

Galaxy Deblending Using Residual Dense Neural Networks

H. Wang, S. Sreejith

To be published in "Physical Review D"

September 2022

Physics Department
Brookhaven National Laboratory

U.S. Department of Energy
USDOE Office of Science (SC), High Energy Physics (HEP) (SC-25)

Notice: This manuscript has been authored by employees of Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy. The publisher by accepting the manuscript for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Galaxy Deblending using Residual Dense Neural network

Hong Wang

*Stony Brook University, Stony Brook, NY 11794 and
Computational Science Initiative, Brookhaven National Laboratory, Upton, NY 11973*

Sreevarsha Sreejith and Anže Slosar

Physics Department, Brookhaven National Laboratory, Upton, NY 11973

Yuewei Lin and Shinjae Yoo

Computational Science Initiative, Brookhaven National Laboratory, Upton, NY 11973

(Dated: July 19, 2022)

We present a new neural network approach for deblending galaxy images in astronomical data using Residual Dense Neural network (RDN) architecture. We train the network on synthetic galaxy images similar to the typical arrangements of field galaxies with a finite point spread function (PSF) and realistic noise levels. The main novelty of our approach is the usage of two distinct neural networks: i) a *deblending network* which isolates a single galaxy postage stamp from the composite and, ii) a *classifier network* which counts the remaining number of galaxies. The deblending proceeds by iteratively peeling one galaxy at a time from the composite until the image contains no further objects as determined by the classifier, or by other stopping criteria. By looking at the consistency in the outputs of the two networks, we can assess the quality of the deblending. We characterize the flux and shape reconstructions in different quality bins and compare our deblender with the industry standard, **SExtractor**. We also discuss possible future extensions for the project with variable PSFs and noise levels.

I. INTRODUCTION

Astronomical survey science is entering a decade which will witness a surfeit of new optical imaging data that will transform astronomical and cosmological research. The new datasets obtained from upcoming ground-based and space-based observing facilities will image vast areas of the sky at unprecedented depths. Among the most important of these will be the Legacy Survey of Space and Time (LSST) from the Vera Rubin Observatory [1], which will complement not only the already existing datasets from Dark Energy Survey (DES, [2]) and Hyper-Suprime Cam (HSC,[3]), but also the data from upcoming space-based surveys like Euclid [4]. Due to the large field of view, most of these surveys will not be able to employ adaptive optics with a point spread function (PSF) that is comparable to the object size (around 0.5-1 arcsecs on the ground and 0.1 arcsecs in space). At the depths of these surveys, this will result in blending, affecting over half of the galaxies in the survey (for example, in [5], the authors estimate that approximately 60% of the objects in the HSC are affected by this problem).

Blending, especially if undetected, can introduce serious systematic errors in survey analysis. These include potential catastrophically wrong inference of fluxes, leading to biased photometric redshifts, biased estimates of shear and local object density dependent systematics, that can interact with galaxy and cluster detections. The main approach to solving the blending problem is two-pronged. On one hand, there is the realization that source separation (deblending) will never be perfect and we should focus on understanding its properties and effects on the analysis. This is done through simulations

and artificial source injections into real data. On the other hand, the better the deblender is, the smaller the corrections needed to make the analysis unbiased. Therefore, we need both good deblenders and good schemes for understanding their imperfections.

In principle, deblending is a well defined problem. The basic model is that the images of individual galaxies are combined on the projected plane, assuming perfect transparency (i.e. intensities add up), and then observed through the atmosphere and telescope with known PSF and noise properties. The most sophisticated deblenders on the market combine machine learning approaches for setting priors on galaxy shapes, and physical modelling for things that we can model explicitly, like the noise and the PSF [6, 7].

Several machine learning based methods have been proposed recently, in order to grapple with the galaxy deblending problem. [8] designed a branched deblender with generative adversarial networks, which can work on the deblending of images with two overlapped galaxies. [9] developed a framework for measuring the photometry of blended galaxies as well as to perform segmentation with a standard convolutional neural network (CNN) and a U-Net. [5] introduced an algorithm where a Variational Auto Encoder(VAE)-like neural network was used for galaxy deblending. Most of the current generation neural network based galaxy deblenders require the galaxy to be located at the center of the image in order to give the network a sign as to which galaxy to recover, which is often impractical and prevents the model from being used iteratively, i.e., it will degrade the model's performance on the image after the first galaxy is removed. Besides, the number of galaxies in the blended image for

the previous methods is always fixed, and to the best of our knowledge, there is no neural network based galaxy deblending framework that can work on images with an arbitrary and unknown number of galaxies. In this paper, we propose an innovative deblending framework with a deblender, and a classifier that can deblend galaxies from a blended image iteratively without any prior information about the number of galaxies. In addition, since the deblender recovers galaxies based on their luminosity, our framework has no constraint on the position of the galaxies. Nevertheless, like other machine learning approaches to deblending, our work remains exploratory as we aim to better understand the applicability of neural networks to astronomical image analysis.

This paper is organized as follows: in Section II, we introduce the architecture of our framework including the deblender and the classifier, and the experimental settings used to train the model. In Section III, we present the experimental results and the comparison with the industry standard deblending method - Source Extractor (SExtractor, [10]). The discussion and conclusion follow, in the last Section IV.

II. NEURAL NETWORK ARCHITECTURE AND TRAINING

A. Proposed Framework

Our goal is to deblend galaxies from astronomical images with an arbitrary and unknown number of overlapped galaxies. The proposed framework consists of two components: a *deblender*, which isolates the image corresponding to a single galaxy from an astronomical image, and a *classifier*, which counts how many galaxies remain in the image. The *deblender* and *classifier* are then used iteratively to separate the scene into its constituent galaxy images. This is illustrated in the FIG. 1 and represented in the following meta-code:

```
while True:
    num_galaxies = classifier(image)
    if num_galaxies == 0:
        exit loop
    deblended_galaxy = deblender(image)
    image -= deblended_galaxy
```

Specifically, given a noisy blended image with multiple galaxies, the deblender will take it as the input and output a noiseless image with a single galaxy as it is trained. This single-galaxy image is then subtracted from the input of the deblender to get the residual image, which is another noisy blended image. The classifier is used to determine if there are further galaxies in the residual image, and if there are, this process is repeated until there is no galaxy left in the residual image.

In an ideal case, the classifier will detect one fewer galaxy at each step and the process will stop when there are no more galaxies left. We call such deblends *High*

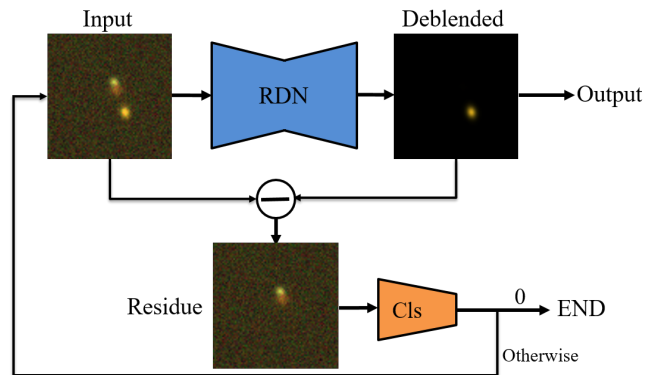


FIG. 1: Figure showing the proposed deblender-classifier framework represented as meta-code in Section II A.

quality deblends. In a non-ideal case, one of the two following scenarios often play out. The number of galaxies in the image does not decrease by one galaxy per iteration, but the process still eventually comes to a halt with zero galaxies detected in the final image. When this happens, we refer to the result as *Medium quality* deblends. The third case is if the process gets stuck in an infinite loop where the *classifier* maintains that there are more than zero galaxies in the image, but the deblender fails to locate them. Then, if the classifier predicts the same non-zero number of galaxies in the residual image for three consecutive iterations we terminate the iteration and call the results *Low quality* deblends.

B. Deblender

Given an astronomical image with multiple overlapped galaxies, the deblender aims to deblend one galaxy from it. A Residual Dense Network (RDN) [11] is trained as the deblender in this framework. RDN shows superior performance in image super-resolution [11] and image restoration [12]. For the deblending task, the RDN will take the noisy blended images with multiple galaxies as input and will give a noiseless image containing the brightest galaxy as the output.

FIG. 2 shows the architecture of the deblender in our framework. It starts with a shallow feature extractor, which includes two convolutional layers. Following this is the main component of the RDN, namely the residual dense blocks (RDBs). Each RDB contains C convolutional layers with a ReLU activation function. The layers in the RDB are densely connected to make better use of the local features. The deblender contains D RDBs in total. An advantage of the original RDN is the design of feature fusion and residual learning in multiple levels, i.e. locally within each RDB and globally among different RDBs, which will extract hierarchical features and provide the full usage of feature information. For the galaxy deblending problem in which the input images are likely to be overlapped blended galaxies with noise, the exist-

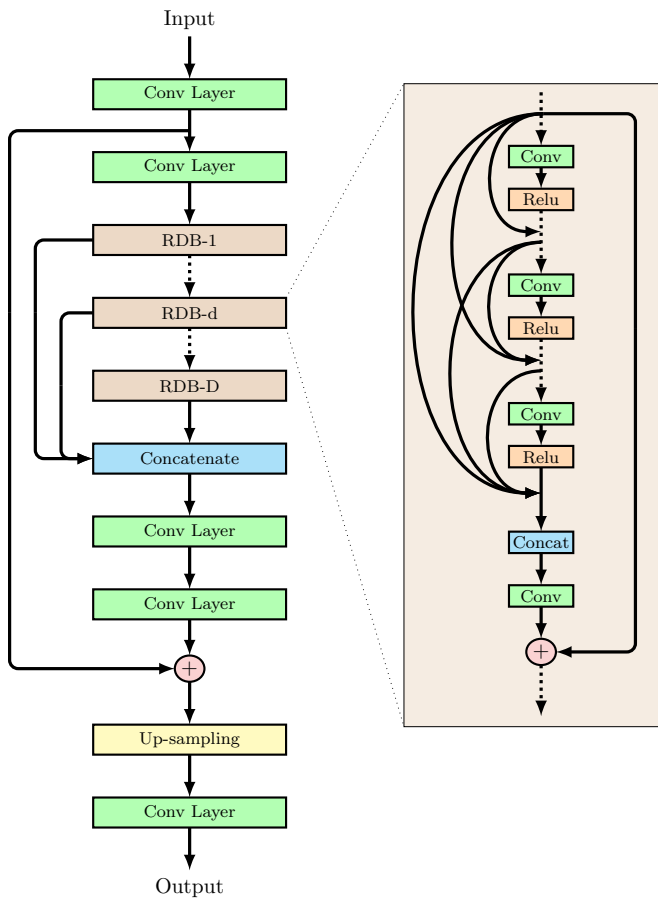


FIG. 2: Figure showing the architecture of the RDN used as the deblender in this work.

tence of those direct and dense connections for the features will help to preserve the detailed morphology and color information of the galaxies during deblending and will lead to recovering better, the pertinent galaxy. The efficient sub-pixel convolutional neural network (ESPCN) followed by a convolutional layer forms the up-sampling net in the RDN. Since the model in our framework is used for deblending, the output image will have the same size as the input image. Therefore, the scale factor in the deblender is set to be 1. During training, the deblender will take the noisy blended images as the input and output noiseless images with a single galaxy. In order to avoid confusing the deblender, we train the model to output the brightest galaxy in the input image all the time. ℓ_1 -norm is used to compute the difference between the prediction of the deblender and the ground truth, thus the loss function for the deblender is written as:

$$l_{\text{RDN}}^{\theta}(I, I_{gt}) = \frac{1}{WH} \sum_{x=1}^W \sum_{y=1}^H |f(I, \theta)_{x,y} - (I_{gt})_{x,y}| \quad (1)$$

where I is the input noisy blended image and $f_{\theta}(I)$ is the output of the RDN with θ referring to the parameters

in the model. I_{GT} denotes the ground truth image, which is a noiseless image containing the brightest galaxy.

C. Classifier

For the classifier portion of the deblending network, we use a VGG-16 network [13]. The VGG-16 contains 13 convolutional layers and 3 fully connected layers. In our framework, it is modified to give four classes as output, i.e., the classifier can tell if the images have 0, 1, 2 or 3 galaxies.

In the first phase of training, the classifier is trained on images with galaxies ranging from 0 to 3. During the second phase in the training process where the deblender and the classifier are trained jointly, the deblended images from the deblender will also be passed through the classifier to see if the RDN has deblended only one galaxy from the blended image. This loss will also be used to update the RDN's parameters. Thus, the training set of the classifier contains both noiseless and noisy images with 0 to 3 galaxies. The cross-entropy (CE) loss is utilized to train the classifier parameterized by ϕ in Eq. 2 where y represents the one-hot ground truth label.

$$l_{CE}^{\phi}(I) = -y^T \cdot \log(g(I, \phi)) \quad (2)$$

D. Training & Test Data

The galaxy images (blended and otherwise) used in this work has been generated using BlendingToolKit [14]. It is a complete framework whose functionalities include generation of images with blended objects as well as providing measurements of the deblending and detection performance. It relies on the industry standard `galsim` [15] packages to make the actual renderings of the galaxies and is built on top of the `WeakLensingDeblending` [16] package.

We first generate noiseless images of single galaxies, which can be considered as the pre-blended ground truth images. We employ the LSST DM galaxy catalog with a span of 1 square deg (`OneDegSq.fits` supplied with the `BlendingToolKit`). This catalog is derived from the `CatSim` galaxy catalog used by the LSST Data Management. It is based on dark matter haloes from the Millennium simulation and a semi-analytic baryon model from [17]. The semi-analytic model was calibrated in luminosity, color and morphology on low redshift galaxies. LSST cosmological catalogs were generated from this model by constructing a lightcone, covering redshifts $0 < z < 6$ by tiling 58 simulation snapshots with halo masses down to $2.5 \times 10^9 M_{\odot}$. The final catalog comprises of a 4.5×4.5 degree footprint on the sky, of which one square degree is sampled in `OneDegSq.fits`. The resulting galaxies are thus somewhat realistic in morphologies, luminosities and

colours, but lack the the full morphological richness of real galaxies.

We use the default options, which include pixel scale of 0.2 arcsec/pixel and impose a magnitude cut on the i -band, $i < 25.3$. This cut corresponds to the LSST gold sample cut and thus represents a set of galaxies that should be measured with reasonable accuracy by the LSST. While it is true that fainter interlopers might negatively affect the deblending procedure, we have ignored this in this exploratory paper. The default option also includes a constant Gaussian PSF with a FWHM of 0.73, 0.70 and 0.67 arcsec in the g , r & i bands respectively. Then, the noiseless blended images are generated by a pixelwise summation between these single galaxy images. Additional random Poisson noise is added to the image to get the noisy blended images corresponding to the 10-year depth survey depth of the LSST. We selected g , r and i bands from the resultant images and converted them into RGB images. The original dimension of the generated images is (120, 120, 3). We crop each image to (80, 80, 3) from the center. The pixel values of the blended images will be normalized to $[0, 1]$ before entering the framework and the pre-blended images will be scaled to $[-1, 1]$ following [8].

Galaxies in the current draft are randomly positioned (with appropriate buffer around the edges of the image), corresponding to the random chance alignments between galaxies, which is cause of the majority of blends. Blending occurring due to physical associations of galaxies would in any case require a different training set.

Unlike the other machine learning based deblending methods where one of the galaxies has to be located at the center of the image [8, 9], there is no constraint on the positions of the galaxies for our deblender, because we train it to output the object with the highest luminosity. This is also the reason why our framework can work iteratively. The use of the pixelwise summation when creating the blended images is also instrumental in making the iterative process possible. During training, the noisy blended images are treated as the input and the noiseless image with the brightest galaxy as the pre-blended ground truth for the deblender. Our goal is to train a deblender that can deblend and denoise at the same time and will always recover the brightest galaxy from the blended images. The classifier is trained with 0 – 3 galaxies and for each of the four classes, half of the training data contains noiseless images while the other half contains noisy images. This is essential in order to train a classifier that can estimate the number of galaxies in both noiseless and noisy images at the same time.

We generate 50,000 2-galaxy blended images as the training set for the deblender using BlendingToolkit. The test set contains 1000 blended images with 2 galaxies in each. There are also several extended test sets for 3 to 7-galaxy blended images for generalization, each containing 1000 images. For the classifier, we generated blended images containing 0, 1, 2 and 3 galaxies, where there are 100,000 images of each class (making it 400,000

in total) - half of these are noiseless images and the other half noisy. The test set for the classifier contains 2000 images of each class and 8000 images in total. The test set for the classifier also has the same ratio of noiseless and noisy images. The classifier gives 4 classes as output indicating the number of galaxies in the input image.

E. Training procedure

The training process can be divided into two phases. In the first phase, the deblender and the classifier are pre-trained separately until they perform reasonably. Then they are fine-tuned jointly to boost the performance of both, while working in the iterative setting. Specifically, we expect that the classifier will force the deblender to output only one galaxy, while the deblender will provide more intermediate images to motivate the classifier to have better discrimination ability.

During the pre-training of the RDN, the noisy 2-galaxy blend images are used as the input and the noiseless images with the brighter single galaxy as the ground truth. Eq. 1 is the objective function during this process. A model trained in this way will have the ability to not only deblend the brighter galaxy, but also to simultaneously denoise. Meanwhile, in the first training phase the classifier is trained on the dataset containing both noisy and noiseless blended images with 0, 1, 2 and 3 galaxies. This pre-training phase makes the trained classifier able to classify both noisy and noiseless images and the trained RDN to determine the end of the whole workflow. They will be used together during the second phase of the training.

During the second phase, the deblender and the classifier will be trained jointly. The framework will operate as it is designed, i.e., the deblender will be run twice because the input contains two galaxies. The blended image, denoted as I_{blend} will be passed through the deblender to get the deblended $I_{\text{deblend-1}}$, which is compared with the ground truth $I_{\text{gt-1}}$ to build the ℓ_1 -norm loss $l_{\text{RDN}}(I_{\text{deblend-1}}, I_{\text{gt-1}})$. The first residual image $I_{\text{res-1}}$ is the subtraction between I_{blend} and $I_{\text{deblend-1}}$. Generally, the residual images are fainter than the blended images and therefore they will be normalized by dividing with the maximum pixel values. Then, the deblender will predict and scale back to recover the second galaxy $I_{\text{deblend-2}}$ from $I_{\text{res-1}}$ and another residual image $I_{\text{res-2}}$ will be calculated. For simplicity, we use $\mathbb{S} = \{I_{\text{blend}}, I_{\text{deblend-1}}, I_{\text{deblend-2}}, I_{\text{res-1}}, I_{\text{res-2}}, I_{\text{gt-1}}, I_{\text{gt-2}}\}$ to denote all the images. In this phase, the classifier will be optimized by all the available images in \mathbb{S} as formulated in Eq. 3. Since \mathbb{S} is class-imbalanced with more 1-galaxy images than 0 and 2-galaxy images, α_I is used as a weight. For 1-galaxy images, $\alpha_I = 0.2$, otherwise $\alpha_I = 1$. For the deblender, in addition to $l_{\text{RDN}}(I_{\text{deblend-1}}, I_{\text{gt-1}})$, it is updated based on a loss from the classifier's prediction on the deblended images. The objective function for the second phase can be

formulated as Eq. 4. λ here is a trade-off coefficient.

$$l_{phase-2}^{\phi} = \sum_{I \in \mathcal{S}} \alpha_I l_{CE}(I) \quad (3)$$

$$l_{phase-2}^{\theta} = l_{\text{RDN}}(I_{\text{deblend-1}}, I_{\text{gt-1}}) + \lambda l_{CE}(I_{\text{deblend}}) \quad (4)$$

III. RESULTS

A. Experimental settings

1. Implementation details

The RDN used in our framework contains $D = 16$ RDBs with $C = 8$ convolutional layers in each block. During the first phase of the training process, the learning rate for the RDN is 10^{-4} , the batch size is 128, and it is trained for 150 epochs with Adam optimizer. For the classifier, the initial learning rate is 0.1, the batch size is 200, and it is trained for 200 epochs. In the second phase, the deblender and the classifier are updated jointly, where the learning rate decays to 10^{-5} for the former and 10^{-6} for the latter. After some experimentation, we chose the trade-off coefficient $\lambda = 10^{-4}$. The batch size is 8 due to GPU limit and the framework is trained for 10 epochs in this phase.

2. Evaluation

We trained the framework on 2-galaxy blended images and tested the trained model on blended images with 2 to 7 galaxies. Testing set for each class contains 1000 noisy blended galaxy images.

We start by showing some example results from 2-galaxy and 3-galaxy blended images in FIG.s 3 and 5. These results show the typical output from a high-quality deblend. The network is using both the morphological as well as the color information to isolate images of individual galaxies which is helpful in recovering images even in difficult cases.

In Table I we present the fraction of deblends sorted by the quality category (see Section II A) for all the testing sets with various numbers of galaxies. Note that this only indicates how well the scheme thinks it is doing, rather than how well it is actually doing. With this caveat, the process shows relatively good results with a large fraction of nominally high-quality deblends.

Somewhat surprisingly, the method continues to work even when the number of galaxies in the blend goes well beyond two as in training. Table I shows a steady decrease in the number of high-quality deblends rather than

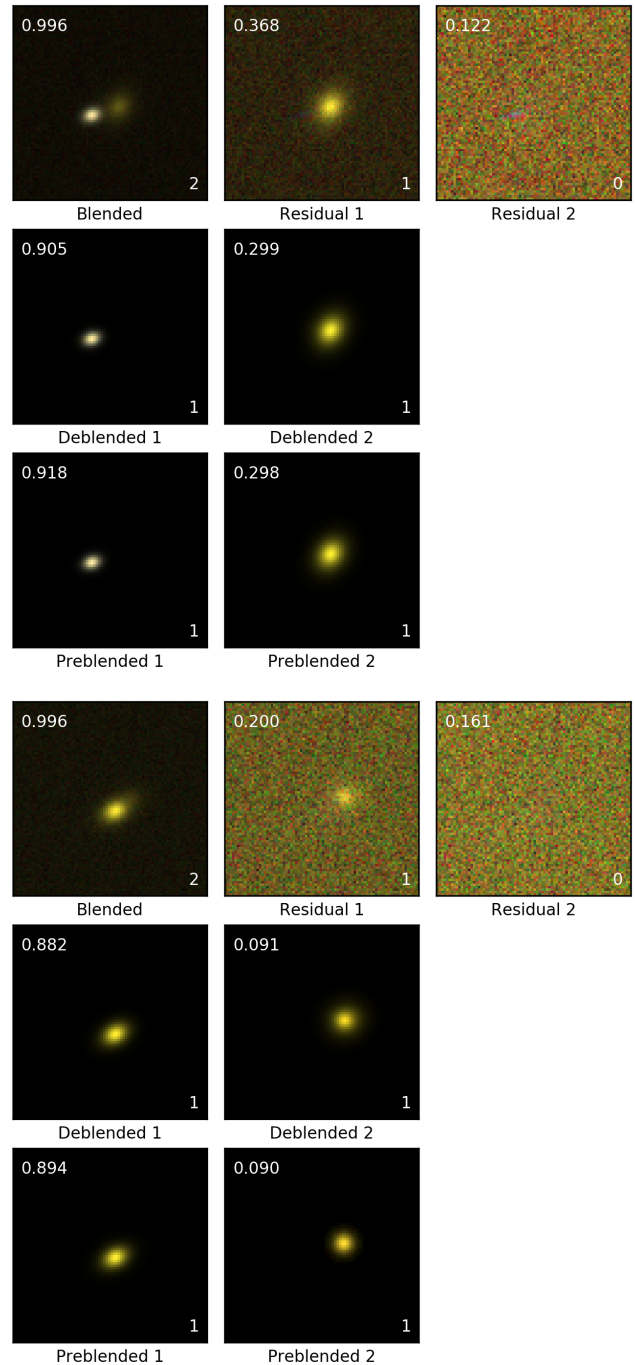


FIG. 3: Figure showing iterative results for 2-galaxy blended images. In each of the two examples shown here, the top row shows the noisy image as it progresses through the iterative process with the input image on the left and the remaining noise image (after the galaxies were subtracted) on the right. The middle row shows the galaxy images that were isolated from the input image by the *deblender*. The bottom row shows the ground truth, which can be visually compared with the deblended images above. All the images were rescaled by its maximum pixel value which is indicated at the top left. The other number at the bottom right corner of each image represents the number of remaining galaxies as determined by the classifier.

	2-galaxy	3-galaxy	4-galaxy	5-galaxy	6-galaxy	7-galaxy
High-quality	77.3%	50.4%	40.6%	27.2%	14.7%	10.5%
Medium-quality	15.3%	36.9%	46.4%	52.8%	62.5%	63.6%
Low-quality	7.4%	12.7%	13.0%	20.0%	22.8%	25.9%

TABLE I: Table showing results of the iterative test on blended images with various number of galaxies.

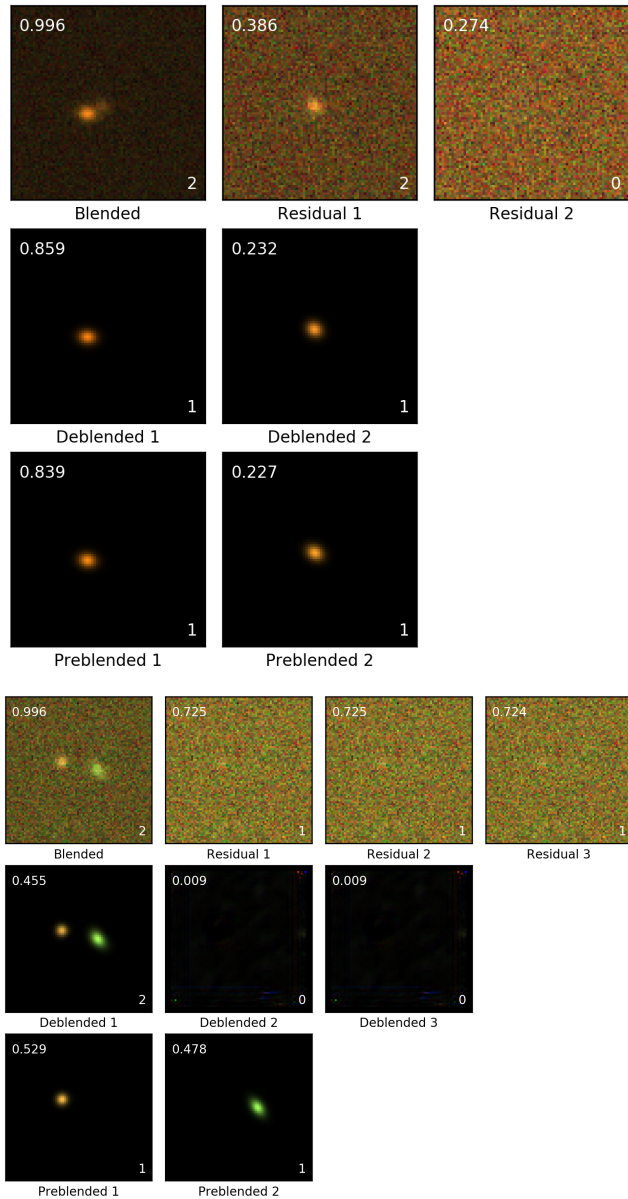


FIG. 4: As FIG. 3, but for medium and low quality samples.

a catastrophic drop in performance. A high-quality deblend for 7-galaxy sample is presented in FIG. 12 in Appendix A. It seems that the network can deal with additional flux around the objects, regardless of whether it is coming from a single object or multiple objects.

To quantify the quality of the deblends, we start with

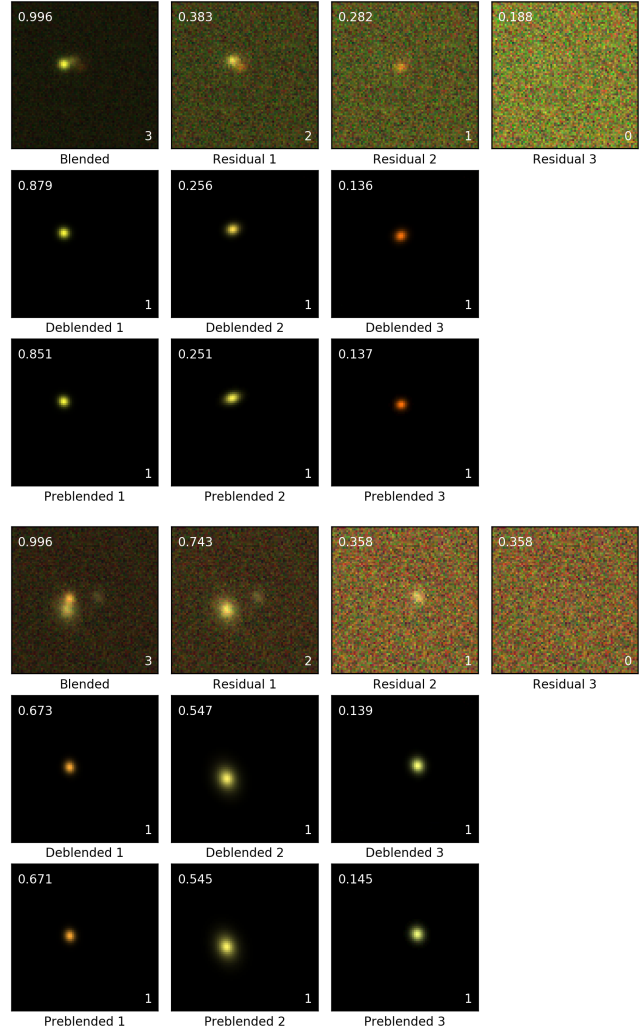


FIG. 5: As FIG. 3, but for 3-galaxy blends.

some standard image analysis metrics. We applied the peak signal-to-noise ratio (PSNR) and the structural similarity index (SSIM) as metrics to evaluate the quality of deblended images from the deblender when compared with the ground truth.

PSNR represents a peak error in the measurement of reconstruction quality in image compression. It computes the logarithm of the ratio between the maximum pixel value (MAX) of the ground truth and the mean squared error (MSE) between the test image and the ground truth in decibels. In our experiment, the ground truth consists of noiseless images with a single galaxy and

	2-galaxy problem				3-galaxy problem			
	PSNR		SSIM		PSNR		SSIM	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Deblended 1	56.51	57.93	0.9973	0.9997	54.66	57.19	0.9954	0.9997
Deblended 2	58.18	59.02	0.9967	0.9997	55.96	58.77	0.9929	0.9996
Deblended 3					57.71	59.29	0.9956	0.9995

TABLE II: Table showing PSNR(dB) and SSIM for 2-galaxy and 3-galaxy deblending problems.

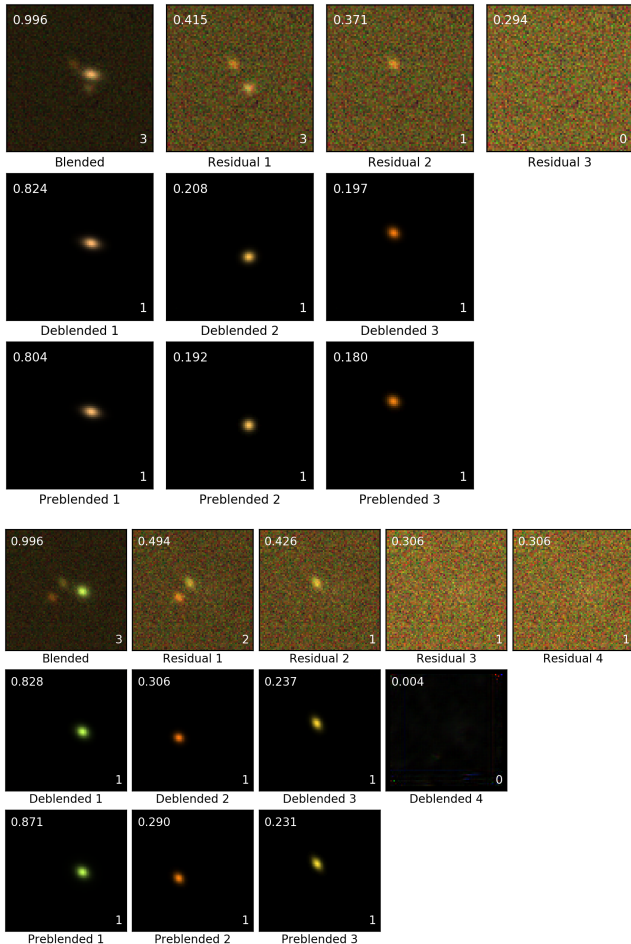


FIG. 6: Figure showing medium and low quality samples for 3-galaxy blends.

the test images are the corresponding deblended images from the RDN. PSNR is formulated as Eq. 5:

$$\text{PSNR(dB)} = 20 \cdot \log_{10}(\text{MAX}) - 10 \cdot \log_{10}(\text{MSE}) \quad (5)$$

SSIM [18], known as the structural similarity index, is commonly used to evaluate the similarity between two images using the means μ_x and μ_y , the variances σ_x and σ_y and the covariance σ_{xy} . In Eq. 6, $c_1 = (k_1 L)^2$ and $c_2 = (k_2 L)^2$ are two small constants to avoid the instability with a weak denominator. We use $k_1 = 0.01$, $k_2 = 0.03$ by default and L is the dynamic range of pixel values.

$$\text{SSIM} = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)} \quad (6)$$

Results are summarized in Table II and Table IV in Appendix A. We do observe a few trends here. First, the PSNR of the second deblended galaxy is higher than that of the first. This somewhat counter-intuitive result comes from the normalization of PSNR, since the second deblend is fainter than the first one. In other words, the PSNR is not telling us that the 2nd deblended galaxy is better deblended, only that its quality is less degraded than expected, given how much fainter it is. This interpretation is confirmed by the SSIM values, which are higher for the first deblend, but only marginally. In both cases we see that the 3-galaxy problem is more difficult than the 2-galaxy problem. We also see that the median values are systematically above the mean values, telling us that the mean is pulled down by a small number of catastrophic outliers. To study these results in language that is more relevant to the astronomy community, we turn to the recovery of the fluxes and moments, described in the next section.

B. Flux and image moment recovery using RDN

An important measure of the effectiveness of any algorithm such as ours is the accurate recovery of physical parameters from the input image. To that end, we have opted to compare the fluxes and the second order image moments (which represent the equivalent ellipse of the image and by extension, the shape and orientation) retrieved by the RDN against that of the truth images. For ease of reference, the brighter galaxies in each field are referred to as x galaxies and the fainter ones as y galaxies.

FIG. 7 shows the RDN recovered fluxes (y -axis) plotted against the true fluxes (X -axis) for the x and y galaxies for the entire test set. The larger panels show the distribution for the entire test set while the smaller panels at the bottom show the distribution separated by quality as detailed in Section 1. The dotted teal line is the $y = x$ expectation.

As is evident from FIG. 7, the RDN is very efficient in recovering the fluxes for the brighter galaxies in our fields, more or less uniformly so for the high, medium and low

sample	quality	α_f	σ_f^2	$ \alpha_e $	$\text{Arg}(\alpha_e)$	σ_e^2
x	high	1.008 ± 0.009	1.1	0.52 ± 0.04	-0.01 ± 0.04	0.034
x	medium	1.029 ± 0.014	6.0	0.56 ± 0.08	-0.08 ± 0.12	0.046
x	low	1.014 ± 0.014	2.2	0.65 ± 0.12	-0.15 ± 0.14	0.049
y	high	1.008 ± 0.010	1.2	0.38 ± 0.04	0.02 ± 0.05	0.024
y	medium	0.959 ± 0.032	2.4	0.50 ± 0.07	0.19 ± 0.10	0.026
y	low	0.52 ± 0.14	6.7	0.18 ± 0.13	1.20 ± 1.09	0.084

TABLE III: Table showing the recovered response factors for flux and galaxy ellipticities split by the deblend type (x or y) and quality flag. Errors on responses were derived by subsampling using bootstrap technique with replacement.

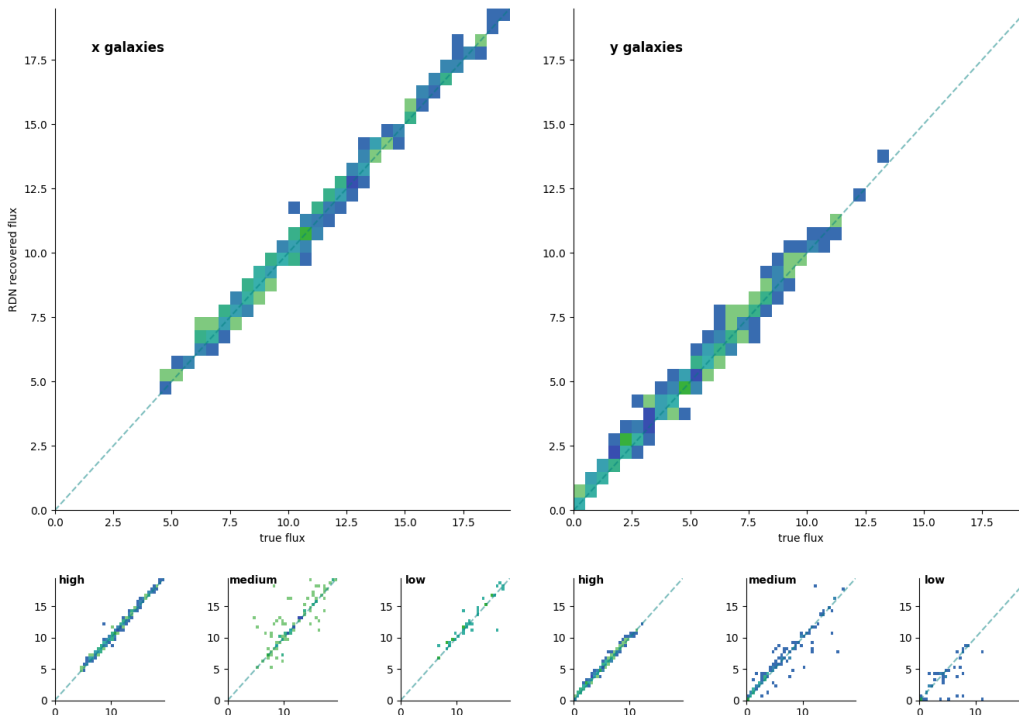


FIG. 7: 2-d histograms showing the comparison of the fluxes recovered from the RDN and the true fluxes for the brighter x galaxies and the fainter y galaxies. The teal line is $y = x$ expectation. The smaller plots show the same relation separated in quality bins of the image data.

quality images. As expected, the lower quality assignments are associated with a higher scatter between real and measured fluxes, although for the brighter galaxy, there is no measurable difference between the medium and low quality marks (given the size of our test catalog). For the fainter galaxies, however, the low quality deblends are essentially uncorrelated with the true values. We find that except for secondary objects in the low-quality bin, our RDN based estimator is essentially unbiased in recovering flux values.

To quantify this we model the recovered flux f_r as

$$f_r = \alpha_f f_t + n_f \quad (7)$$

where α_f is a proportionality constant and n_f is flux noise which has zero mean and variance σ_f^2 . The maxi-

mum likelihood estimate for α_f and σ_f^2 are given by

$$\alpha_f = \frac{\sum f_r f_t}{\sum f_t^2} \quad (8)$$

$$\sigma_n^2 = \frac{1}{N} \sum (\alpha_f f_r - f_t)^2 \quad (9)$$

In Table III we show the values of the recovered parameters. We see that in the case of flux, α_f is consistent with unity for both x and y galaxies indicating unbiased flux measurements, except for the low quality y regime, where deblending suffers from bad systematic errors. We also observe from the table that the variance/noise σ_f^2 increases as the quality of the blends decreases, except in the case of medium quality x galaxies, which is likely because of some outliers (as evidenced by the scatter for the respective quality type in FIG. 7).

Similar to the fluxes, the shapes and orientations of the galaxies, represented by their second order image moments, are crucial characteristics that we aimed to recover using the RDN, in particular with the weak lensing application in mind. We define second moments as

$$\mu_{ij} = \frac{\sum_{xy} x^i y^j I(x, y)}{\sum_{x, y} I(x, y)} \quad \text{for } i + j = 1$$

$$\mu_{ij} = \frac{\sum_{xy} (x - \mu_{10})^i (y - \mu_{01})^j I(x, y)}{\sum_{xy} I(x, y)} \quad \text{for } i + j = 2$$

where the sum goes over image pixels, I is the pixel intensity and \bar{x} and \bar{y} are the mean values. The first moments measure the central position of the stars, while the second moments measures the shape. Commonly used shape parameters are ellipticities, given by

$$e_1 = \frac{\mu_{20} - \mu_{02}}{\mu_{20} + \mu_{02}} \quad (10)$$

$$e_2 = \frac{2\mu_{11}}{\mu_{20} + \mu_{02}} \quad (11)$$

The ellipticities so defined are dimensionless and are invariant both under scaling in flux and size. We are calculating these quantities for true and neural network recovered images, which are noiseless (i.e. flux is zero outside object's extent) and so there is no need for more advanced techniques such as adaptive moments.

In order to plot a two dimensional quantity such as ellipticity against the truth, we plot the value of $e_p e_t^*$ on the complex plane in the Figure 8. This in turn "rotates" ellipticity by the inverse of its true argument, and in the ideal case, the results should scatter around the positive real axis. Additionally, dividing by e_t blows up the value for galaxies with small ellipticities making the plot unwieldy. We see that with the exception of a few catastrophic outliers, the galaxies indeed do scatter around the real axis. The y galaxies are noisier, but neither is biased.

In order to quantify the shape recovery we fit a similar model as in the case of flux:

$$f_e = \alpha_e e_t + n_e \quad (12)$$

Note that since e_t is now a complex quantity, α is also a complex quantity. Results of this exercise are shown in the Table III. For an ideal recovery, $\alpha_e = 1$ and real. Since the noise is more likely to scatter object into a rounder shape, $\alpha_e < 1$. We find $\alpha_e \sim 0.55$ for the x galaxies seemingly independent of the quality flag. Average α_e for y galaxies is even lower. The argument of α_e is consistent with zero except for the low quality y galaxies. This indicates that the recovery is biased in amplitude, as is expected for quantities that are non-linear in image fluxes, but not in orientation. Only in the case of low quality y galaxies, the results indicate that the output shapes are uncorrelated with the input shapes.

We also find that the value of noise, σ_e^2 increases with the quality flag as expected.

In an earlier version of this work we have found cases where the recovered galaxies would be drastically different than input cases: the neural network would sometimes turn round objects into highly elliptical ones or vice-versa. We have manually looked at several of these cases in the Appendix B without identifying a clear regularity.

C. Comparison with SExtractor

In this section, we compare our deblending strategy with what is widely considered as the industry standard, Source Extractor (SExtractor, [10]), which has been the baseline detection, deblending and image extraction software in astronomy for over two decades. It returns a set of user specified parameters from a more extensive default parameter file, by following the specific configurations defined by the user. In this work, we have used the Python formulation of SExtractor, *sep* [19] with the settings employed for DES[20].

We note that this is a fundamentally unfair comparison for several reasons. First, the RDN has been trained for this particular set of galaxies and thus internally employs a correctly tuned prior for the distribution of morphologies, fluxes and ellipticities, while SExtractor employs a general algorithm and is thus intrinsically more robust. While the RDN's approach to deblending is peeling off the individual objects, SExtractor detects objects by creating segmentation maps. Moreover, following DES configuration, SExtractor is run on the co-added image (R+G+B) and is thus missing the color information. It is also important to note that although SExtractor outputs a number of photometric features, it is primarily used to segment image while a detailed analysis is left to other codes. Nevertheless, it is an appropriate sanity check to measure what kind of improvements can be brought about by employing more sophisticated methods.

FIG. 9 shows the comparison of the flux recovery by the RDN and SExtractor for both the x and y galaxies. The top left and right panels feature the RDN recovered fluxes on the y -axis and the true fluxes on the X -axis for the x and the y galaxies respectively. The bottom left and right panels are the same representations for the SExtractor fluxes. The teal line is a $y = x$ fit. As is evident from the plots, the RDN does a better job of recovering the object fluxes for both the x and y galaxies. SExtractor has a tendency to put a disproportionate amount of flux into the first detected objects, and hence the x images are biased high, while the y images are biased low.

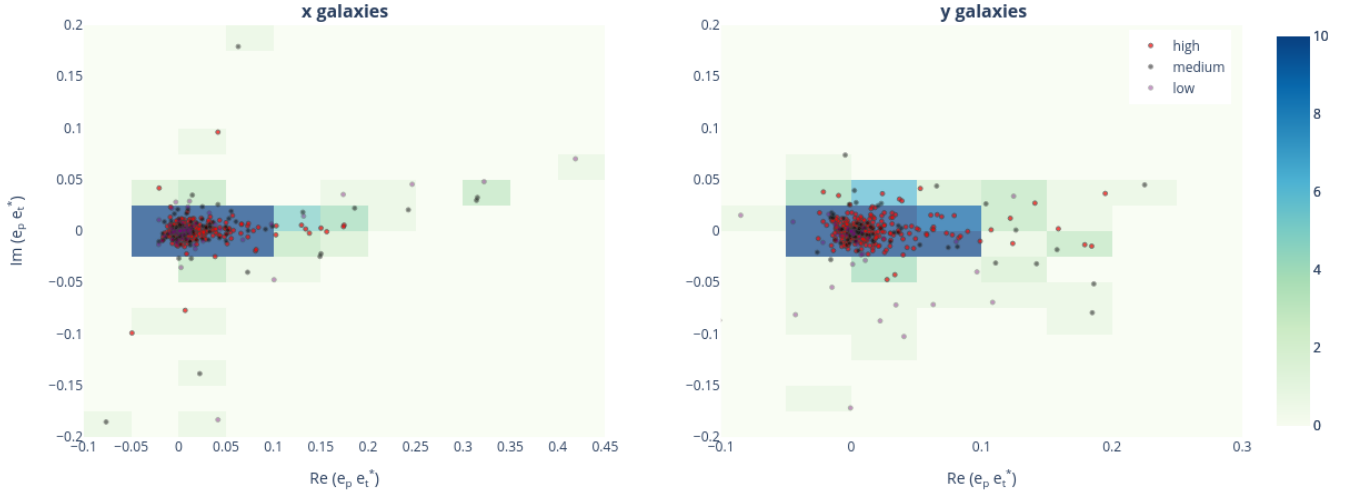


FIG. 8: 2-d histograms showing the quantity $e_p e_t^*$ plotted on the complex plane. For a noiseless recovery, the points would scatter on the positive x -axis line weighted by the square of true ellipticity. The left plot is for the x galaxies and the right plot is for the y galaxies. Different quality flags are represented by different colored points as per the legend. The scatter around this line represents the level of the noise in moment recovery. See table III for a quantification of the deblending response.

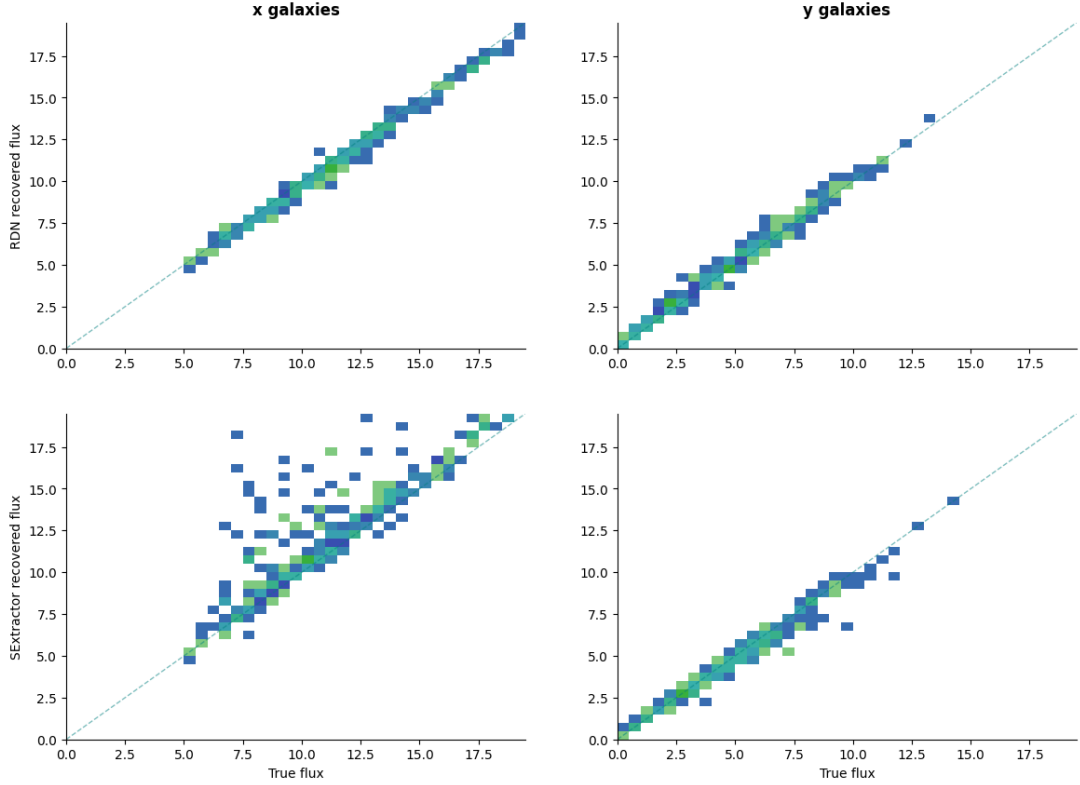


FIG. 9: 2-d histograms showing the flux recovery from RDN and SExtractor as compared to the true fluxes. The top left and right panels show the true fluxes on the X -axis and the RDN fluxes on the y -axis respectively for x & y galaxies. The teal line is the $y = x$ expectation. The bottom left and right panels are the same representations for SExtractor.

D. Failure modes

In this section we investigate the most common failure modes in the two approaches by cherry-picking cases in which one method performs correctly while the other fails.

We manually looked at all the cases where the RDN detects just a single galaxy. In a majority of those, the secondary object was very faint and overlapped almost perfectly with the primary, leading to the deblending being suspended because the *classifier* determined that there are no further objects. However, for 12 out of the 1000 objects in the test set, for the deblended x galaxies, the result clearly shows two objects, examples of which are shown in FIG. 10. These problems are trivial for the SExtractor approach. A visual inspection of all these images was conducted and it was observed that, both the objects in 7 of these fields are very similar in shape and, in 5 fields, they had similar brightnesses. It stands to reason that the RDN has some difficulty in deblending objects with similar structural parameters, most likely because it cannot decide which the “brighter” object is, especially in the presence of noise. In fact, when presented with such examples in training, it simply pays a large loss penalty whenever it starts with the “wrong” object, without a clear rule about how to pick one of the two objects if the noise hides the identity of the bright ones. We will discuss this further in Section IV.

When SExtractor is run on the test set, it returns photometric parameters for 979 x galaxies and 729 y galaxies (the x counterparts of all y galaxies are present in the x set). As seen in FIG. 10, this includes 9 instances where the RDN fails to deblend the different galaxies in the field. However, overall, there are more instances where the RDN successfully deblends the x and y galaxies, while SExtractor either doesn’t detect either galaxy, or only detects the x galaxy, but not the y galaxy, a few examples of which are illustrated by FIG. 11. We note that these are traditional blend merges, where the deblender merges two distinct objects into a single one. Since human eye is pretty good at detecting these, we could possibly improve upon this by fine-tuning the SExtractor settings, potentially at the cost of artificially shredding objects with substructure. These are also examples where color information is most helpful.

IV. SUMMARY AND CONCLUSIONS

In this work we present a new approach to deblending using a Residual Dense Network, which was trained on blended galaxy images with a realistic PSF and noise levels, and which has performed decently in deblending and recovering object fluxes and shapes. Compared to previous works, our set up does not assume that the object to be deblended is located at the center of the image. We also do not need to assume in advance, the number of objects to be deblended. By using two neural networks,

one to perform deblending and one to determine the number of remaining objects in the field, we can classify the quality of deblends. We have shown that deblends that are designed to be higher quality have indeed, less noisy fluxes and shapes.

In the most current deblending networks found in literature, the network is trained to deblend the central galaxy. Given that centering in the presence of noise presents its own problems, we have trained our network to deblend the brightest galaxy remaining in the image. This works very well, where the choice of the brightest is clear, but leads to a particular failure mode where the network cannot decide where to start, when two objects are approximately equally bright (see FIG. 10). This might be a generic feature of any kind of deblender that proceeds by peeling off one galaxy at a time. For example, even if we choose the galaxy closer to the center, there would likely be an equivalent failure mode when two galaxies are equally close to the center. The same thing could happen for other initial assumptions that attempt to distinguish between the galaxies. There are several approaches to solve this. One possibility would be to use a symmetrized loss function that would train the network to deblend either galaxy. Another possibility would be to have a more sophisticated network that deblends multiple galaxies at once, perhaps again with an appropriately symmetrized loss function.

Despite the RDN deblender being trained on two-galaxy blends, it performs surprisingly stably all the way to the maximum 7-galaxy blends, where it still correctly deblends 7 galaxies in 10% of cases.

We have found that our method outperforms the industry standard SExtractor on flux recovery, subject to the strong caveats mentioned in Section III C. However, we found that the RDN recovered shapes are significantly worse and occasionally catastrophically wrong. The very good results could partially be the result of the limited morphological variety of the presented images; however, since we do not attempt to deconvolve the image, this might not be a major issue. Our work reiterates concerns with using a deeply non-linear deblender based on NN for work requiring precision shapes, such as weak lensing. At the same time, such work could lead to potentially interesting new work on galaxy-star separation and morphological classification of galaxies.

Similar to the majority of other neural network approaches, our deblender cannot deal optimally with a PSF that is variable and is different for the different input channels (for e.g. for different bands), nor can it deal with a pixel mask. This is a significant deficiency that is usually swept under the rug under the assumption that the network can be retrained for a set of different PSFs. However, in practice, we have a different PSF for different bands and with the typical astronomical observations in five or six bands, the number of possible PSF combinations becomes unmanageable. The PSF shape in each band thus needs to be a part of the input and network training. We are leaving these important aspects

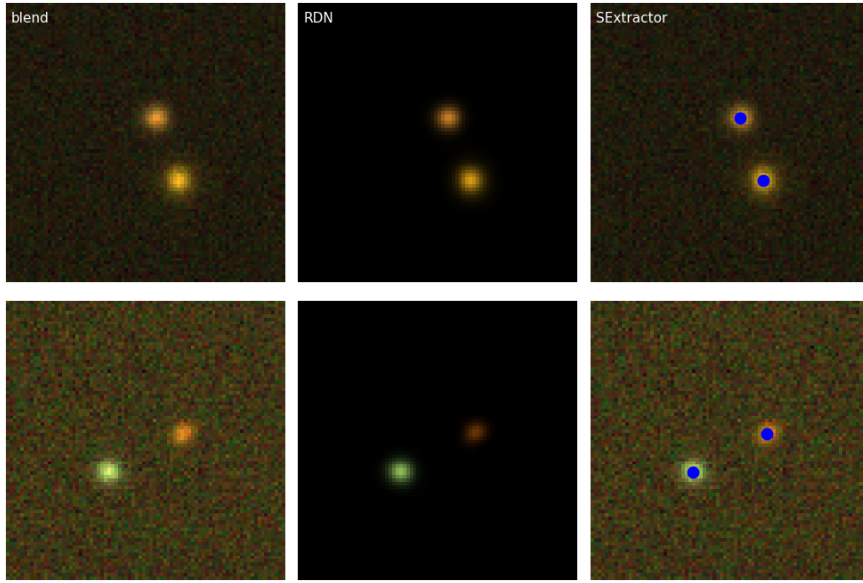


FIG. 10: Figure showing instances where the RDN fails, while SExtractor succeeds in detecting and deblending the objects in the field. On the left is the original blend, in the centre is the output of the RDN and on the right is the blend with the SExtractor detections superimposed in blue.

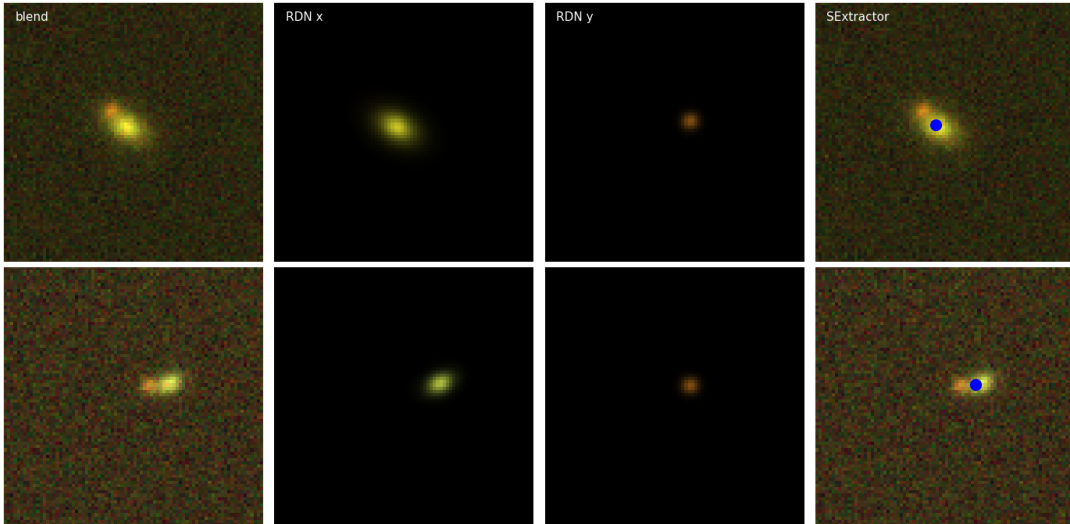


FIG. 11: Figure showing instances where SExtractor fails to detect and deblend, while the RDN succeeds. On the far left is the original blend, the RDN result is represented in the two central panels for the x and y galaxies, and the blend with the SExtractor detection superimposed in blue is on the far right panel.

for future work.

One of the main advantages of most ML approaches to deblending is that they also automatically denoise and deconvolve the image and can in principle fill in the missing information in case of pixel masks. For a sufficiently sophisticated neural network, this would trivialize many typical operations such as flux and shear estimation on the resulting truth images. However, it would also require us to have a more advanced way of propagation of both statistical and systematic uncertainties.

This work has demonstrated the feasibility of using

RDN in astronomical image analysis and has highlighted some of the unresolved issues with the current machine learning approaches to the deblending problem in the case of realistic astronomical images.

ACKNOWLEDGEMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research through the

SciDAC grant “Accelerating HEP Science: Inference and Machine Learning at Extreme Scales” under DOE Con-

tract Number DE-SC-0012704. We thank our collaborators for useful input.

-
- [1] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, and et al., LSST Science Book, Version 2.0, arXiv e-prints , arXiv:0912.0201 (2009), arXiv:0912.0201 [astro-ph.IM].
- [2] T. M. C. Abbott, F. B. Abdalla, A. Alarcon, J. Aleksic, S. Allam, S. Allen, A. Amara, J. Annis, J. Asorey, S. Avila, and et al., Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing, *Phys. Rev. D* **98**, 043526 (2018), arXiv:1708.01530 [astro-ph.CO].
- [3] C. Hikage, M. Oguri, T. Hamana, S. More, R. Mandelbaum, M. Takada, F. Köhlinger, H. Miyatake, A. J. Nishizawa, H. Aihara, R. Armstrong, J. Bosch, J. Coupon, A. Ducout, P. Ho, B.-C. Hsieh, Y. Komiyama, F. Lanusse, A. Leauthaud, R. H. Lupton, E. Medezinski, S. Mineo, S. Miyama, S. Miyazaki, R. Murata, H. Murayama, M. Shirasaki, C. Sifón, M. Simet, J. Speagle, D. N. Spergel, M. A. Strauss, N. Sugiyama, M. Tanaka, Y. Utsumi, S.-Y. Wang, and Y. Yamada, Cosmology from cosmic shear power spectra with Subaru Hyper Suprime-Cam first-year data, *PASJ* **71**, 43 (2019), arXiv:1809.09148 [astro-ph.CO].
- [4] A. Refregier, A. Amara, T. D. Kitching, A. Rassat, R. Scaramella, and J. Weller, Euclid Imaging Consortium Science Book, arXiv e-prints , arXiv:1001.0061 (2010), arXiv:1001.0061 [astro-ph.IM].
- [5] 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, *MNRAS* **500**, 531 (2021), arXiv:2005.12039 [astro-ph.IM].
- [6] P. Melchior, F. Moolekamp, M. Jerdee, R. Armstrong, A. L. Sun, J. Bosch, and R. Lupton, SCARLET: Source separation in multi-band images by Constrained Matrix Factorization, *Astronomy and Computing* **24**, 129 (2018), arXiv:1802.10157 [astro-ph.IM].
- [7] F. Lanusse, P. Melchior, and F. Moolekamp, Hybrid Physical-Deep Learning Model for Astronomical Inverse Problems, arXiv e-prints , arXiv:1912.03980 (2019), arXiv:1912.03980 [astro-ph.IM].
- [8] D. M. Reiman and B. E. Göhre, Deblending galaxy superpositions with branched generative adversarial networks, *Monthly Notices of the Royal Astronomical Society* **485** (2019).
- [9] 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**, 2481 (2020).
- [10] E. Bertin and S. Arnouts, SExtractor: Software for source extraction., *A&AS* **117**, 393 (1996).
- [11] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu, Residual dense network for image super-resolution, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2018) pp. 2472–2481.
- [12] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu, Residual dense network for image restoration, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020).
- [13] K. Simonyan and A. Zisserman, Very deep convolutional networks for large-scale image recognition, arXiv preprint arXiv:1409.1556 (2014).
- [14] <https://github.com/LSSTDESC/BlendingToolKit>.
- [15] B. T. P. Rowe, M. Jarvis, R. Mandelbaum, G. M. Bernstein, J. Bosch, M. Simet, J. E. Meyers, T. Kacprzak, R. Nakajima, J. Zuntz, H. Miyatake, J. P. Dietrich, R. Armstrong, P. Melchior, and M. S. S. Gill, GALSIM: The modular galaxy image simulation toolkit, *Astronomy and Computing* **10**, 121 (2015), arXiv:1407.7676 [astro-ph.IM].
- [16] <https://github.com/LSSTDESC/WeakLensingDeblending>.
- [17] G. De Lucia, V. Springel, S. D. M. White, D. Croton, and G. Kauffmann, The formation history of elliptical galaxies, *MNRAS* **366**, 499 (2006), arXiv:astro-ph/0509725 [astro-ph].
- [18] 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**, 600 (2004).
- [19] K. Barbary, K. Boone, and C. Deil, Sep: V0.3.0 (2015).
- [20] <https://github.com/esheldon/sxdes>.

Appendix A: Iterative results

FIG. 12 visualizes a typical high-quality deblend for the 7-galaxy blended images. It shows good recovery of the morphological and color information for all 7 galaxies. The rest of the galaxies are much fainter compared with the brightest one for the original blended image, but after several iterations, all the residual galaxies emerge for the network to deblend. As the classifier was trained for up to 3-galaxy blended images, the predictions for images with more than 3 galaxies are skipped.

PSNR and SSIM results for 4 to 7-galaxy blended images are in Table IV. There are some similar trends as in Table II, where the fainter galaxies have higher PSNR and SSIM. This could be interpreted by the normalization of PSNR as in Section III A 2.

Appendix B: RDN recovery

FIG. 13 shows a few examples of where the RDN recovered images differ drastically from the truth images. In the first example from the bottom, the true x image is significantly rounder than the RDN recovered x image, while the opposite is true for the respective y images. In the second example, the objects are not deblended at

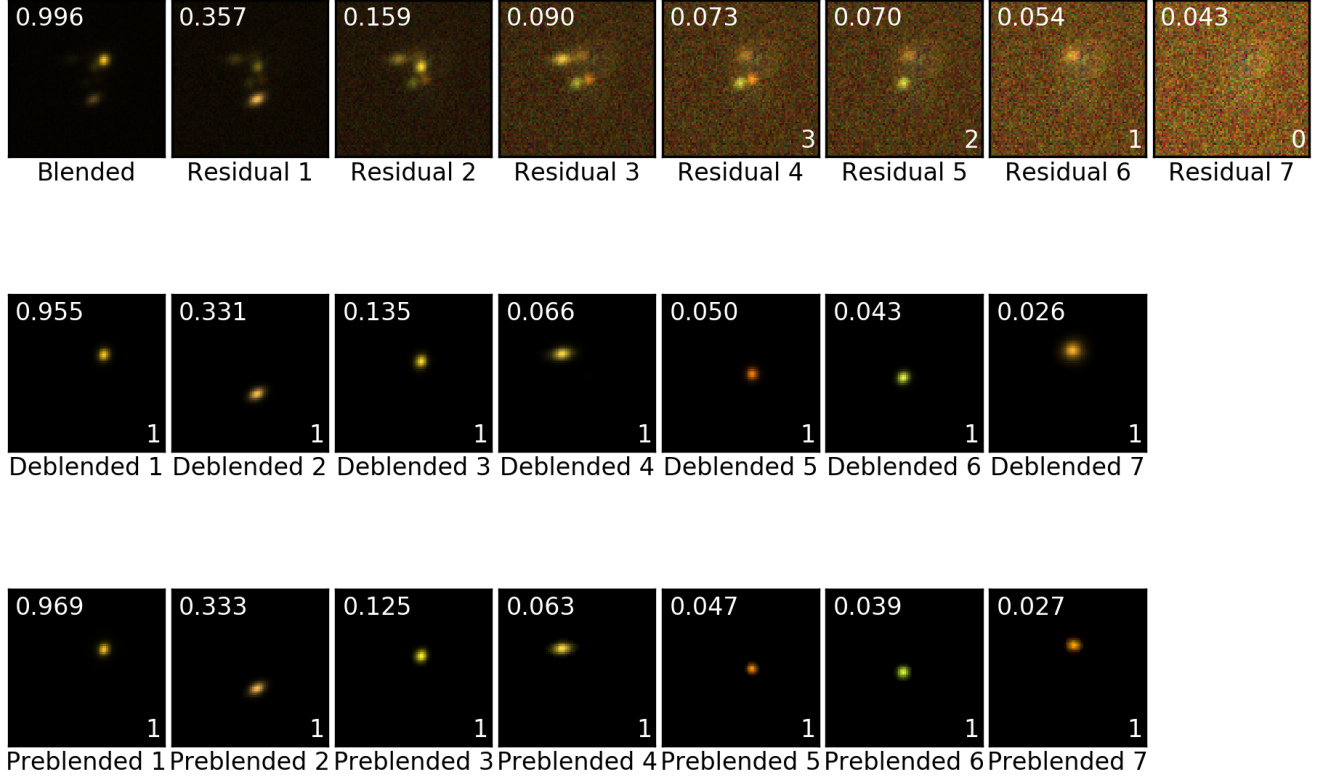


FIG. 12: Figure showing iterative results for 7-galaxy blended images. The classifier was trained for 0 to 3-galaxy blended images, therefore the prediction for some of the residual images were skipped.

	4-galaxy problem		5-galaxy problem		6-galaxy problem		7-galaxy problem									
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM								
	Mean	Median	Mean	Median	Mean	Median	Mean	Median								
1	53.74	56.52	0.9946	0.9996	51.91	54.94	0.9927	0.9994	50.57	52.90	0.9921	0.9991	49.61	51.72	0.9907	0.9987
2	54.17	57.82	0.9906	0.9996	52.21	55.14	0.9879	0.9991	50.40	53.09	0.9856	0.9983	49.60	51.00	0.9846	0.9981
3	54.50	57.80	0.9902	0.9994	52.22	54.35	0.9867	0.9986	50.35	50.35	0.9843	0.9962	49.68	49.25	0.9837	0.9947
4	56.70	59.48	0.9948	0.9994	53.15	54.12	0.9880	0.9977	50.67	49.96	0.9849	0.9926	50.02	48.19	0.9847	0.9910
5					56.89	58.16	0.9934	0.9990	52.25	51.77	0.9876	0.9940	50.52	47.72	0.9855	0.9895
6									56.57	56.90	0.9924	0.9980	52.32	50.51	0.9890	0.9923
7													56.28	55.14	0.9925	0.9967

TABLE IV: Table showing PSNR(dB) and SSIM for 4 to 7-galaxy deblending problems.

all, as seen in the RDN x image which is identical to the blended image and the extremely faint object in the RDN y image. A possible reason for this is the similarity in the shapes (and to some extent the orientation and color) of both the galaxies in the blended image. In the third case, which is a particularly intertwined blend, the RDN recovered x image is slightly brighter and larger than the truth image, while the y image is the opposite.

The fourth case is similar to the second in the sense that both the RDN x and the blended image are identical, but in this case, the RDN y image is blank - however, the two objects in the field do not share a similar shape, orientation or color, and is therefore more mysterious. The fifth example also features a particularly closely associated blend like the third, but in this case, the RDN recovered images for both galaxies is rounder than the truth image.

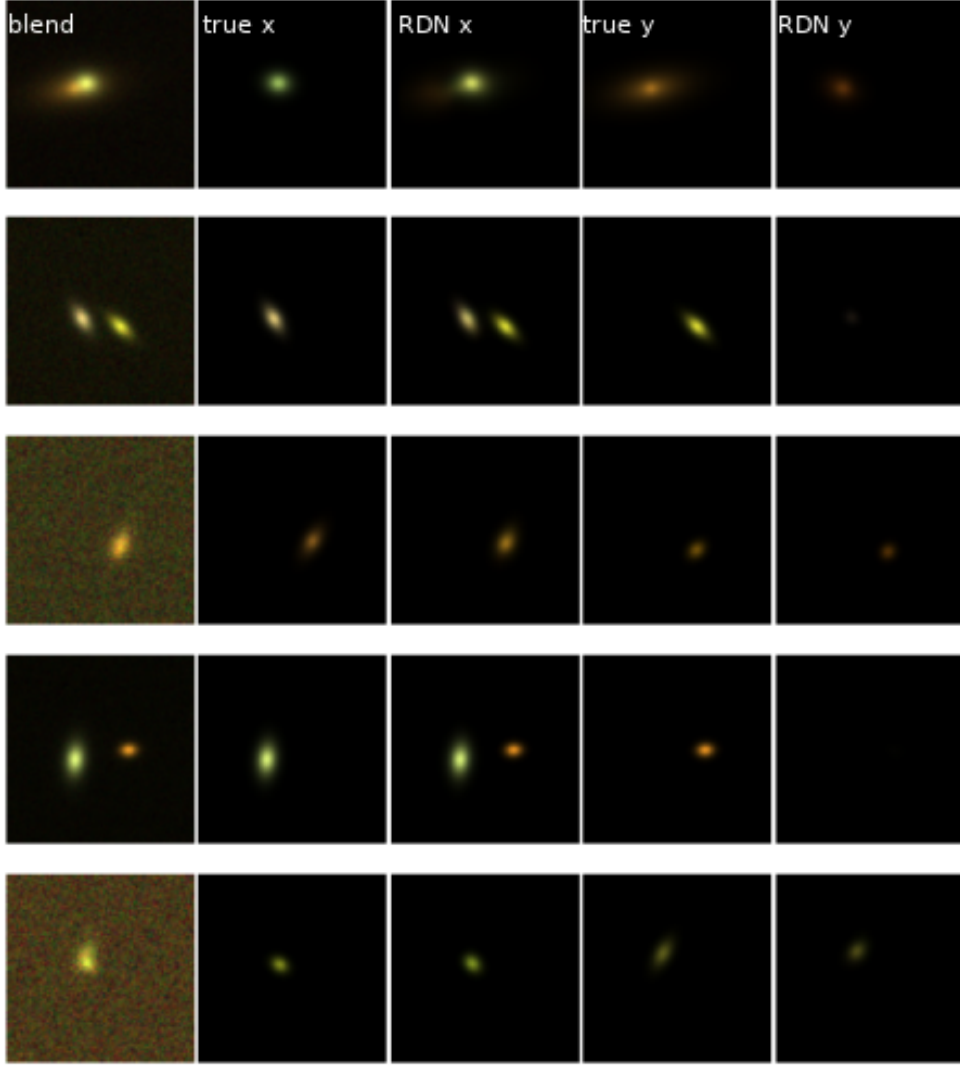


FIG. 13: Figure showing some specific failures of the RDN. Row-wise left to right are the blended image, the truth image for the x galaxy, the RDN recovered image for the x galaxy, the truth image for the y galaxy, and the RDN recovered image for the y galaxy. A detailed explanation is given in Appendix B.