More, R. Mandelbaum, and et al. Cosmology from cosmic shear power spectra with Subaru Hyper Suprime-Cam first-year data. , 71(2):43, Apr. 2019.
- [8] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [9] O. Kupyn, V. Budzan, M. Mykhailych, D. Mishkin, and J. Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8183–8192, 2018.
- [10] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. *arXiv e-prints*, page arXiv:1612.01474, Dec. 2016.
- [11] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, and et al. LSST Science Book, Version 2.0. *arXiv e-prints*, page arXiv:0912.0201, Dec. 2009.
- [12] R. Mandelbaum. Weak Lensing for Precision Cosmology. , 56:393–433, Sept. 2018.
- [13] R. M. Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [14] A. Refregier, A. Amara, T. D. Kitching, A. Rassat, R. Scaramella, and J. Weller. Euclid Imaging Consortium Science Book. *arXiv e-prints*, page arXiv:1001.0061, Jan. 2010.
- [15] D. M. Reiman and B. E. Göhre. Deblending galaxy superpositions with branched generative adversarial networks. *Monthly Notices of the Royal Astronomical Society*, 485(2), 2019.
- [16] B. Rowe, M. Jarvis, R. Mandelbaum, G. Bernstein, J. Bosch, M. Simet, J. Meyers, T. Kacprzak, R. Nakajima, J. Zuntz, H. Miyatake, J. Dietrich, R. Armstrong, P. Melchior, and M. Gill. Galsim: The modular galaxy image simulation toolkit. *Astronomy and Computing*, 10:121–150, 2015.
- [17] C. J. Schuler, M. Hirsch, S. Harmeling, and B. Schölkopf. Learning to deblur. *IEEE transactions on pattern analysis and machine intelligence*, 38(7):1439–1451, 2015.
- [18] H. Wang, S. Sreejith, A. Slosar, Y. Lin, and S. Yoo. Galaxy Deblending using Residual Dense Neural networks. *arXiv e-prints*, page arXiv:2109.09550, Sept. 2021.
- [19] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

A Additional results

Difference of moments between predicted images and ground truth: Table 2 shows the difference between the model’s output and the true image for astronomical quantities of interest. All four comparison models in the text are included.

	Accurate PSF			Average PSF		
	Mean	Variance	RMS	Mean	Variance	RMS
M_{00}	0.015	131.416	11.464	-0.053	129.831	11.394
$\langle x \rangle$	0.011	0.046	0.214	0.003	0.045	0.212
$\langle y \rangle$	0.009	0.045	0.213	0.017	0.044	0.211
μ_{20}	-0.038	2.577	1.606	-0.160	2.867	1.701
μ_{02}	-0.096	2.712	1.650	-0.169	2.977	1.734
μ_{11}	-0.017	0.413	0.643	-0.007	0.491	0.701
e_1	0.007	0.020	0.140	-0.003	0.024	0.154
e_2	-0.002	0.020	0.142	0.005	0.039	0.198
$\ e\ $	0.155	0.016	0.152	0.179	0.031	0.202

(a) Results for model with accurate PSF and average PSF, both with moment loss

	Accurate PSF (ℓ_1 loss)			Gaussian regularization		
	Mean	Variance	RMS	Mean	Variance	RMS
M_{00}	-1.142	259.151	16.139	-0.484	124.569	11.172
$\langle x \rangle$	0.005	0.050	0.224	0.002	0.050	0.225
$\langle y \rangle$	0.005	0.050	0.223	0.011	0.049	0.223
μ_{20}	-0.222	2.555	1.614	0.103	2.469	1.574
μ_{02}	-0.255	2.655	1.649	0.072	2.566	1.604
μ_{11}	-0.016	0.466	0.683	-0.001	0.440	0.664
e_1	0.003	0.023	0.151	0.003	0.018	0.134
e_2	-0.003	0.024	0.154	0.004	0.019	0.137
$\ e\ $	0.171	0.017	0.138	0.153	0.013	0.150

(b) Results for model with ℓ_1 loss only and model with Gaussian regularization

Table 2: Difference of moments between predicted images and ground truth.

Recovery of ellipticity: Figure 5 is the same as Figure 4 but for e_2 . Similar conclusion can be drawn that the model recovers the shape high signal-to-noise objects without bias and model trained with average PSF lost some information with a increased scatter around the mean compared with the fiducial model.

Recovery of shapes and its biases: We also study the performance of our network in removing the subtle biases that remain in the prediction. In particular, we want to study the remaining contamination by the PSF shape. Our main motivation is to understand whether such methods could ever be used for weak lensing application, where residual contamination with the PSF leads to to the so-called additive bias (see [12] and references there-in for an in-depth review). In short if the shape of galaxy remains correlated with the PSF after the convolution, then this PSF shape will leak into shear estimation and lead to shape correlations that can mimic real weak gravitational signal.

The true galaxy ellipticities are uncorrelated with PSF shapes, but since the observed image is a convolution of the PSF with true shape, the resulting predicted image ellipticities might still correlate with the PSF shape. On the other hand, if a measurement is perfect, the predicted ellipticities should correlated perfectly with true ones. To quantify this effect, we define the following correlations

$$\xi_{\text{PSF}} = \frac{\langle e_{\text{PSF}} \cdot e_{\text{pred}} \rangle}{\langle |e_{\text{PSF}}| |e_{\text{pred}}| \rangle}, \quad (12)$$

$$\xi_{\text{true}} = \frac{\langle e_{\text{true}} \cdot e_{\text{pred}} \rangle}{\langle |e_{\text{true}}| |e_{\text{pred}}| \rangle}, \quad (13)$$

where averages are calculated by simply averaging data in the testing dataset. Note that since individual ellipticities of either PSF, true galaxy image or predicted galaxy image can reach zero, we average the

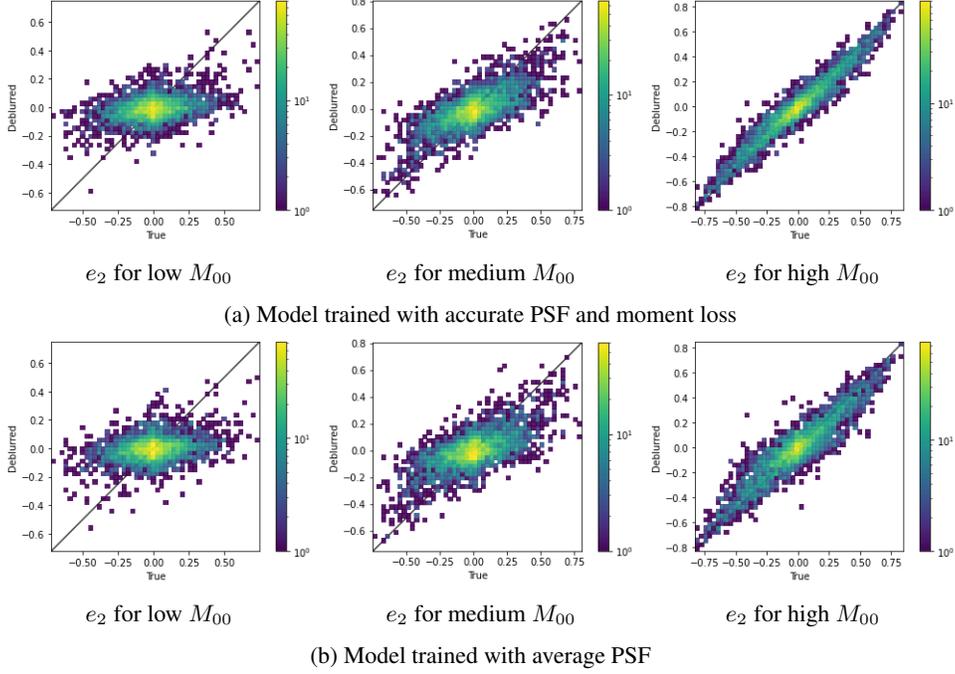


Figure 5: 2D distributions of e_2 recovery split by signal-to-noise.

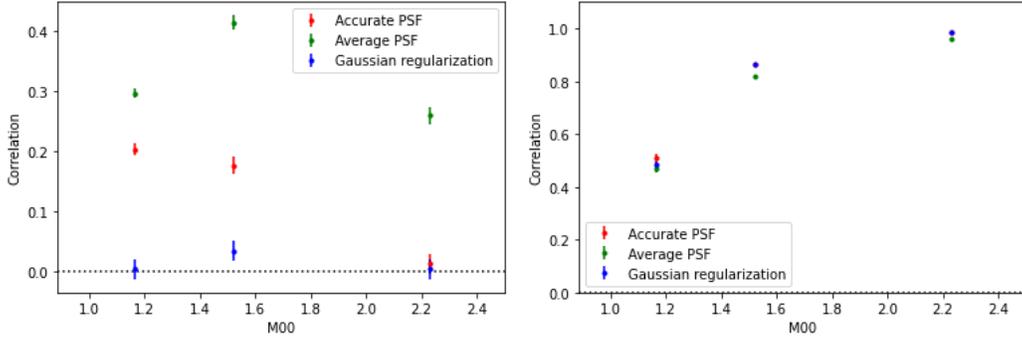


Figure 6: Cross correlation of the network output ellipticities with the PSF (left) and with the true image ellipticities (right).

numerator and denominator separately. In ideal case, we expect the predicted image to remove any trace of PSF from the resulting image, so $\xi_{\text{PSF}} = 0$. However, this could also be achieved by making all objects round (i.e. $e_{\text{pred}} = 0$ for every objects). Therefore we also look at the correlation with the true value, which should be, in the ideal case perfect, i.e $\xi_{\text{true}} = 1$.

Results of this exercise are plotted in the Figure 6. We see that our fiducial case (denoted "Accurate PSF" and plotted red) performs better than the "Average PSF" case across all signal-to-noise levels. The cross-correlation with PSF is significantly lower and the cross-correlation with the truth significantly higher. We conclude that the deconvolution is taking place, albeit imperfectly: although the network does successfully decorrelate PSF in the high SNR, but it still contains significant contamination in the lower SNR region.

In order to improve upon this, we explore another version of deconvolution in which we replace the Wiener filter Equation 2 with a modified term, which we refer to as "Gaussian Regularization":

$$\text{WD}^{\text{GaussReg}} = \frac{P_s(k)}{P_s(k) + P_n(k)G^{-2}(k)} \text{PSF}^{-1}(k), \quad (14)$$

In other words, the PSF is linearly deconvolved from the noisy image, but the resulting image is then regularized with a circular Gaussian roughly the size of the PSF. In the original formulation of the Wiener filter, modes are anisotropically regularized, taking into account that anisotropic PSF has destroyed more information in some directions than the others. However, the result is that the PSF shape can sneak back into the resulting image. The Gaussian regularization tries to prevent this at the expense of formally decreased optimality, i.e. the resulting process is now manifestly worse than minimum variance solution. We choose a Gaussian whose size is close to typical PSF size and have explicitly demonstrated that the results are largely insensitive to precise choice of this kernel, if the network is retrained.

In our tests, however, this method seems to perform pretty well. In Figure 6, we see that the Gaussian regularization model performs significantly better than our fiducial network with no evidence of PSF correlation at no loss in correlation with the truth. In Tables 2, we see that the Gaussian regularization performs essentially as well as the fiducial model within the noise and that while individual RMS values scatter favouring one or the other method, there is no clear winner.