



**HAL**  
open science

## Photometry of high-redshift blended galaxies using deep learning

Alexandre Boucaud, Marc Huertas-Company, Caroline Heneka, Emille Ishida, Nima Sedaghat, Rafael De souza, Ben Moews, Hervé Dole, Marco Castellano, Emiliano Merlin, et al.

► **To cite this version:**

Alexandre Boucaud, Marc Huertas-Company, Caroline Heneka, Emille Ishida, Nima Sedaghat, et al.. Photometry of high-redshift blended galaxies using deep learning. Monthly Notices of the Royal Astronomical Society, 2019, 491 (2), pp.2481-2495. 10.1093/mnras/stz3056 . hal-02447731

**HAL Id: hal-02447731**

**<https://hal.science/hal-02447731>**

Submitted on 26 Jun 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Photometry of high-redshift blended galaxies using deep learning

Alexandre Boucaud<sup>1</sup>,<sup>1</sup>★ Marc Huertas-Company,<sup>2,3</sup> Caroline Heneka,<sup>4</sup>  
Emille E. O. Ishida,<sup>5</sup> Nima Sedaghat<sup>6</sup>,<sup>6</sup> Rafael S. de Souza,<sup>7</sup> Ben Moews<sup>8</sup>,<sup>8</sup>  
Hervé Dole,<sup>9</sup> Marco Castellano,<sup>10</sup> Emiliano Merlin,<sup>10</sup> Valerio Roscani,<sup>10</sup>  
Andrea Tramacere,<sup>11</sup> Madhura Killedar<sup>12</sup> and Arlindo M. M. Trindade<sup>13</sup>

## Collaboration for the COIN

<sup>1</sup>APC, Astroparticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/Irfu, Observatoire de Paris, Sorbonne Paris Cité, 10 rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France

<sup>2</sup>Instituto de Astrofísica de Canarias (IAC), Departamento de Astrofísica, Universidad de La Laguna (ULL), E-38200, La Laguna, Spain

<sup>3</sup>LERMA, Observatoire de Paris, PSL Research University, CNRS, Sorbonne Universités, Université Paris Diderot, F-75014 Paris, France

<sup>4</sup>Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56126 Pisa, Italy

<sup>5</sup>Université Clermont Auvergne, CNRS/IN2P3, LPC, F-63000 Clermont-Ferrand, France

<sup>6</sup>Department of Computer Science, University of Freiburg, Georges-Koehler-Allee 052, D-79110 Freiburg, Germany

<sup>7</sup>Department of Physics and Astronomy, University of North Carolina at Chapel Hill, NC 27599-3255, USA

<sup>8</sup>Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ, UK

<sup>9</sup>Institut d'Astrophysique Spatiale, CNRS, Univ. Paris-Sud, Université Paris-Saclay, F-91400 Orsay, France

<sup>10</sup>INAF – Osservatorio Astronomico di Roma, Via Frascati 33, I-00078, Monte Porzio Catone, Italy

<sup>11</sup>Department of Astronomy, University of Geneva, 24 rue du Général-Dufour, CH-1211 Genève, Switzerland

<sup>12</sup>Sydney Informatics Hub, The University of Sydney, Sydney, NSW 2008, Australia

<sup>13</sup>Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto, CAUP, Rua das Estrelas, P-4150-762 Porto, Portugal

Accepted 2019 October 10. Received 2019 September 18; in original form 2019 April 30

## ABSTRACT

The new generation of deep photometric surveys requires unprecedentedly precise shape and photometry measurements of billions of galaxies to achieve their main science goals. At such depths, one major limiting factor is the blending of galaxies due to line-of-sight projection, with an expected fraction of blended galaxies of up to 50 per cent. This proof-of-concept work explores for the first time the use of deep neural networks to estimate the photometry of blended pairs of galaxies in space-based monochrome images similar to the ones that will be delivered by the *Euclid* space telescope under simplified idealized conditions. Using a clean sample of isolated galaxies from the CANDELS survey, we artificially blend them and train two different network models to recover the photometry of the two galaxies. We show that our approach can recover the original photometry of the galaxies before being blended with  $\sim 7$  per cent mean absolute percentage error on flux estimations without any human intervention and without any assumption on the galaxy shape. This represents an improvement of at least a factor of 4 compared to the classical SExtractor approach. We also show that, forcing the network to simultaneously estimate fractional segmentation maps results in a slightly improved photometry. All data products and codes have been made public to ease the comparison with other approaches on a common data set. See <https://github.com/aboucaud/coinblend>.

**Key words:** methods: data analysis – methods: statistical – techniques: image processing – galaxies: general – galaxies: photometry.

## 1 INTRODUCTION

The upcoming years will be marked by the arrival of a new generation of deep and wide galaxy surveys from ground [e.g. Large Synoptic Survey Telescope (LSST), Ivezić et al. 2008], and space

(e.g. *Euclid*, Racca et al. 2016). Under this new paradigm of *big-data* surveys, the community aims to achieve an unprecedented level of accuracy and precision both in terms of photometry (e.g. photometric redshifts, Krone-Martins, Ishida & de Souza 2014; Elliott et al. 2015; Beck et al. 2017; Salvato, Ilbert & Hoyle 2018), and shear measurements (e.g. Kilbinger et al. 2017; Kitching et al. 2017) for an unprecedented number of objects. This requires to revisit most

\* E-mail: alexandre.boucaud@apc.in2p3.fr

of the commonly used procedures to extract measurements from images, in order to reduce as far as possible all the systematic effects and reach the requirements. One particular important source of error is the blending of sources. As surveys become deeper and deeper, we expect an increasing fraction of overlapping galaxies which could bias the measurements at levels beyond requirements (Dawson et al. 2016). For example, the estimates for LSST say that  $\sim 45$  to 66 percent of the sources are expected to overlap to a degree of being problematic for a number of methods, with  $\sim 75$  percent of blends probably composed of only two objects (Dawson & Schneider 2014; Dawson et al. 2016). If galaxies are not properly separated, their photometry is biased, which has a direct impact on the derived redshift (and all other physical properties). Efficient algorithms to automatically separate (deblend) detected sources are crucial and will be a key ingredient in the processing pipelines of the next generation surveys. However, there is currently no standard solution in the literature and deblending remains an open issue among the community.

The most widely used software for detecting and separating objects in large fields is *SEXTRACTOR*<sup>1</sup> (Bertin & Arnouts 1996) but its use is far from being an optimal solution. In a nutshell, this software looks for saddle points in the luminosity profiles of galaxies using multiple thresholds. The main problem with such an approach is that it is very sensitive to the configuration parameters and it is difficult to configure so that it works in a wide variety of cases. The fraction of blended sources which are not identified as such can reach significant fractions (Laidler et al. 2006). An alternative way is to simultaneously fit a parametric model to all galaxies in the image and use the best-fitting models to estimate the photometry (Pignatelli, Fasano & Cassata 2006; Mancone et al. 2013; Safarzadeh et al. 2015). This approach typically reaches better photometric accuracy but still requires to properly identify the centroids of all the different objects. It also assumes simplistic models for the galaxy surface brightness distribution which do not encapsulate all the diversity of galaxy morphologies, especially in the more distant Universe. It is also expensive in terms of computing time. The classical deblending approaches are therefore insufficient to reach the level of requirements on measurements of galaxy properties for upcoming surveys. It is thus timely to investigate and compare different approaches.

Several groups are working on alternative solutions more adapted to large volumes of data (Joseph, Courbin & Starck 2016; Tramacere et al. 2016; Ivezić, Connolly & Jurić 2017). For example, recent works by the LSST collaboration (Melchior et al. 2018) have started to develop more global approaches based on non-negative matrix factorization, that can achieve a more efficient source separation and enable to put flexible constraints or priors on the shape of the signals. This approach is however optimized for ground-based data in which galaxies have little resolved structures and also takes full advantage of the multiwavelength nature of LSST data. It is less well adapted for monochrome space data such as the images delivered by the *Euclid* space telescope (Jones & Heavens 2019).

The goal of this paper is to explore if machine learning and more precisely deep learning is an approach worth investigating for segmenting blended galaxies and estimating their photometry. During recent years, the use of deep learning approaches for tasks related to galaxy images has become a burgeoning field of research in astronomy. One of the earliest and most pervasive area of application is the classification of galaxy morphologies

(e.g. Dieleman, Willett & Dambre 2015; Barchi et al. 2017; Domínguez Sánchez et al. 2018; Huertas-Company et al. 2018; Khalifa et al. 2018). More recent research includes the recovery of galaxy features in noisy images by Schawinski et al. (2017), the finding of galaxy–galaxy strong lensing effects by Lanusse et al. (2018), and the generation of physically realistic synthetic galaxy images to augment existing data sets and consequently provide the aforementioned deep learning approaches with larger training sets (Ravanbakhsh et al. 2017; Fussell & Moews 2019).

In a recent work, Reiman & Göhre (2019) used deep learning for the first time to deblend Sloan Digital Sky Survey (SDSS) galaxies. They introduce a modified Generative Adversarial Network (GAN, Goodfellow et al. 2014) to separate blended galaxies, combining aspects of the super-resolution GAN (SRGAN) by Ledig et al. (2017) and the deep residual learning framework by He et al. (2016). With the generator as a modified residual network that features two branches, each branch generates one of the two blended galaxies. They show promising results. However, their procedure to generate blended images for training based on the pixelwise maximum of the two individual stamps does not reflect the true process resulting in line-of-sight blending which sums the photons coming from both sources.

In this paper, we further explore the use of machine learning to both segment and measure the photometry of blended pairs of galaxies. The approach presented here is designed having *Euclid* data in mind as the main target of application (i.e. monochrome space-based data). The goal is thus to obtain a neural network optimized to predict the photometry of pairs of galaxies observed with fairly high spatial resolution in one single band.

This paper is organized as follows: in Section 2, we describe the realistic image data set of blended galaxy pairs we created. We then detail the reasons behind the choice of deep learning methods for this paper and carefully unroll the methodology used to set up our networks in Section 3. These methods are applied to our emulated data set in Section 4, where we compare the results with *SEXTRACTOR*, before discussing the pros and cons in the final Section 5.

## 2 DATA SET OF ARTIFICIALLY BLENDED GALAXIES

Galaxy blending is a confusion effect created by the projection of photons from galaxies on a given line of sight, to the 2D plane. As telescopes get more sensitive, we have access to a higher number of galaxies and thus to a higher chance of finding multiple objects in the same line of sight (Dawson & Schneider 2014).

The quantification of the effects of blending on the derived galaxy properties is a difficult task by nature, due to the loss of spatial information from pixelation in the sensors and convolution with the point spread function. Most existing methods require additional knowledge (several wavelength bands), or a priori knowledge, like parametric models, of the galaxy profiles, symmetries, etc. Moreover, to assess the accuracy of such methods, we are often left with bottom-up approaches like the simulation of galaxy blending using software like *GALSIM* (Rowe et al. 2015), for which we have access to the true light distribution of each object in the image. But as realistic as they can be, simulated images often show their limits when compared with the diversity and the singularity of real data images (Haussler et al. 2007). This is particularly critical for machine learning which implicitly assumes that the training sets are fully representative of the real data.

<sup>1</sup><https://www.astromatic.net/software/SEXtractor>

In order to get a realistic representation of observations, for this work we decide to simulate blended objects from real observations. Although this approach eventually propagates the biases and errors existing in the observations, it has the advantage of including fully realistic morphologies. We describe in the next paragraph the methodology we follow to generate our galaxy sample.

## 2.1 Parent sample

The parent sample used in this paper is the  $H$ -band selected catalogue from the Cosmic Assembly Near-infrared Deep Extragalactic Legacy (CANDELS) survey, presented in Dimauro et al. (2018). The catalogue contains galaxies with  $F_{160W} < 23.5$ , for which both visual morphologies and parametric bulge–disc decomposition are performed. From this parent data set, we first define a clean sub-sample of isolated galaxies with unambiguous morphologies that are then used to perform the blends. More precisely, we use the neural-network-based morphological classification published in Huertas-Company et al. (2015) and select galaxies with four different morphological types:

- (i) pure bulges:  $P_{\text{SPH}} > 0.8$ ,
- (ii) pure discs:  $P_{\text{DISC}} > 0.8$ ,
- (iii) two component bulge+disc:  $P_{\text{SPH}} > 0.8$  and  $P_{\text{DISC}} > 0.8$ ,
- (iv) irregular galaxies:  $P_{\text{IRR}} > 0.8$ .

Note that the purpose of this selection is not to have a complete sample of galaxies, but to have a clean data set of isolated galaxies with different morphologies for which we can reasonably trust the segmentation procedure. By selecting galaxies with very large probabilities of being in a given morphological type we can be reasonably certain that we remove originally blended systems or complex structures such as mergers. We also notice that the sample is built to be representative of the general field population at intermediate redshifts, which is expected to be the dominant population in the Euclid Survey. However, extreme environments such as massive clusters or nearby very resolved galaxies are not considered here.

From this initial sample, we generate  $128 \times 128$  pixel stamps centred on the objects. We then remove all other objects present in the stamps. To that purpose, we apply a morphological dilation to the original segmentation obtained with SExtractor and replace all distinct regions with random pixels sampled from empty regions in the background. This process is illustrated in Fig. 1.

In order to further clean the sample, we visually inspect the selected galaxies and remove the ones which still present anomalies such as originally blended systems not detected by SExtractor, or the ones for which the removal of companions created some visual artefacts in the images. The final sample results in  $\sim 2000$  galaxies whose types are summarized in Table 1. Fig. 1 shows a selection of these galaxy stamps, along with the stamp after the removal of neighbouring objects and the associated SExtractor segmentation map of the central isolated galaxy.

## 2.2 Blending

To create the artificially blended systems, we combine the galaxies of the clean sample we just obtained using the following procedure. First, we randomly select one galaxy, referred to as the *central galaxy*, with a magnitude and an effective radius respectively denoted  $\text{mag}_{\text{cen}}$  and  $R_{\text{cen}}$ .  $R_{\text{cen}}$  is the semimajor axis of the best Sérsic fit model from the catalogue by Dimauro et al. (2018). Second, we pick a second galaxy in the catalogue, referred to as

the *companion galaxy* with properties  $\text{mag}_{\text{comp}}$  and  $R_{\text{comp}}$ , so that it satisfies  $\text{mag}_{\text{cen}} - 2 < \text{mag}_{\text{comp}} < \text{mag}_{\text{cen}} + 2$ . Then, we set  $R = \max(R_{\text{cen}}, R_{\text{comp}})$  as the biggest effective radius between the two galaxies and randomly select a couple of shifts  $(\Delta x, \Delta y)$  from a uniform distribution ranging from  $0.5 \cdot R$  and half of the image size. We use these shifts to apply a translation to the stamp of the *companion galaxy*. Finally, the blend is created by adding up the pixels of two stamps.

Note that the blending process contains two oversimplifications as compared to real observed blends. First, we avoid overlap in the very inner parts of the central galaxy ( $< 0.5R_{\text{c}}$ ) and secondly, the central galaxy is always placed at the centre of the stamp. We are fully aware of these simplifications but consider this enough complexity for our blends in a first proof-of-concept work.

We repeat this process to build up a sample of 30 000 blend images, which necessarily contains some redundancy because each galaxy appears in multiple stamps. However this redundancy should not be considered problematic due to the strict separation between training and test galaxies (see Section 2.3). It allows us to build a large enough sample to train the networks as described in the following. We show in Fig. 2 some examples of blended pairs with different magnitude differences and distances between the two galaxies.

To summarize, at the end of this process, we have for every generated blend system:


- (i) the original CANDELS cut-outs of the *central* and *companion galaxy*,
- (ii) the associated SExtractor segmentation maps,
- (iii) the associated SExtractor photometry (FLUX\_AUTO),
- (iv) the generated blended stamp.

With the purpose of triggering the comparison with other approaches, the software used to generate the blends as described above has been publicly released as a package called `candels-blender`.<sup>2</sup>

## 2.3 Training, validation, and test data sets

As explained in the previous sections, the blend stamps contain some level of redundancy since the same galaxy can appear in several of them. This could artificially improve the results evaluated in the test set because the network might have seen already the same galaxy in the training phase. To avoid this potential bias, we adopt a specific procedure. Following a standard approach in machine learning (e.g. Bengio 2012), we split the data set into three subcategories: training, validation, and test, respectively 60 per cent, 20 per cent, and 20 per cent of the full data set. During the training, the model loss (i.e. cost function) is periodically computed on the validation sample to ensure it is not diverging from the training, which would indicate overfitting or a bad convergence of the network. Training and validation samples can be randomly selected from the same data set, however the test sample, on which the metrics are computed, must be carefully chosen to be both distinct from and representative of the sample used for the training and validation. To achieve this feature and obtain meaningful results, we isolate the sample of galaxy stamps used for the test data set at the very beginning by randomly picking them out of the catalogue. This way, all the galaxies used to construct the blends for the training and validation are never to be found in the test sample of blends, and vice versa.

<sup>2</sup><https://github.com/aboucaud/candels-blender>



**Figure 1.** Selection of CANDELS cut-outs displaying at their centre galaxies with different morphologies. The left-hand column shows the original CANDELS image. The middle one shows the same cut-out after the neighbours removal procedure, leaving the central galaxy fully isolated. The rightmost column shows the SECTRATOR segmentation map for that isolated galaxy. The galaxy images have been asinh-stretched to enhance the details.

**Table 1.** Morphological mix of the final data set used to generate the blended systems.

Galaxy type	Number before inspection	Number after inspection
Bulge	386	352
Disc	473	433
Bulge+disc	884	702
Irregular	875	432
Total	2618	1919

In the end, we have a training/validation set composed of 25 000 blends and a test set of 5000 blends. This generated data set is used to train several deep neural network architectures as described in the next section.

### 3 METHODS

Our goal is to recover, with deep learning, the photometry of the two galaxies before the blending process. The sample being made of real galaxies, we make the assumption that the ground truth (also referred to as *the target* in supervised learning) is the flux of the isolated galaxy computed by SEXTRACTOR on the original CANDELS cut-out. We also assume that the segmentation mask provided by SEXTRACTOR for the isolated galaxies is correct. We understand that these are strong assumptions. However, the main purpose of the work does not depend on the absolute accuracy of the training sample. The main objective is to calibrate how well we can recover the photometry of blended galaxies relative to the accuracy obtained on the same galaxies when they are isolated. In that respect, the ground truth can be replaced with any other measurement.

We perform two different experiments. In the first one we use a standard Convolutional Neural Network (CNN, LeCun et al. 1989; Schmidhuber 2014; Sze et al. 2017) to directly compute the fluxes of the two galaxies from the blend image. We call this configuration `blend2flux`. In the second experiment, we recover with a unique architecture, the fractional segmentation map, which is an image in which pixel values between 0 and 1, depending on the fraction of flux belonging to a given galaxy; as well as the flux, for each of the two galaxies. The idea is to calibrate whether having information on the fractional segmentation map helps the network to obtain a more reliable photometry. We call this second experiment `blend2mask2flux`.

The networks are implemented, trained, and evaluated using the PYTHON API Keras,<sup>3</sup> which runs on top of TensorFlow.<sup>4</sup> The source code needed to reproduce the results of this paper as well as all the plots will be publicly released upon acceptance.

#### 3.1 Configuration 1: `blend2flux`

Experience with deep learning has proven that reducing pre-processing to a minimum often results in better results (Liang & Hu 2015). We thus start off with a deep neural network model that predicts fluxes directly from the blended images without any intermediate step. We use to that purpose a standard CNN configuration including a feature extraction convolutional part followed by a fully connected (or dense) network. The input of the network is thus a

one-channel image with two blended galaxies and the output is a vector of two floating numbers corresponding to the fluxes of each galaxy.

We build a modular version of this sequential network, where the number of layers of both the convolutional and the dense network, as well as their filter size are adjustable. The architecture whose results are shown in this paper is sketched in detail in Fig. 3. The CNN part is made of five convolutional layers activated using a ReLU function and using convolution kernels of size  $3 \times 3$  only. Max-pooling layers are inserted in between each convolution layer to downsample the images. The first layer starts with a filter size of 256, and doubles this filter size every other layer. After the fifth convolutional layer, the data are flattened to be fed to a three-layer classical neural network, finally yielding a vector of size two with the fluxes. Given that our network is aiming at correct relative flux measurements, we choose to use the mean absolute percentage error (MAPE, see equation 1) as our loss function. To adjust the weights during training, we select the Adam algorithm, a popular optimizer for deep learning due to its fast and effective learning (Kingma & Ba 2014). Adam is an extended stochastic gradient descent algorithm, meaning it iteratively updates network weights with individual adaptive learning rates based on both first and second moments of the gradients.

$$\text{MAPE}(y_{\text{meas}}, y_{\text{true}}) = \frac{100}{n} \sum_n \left| \frac{y_{\text{true}} - y_{\text{meas}}}{y_{\text{true}}} \right| \quad (1)$$

This `blend2flux` network, which has about 25.7 million free parameters, is then trained from scratch using the training set of 25 000 images. We consider the network as having converged after the validation loss, computed on the validation part of the training sample, stays on a plateau for a full 10 consecutive epochs after having decreasing the learning rate several times (Yao, Rosasco & Caponnetto 2007). For this network, it happened after 70 epochs which took less than 5 h of training on an Nvidia K80 GPU.


The network built being modular, we trained a few variations around the fiducial network presented above to compare their relative performance. The results of the various network models as a function of the number of filters for the first convolutional layer are summarized in Table 2 with the fiducial results in the middle column. The table shows that doubling the initial filter size (right-hand column) only slightly increases the performance on the validation set regarding the estimated fluxes in Section 4, at the expense of quadrupling the number of parameters (hence the training time and computation cost). Using instead a smaller network with an initial filter size of 64 (left-hand column) reduces the number of parameters to about 1.6 million, which has a higher impact on the performance ( $\sim 1$  per cent worse). The network still reaches though a precision below 10 per cent on estimated fluxes, despite being significantly reduced in size. We therefore want to stress here that smaller and simpler networks than our fiducial one still outperform SEXTRACTOR results.

#### 3.2 Configuration 2: `blend2mask2flux`

In a second experiment, we aim at recovering the individual segmentation maps for the two galaxies in addition to the photometry. The objective of this exercise is to quantify if a fraction segmentation map contains additional information that the networks can use to improve the photometry. We achieve this objective using a concatenation of two different networks, one to produce the fractional segmentation maps, and a second to predict the fluxes from these maps and the blend image. We call that composite

<sup>3</sup><https://keras.io>

<sup>4</sup><https://github.com/tensorflow/tensorflow>




**Figure 2.** Selection of blended systems created and used in this work. The stamps are ordered vertically by the distance in pixels between the galaxy centres, and horizontally with respect to the magnitude difference between the galaxies. The images have all been asinh-stretched for visualization purposes.

network `blend2mask2flux`. One important constrain when building this network was to ensure it had approximately the same number of free parameters as the fiducial `blend2flux`.

To produce the probabilistic maps, we use a suitable deep network architecture from the literature called a U-Net (Ronneberger, Fischer & Brox 2015). U-Net was designed to perform biomedical image segmentation and has already proven useful to detect and segment overlapping chromosomes. The network architecture is quite unique and characterized by an ability to capture both fine and large-scale information of the input image by keeping a copy of each downsampling step (convolution + *max-pooling*) and concatenating it at the upsampling step. For our purpose, we create a modular version of the original U-Net architecture made of blocks of two convolutional layers activated with ReLU, followed by either a downsizing or upsizing layer (respectively, *max-pooling* and up convolution layers). Because the output images are of the same shape as the input blend, each downsizing block is associated with

an upsizing one in the network, and the model can therefore be parametrized by the number of consecutive downsizing blocks, as well as the size of the filters (number of convolution kernels). After some tests and with a range of these parameters, we selected a U-Net with a depth of 5 and an initial filter size of 32, which we also refer to as the *fiducial* model. The exact architecture of this network is depicted on Fig. 4. The last activation of the model is a sigmoid function that creates output images with pixels in the range [0, 1]. We use a binary cross-entropy loss to train the model. To compare with the input results, we convert the fractional segmentation maps to traditional binary segmentation maps using a threshold at 0.5. Further results of this pure segmentation stage will be discussed in a specific Section 4.4 at the end of this paper.

The second part of this composite model is the retrieval of the photometry using the blend image and the fractional segmentation maps obtained with the U-Net. For this part, we use an architecture similar to the `blend2flux` model shown in Section 3.1 with a



**Figure 3.** Schematic representation of the fiducial `blend2flux` network. The network takes as input an image of a blended system and outputs the fluxes of the two galaxies. The blue boxes correspond to the convolutional part of the network. The yellow part is the fully connected section. The sizes of the different layers and convolutions are also indicated.

reduced number of free parameters, and changing the input to a three-channel input – the concatenation of the blend image, the fractional segmentation maps of the central galaxy, and of the companion galaxy – (instead of one channel – the blend image – in the original `blend2flux` network). Like the `blend2flux` model, the output of the network is evaluated using the MAPE loss.

The composite `blend2mask2flux` network is trained following a particular process. First, the U-Net is trained alone to produce accurate segmentation maps of the two galaxies. Then, we load the pre-trained weights of the U-Net into the `blend2mask2flux` network, and train the network end-to-end with respect to the flux retrieval, i.e. using the MAPE loss on the photometry. Note that we still keep the loss on the U-Net but with a weight of 0.1 compared to the photometry loss. This last optimization step, during which we optimize the network with respect to both the segmentation and the photometry loss, also fine-tunes the fractional segmentation maps for flux measurement. A more detailed discussion on this aspect can be found in Section 4.4.

This `blend2mask2flux` network presented above has 18.5 million free parameters, a number very close and even inferior to the *fiducial* `blend2flux` model. The U-Net part is trained from scratch on the 25 000 image training set for about 50 epochs. Then the end-to-end `blend2mask2flux` network is trained during a few hundred epochs with a small learning rate. This full process takes about 50 h of training on a Nvidia K80 GPU, much longer than that of the `blend2flux` network. Both the model complexity, and the training process (reduced batch size for the U-Net training) are accountable for that order of magnitude time difference.

### 3.3 Baseline: SEXTRACTOR

In order to have a baseline to compare with, we also run a classical SEXTRACTOR segmentation procedure on the blended systems. We highlight that the comparison is not completely fair since SEXTRACTOR does not only measure photometry but also detects the objects without any prior on the number of existing objects. However, the two deep learning approaches implicitly incorporate a prior on the number of blended galaxies through the training set (networks are trained only with images containing two objects).

In order to minimize that effect, and help SEXTRACTOR as much as possible, we adapted the procedure reported by Rix et al. (2004), Galametz et al. (2013), where SEXTRACTOR is first ran in a *cold*

**Table 2.** `blend2flux` network performance computed on the entire test set using MAPE.

Initial filter size	64	256	512
Number of parameters (Mio.)	1.6	25.7	102.7
Flux error central (per cent)	9.33	8.39	8.25
Flux error companion (per cent)	8.79	8.01	7.98
Total flux error (per cent)	9.06	8.20	8.12

mode, aiming to select the larger elements in a blended image followed by a second round where it is ran in a *hot* mode – which is more sensitive to small structures. In our particular case, where the data are known to have only two elements, the *cold* mode was used to scan all the images and a subsequent run with the *hot* mode was restricted to those images for which SEXTRACTOR identified only one object. Our code used the PYTHON package `sep` (Barbary 2016) and the parameters used for both modes are described in Table 3.

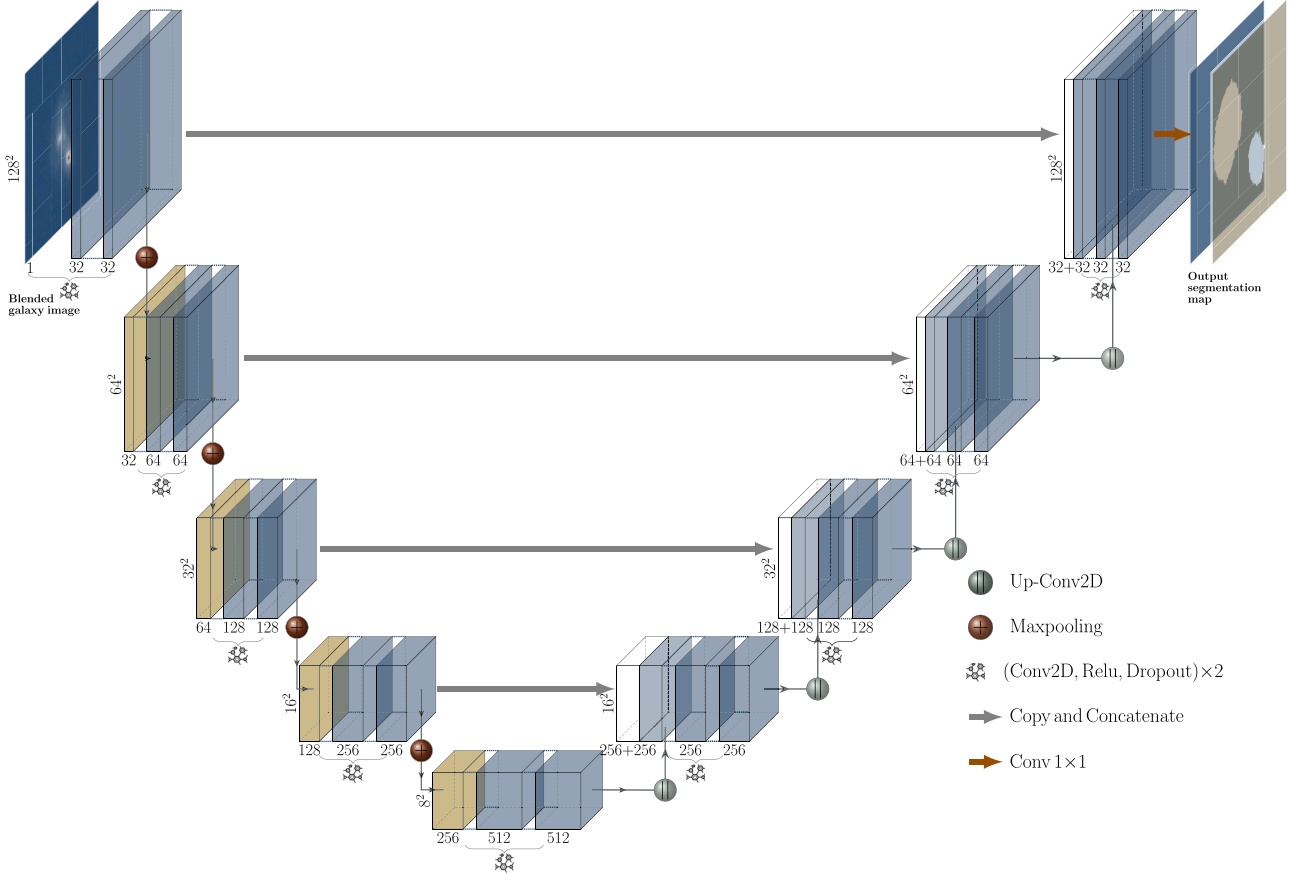
Following this procedure, results can be divided in three cases:

- (i) SEXTRACTOR detects exactly two galaxies (75 per cent): fluxes were associated with central or companion galaxy based on the closest detection.
- (ii) SEXTRACTOR detects only a single object, meaning it is not able to deblend the pair (22 per cent): detected object was associated with the central galaxy if its measured centroid is located within  $0.5R_{\text{cen}}$  from the centre of the image. Otherwise, detection was associated with the companion galaxy.
- (iii) SEXTRACTOR overdeblends and detects more than two objects (3 per cent): the two brightest detections were considered – others were ignored.

## 4 RESULTS

In this section, we evaluate the results of the three experiments described previously. The main objective is to test the photometric accuracy of blended objects as compared with the photometry obtained on the same objects when they are isolated. We use the magnitude difference as the main indicator of accuracy and explore the results as a function of two main parameters: the magnitude difference between the two galaxies and the distance between the two galaxy centroids.





**Figure 4.** Schematic representation of the U-Net part of the `blend2mask2flux` network. The network takes as input an image of blended system and outputs two fractional segmentation maps, each pixel is assigned a float value between 0 and 1 corresponding to the fraction of flux belonging to a given galaxy. The lines indicate the connections among the different layers.

#### 4.1 Overall photometric accuracy

Figs 5 and 6 show the recovered magnitude in the blended systems (hereafter output magnitude) for the three different methods, the `blend2flux` and `blend2mask2flux` networks and SEXTRACTOR, as a function of the magnitude measured on the same isolated galaxies (hereafter input magnitude). On Fig. 5, we focus the results on the central (top) and the companion (bottom) galaxy using the blends for which SEXTRACTOR detects them. On Fig. 6, we aggregate the results on both galaxies, and distinguish the cases for which SEXTRACTOR detects the pair (top) and over- or underblends (bottom).

On both figures, the deep learning architectures behave very similarly. The relation between the two quantities is centred on the one-to-one line and the typical scatter is  $\sim 0.1$  mag. The scatter is roughly constant over all the luminosity range explored which means that the photometry can be recovered with similar accuracy for bright and faint objects in our sample. This is clearly not the case for the SEXTRACTOR results which present a noticeable increase of the scatter at the faint end. This difference highlights an important advantage of a machine learning approach. The loss function used (MAPE) does penalize errors regardless of the flux, which helps producing an unbiased estimator.

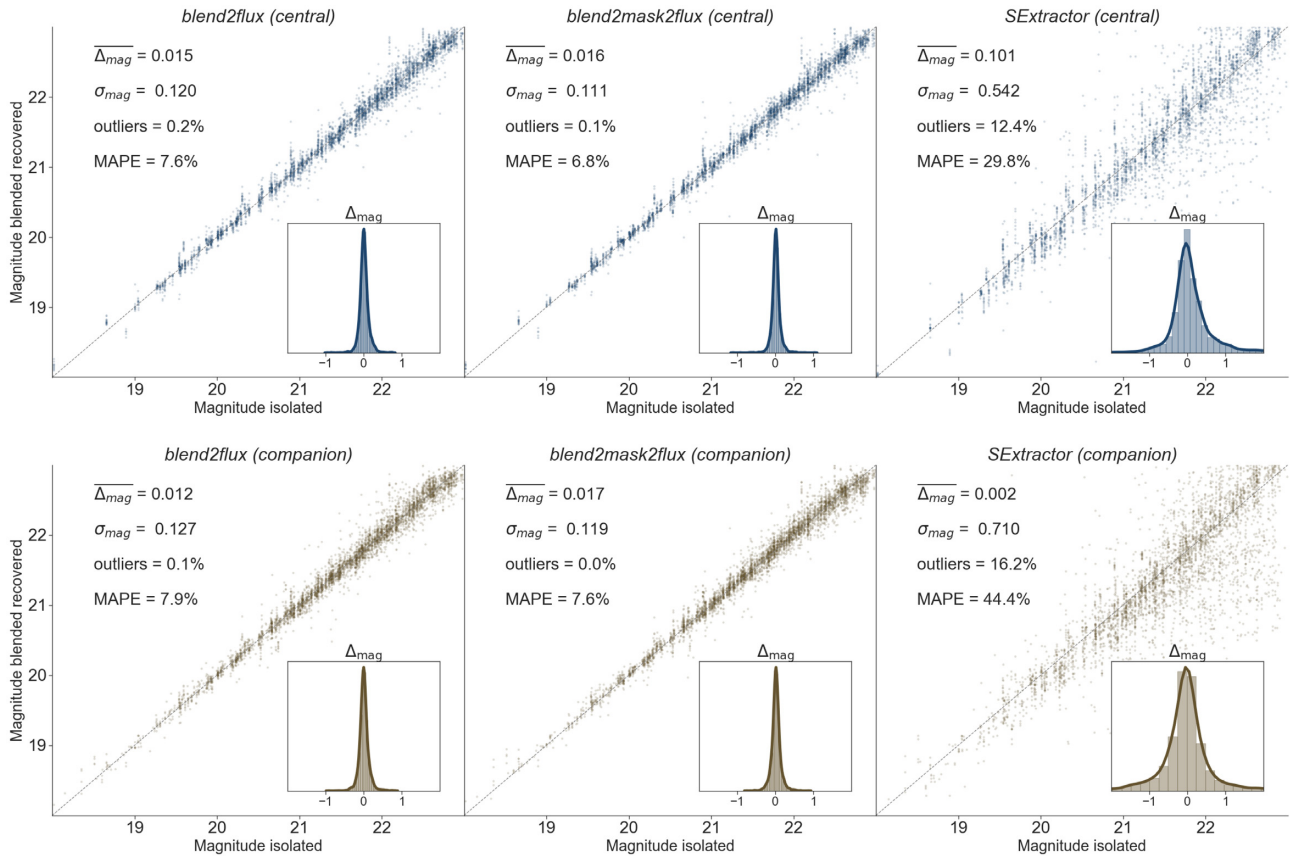
In each panel of Figs 5 and 6, we quantify in more detail the bias and scatter on the recovered photometry. The embedded histograms show the distribution of photometric error  $\Delta_{\text{mag}} =$

**Table 3.** SEXTRACTOR parameters for *hot* and *cold* modes.

Parameter	Hot	Cold
Detection threshold	4	5
Minimum pixel area per object	6	10
Minimum contrast ratio	0.0001	0.01
Number of thresholds for deblending	64	64

$\text{mag}_{\text{blend}} - \text{mag}_{\text{isolated}}$  between the output and input magnitudes. The distributions for both the central and the companion galaxy are generally well centred around zero for the three codes, which indicates that the estimators are globally unbiased. We note that the SEXTRACTOR panels present a slightly skewed histogram and positive bias of 0.1 mag for the central galaxy. We explain this slight bias by looking at the selection process of the companion galaxy described in Section 2, which is skewed a bit towards selecting fainter galaxies than the central ones.

The visible difference between the methods are shown in the scatter. Both deep learning approaches present a very low scatter of  $\sim 0.1$  mag compared to the  $\sim 0.5$ – $0.7$  mag scatter of SEXTRACTOR. Another good indicator of the model performance, used for training the models, is the MAPE (see equation 1), computed here on the magnitude. Again, both network model show good and similar performance, with always a slight advantage for the `blend2mask2flux`, whereas SEXTRACTOR is



**Figure 5.** Magnitude measured on the blend systems as a function of the magnitude measured by SExtractor on the same isolated galaxies (isolated magnitude). The top row shows the results for the central galaxy using the blends for which SExtractor detected either the two galaxies or only the central one. The bottom row shows the results for the companion galaxy using the blends for which SExtractor detected either the two galaxies or only the companion. The columns refer to different codes or models applied to the blend images, respectively from left to right `blend2flux`, `blend2mask2flux` and SExtractor. The dashed line denotes identical estimation from blended and isolated galaxy images to guide the eye. The inner panels show the histograms of photometric errors ( $\Delta_{mag} = \text{mag}_{\text{blend}} - \text{mag}_{\text{isolated}}$ ). The numbers in each panel indicate the average photometric error  $\overline{\Delta_{mag}}$ , the dispersion  $\sigma_{mag}$ , the fraction of outliers, defined as  $|\Delta_{mag}| > 0.75$ , and the MAPE on the flux.

distanced. These two indicators show an overall improvement of the measured photometry of a factor 4 using the deep learning models compared to using SExtractor on these blended galaxies.

Another important difference between the methods is the fraction of catastrophic errors, i.e. cases in which the estimated photometry in the blended systems significantly differs from the input value. We arbitrarily define catastrophic errors as  $|\Delta_{mag}| > 0.75$ , which corresponds to an error of a factor of 2 in flux. The fraction of outliers defined that way is two orders of magnitude smaller with the deep learning methods compared to SExtractor. Both network architectures achieve a comparable fraction of  $\sim 0.1$  per cent outliers, whereas the SExtractor fraction is of the order of  $\sim 10$  per cent, even when restricting the results to the cases where SExtractor detects both objects (see top panel of Fig. 6).

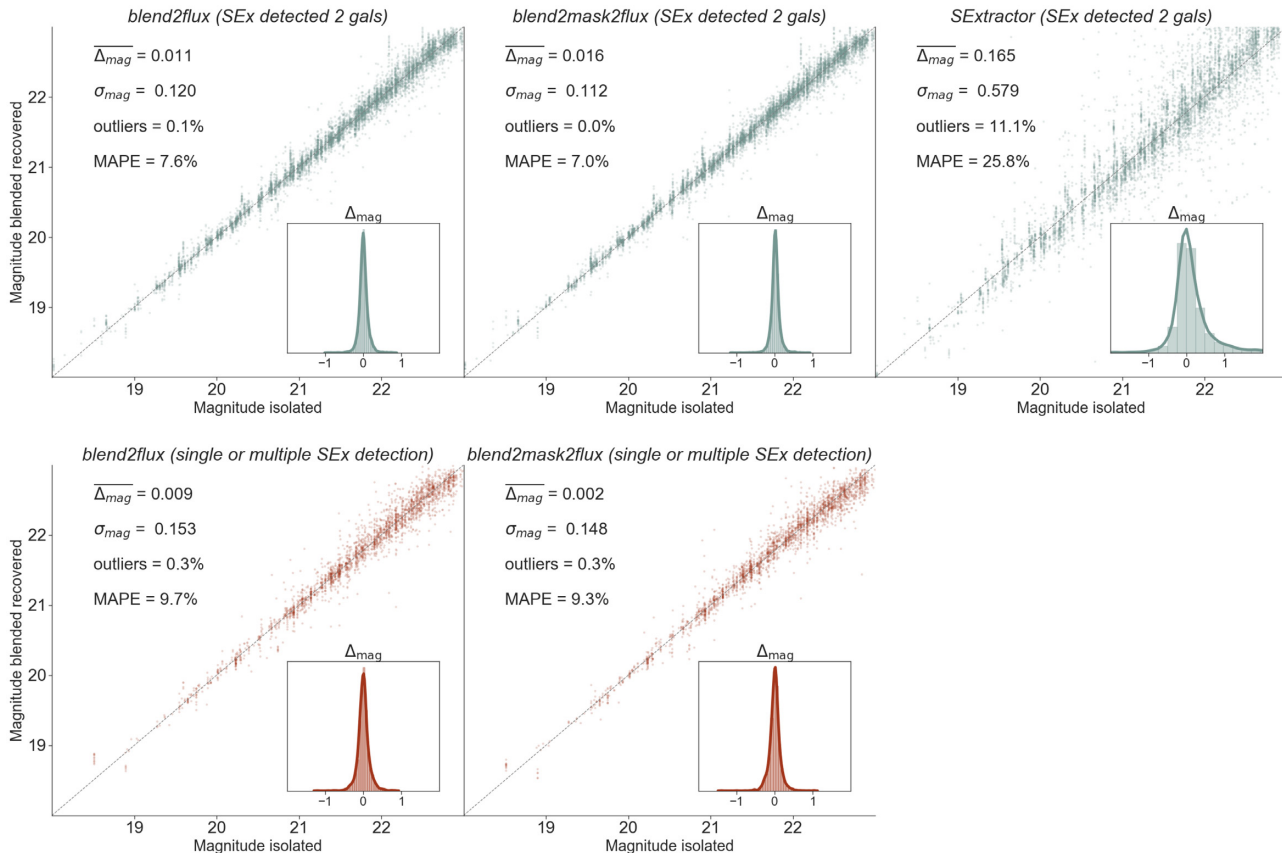
Lastly, as shown on the bottom panels of Fig. 6, the performance of both `blend2flux` and `blend2mask2flux` models on the galaxies that SExtractor did not manage to accurately deblend (25 per cent) gets affected compared with the well deblended cases (top panels) but remains unbiased with a low scatter and an outlier rate below 0.4 per cent.

## 4.2 Photometric accuracy versus blend properties

Aiming for an unbiased performance for a range of blend properties, we report results as a function of the magnitude difference between blended galaxies and the distance between the two objects.

In Fig. 7, we show the binned photometric error (the bias and scatter in our magnitude estimate) as a function of the difference in magnitude between the two blended galaxies. We observe that the two deep learning approaches present a very stable behaviour across the whole range of magnitude difference. As expected, the bias slightly increases when one of the galaxies in the pair is significantly brighter. However, it remains below  $\sim 0.06$ . Overall the bias remains always lower than the SExtractor-based estimates. The deep learning results are also very stable in terms of scatter which is of the order of  $\sim 0.1$  mag. Here, the scatter for SExtractor-based estimates is always significantly larger ( $\sim 0.25$  mag) than for the networks, and also shows a strong increase with magnitude difference between central galaxy and companion. For both networks this trend is only slightly visible.

We look at the the bias and scatter with respect to the normalized distance between the two galaxies (Fig. 8). While the deep learning results display little photometric dependence with distance for both



**Figure 6.** Magnitude measured on the blend systems as a function of the magnitude measured by SExtractor on the same isolated galaxies (isolated magnitude). The top row shows the results for the central and companion galaxies on the blends for which SExtractor detected exactly two galaxies, while the bottom row show the results on the blends for which SExtractor detected either one or more than two galaxies (under- or overdeblending). The different columns indicate different codes or models applied to the blend images, from left to right `blend2flux`, `blend2mask2flux`, and SExtractor. The dashed line denotes identical estimation from blended and isolated galaxy images to guide the eye. The inner panels show the histograms of photometric errors ( $\Delta_{mag} = \text{mag}_{\text{blend}} - \text{mag}_{\text{isolated}}$ ). The numbers in each panel indicate the average photometric error  $\overline{\Delta_{mag}}$ , the dispersion  $\sigma_{mag}$ , the fraction of outliers, defined as  $|\Delta_{mag}| > 0.75$ , and the MAPE on the flux.

central galaxy and companion, the SExtractor results show a fluxes underestimated by up to 1 mag for close objects ( $< R_e$ ).

In both plots, the SExtractor measurements are systematically more biased and more scattered than the machine learning-based estimates across the whole range of parameters. We also notice that both neural net architectures behave very similarly.

### 4.3 Photometric accuracy and morphology

The galaxies in our sample are classified into four morphological types (pure bulge, pure disc, two component bulge+disc, irregular) and are distributed as was shown in Table 1. One major property of machine learning methods is that they do not assume any external prior on the galaxy shape (as opposed to model fitting techniques). Alternatively, the prior is inferred from the data during training. We explore in Fig. 9 the dependence of the photometric accuracy on the morphological type. We plot the median bias and scatter in bins of magnitude and distance now divided by morphological type.

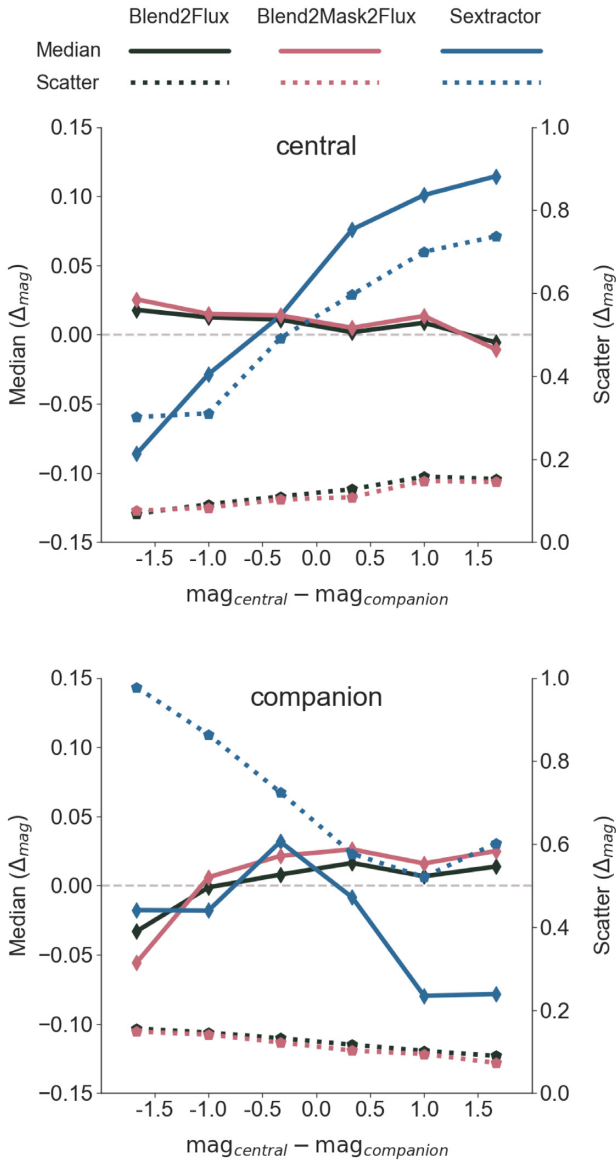
Surprisingly, spheroidal galaxies tend to present larger errors when these galaxies are fainter than the other galaxy in the blended system ( $\Delta_{mag} > 0$ ). This behaviour seems to be present in both codes. The reason for this is unclear. One possible explanation is that the outskirts of the spheroids are too faint to be detected. Since these objects typically have large Sérsic indices (i.e. steep luminosity

profiles), the fraction of light in the outskirts is not negligible. Another possible explanation is that, as can be seen in Table 1, spheroids are under represented as compared to other classes. It might be that the network did not see enough examples. The behaviour of SExtractor is different, in the sense, that irregulars clearly present a larger bias than the other morphological types. This is expected, since irregular light profiles are more difficult to capture. The machine learning approaches perform better since the priors were learned during the training phase.

As can be seen in Fig. 9 the photometric accuracy (magnitude scatter) overall is considerably lower for our two deep learning algorithms than for SExtractor results.

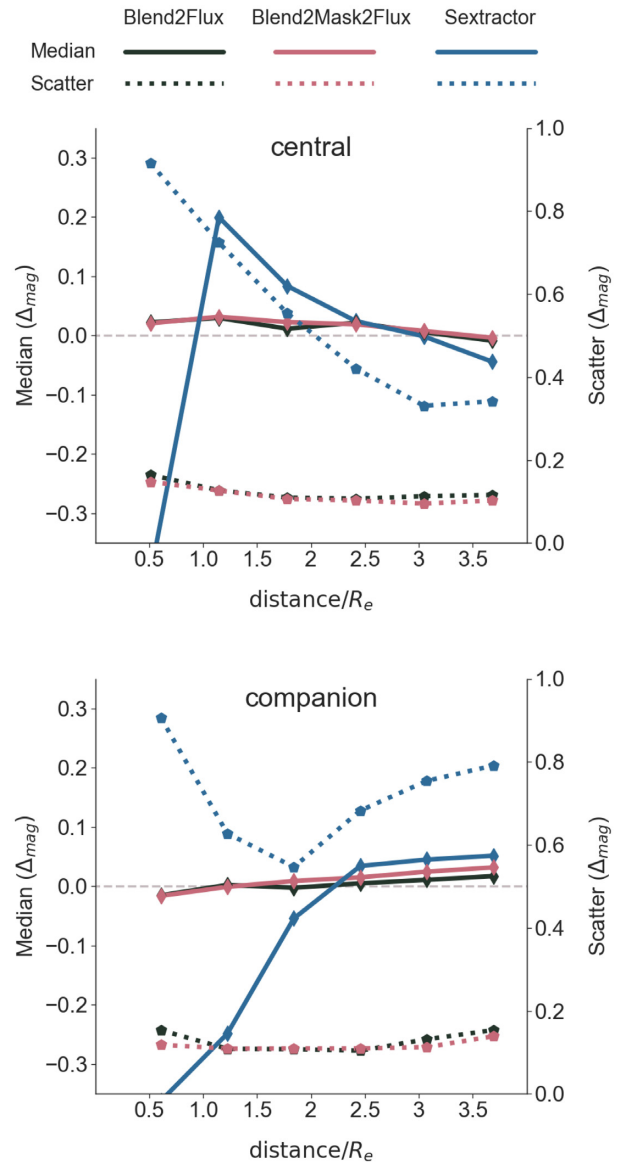
### 4.4 Segmentation maps

Throughout the paper, segmentation maps have been considered as a by-product of both SExtractor and the `blend2mask2flux` network, possibly improving the photometry. In this subsection, we focus on the recovery of the segmentation maps of blended galaxies from the deep learning architecture, as well as the comparison between the results of the initial training of the U-Net alone and the ones after the training of the full `blend2mask2flux` network, which is characterized by the tuning of the fractional segmentation maps for photometry.



**Figure 7.** Binned photometric errors ( $\Delta_{\text{mag}} = \text{mag}_{\text{blend}} - \text{mag}_{\text{isolated}}$ ) as a function of the relative magnitude difference in the blended system. Solid lines and left-hand axis denote median values and dotted lines and right-hand axis refer to scatter values. Colours indicate different codes: `blend2flux` (dark green), `blend2mask2flux` (pink), and `SEXTRACTOR` (blue). The top plot shows the values for the central galaxy and the bottom plot for the companion galaxy. In both plots, only the results where the selected galaxy is detected with the `SEXTRACTOR` method are used.

In Section 3.2, we describe the `blend2mask2flux` model as a hybrid network made of a U-Net whose output is fed to a modified `blend2flux` network. The U-Net is in charge of reproducing the two `SEXTRACTOR` segmentation maps of the original CANDELS galaxy cutouts from the blend image. In other words, its task is to take as input the full  $128 \times 128$  blend image and produce two binary  $128 \times 128$  images that correspond to the masks of the central and companion galaxy; this can be seen for a selection of four blends in Fig. 10. For better assessment of the accuracy of the method, we trained the network to output the fractional segmentation maps in a specific order, central galaxy first, and then the companion. The cost function (loss) used to train the modified



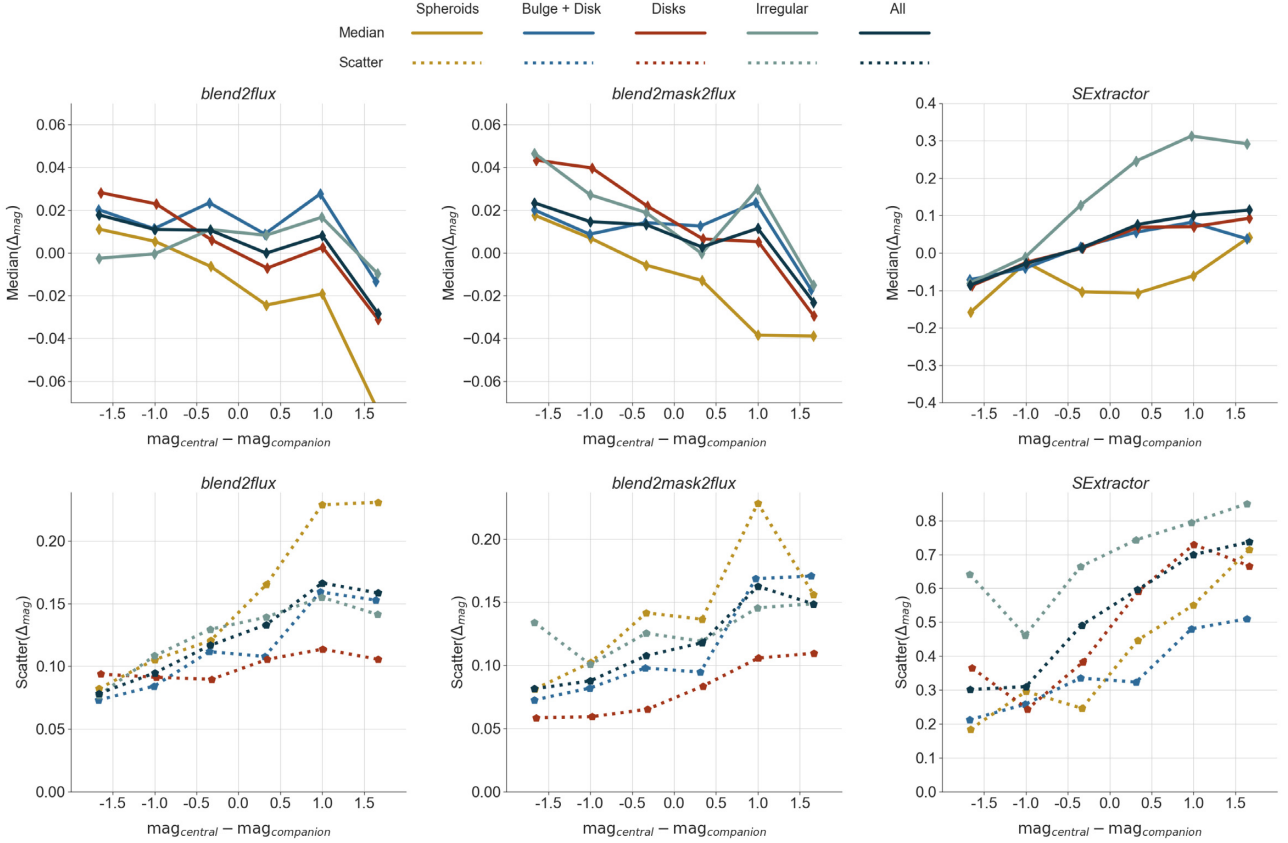
**Figure 8.** Binned photometric errors ( $\Delta_{\text{mag}} = \text{mag}_{\text{blend}} - \text{mag}_{\text{isolated}}$ ) as a function of the distance between the two galaxies, normalized by the effective radius of the selected galaxy. The notations and data selection are the same as the one in Fig. 7.

U-Net is a *binary-crossentropy*, which performs well for a pixel-by-pixel binary classification as needed for our segmentation maps.

To score the results, the *binary-crossentropy* loss is not very informative since every pixel rightfully classified as *background* adds up to the accuracy, while we would like to assess the similarity to the target `SEXTRACTOR` segmentation map. For that purpose, we use a metric called Intersection over Union (IoU) also known as Jaccard index (Jaccard 1901) of set A and B

$$\text{IoU}(A, B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|}. \quad (2)$$

It is usually defined in computer vision for bounding boxes, but can be adapted to any shape. This metric has the advantage of decreasing very rapidly to zero in case of a mismatch between A and B in terms of location or morphology. Therefore a score superior to 0.5 is considered a good score.



**Figure 9.** Dependence of photometric bias (solid lines, top panels) and scatter (dotted lines, bottom panels) on the morphological type for the three codes considered in this work as a function of the magnitude difference. From left to right, the different panels show the results for `blend2flux`, `blend2mask2flux`, and `SExtractor`, respectively. The different colours indicate the morphological type: spheroids (yellow), disc + spheroids (blue), and discs (red) and irregulars (light green). The dark blue lines show the results for all galaxies.

After training, the U-Net with the setup described in Section 3.2 obtains an average IoU score of

$$\text{IoU}_{\text{U-Net}} = 0.82$$

on the test images, which is an indication of a very good recovery. However, once the `blend2mask2flux` is trained end-to-end to recover the photometry, thus allowing the parameters of the U-Net section to vary, the average IoU score on the test data set drops to

$$\text{IoU}_{\text{blend2mask2flux}} = 0.70.$$

The outcome and evolution from the pure segmentation objective to the photometry objective can be seen in Fig. 10. We observe that the masks predicted by the two networks differ. The `blend2mask2flux` network has therefore modified the segmentation maps to improve photometric accuracy.

To further quantify the evolution of the segmentation when tuning the network for the photometry, we show the histograms of the IoU score for the two models in Fig. 11. We see that the pure segmentation network has a very small dispersion, which broadens and becomes worse when it is optimized for the photometry. To better understand this difference, we compare directly the area of the overlapping region in the predicted masks to the ground truth. We compute the area by summing the pixels belonging to the segmentation maps of both galaxies  $n_{\text{pred}}$  and compare with the sum on the initial `SExtractor` segmentation maps  $n_{\text{true}}$  via a simple estimator averaged over the entire test

data set


$$\Delta n = \frac{\langle n_{\text{true}} - n_{\text{pred}} \rangle}{\langle n_{\text{true}} \rangle}.$$

We measure  $\Delta n \simeq -0.05$  for the U-Net and  $\Delta n \simeq 0.31$  for the `blend2mask2flux` model, which means that the number of pixels belonging to the overlapping region is slightly overestimated by the former method and well underestimated by the latter. The interpretation is that the `blend2mask2flux` network tends to ignore the regions of overlap between galaxies since inferring the flux from these regions is more uncertain. This most probably results in a small positive bias reported in Fig. 7.

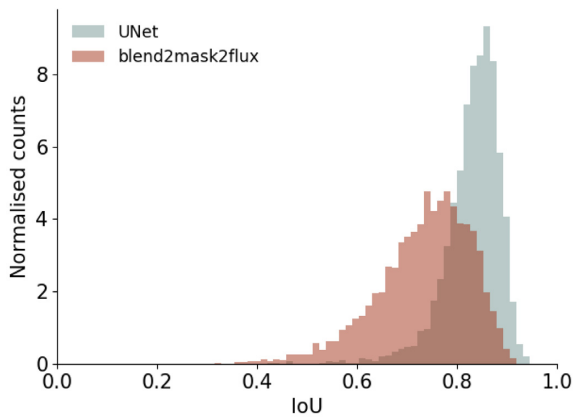
## 5 SUMMARY AND CONCLUSIONS

We have presented a first test of a deep learning method to measure the photometry of blended systems in monochrome space-based images of the distant Universe under simplified conditions.

First, we built a training set out of observed high-redshift galaxies from the CANDELS survey which have been artificially blended. The data set covers a representative range of morphologies (bulges, discs, and irregulars), magnitude differences ( $-2 < \Delta_{\text{mag}} < 2$ ), and distances ( $0.5R_e < D < 4R_e$ ) between the pairs. The data set of blended pairs is made public with this work to promote comparisons with other approaches.



**Figure 10.** Selection of four simulated blends from the test data set and the recovery of the individual galaxy masks through the U-Net and `blend2mask2flux` networks. The leftmost column shows the stamp of blended galaxies that is input of the network. The next two columns are the segmentation masks obtained on the individual galaxy images with SEXTRACTOR (considered ground truth in this work), blue refer to the central galaxy and brown to the companion one. The last four columns show the predicted maps after threshold, yielded respectively by the U-Net architecture and the `blend2mask2flux` model as labelled. For each recovered galaxy mask, the segmentation score (IoU) as compared the input mask is indicated in the lower right corner.



**Figure 11.** Histogram of the distribution of segmentation score (IoU) comparing the segmentation maps of the individual galaxies with the one obtained on the blended image with the U-Net and the `blend2mask2flux` model, for all the blended galaxies in the test set. The U-Net model being the starting point of the `blend2mask2flux` model, we can consider the difference between the orange histograms and bars to the blue ones as an indicator of the impact on masks when optimizing for the photometry.

We have tested two different neural network architectures. The first one measures the fluxes of the two galaxies directly from the images. The second approach, more complex, also estimates the fractional segmentation maps. The networks are trained with a sam-

ple of 25 000 galaxy pairs and tested on an independent sample of 5000 pairs. The results are compared to the standard SEXTRACTOR approach on the same test set. Our main results are:

(i) Both deep learning approaches result in an unbiased photometric estimate with a typical uncertainty of  $\sim 0.1$  mag. This represents an improvement of at least a factor of 4 in flux error as compared with SEXTRACTOR even if the comparison is restricted to the cases where SEXTRACTOR detects exactly two objects as expected.

(ii) The fraction of galaxies for which the photometric error exceeds 0.75 mag is as low as  $\sim 1$  per cent in the two machine learning approaches. This value reaches  $\sim 12$  per cent for SEXTRACTOR. Our deep learning methods also reach an excellent photometric accuracy in these cases where SEXTRACTOR over- or underdeblends (i.e. finds more or less than two galaxies per image).

(iii) The photometric accuracy obtained with the two deep learning approaches is very stable across all magnitude differences and distances explored in this work. Even for large magnitude differences between the two galaxies (factor  $\sim 2$ ), the photometric uncertainty stays close to  $\sim 0.2$  mag. This represents a major improvement as compared to SEXTRACTOR whose performance strongly depends on the properties of the blended system; at large magnitude differences its photometric uncertainty can reach 1 mag.

(iv) The presented method does not assume any pre-defined model for the shape of galaxies. This is translated into a comparable photometric accuracy for all the morphological types explored in this work (discs, bulges, and irregulars).

(v) Estimating and using the segmentation maps to estimate photometry results in a marginal gain in terms of photometric accuracy. The more complex network reaches slightly lower photometric errors and a smaller fraction of outliers at the expense of significantly larger training times. However, it has the advantage that the fractional segmentation maps can be used to estimate uncertainties. We will explore this in future work.

(vi) When a network is asked to optimize both for segmentation masks and photometry simultaneously, the recovered masks are usually tighter than the original ones derived by SEXTRACTOR on the isolated galaxies. This is especially true for irregular and bulgy galaxies.

This proof-of-concept work shows that machine learning can be used as a powerful tool on large imaging data sets, to measure the photometry. Despite the simplistic constraints we imposed on our data set: two galaxies per stamp, one galaxy pinned at the image centre and no blends with completely overlapping galaxy centroids (also referred to as *unrecognized blends*), our photometric measurement networks have demonstrated that on monochromatic images, they outperform the traditional SEXTRACTOR approach with respect to photometric accuracy, precision, outliers fraction and stability towards different morphological types.

On top of that, we trained a network to also produce fractional segmentation maps. With a lower number of free parameters, the network using these maps systematically achieved better results on the photometry than the direct mapping between the blend image to the flux measurement. These probabilistic maps – that once thresholded are called segmentation maps – may well be used as a starting point by other software to guide the modelling of the blend galaxies.

Future work will focus on generalizing the approach presented here to a more realistic situation, including multiple ( $>2$ ) galaxies and more complex blend configurations as well as estimating full posterior of fluxes instead of single point estimates.

## ACKNOWLEDGEMENTS

We thank the anonymous referee for the constructive report that helped improve the paper.

We acknowledge the support from CC-IN2P3<sup>5</sup> which provided the computing resources needed to train the deep learning models. Figs 3 and 4 made use of modified versions of the LATEX package PLOTNEURALNET (Iqbal 2018) and several code snippets openly available on stack overflow. We deeply thank everyone who contributes to open forums and learning platforms.

This work was partially developed during the Fifth COIN Residence Program<sup>6</sup> (CRP#5) held in Chania, Greece in 2018 September. We thank Vassilis Charmandaris for encouraging the accomplishment of this event. COIN and EEOI are financially supported by CNRS as part of its MOMENTUM programme over the 2018–2020 period. RSS acknowledges the support from NASA under the Astrophysics Theory Program Grant 14-ATP14-0007.

## REFERENCES

Barbary K., 2016, *J. Open Source Softw.*, 1, 58

<sup>5</sup><https://cc.in2p3.fr>

<sup>6</sup><https://cosmostatistics-initiative.org/residence-programs/coin-residence-program-5-kania-greece/>

- Barchi P. H., da Costa F. G., Sautter R., Moura T. C., Stalder D. H., Rosa R. R., de Carvalho R. R., 2017, *J. Comput. Interdiscip. Sci.*, 7, 114
- Beck R., Lin C.-A., Ishida E. E. O., Gieseke F., de Souza R. S., Costa-Duarte M. V., Hattab M. W., Krone-Martins A., 2017, *MNRAS*, 468, 4323
- Bengio Y., 2012, *Practical Recommendations for Gradient-based Training of Deep Architectures*. Springer, Berlin, Heidelberg, p. 437
- Bertin E., Arnouts S., 1996, *A&AS*, 117, 393
- Dawson W. A., Schneider M. D., 2014, doi:10.2172/1122197
- Dawson W. A., Schneider M. D., Tyson J. A., Jee M. J., 2016, *ApJ*, 816, 11
- Dieleman S., Willett K. W., Dambre J., 2015, *MNRAS*, 450, 1441
- Dimauro P. et al., 2018, *MNRAS*, 478, 5410
- Domínguez Sánchez H., Huertas-Company M., Bernardi M., Tuccillo D., Fischer J. L., 2018, *MNRAS*, 476, 3661
- Elliott J., de Souza R. S., Krone-Martins A., Cameron E., Ishida E. E. O., Hilbe J., COIN Collaboration, 2015, *Astron. Comput.*, 10, 61
- Fussell L., Moews B., 2019, *MNRAS*, 485, 3203
- Galametz A. et al., 2013, *ApJS*, 206, 10
- Goodfellow I., Pouget-Abadie J., Mirza M., Xu B., Warde-Farley D., Ozair S., Courville A., Bengio Y., 2014, in *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc., p. 8
- Hausser B. et al., 2007, *ApJS*, 172, 615
- He K., Zhang X., Ren S., Sun J., 2016, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Las Vegas, USA, p. 770
- Huertas-Company M. et al., 2015, *ApJS*, 221, 8
- Huertas-Company M. et al., 2018, *ApJ*, 858, 114
- Iqbal H., 2018, *HarisIqbal88/PlotNeuralNet v1.0.0*. Available at: <https://doi.org/10.5281/zenodo.2526396>, Nov 21st 2019
- Ivezic Z. et al., 2008, *Serb. Astron. J.*, 176, 1
- Ivezic Ž., Connolly A. J., Jurić M., 2017, in *Brescia M., Djorgovski S. G., Feigelson E. D., Longo G., Cavaoti S., eds, IAU Symp. Vol. 325, Astrominformatics*. Kluwer, Dordrecht, p. 93
- Jaccard P., 1901, *Bull. Soc. Vaudoise. Sci. Nat.*, 37, 547
- Jones D. M., Heavens A. F., 2019, *MNRAS*, 483, 2487
- Joseph R., Courbin F., Starck J. L., 2016, *A&A*, 589, A2
- Khalifa N. E., Hamed Tah M., Hassanien A. E., Selim I., 2018, in *2018 Int. Conf. Computing Sciences and Engineering (ICCSE)*. IEEE, Kuwait City, Kuwait
- Kilbinger M. et al., 2017, *MNRAS*, 472, 2126
- Kingma D. P., Ba J., 2015, *3rd International Conference for Learning Representations (ICLR)*, San Diego
- Kitching T. D., Alsing J., Heavens A. F., Jimenez R., McEwen J. D., Verde L., 2017, *MNRAS*, 469, 2737
- Krone-Martins A., Ishida E. E. O., de Souza R. S., 2014, *MNRAS*, 443, L34
- Laidler V. G. et al., 2006, in *Gabriel C., Arviset C., Ponz D., Enrique S., eds, ASP Conf. Ser. Vol. 351, Astronomical Data Analysis Software and Systems XV*. Astron. Soc. Pac., San Francisco, p. 228
- Lanusse F., Ma Q., Li N., Collett T. E., Li C.-L., Ravanbakhsh S., Mandelbaum R., Póczos B., 2018, *MNRAS*, 473, 3895
- LeCun Y., Boser B., Denker J. S., Henderson D., Howard R. E., Hubbard W., Jackel L. D., 1989, *Neural Comput.*, 1, 541
- Ledig C. et al., 2017, in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, p. 105
- Liang M., Hu X., 2015, in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, p. 3367
- Mancone C. L., Gonzalez A. H., Moustakas L. A., Price A., 2013, *PASP*, 125, 1514
- Melchior P., Moolekamp F., Jerdee M., Armstrong R., Sun A.-L., Bosch J., Lupton R., 2018, *Astron. Comput.*, 24, 129
- Pignatelli E., Fasano G., Cassata P., 2006, *A&A*, 446, 373
- Racca G. D. et al., 2016, in *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, p. 990400 preprint ([arXiv:1610.05508](https://arxiv.org/abs/1610.05508))
- Ravanbakhsh S., Lanusse F., Mandelbaum R., Schneider J., Poczos B., 2017, in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*, p. 1488
- Reiman D. M., Göhre B. E., 2019, *MNRAS*, 485, 2617
- Rix H.-W. et al., 2004, *ApJS*, 152, 163

Ronneberger O., Fischer P., Brox T., 2015, in Navab N., Hornegger J., Wells W. M., Frangi A. F., eds, *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015)*. Springer International Publishing, Cham, p. 234

Rowe B. et al., 2015, *Astron. Comput.*, 10, 121

Safarzadeh M., Ferguson H. C., Lu Y., Inami H., Somerville R. S., 2015, *ApJ*, 798, 91

Salvato M., Ilbert O., Hoyle B., 2019, *Nat. Astron.*, 3, 212

Schawinski K., Zhang C., Zhang H., Fowler L., Santhanam G. K., 2017, *MNRAS*, 467, L110

Schmidhuber J., 2015, *Neural Networks*, 61, 85

Sze V., Chen Y., Yang T., Emer J. S., 2017, *Proc. IEEE*, 105, 2295

Tramacere A., Paraficz D., Dubath P., Kneib J. P., Courbin F., 2016, *MNRAS*, 463, 2939

Yao Y., Rosasco L., Caponnetto A., 2007, *Const. Approx.*, 26, 289

This paper has been typeset from a  $\text{\TeX}/\text{\LaTeX}$  file prepared by the author.