



**HAL**  
open science

## CNN photometric redshifts in the SDSS at $r \leq 20$

M Treyer, R Ait Ouahmed, J Pasquet, S Arnouts, E Bertin, D Fouchez

► **To cite this version:**

M Treyer, R Ait Ouahmed, J Pasquet, S Arnouts, E Bertin, et al.. CNN photometric redshifts in the SDSS at  $r \leq 20$ . Monthly Notices of the Royal Astronomical Society, 2023, 527 (1), pp.651-671. 10.1093/mnras/stad3171 . hal-04237425

**HAL Id: hal-04237425**

**<https://hal.science/hal-04237425v1>**

Submitted on 21 Nov 2023

**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.



Distributed under a Creative Commons Attribution 4.0 International License

# CNN photometric redshifts in the SDSS at $r \leq 20$

M. Treyer<sup>1</sup>,<sup>\*</sup> R. Ait Ouahmed,<sup>1</sup> J. Pasquet,<sup>2,3</sup> S. Arnouts,<sup>1</sup> E. Bertin<sup>4,5</sup> and D. Fouchez<sup>6</sup>

<sup>1</sup>Aix Marseille Université, CNRS, CNES, LAM, F-13388 Marseille, France

<sup>2</sup>AMIS – Université Paul-Valéry – Montpellier 3, F-34199 Montpellier, France

<sup>3</sup>UMR TETIS – Inrae, AgroParisTech, Cirad, CNRS, Univ. Montpellier, F-34398 Montpellier, France

<sup>4</sup>Sorbonne Université, CNRS, IAP, F-75014 Paris, France

<sup>5</sup>CFHT, Kamuela, HI 96743, USA

<sup>6</sup>Aix Marseille Univ., CNRS/IN2P3, CPPM, F-13288 Marseille, France

Accepted 2023 October 12. Received 2023 September 12; in original form 2023 June 5

## ABSTRACT

We release photometric redshifts, reaching  $\sim 0.7$ , for  $\sim 14$ M galaxies at  $r \leq 20$  in the  $11\,500\text{ deg}^2$  of the SDSS north and south Galactic caps. These estimates were inferred from a convolution neural network (CNN) trained on *ugriz* stamp images of galaxies labelled with a spectroscopic redshift from the SDSS, GAMA, and BOSS surveys. Representative training sets of  $\sim 370$ k galaxies were constructed from the much larger combined spectroscopic data to limit biases, particularly those arising from the over-representation of luminous red galaxies. The CNN outputs a redshift classification that offers all the benefits of a well-behaved PDF, with a width efficiently signalling unreliable estimates due to poor photometry or stellar sources. The dispersion, mean bias, and rate of catastrophic failures of the median point estimate are of order  $\sigma_{\text{MAD}} = 0.014$ ,  $\langle \Delta z_{\text{norm}} \rangle = 0.0015$ ,  $\eta(|\Delta z_{\text{norm}}| > 0.05) = 4$  per cent on a representative test sample at  $r < 19.8$ , outperforming currently published estimates. The distributions in narrow intervals of magnitudes of the redshifts inferred for the photometric sample are in good agreement with the results of tomographic analyses. The inferred redshifts also match the photometric redshifts of the redMaPPer galaxy clusters for the probable cluster members.

**Key words:** methods: data analysis – techniques: image processing – catalogues – surveys – galaxies: distances and redshifts.

## 1 INTRODUCTION

As photometric redshifts have become inescapable in most cosmological endeavours, so have machine learning techniques to predict them. Spectroscopy alone can no longer fulfil the task of measuring the distances to the millions of sources detected in current photometric sky surveys, e.g. DES (Dark Energy Survey Collaboration 2016) and KIDS (de Jong et al. 2013), let alone future ones such as Euclid (Laureijs et al. 2011) and Vera Rubin/LSST (Ivezić et al. 2019). Spectral energy distribution (SED) template fitting techniques have been widely used for several decades to estimate the redshifts of galaxies from multiband photometry, the so-called photometric redshifts. This technique relies on a set of observed or modelled SEDs assumed to represent the diversity of observed galaxies (e.g. Arnouts et al. 1999; Ilbert et al. 2006; Brammer, van Dokkum & Coppi 2008). It allows physical parameters to be derived in addition to redshift probability density functions (PDFs). The first neural networks for the estimation of photometric redshifts emerged in the early 2000s (Firth, Lahav & Somerville 2003; Tagliaferri et al. 2003; Collister & Lahav 2004). Since then, machine learning progressed enormously, helped by the growing wealth of data and computing capabilities. A machine learning algorithm learns to map the multidimensional photometric information using labelled and/or unlabelled data, i.e. data with or without known spectroscopic redshifts (supervised and unsupervised

training methods respectively, which can be combined). The accuracy of the photometric redshifts derived from such optimized mapping is much higher than via SED fitting provided the galaxies span the same parameter space as the training sample (see Brescia et al. (2021) for a review and Henghes et al. (2022) for a comparison of several such methods). Unlike SED fitting however, most machine learning methods have only provided point estimates. Few studies have had a probabilistic approach able to estimate uncertainties, Sadeh (2014) and Sadeh, Abdalla & Lahav (2016) being the first to provide redshift PDFs via a classifier. Jones et al. (2023) recently proposed Bayesian neural networks assuming Gaussian’s PDFs ‘as a promising way to provide accurate predictions with uncertainty estimates’.

‘Deep learning’ is the latest step forward in the pursuit of photometric redshifts. Thanks to the development of convolutional neural networks (CNN, LeCun et al. 1998), and with the help of graphics processing units (GPUs), (regularly sampled) images may now be used directly instead of, or sometimes in addition to extracted features (magnitudes, colours, etc.), which only transmit a fraction of the available photometric information, with variable reliability. Deep neural networks were designed to handle the much larger amount of information contained in the image pixels (Hoyle 2016; D’Isanto & Polsterer 2018). They consist of successive layers of artificial neurons, each performing a linear transformation of the input followed by a nonlinear ‘activation function’. Weights are updated as the network processes (learns from) batches of the training data, until a suitable solution is found (a loss function is minimized).

\* E-mail: [marie.treyer@lam.fr](mailto:marie.treyer@lam.fr)

**Table 1.** The photometric and spectroscopic data (Sections 2.2 and 2.1).

Survey	Magnitude	Size	Spectra
SDSS	$10 \leq r \leq 20$	13.8M	No
SDSS	$r \leq 17.8$	660k	Yes
SDSS	$17.8 \leq r \leq 20$	162k	Yes
GAMA	$r \leq 20$	210k	Yes
BOSS	$r \leq 20$	486k	Yes

Pasquet et al. (2019) (hereafter P19) presented a CNN to estimate photometric redshifts straight from multi-band stamp images of galaxies, without any feature extraction nor colour images. The network was designed as a classifier into small contiguous redshift bins, the output of which was normalized to produce PDFs. As a proof of concept, it was applied to the flux-limited,  $r < 17.8$ , spectroscopic Main Galaxy Sample of the SDSS (York et al. 2000), using *ugriz* stamp images and the galactic reddening values along the lines of sight as input data, with the spectroscopic redshifts as labels in the context of supervised learning. The weighted mean values of the so-called PDFs were found to be photometric redshifts of unprecedented accuracy in the limited redshift range of interest ( $z < 0.4$ ). Other methods exploiting galaxy images have since been proposed (Dey et al. 2022; Hayat et al. 2021; Schuldt et al. 2021; Henghes et al. 2022).

Here we use a more complex CNN architecture to estimate photometric redshifts for the  $\sim 14$  million galaxies at  $r \leq 20$  without spectroscopy in the SDSS footprint. The photometric and spectroscopic data are presented in Section 2 (and Appendix A). The architecture, input and output of the network are described in Section 3 (and Appendix F). Training experiments are described in Section 4. The performance of the final experiment is tested in Section 5. Its inference on the photometric sample is presented in Section 6 (and Appendices C and D). We conclude this work in Section 7. Additionally, a recipe for classifying galaxies into blue/star-forming or red/passive types is given in Appendix B and alternative training strategies are explored in Appendix E. The CNN input and output are available at: <https://deepdip.iap.fr/treyer+2023>.


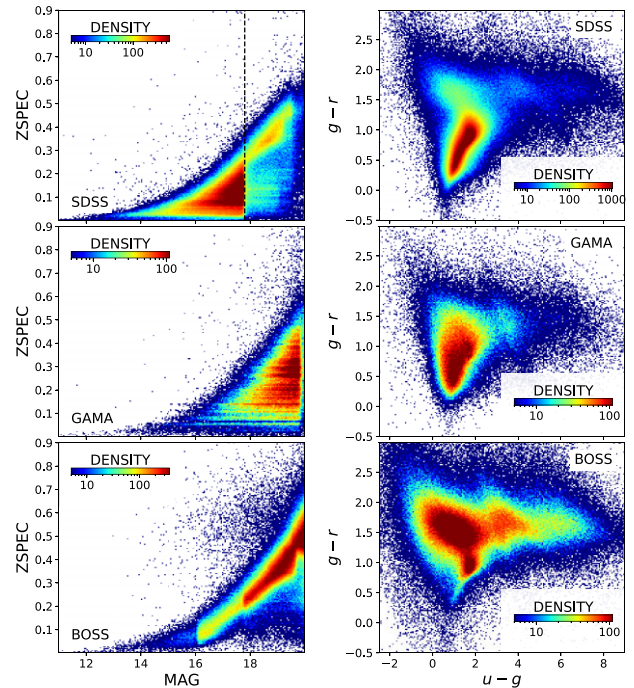
## 2 THE DATA

The data detailed below are summarized in Table 1.

### 2.1 The photometric data

Our catalogue is drawn from the SDSS data release 16 (DR16, Ahumada et al. 2020). The SDSS is a multi-band imaging and spectroscopic redshift survey that was conducted on a dedicated 2.5m telescope at Apache Point Observatory in New Mexico. It provides photometry in the *ugriz* passbands over  $\sim 11,500$  deg<sup>2</sup> of the North and South galactic caps to a limiting magnitude of  $r = 22.5$ . Via the SDSS CasJob web service, we retrieved  $\sim 15.3$ M catalogue entries of non-point-like sources (**type** = 3) with dereddened petrosian magnitudes  $r \leq 20$ , of which  $\sim 1.5$ M have spectroscopic redshifts. Thus the final number of purely photometric sources for which we infer redshifts is  $\sim 13.8$ M. A sky map of this data set is shown in Appendix A.

Photometric redshifts by Beck et al. (2016) (hereafter B16) are available for nearly all these sources. They were computed using a  $k$ -nearest neighbour algorithm (kNN, Csabai et al. 2007) with five dimensions (the  $r$ -band magnitude and 4 colours:  $(u - g)$ ,  $(g - r)$ ,  $(r - i)$ ,  $(i - z)$ ). The training data included deep, high-redshift

**Figure 1.** The redshift distributions of the SDSS, GAMA, and BOSS surveys presented in Section 2.2, totalling 1.5M galaxies.**Figure 2.** Magnitude/redshift and  $(u - g)/(g - r)$  distributions of the SDSS, GAMA, and BOSS surveys (Section 2.2).

spectroscopic surveys in addition to the SDSS. A 3D error map ( $r$ ,  $g - r$ ,  $r - i$ ) built on the uncertainties measured for spectroscopic galaxies helps to identify insecure estimates based on the position of a galaxy in this grid. The accuracy of these photometric redshifts makes them a reference in machine learning based on photometric measurements. Last but not least, they are still the only ones available for comparison purposes.

### 2.2 The spectroscopic data

The  $\sim 1.5$ M spectroscopic redshifts, used as training labels, come from the SDSS, GAMA and BOSS spectroscopic surveys described below, matched to the SDSS DR16 photometric catalogue described above. Fig. 1 shows the redshift distributions of these three surveys and Fig. 2 their respective magnitude/redshift and  $(u - g)/(g - r)$  colour distributions. A sky map of the spectroscopic data is shown in Appendix A.

### 2.2.1 The SDSS survey

The SDSS spectroscopy is nearly complete to  $r = 17.8$ , totalling  $\sim 660\text{k}$  galaxies, but reaches fainter magnitude to much lower completeness with targeted populations, adding  $\sim 162\text{k}$  galaxies with  $17.8 < r < 20$ . We use the specific star formation rates (sSFRs) derived from the SDSS spectra by Brinchmann et al. (2004) in the data release 12 (DR12; Alam et al. 2015) to draw an empirical separation between blue, star-forming galaxies and red, passive galaxies in the redshift/( $u - g$ )/( $g - r$ ) space (Appendix B). This separation is used to balance the training samples and evaluate the performance of the training experiments in the two populations.

### 2.2.2 The GAMA survey

The GAMA survey (Driver et al. 2009, 2011; Liske et al. 2015) is a joint European-Australian spectroscopic survey combining UV to FIR photometric data from several ground- and space-based programs, including SDSS. The spectroscopy was carried out using the 2dF/AAOmega multi-object spectrograph on the Anglo-Australian Telescope, building on previous spectroscopic surveys such as SDSS, the 2dF Galaxy Redshift Survey and the Millennium Galaxy Catalogue (MGC). We use the 4 equatorial fields (G02, G09, G12, and G15) available in the data release 3 and 4 (DR4, Driver et al. 2022), covering a total of  $\sim 235 \text{ deg}^2$ . The spectroscopy is 98 per cent complete to  $r = 19.8$  (Liske et al. 2015), except for the G02 field where only the region north of Dec  $\sim 6\text{deg}$  was observed to high completeness. This provides us with a sample of  $\sim 210\text{k}$  spectroscopic galaxies at  $r < 20$  (90 per cent of which at  $r > 17.8$ ) matched to the photometric catalogue. GAMA constitutes the main component of our training set as its completeness makes it most representative of the photometric data set.

### 2.2.3 The BOSS survey

We retrieved an additional 486k spectroscopic sources from, essentially (98.2 per cent), the BOSS survey (Dawson et al. 2013). These are dominated by Luminous Red Galaxies (LRG). We will refer to this sample as ‘BOSS’, although a small contribution ( $\sim 8538$  galaxies) comes from other deep redshift surveys, namely: VVDS Wide and Deep (Le Fèvre et al. 2013), DEEP2 (DR4, Newman et al. 2013), VIPERS (DR2, Scodreggio et al. 2018), UDSz (Bradshaw et al. 2013; McLure et al. 2013), zCOSMOS-bright (Lilly et al. 2007).

## 3 THE CNN

### 3.1 Architecture

The present network is a more complex version of the P19 CNN, intended for a more complex data set. Its architecture is diagrammed in Figs 3 and 4. Fig. F1 in Appendix F lists all the layers with their type, shape, number of parameters and the layer(s) they are connected to upstream. We refer to P19 for a pedagogical description of the role played by the different types of layers. Their fig. 4 can also be compared to our Fig. 3.

As in P19, the input data consist of images of galaxies in the five *ugriz* bands of the SDSS surveys, with the galactic extinction along the line of sight added downstream before the fully connected layers (Section 3.2). The training labels are spectroscopic redshifts.

Here too, we choose to handle the redshift estimation task by means of a classification rather than, but aided by, a non-linear regression. The gain of using a classification rather than a regression or of adding a regression to the classification is negligible with a

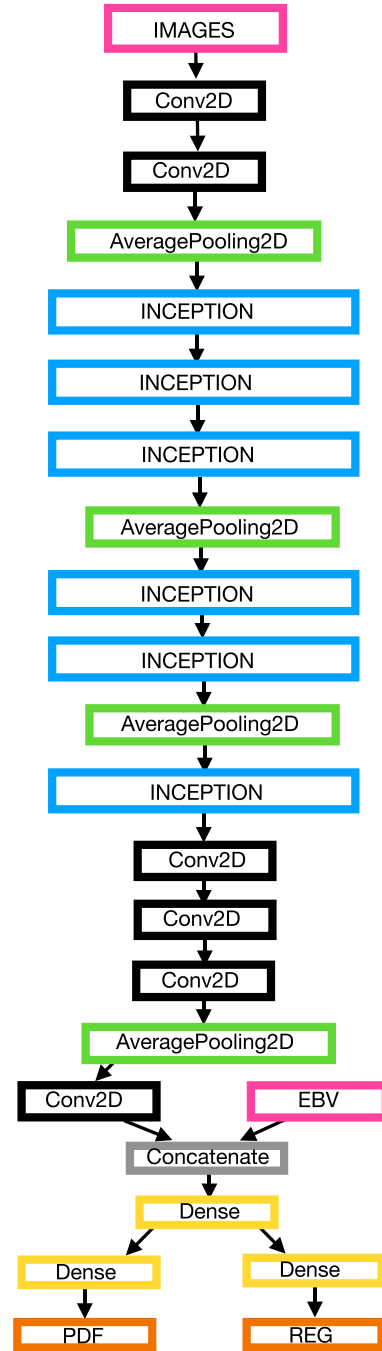


Figure 3. The CNN architecture. The inception module is detailed in Fig. 4.




Figure 4. Inception module.

rich data set such as the bright SDSS used by P19, but it proves more significant with sparser training sets such as the present data (Appendix E1) and even more so with high-redshift data (Ait Ouahmed et al. 2023).

The classes correspond to narrow, mutually exclusive, redshift bins, i.e. each training galaxy belongs to a single class (one-hot encoding of the spectroscopic redshifts). The classifier is trained using the softmax cross-entropy loss function (Baum & Wilczek 1987; Solla, Levin & Fleisher 1988) and the regression with the root mean square error. The two loss functions are simply added (after trying different weighted sums).

A softmax activation function (Bridle 1990) applied to the output layer of the classifier normalizes the outputs to 1. Such outputs were shown, both theoretically and experimentally, to provide good estimates of the posterior probability of classes in the input space (Richard & Lippmann 1991; Rojas 1996) provided the network is sufficiently complex and properly trained. Whether that is the case here may be questioned but we find the classification outputs to be very useful probability density function (PDF) proxies. We will call them ‘PDF’.

Compared to P19, the present network has 1 additional convolutional layer upstream, 6 inception blocks of similar complexity instead of 4 + 1 simpler one, and 3 additional convolutional layers following the last inception module. These have no padding and are followed by an average pooling layer, which reduces the number of trainable parameters to  $\sim 7\text{M}$  compared to  $\sim 27\text{M}$  in P19. Apart from the softmax activation function used in the last dense layer to produce the ‘PDFs’, all but two of the non-linear activation functions (introducing non-linearity into the network) are the commonly used ReLU (Rectified Linear Unit, Fukushima 1969). The second convolutional layer and the last dense layer before the regression output use a hyperbolic tangent activation function, which clips the dynamic range. The network is trained with the Adam optimizer (Kingma & Ba 2015), a stochastic gradient descent method based on the adaptive estimation of first-order and second-order moments.

We note that the results presented below are quite robust to the specifics of the CNN architecture. Replacing the inception blocks by simple convolutions while retaining the same depth and number of trainable parameters, only slightly degrades the metrics. The gain from the inception modules is of the order of that found with averaging a large number of models (Fig. 8): it takes averaging a large number of trained networks without inception blocks to achieve results similar to no averaging using inception blocks.

### 3.2 Training input

The CNN is fed  $64 \times 64$  pixel image cutouts centered on the galaxy coordinates in the five *ugriz* SDSS bands, to which we add the Schlegel, Finkbeiner & Davis (1998) galactic extinction value along the line of sight. The much larger number of galaxies and the different samples involved in the analysis compared to P19 requires a more efficient approach for generating the cutouts. Instead of extracting and resampling individual SDSS frames for every galaxy, we adopt the following procedure:

(i) Using SWARP (Bertin et al. 2002), we first re-project at once the whole SDSS imaging survey on a grid of 27 070 overlapping tiles covering the survey footprint in the five filters, relying on the WCS parameters of the input image headers (Calabretta & Greisen 2002) for the astrometry. The number of input frames contributing to a given output pixel ranges from one or two for ‘regular’ SDSS images, to 64 for some of the galaxies in Stripe 82. Each output tile

is  $18\,192 \times 18\,192$  pixels wide ( $2^\circ \times 2^\circ$  with  $0.396''$  pixels), and is aligned with the local North-South axis using the ZEA (zenithal equal area) projection. 5 arcmin overlaps between nearest neighbours at mid-width/height guarantee that any cutout is entirely contained in at least one of the tiles.

(ii) Galaxy sample requests are organized by tile, and  $64 \times 64$  pixel cutouts are extracted without resampling around every projected galaxy position.

While the new procedure is more than two orders of magnitude faster than that of P19, the generated cutouts are not centered as precisely (up to half-a-pixel only), and the direction to the local north is not always as perfectly aligned with the vertical axis of the pixel grid. However, we did not find these changes to have a measurable impact on the quality of the inferred redshifts. A (64,64,5) data cube is thus produced for the  $\sim 15.3\text{M}$  sources in the DR16 photometric catalogue.

### 3.3 Output assessment

#### 3.3.1 ‘PDF’

We use several tests and quantities designed to evaluate PDFs to assess the behaviour of our ‘PDFs’:

(i) The Probability Integral Transform statistic (PIT, Dawid 1984) is based on the histogram of the cumulative probabilities at the true value ( $\text{CDF}_i = \sum_{z \leq z_i} \text{PDF}_i(z)$  for galaxy  $i$  at spectroscopic redshift  $z_i$ ). A flat PIT distribution is expected from well calibrated PDFs, whereas convex or concave distributions point to over or underconfident PDFs (Polsterer, D’Isanto & Lerch 2022). Indeed excessively narrow (overconfident) PDFs will miss the target too often, overproducing PIT values close to 0 or 1, whereas PDFs that are too wide (underconfident) will encompass the true redshifts more often than not, overproducing intermediate PIT values.

(ii) The credibility test proposed by Wittman, Bhaskar & Tobin (2016) (hereafter WBT) is based on the cumulative distribution of the ‘threshold credibilities’, defined as the cumulative probabilities equal to or above the probability at the true value ( $c_i = \sum_{\text{PDF}_i(z) \geq \text{PDF}_i(z_i)} \text{PDF}_i(z)$ ), i.e. the smallest credible interval (CI) in which the spectroscopic redshift of a galaxy lies. With well calibrated PDFs, 1 per cent of the galaxies have their spectroscopic redshift within their 1 per cent CI, 2 per cent within their 2 per cent CI, etc., which translates into the cumulative distribution of  $c_i \leq c$  being equal to  $c$ .

(iii) The Continuous Ranked Probability Score (CRPS, well known in meteorological predictions, Hersbach 2000) is a quadratic measure of the difference between the forecast cumulative distribution function (CDF) and the empirical CDF of the observation (Zamo & Naveau 2018), here a unit step function around the spectroscopic redshift ( $\text{CRPS}_i = \int_{-\infty}^{z_i} \text{CDF}_i(z)^2 dz + \int_{z_i}^{+\infty} (\text{CDF}_i(z) - 1)^2 dz$ ). It can be viewed as a generalization of the MAE to distributional predictions.

(iv) We quantify the uncertainty of the ‘PDFs’ by the width of the 68 per cent central credible interval, i.e. the redshift width encompassing 68 per cent of the distribution after chopping off the left and right wings in equal measure.

(v) Other important aspects of the redshift estimate contained in the full shape of the ‘PDF’ can also be estimated, e.g. skewness and multi-modality.

Although passing one test does not ensure PDF quality (Amaro et al. 2019) nor, for that matter, PDF status, we expect several

successful tests combined with measures of the accuracy of the point estimates (next section) to provide some level of photometric redshift reliability.

### 3.3.2 Point estimates

Although redshift PDFs may be directly incorporated into Bayesian schemes in certain cosmological studies (e.g. inferring cosmological parameters), point estimates are required for many others (especially if the ‘PDFs’ are not true PDFs). We consider the following 4 photometric redshift estimators: i/ the weighted mean ( $z_{\text{mean}}$ ), ii/ the median value ( $z_{\text{med}}$ ) and iii/ the peak value ( $z_{\text{peak}}$ ) of the ‘PDF’, and iv/ the regression output ( $z_{\text{reg}}$ ). The metrics we use to assess the accuracy of these point estimates are identical to P19. We define:

- (i) the **normalized residuals**  $\Delta z_{\text{norm}} = (z_{\text{cnn}} - z_{\text{spec}})/(1 + z_{\text{spec}})$
- (ii) the **prediction bias**  $\langle \Delta z_{\text{norm}} \rangle$  (mean of the residuals)
- (iii) the **deviation**  $\sigma_{\text{MAD}} = 1.4826 \times \text{MAD}$ , where MAD (Median Absolute Deviation) is the median of  $|\Delta z_{\text{norm}} - \text{Median}(\Delta z_{\text{norm}})|$
- (iv) the **fraction  $\eta$  of outliers** with  $|\Delta z_{\text{norm}}| > 0.05$

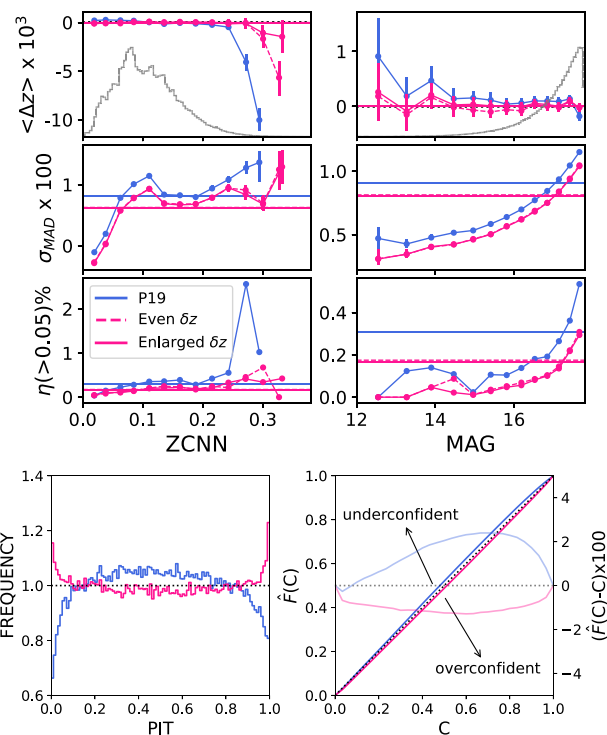
## 4 TRAINING EXPERIMENTS

### 4.1 SDSS at $r \leq 17.8$

For the purpose of comparison with P19, we first train the CNN with the same data set:  $\sim 510\text{k}$  SDSS galaxies at  $r < 17.8$ , the same binning: 180 redshift classes in the range  $0 < z < 0.4$  with constant width  $\delta z = 0.4/180$ , and the same protocol as P19: the database is split into 5 cross-validation samples, each one used in turn for testing while the remaining 80 per cent, augmented with randomly flipped and rotated images.<sup>1</sup>, is used for training (other training parameters are given in Appendix F). This operation is repeated 5 times with randomly initialized weights. The final ‘PDFs’ are the average of the 5 classification outputs.

Fig. 5 shows how this network compares with P19 (pink versus blue lines). The top panels show the bias,  $\sigma_{\text{MAD}}$  and outlier fraction  $\eta$  as a function of CNN  $z_{\text{mean}}$  and magnitude. All metrics are significantly reduced. However, the dashed pink line in the top left panel shows that the new CNN is plagued with a redshift ceiling effect similar to P19, manifest as a steep drop in the bias at  $z_{\text{mean}} \sim 0.3$ , above which no  $z_{\text{mean}}$  is predicted despite spectroscopic redshifts reaching higher values. We found that this effect could be mitigated by enlarging the bins at the highest, underpopulated redshifts, so that each class contains at least 20 training galaxies ( $\delta z = 0.0022$  at  $z < 0.33$ , increasingly larger at  $z > 0.33$ , adding up to 158 bins). The result is shown as the solid pink line. Sparse sampling may not be the only reason for this issue. The  $z \gtrsim 0.3$  tail of the SDSS data at  $r < 17.8$  is entirely populated by red galaxies and these are affected by a colour degeneracy at this particular redshift, which distorts their CNN redshift distribution (see Section 5.1). The lower panels of Fig. 5 show the PIT and WBT tests. According to both, the new ‘PDFs’ are slightly overconfident where P19’s were underconfident. Their mean CRPS is 0.0060, versus 0.0067 for P19.

The performance of the various redshift estimators are shown in Table 2, with the P19 results in parenthesis. P19 had used  $z_{\text{mean}}$  as optimal point estimate as it minimizes the bias and the rate of



**Figure 5.** Comparison between the P19 results (in blue) and the present work (in pink) for the SDSS sample at  $r \leq 17.8$ . Top panels: bias,  $\sigma_{\text{MAD}}$ , and catastrophic failures as a function of  $z_{\text{mean}}$  and magnitude. The dashed and solid pink lines show the present metrics before and after expanding the binning at  $z \gtrsim 0.33$ . Bottom panels: the PIT and WBT tests. The PIT distribution is expected to be flat, the WBT test is expected to follow the unity line. The departures from the unity line are overlaid in faded colours, with units along the right-hand side y-axis ( $10^{-2}$ ).

**Table 2.** Performance comparison between the four redshift estimators for the SDSS trainings at  $r < 17.8$ . The P19 statistics are in parenthesis.

	$z_{\text{mean}}$	$z_{\text{med}}$	$z_{\text{peak}}$	$z_{\text{reg}}$	unit
$\sigma_{\text{MAD}}$	808 (908)	800 (902)	818 (918)	810	$10^{-5}$
$\langle \Delta z_{\text{norm}} \rangle$	3 (4)	-31 (-43)	-87 (-125)	-27	$10^{-5}$
$\eta(> 0.5)$	0.17 (0.31)	0.18 (0.31)	0.27 (0.39)	0.17	%

catastrophic failures without significantly degrading  $\sigma_{\text{MAD}}$ . However, each point estimate has pros and cons:  $z_{\text{med}}$  tends to minimize the dispersion at the expense of the bias,  $z_{\text{peak}}$  also optimizes the dispersion but maximizes the catastrophic failures and generates very noisy redshift distributions, while  $z_{\text{reg}}$  is a slightly degraded  $z_{\text{mean}}$ .

These statistics marginally outperform the two recent attempts at improving the P19 performance. Hayat et al. (2021) proposed a self-supervised representation learning method in which a network was pre-trained with 1.2M unlabelled SDSS galaxies, then fine-tuned with the labelled data. While the great potential of such methods is undeniable, in particular at high redshifts where spectroscopy is very sparse, the present supervised training technique remains competitive with the SDSS data. Dey et al. (2021) presented a deep capsule network that jointly estimates the redshift and the basic morphological type of galaxies (spiral/elliptical). Its backbone consists of a primary convolutional layer followed by Conv-Caps (‘capsule’) layers, composed of multiple neurons that compute not

<sup>1</sup>Note that these data augmentation processes can have an impact on the training if the image transfer function is not isotropic, e.g. if differential chromatic refraction is not negligible in the data.

only the presence or absence of a feature, but also its properties such as rotation, size, velocity, or colour. Although this design is robust and invariant to image orientation, data augmentation techniques such as rotation and flipping were used during training. The latent space of the network has only 16 dimensions, which helps to study its interpretability. However, compared to classical CNNs, capsule networks are trickier to train and can hardly be adapted to deeper architectures for more complex tasks.

#### 4.2 SDSS+GAMA + BOSS at $r \leq 20$


We expand the training to  $r = 20$  using the deeper spectroscopy (Section 2.2). Randomly splitting the full spectroscopic data for training/validation as was done for the SDSS at  $r < 17.8$  (e.g. 80 per cent/20 per cent) is not appropriate here, as the sample is not at all representative of the galaxies expected to populate the photometric sample. While SDSS and GAMA are representative sets of galaxies at  $r \leq 17.8$  and  $17.8 \leq r \leq 19.8$ , respectively, given their large redshift completeness, the bright SDSS data is over-represented compared to the faint GAMA data. Both surveys also feature strong local structures in their redshift distributions (Fig. 1). Last but not least, the large BOSS sample is overpopulated with LRGs at  $z_{\text{spec}} \gtrsim 0.2$ , which really are quite rare compared to ‘normal’ red galaxies at  $r < 20$ . Over-representing a specific population in a specific redshift interval will bias the predictions in favour of those redshifts. The results will be deceptively good on the validation sample as it matches the training sample by design, but suboptimal on a different population, which the photometric sample is expected to be. These effects are quantified in Appendix E3. For these reasons we attempt to create training and testing samples at least roughly representative of the general population at  $r < 20$ .

##### 4.2.1 Test and training samples

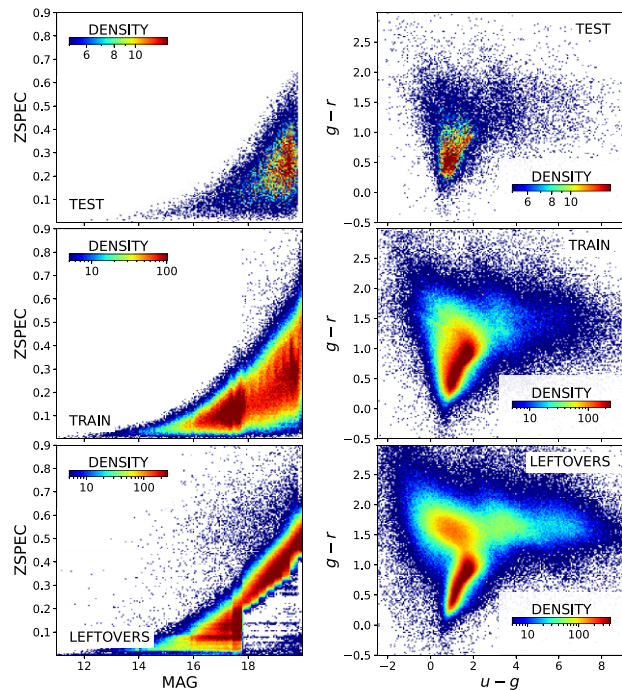
We use SDSS and GAMA as models of the Universe in their respective magnitude range of completeness to create mock samples out of the full spectroscopic data. To do so, we model the redshift distributions of blue and red galaxies in bins of magnitude with smooth ad hoc functions; in each magnitude bin, we randomly extract from the full spectroscopic sample subsets of blue and red galaxies with redshift distributions matching these smooth distributions. The magnitude and redshift bin widths are chosen to strike a compromise between mock resemblance and size.

We first create a GAMA-like test sample using the above method but avoiding GAMA itself in the random extraction in order to sample the full sky coverage. We match the blue/red ratio and the magnitude distribution of GAMA by randomly extracting from a blue and a red redshift-matched sample. The result is a test sample of 25 856 galaxies at  $r \leq 19.8$  resembling GAMA but with much reduced redshift structures and spanning the entire sky. Its magnitude, redshift, and colour distributions are shown in Fig. 6 and in the top panels of Fig. 7. Galaxy properties are not limited to magnitude, redshift and red/blue type, but more sophisticated methods taking more parameters into account (e.g. SOM) would reduce the size of the mock sample too significantly, defeating its purpose.

We subtract this test sample from the spectroscopic sample and proceed to extract training samples using the above method without excluding GAMA and adding galaxies at  $19.8 < r < 20$ , also based on GAMA but more loosely as it is less representative in this range. The SDSS and GAMA-like subsets are concatenated without matching the GAMA number counts, which would deplete the bright end too



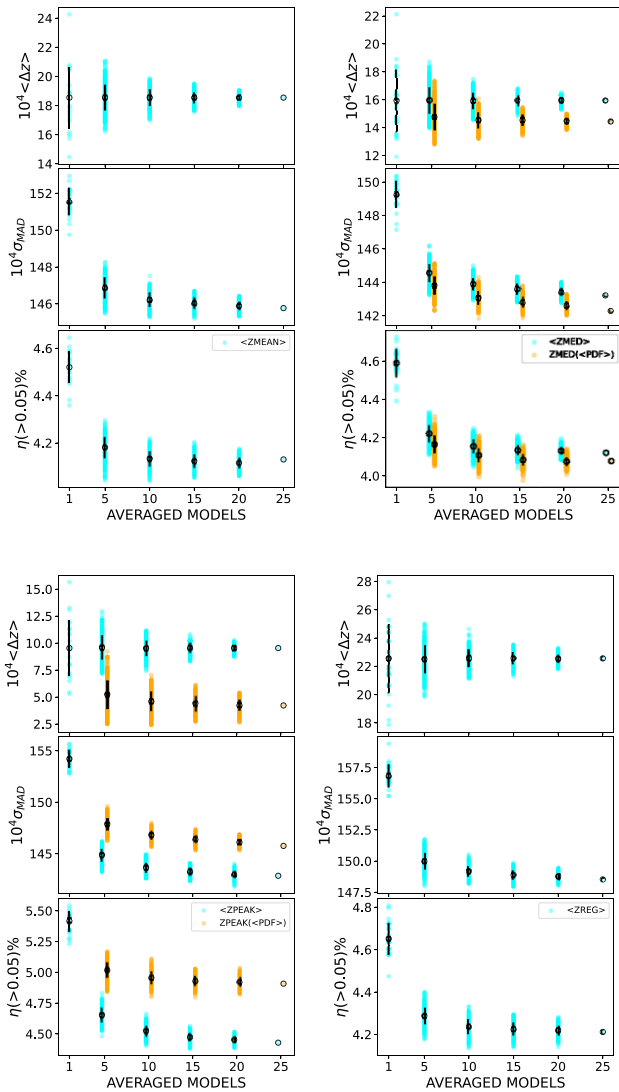
**Figure 6.** The normalized redshift and magnitude (log scale) distributions of the GAMA sample (grey), of the full spectroscopic sample (orange), and of the GAMA-like test sample (pink) for blue and red galaxies (top and bottom panels respectively).



**Figure 7.** Magnitude/redshift distributions (left panels) and  $(u - g)/(g - r)$  distributions (right panels) of the test, training, and leftover samples described in Section 4.2.1.

drastically. We create in this way 25 training samples of  $\sim 370\text{k}$  galaxies, significantly smaller than throwing in the full GAMA and SDSS samples but doing so propagates unwanted training features in the predicted redshift distributions. The 25 samples total  $\sim 910\text{k}$  unique spectroscopic galaxies and leave out a sample of  $\sim 580\text{k}$  galaxies, half of them at  $r < 17.8$ , the other half dominated by LRGs (Section 5.5). The intersection between any 2 training samples is between 45 and 80 per cent. The intersection of all 25 samples amounts to 78 682 sources, most of them in the troughs of the GAMA redshift distribution, whose features reappear when combining any 2 samples. Fig. 7 shows the magnitude/redshift and  $(u - g)/(g - r)$  distributions of the test, training and leftover samples.

Another strategy, leaving a larger sample of leftovers available for testing, would be to train just a few training samples several times. Using just one training sample for instance, picked at random, leaves nearly twice as many galaxies for testing. These however, are, by




**Figure 8.** Bias, deviation, and rate of catastrophic failure of the averaged point estimates (blue) and of the corresponding point estimate of the averaged PDFs when relevant (orange), as a function of the number of CNN outputs being averaged for the test sample. From top left to bottom right:  $z_{\text{mean}}$ ,  $z_{\text{med}}$ ,  $z_{\text{peak}}$ , and  $z_{\text{reg}}$ . The black circles and vertical lines show the mean and standard deviation of the distributions.

design, as un-representative as the smaller sample (bright galaxies and faint red ones) and thus do not provide a more informative testing opportunity. The strategy is explored in Appendix E2.

#### 4.2.2 Averaging

We train the CNN with each of the 25 samples in full (without cross-correlations) using 400 redshift bins between  $z = 0$  and 0.9, corresponding to a bin width  $\delta z = 0.00225$  similar to the SDSS training at  $r < 17.8$ . Other training parameters are given in Appendix F. From the 25 trained networks, we infer ‘PDFs’ and regression values for the test sample.

The benefits of averaging the outputs of several networks are shown in Fig. 8, where the bias, deviation and rate of catastrophic failure of the averaged  $z_{\text{mean}}$ ,  $z_{\text{med}}$ ,  $z_{\text{peak}}$  and  $z_{\text{reg}}$  (in blue) and of the point estimates of the averaged ‘PDFs’ when relevant (in orange),



**Figure 9.** A random sample of outputs in increasing order of spectroscopic redshift. The outputs of the 15 models are shown in grey, with their respective median values marked as coloured vertical lines. The averaged outputs (the final ‘PDF’) are shown in black.

are plotted as a function of the number of models being averaged. The 25 points at  $N = 1$  illustrate the variance between the 25 models. In between these points and the final averages of the 25 models, we randomly picked 400 combinations of  $N$  different models among the tens of thousands of possibilities (several million between 10 and 15). The black circles and vertical lines show the mean and standard deviation of the coloured points.

For all the point estimates,  $\sigma_{\text{MAD}}$  gains the most from averaging, while the bias is quite insensitive to it. The median of the averaged ‘PDFs’ does consistently better than the averaged  $z_{\text{med}}$ . On the contrary the averaged  $z_{\text{peak}}$  do better for  $\sigma_{\text{MAD}}$  and  $\eta$  than the mode of the averaged ‘PDFs’, but a factor of 2 worse for the bias, making the latter method preferable in this case too. For all the point estimates, the gain from using a single model to averaging 5 is significant for  $\sigma_{\text{MAD}}$  and  $\eta$ : averaging any 5 models does better than the best model among the 25. Then the overlap becomes large and it becomes possible to do better with certain combinations of, say, 10 than of more models, including the final average of 25. However, the means and standard deviations show that it is unlikely, and also that 25 is an overkill.

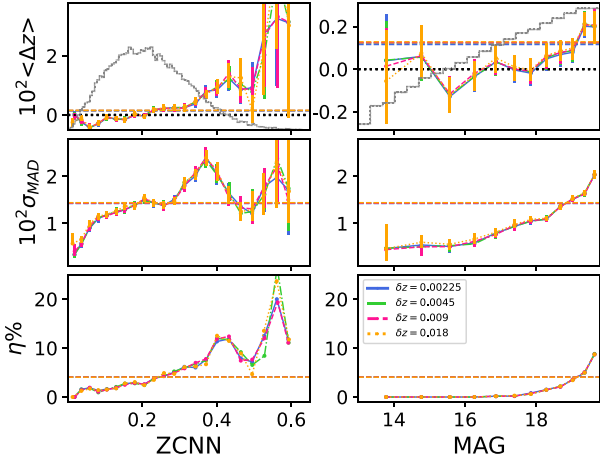
In all that follows, we use 15 training samples randomly selected among the 25, a reasonable compromise between performance and computing time. This choice leaves unchanged the total number of spectroscopic galaxies used for training and left over. We compute the point estimates of the averaged ‘PDFs’ rather the average of the point estimates. The uncertainty on a given metric may be estimated from the standard deviations: they are of order  $3 \times 10^{-5}$  for  $\sigma_{\text{MAD}}$  and the bias, and of order 0.03 for  $\eta$ .

The benefit of averaging ‘PDFs’ is further illustrated in Fig. 9, which shows random examples of outputs in increasing order of spectroscopic redshift. The 15 CNN outputs are shown in grey, with their respective  $z_{\text{med}}$  marked as coloured vertical lines. The averages of the 15 outputs are plotted in black. We interpret the fact that these consistently provide better redshift estimations than the individual ones, which are themselves consistent with one another, as convergence towards true PDFs. At any rate, we remove the quotes to alleviate the text.


#### 4.2.3 Redshift binning

In addition to the binning used above with  $\delta z = 0.00225$ , we train the 15 samples using 200, 100 and 50 bins between  $z = 0$  and 0.9, corresponding to  $\delta z = 0.0045, 0.009$  and 0.018. Fig. 10 shows how little sensitive the point estimates ( $z_{\text{med}}$  here) are to the PDF resolution. The bias,  $\sigma_{\text{MAD}}$ , and rate of catastrophic failure, plotted as a function of  $z_{\text{med}}$  and magnitude, are undistinguishable from the





**Figure 10.** Test sample metrics as a function of  $z_{\text{med}}$  and magnitude for the four classification schemes. The grey histograms in the top panels are the shapes of the spectroscopic redshift and magnitude distributions, in log scale for the latter. The metrics are very little sensitive to the choice of classification binning within the explored range.




**Figure 11.** The DIP score distribution of the test sample for the four binning experiments.


smallest to the largest binning, suggesting that too fine a binning is superfluous.

We use the DIP test (Hartigan & Hartigan 1985) to detect multimodality in the different sets of PDFs. It consists of measuring the maximum distance at any point between the CDF and the closest uni-modal CDF, a uni-modal distribution having a score of 0 by definition. Fig. 11 shows the significant impact of the bin width on the distribution of DIP scores in the test sample. Since the bin width has, on the contrary, very little impact on the point estimates in the range we tested, we conclude that the many spikes generated by the small bins do not represent meaningful multimodalities but simply the incapacity of the CNN to classify redshifts in such fine a grid. Increasing the bin width makes the PDFs increasingly uni-modal but would eventually degrade the predictions significantly, to the point of no prediction at all in the extreme case of only 1 bin. There is therefore an optimal resolution for the classifier, given the training data.

Fig. 12 shows the PIT and WBT tests for the 4 binning scenarios. Both are nearly perfect but the latter is more discriminatory. It shows that  $\delta z = 0.009$  is close to the optimal resolution. Smaller bins produce ‘overconfident’, overly spiky PDFs, larger bins produce ‘underconfident’, underinformative ones. The  $\delta z = 0.009$  binning is consequently our final choice in the rest of this work. The mean CRPS for this binning and the two smaller ones is  $\sim 0.013$ . It is slightly higher ( $\sim 0.015$ ) for the largest binning.



**Figure 12.** PIT and WBT statistics for the four binning scenarios. The departures from the unity line in the WBT test are overlaid in faded colours, with units along the right-hand side y-axis ( $10^{-2}$ ).




**Figure 13.** The point estimate distributions of blue and red test galaxies, with their stacked PDFs (the missing regression value is very similar to  $z_{\text{mean}}$ ). The red galaxy distribution is strongly distorted around  $\sim 0.35$ , to a variable degree depending on the point estimate. This is attributed to a colour degeneracy (see Section 5.1).

## 5 TEST RESULTS


### 5.1 Colour dependence

Fig. 13 shows the distribution of the different point estimates for the blue and red galaxies in the test sample, as well as the stacked PDF, the spectroscopic redshift distribution and the B16 distributions (the regression value is not displayed to limit the clutter and because it is very similar to  $z_{\text{mean}}$ ). The strong distortion around  $z_{\text{spec}} \sim 0.35$  for the red galaxies can be attributed to a degeneracy in their optical colours. It is more or less severe depending on the point estimate:  $z_{\text{mean}}$  (and  $z_{\text{reg}}$ ) generate the strongest distortion within the smoothest distributions,  $z_{\text{peak}}$  minimizes the distortion but generates the noisiest (discretized) distributions,  $z_{\text{med}}$  is the best compromise considering both the blue and red populations. The stacked PDF best fits the  $z_{\text{spec}}$  distribution of both blue and red galaxies. The B16 redshifts are significantly more distorted, probably because they were computed from the measured colours while the CNN captures more information from the full images.

The left panel of Fig. 14 shows the  $(g - i)$  colour versus  $z_{\text{spec}}$  distribution of red galaxies in the spectroscopic sample. While colour and redshift are well correlated at low redshift, the relation flattens out at  $z_{\text{spec}} \gtrsim 0.31$ , making it a harder task to predict redshifts in this range. The right panel displays the  $(g - i)$  colour versus  $z_{\text{med}}$  distribution of red galaxies in the test sample, colour-coded by the mean PDF width. The overlaid black line is the shape of the  $z_{\text{med}}$  distribution (also displayed in Fig. 13). The strong distortion around  $\sim 0.35$  coincides with increased PDF widths, i.e. increased uncertainties in the classification. As shown in Fig. 13, stacking these wider PDFs nearly suppresses the point estimate distortion.



**Figure 14.** Left:  $(g - i)$  colour against  $z_{\text{spec}}$  for the red galaxies in the full spectroscopic sample. The colour clearly becomes indiscriminate at  $z_{\text{spec}} \gtrsim 0.31$ . Right: the  $(g - i)$  versus  $z_{\text{med}}$  distribution of the red galaxies in the test sample colour-coded by the PDF width. The black profile is their  $z_{\text{med}}$  distribution. The mislocated galaxies in the degenerate interval have wider than average PDFs, reflecting the greater uncertainty in their classification.



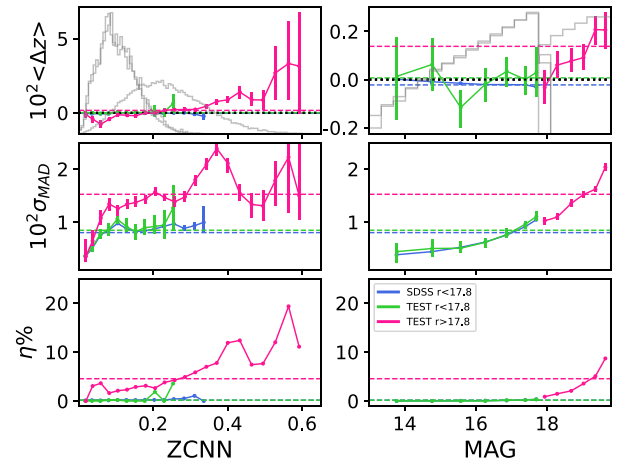
**Figure 15.** PIT and WBT statistics for the blue and red subpopulations in the  $\delta z = 0.009$  scenario. The PIT distribution highlights the red population bias. However, both the PIT distribution and WBT test for the full sample are near perfect. The departures from the unity line in the WBT test are overlaid in faded colours, with units along the right-hand side y-axis ( $10^{-2}$ ).

Fig. 15 shows the PIT and WBT tests for the blue and red subpopulations and the full sample. The PIT distribution reveals a large positive bias for the red galaxies, whose PDFs are excessively to the right of their true redshift. This is indeed the case for the many galaxies shoved into the CNN redshift distortion from lower  $z_{\text{spec}}$ . The WBT test is very close to the unity line for both galaxy types, and closer still for the whole population. For comparison with traditional  $1\sigma$  coverage tests: 67.3 per cent of the test galaxies have their smallest credible interval defined by their spectroscopic redshift smaller than or equal to 68 per cent, while 64.6 per cent of the galaxies have their spectroscopic redshift within the 68 per cent central credible interval defining the PDF width independently of the spectroscopic redshift. The difference points to the fact that the PDFs are non-Gaussian.


## 5.2 Magnitude dependence

Fig. 16 shows how the performance on the test sample degrades, predictably, from  $r < 17.8$  to  $r > 17.8$ . The  $r < 17.8$  regime remains competitive with the bright SDSS training described in Section 4.1, despite the much lower number of galaxies in this magnitude range in the present training samples ( $\sim 132\text{k}$  versus  $\sim 414\text{k}$ ). This is expected from P19 who found that the performance remained virtually unchanged when the training sample size was reduced to  $\sim 100\text{k}$ .

P19 also found that performance could be improved by discarding sources with the largest PDF widths, reporting the result of rejecting the largest 10 per cent and 20 per cent. Here we design a more conservative, magnitude-dependent cut in PDF width meant to



**Figure 16.** Test sample metrics as a function of  $z_{\text{med}}$  and magnitude at  $r < 17.8$  and  $r > 17.8$ , with the SDSS MGS training performance for comparison (Section 4.1). The performance degrades significantly from bright to faint. The bright regime is similar to the SDSS MGS training despite the much lower number of training galaxies.



**Figure 17.** Magnitude-dependent threshold designed to exclude the largest 5 per cent of PDF widths in the test sample (red crosses). The dashed line is the analytical fit (equation 1), plateaued out below  $r = 15$  at a PDF width of  $\sim 0.02$ .

exclude the worst 5 per cent of PDF widths at a given  $r$ –band magnitude in the test sample. The computed threshold is shown in Fig. 17 as red crosses. We fit the trend with a 4<sup>th</sup> degree polynomial,  $\sum_{i=0}^4 a_{5 \text{ per cent}}[i] r^i$ , where:


$$a_{5 \text{ per cent}} = [1.26761997 \times 10^1, -3.30587750, 3.22381178 \times 10^{-1}, -1.39350947 \times 10^{-2}, 2.25736799 \times 10^{-4}] \quad (1)$$

capping it at bright magnitude at the  $r = 15$  value of  $\sim 0.02$ . The final threshold is the red dashed lines in Fig. 17. As shown in Table 3 (see the next section), excluding galaxies with PDF width above this threshold improves redshift quality at a minor cost (4.6 per cent of the test sample), while allowing for the expected increase in uncertainty with magnitude. Also expected is the underlying dependence with spectroscopic redshift. High-redshift galaxies not only tend to be fainter, they are also sparse in the training samples, hence their poorer outcome. The fraction of galaxies excluded by the threshold increases from 2 to 16 per cent between  $z = 0$  and 0.6.

We also note that, despite our attempt at matching representative magnitude-dependent redshift distributions, this limit inevitably carries biases. But by enveloping the spread of PDF widths of a sample of spectroscopic galaxies whose image quality is on average more reliable than in the full photometric SDSS survey (Section 6), it is nevertheless a tool of quality control.

**Table 3.** The CNN performance on the GAMA-like test sample, in units of  $10^{-5}$  for  $\sigma_{\text{MAD}}$  and  $\langle \Delta z_{\text{norm}} \rangle$ , and in per cent for  $\eta (> 0.05)$ . The **B16** performance is reported under  $z_{\text{B16}}$ . The numbers and second lines in parentheses are the results of excluding galaxies above the PDF width threshold (equation 1).

	Blue					Red					All				
N	15 120 (14 463)					10 736 (10 192)					25 856 (24 655)				
$z_{\text{phot}}$	$z_{\text{mean}}$	$z_{\text{med}}$	$z_{\text{peak}}$	$z_{\text{reg}}$	$z_{\text{B16}}$	$z_{\text{mean}}$	$z_{\text{med}}$	$z_{\text{peak}}$	$z_{\text{reg}}$	$z_{\text{B16}}$	$z_{\text{mean}}$	$z_{\text{med}}$	$z_{\text{peak}}$	$z_{\text{reg}}$	$z_{\text{B16}}$
$\sigma_{\text{MAD}}$	1446	<b>1403</b>	1455	1474	3058	1474	<b>1443</b>	1466	1503	2853	1466	<b>1421</b>	1460	1481	3062
	(1392)	<b>1355</b>	1403	1413	3002	1400	<b>1382</b>	1415	1421	2734	1402	<b>1367</b>	1406	1422	2987)
$\langle \Delta z \rangle$	-46	-78	-163	<b>-22</b>	-59	559	479	<b>339</b>	581	2706	205	153	<b>45</b>	228	1089
	(-77)	(-99)	(-128)	<b>(-66)</b>	(-59)	443	373	<b>281</b>	459	2572	138	97	<b>41</b>	151	1031)
$\eta(\%)$	<b>3.7</b>	<b>3.7</b>	4.76	3.81	15.97	4.66	<b>4.41</b>	4.96	4.81	21.36	4.1	<b>3.99</b>	4.84	4.22	18.2
	<b>(2.75)</b>	2.8	3.56	2.77	15.3	<b>3.</b>	3.01	3.61	3.08	19.85	<b>2.86</b>	2.88	3.58	2.9	17.18)



**Figure 18.** Comparison of the point estimates (in percentages):  $z_{\text{mean}}$  is very similar to the regression value but positively skewed with respect to  $z_{\text{med}}$ . The mean difference between any two point estimates is less than 0.5 per cent.


### 5.3 Point estimate summary

Fig. 18 shows the distributions of the discrepancies between different point estimates. The PDF mean tends to be larger than the mode, which tends to be larger than the median (positive skewness), however, the mean difference between any 2 point estimates is less than 0.5 per cent. The regression value is very similar to the mean value.

The performance of the different point estimates for the blue, red and combined galaxies in the test sample are reported in Table 3, with the best scores in bold face. The metrics resulting from applying the PDF width threshold are reported in the second lines in parenthesis. The best dispersions are consistently achieved with  $z_{\text{med}}$ . The mean bias is a more shifting quantity since it may take negative values. Good scores are of the order of  $10^{-4}$ , poor ones above  $10^{-3}$ , as is the case for red galaxies. The rate of catastrophic failures remains below 5 per cent in all cases. Compared to **B16**, the precision is improved by a factor of  $\sim 2$  or more, the red galaxy bias and the rate of catastrophic failures by a factor of  $\sim 5$ .

### 5.4 Outliers

Fig. 16 shows that the fraction of catastrophic failures increases with magnitude and predicted redshift (it increases similarly with spectroscopic redshift), like all metrics. Table 3 shows that applying the PDF width threshold significantly reduces their fraction, however about two thirds remain within the range of other galaxies. Fig. 19, which displays the PIT intervals against  $|\Delta z_{\text{norm}}|$ , assuming  $z_{\text{med}}$ , reveals that most outliers, defined as being above the green dashed line ( $|\Delta z_{\text{norm}}| > 0.05$ ), tend to have PIT values close to 0 and 1 (equivalently, **WBT** credibility intervals close to 1). This means that their spectroscopic redshift tends to lie to the left or right of their PDF. Those with intermediate PIT values would be largely excluded by the PDF width threshold (excluded outliers are crossed




**Figure 19.** PIT intervals against  $|\Delta z_{\text{norm}}|$  assuming  $z_{\text{med}}$  (the convergence at  $\text{PIT} = 0.5$  and  $\Delta z_{\text{norm}} = 0$  arises from the very definition of  $z_{\text{med}}$ ). The grey scale is inverted to highlight the low-density outlier region above the green dashed line ( $|\Delta z_{\text{norm}}| > 0.05$ ). Outliers excluded by the PDF width threshold are crossed in red. All the others have PIT values close to 0 and 1, i.e. narrow PDFs missing the spectroscopic redshifts.

in red). Let's note that most cases of extreme PIT value are not outliers, they are expected from a flat PIT distribution. Outliers with narrow PDFs missing the spectroscopic redshifts cannot be easily identified. Their visual inspection in the five bands does not reveal any specific photometric defect, nor is there anything noteworthy in their spatial distribution. They simply are the tail of a continuous  $|\Delta z_{\text{norm}}|$  degradation occurring with increasing redshift and magnitude. Reducing it is likely to be difficult without a richer training set at high redshift.

### 5.5 Spectroscopic leftovers

The  $\sim 583k$  galaxies in the spectroscopic sample that belong neither to the test sample nor to any of the 15 training samples have a very bimodal distribution (see Fig. 7): one half at  $r < 17.8$  is a mix of red and blue galaxies left over from, mostly, the bright SDSS catalogue, the other half at  $17.8 < r < 20$  is mainly red galaxies. Fig. 20 shows the redshift distributions of these 2 subsets, with the  $z_{\text{med}}$  metrics before and after applying the PDF width threshold. In the bright interval, the excess of galaxies around  $z_{\text{spec}} \sim 0.08$  is significantly larger than in the full SDSS catalogue as the training samples were designed to avoid it. The CNN predictions are overly smooth compared to the true distribution, probably a counter bias from smoothing the training redshift distribution in order to avoid biases from such local structures. Nevertheless, the metrics remain close to the values derived in Table 2.

In the faint magnitude panel, the situation is less favourable. The sample is dominated by high-redshift LRGs that we chose to avoid in the training samples, and by red galaxies in the redshift interval of the colour degeneracy (Section 5.1). The distortion induced



**Figure 20.** The redshift distributions of leftover galaxies at  $r < 17.8$  and  $17.8 < r < 20$  (left and right panels, respectively). The grey shaded histograms are the spectroscopic redshift distributions;  $z_{\text{mean}}$  and  $z_{\text{med}}$  are in orange and green, respectively. The faint interval is dominated by LRGs in the BOSS sample, hence the strong CNN redshift distortion around  $\sim 0.35$ , here emphasizing an existing feature. The metrics reported in the top right corners are for  $z_{\text{med}}$ .

for these galaxies around  $\sim 0.35$  happens to emphasize an actual feature, also a leftover from the creation of the smooth training catalogues. Although this population largely differs from the test sample, the deviation and rate of catastrophic failures remain within the range of values derived in Table 3. However, the bias is much larger, and negative. The redshifts are visibly underestimated. The training samples were designed to represent ‘normal’ red galaxies within the magnitude and redshift range of the GAMA sample. The performance is poorer for these galaxies due to the colour degeneracy, but poorer still for LRGs that are purposely under-represented in the training samples compared to their overwhelming presence in the spectroscopic sample. Adding high-redshift LRGs to the training samples reduces the present bias but at the cost of increasing it in the test sample. This is shown in Appendix E3.

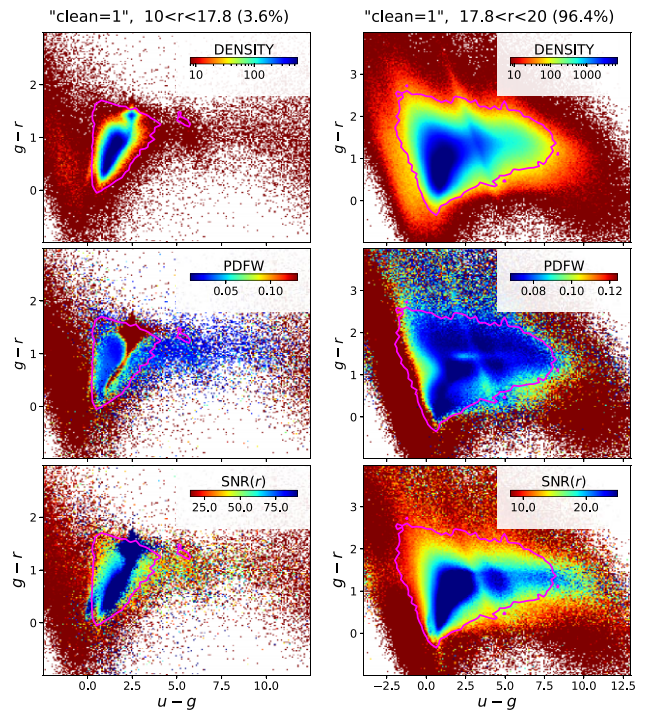
## 6 INFERENCE RESULTS

We infer the PDFs of the  $\sim 13.8\text{M}$  sources without spectroscopic redshift in the photometric sample (Section 2.1), split in two according to the SDSS keyword ‘clean’ referring to photometric quality.

The results are presented below for the ‘clean’ (**clean** = 1) sources ( $\sim 81.7$  per cent), and in Appendix C for the ‘dirty’ (**clean** = 0) sources.

### 6.1 The ‘clean’ sample

Fig. 21 shows the  $(u - g)/(g - r)$  colour distributions of the ‘clean’ sources at  $10 < r < 17.8$  and  $17.8 < r < 20$  (left and right panels, respectively). The colour-code indicates the density in the top panels, the mean PDF width in the middle panels and the mean  $r$ -band signal-to-noise ratio (SNR) in the bottom panels. The star sequence is conspicuous at  $r < 17.8$  with very poor PDFs, unsurprisingly since stars were not included in the training. Also unsurprisingly, the PDFs are very inconclusive in regions of the colour/colour plots not or ill represented in the training samples, the contours of which are shown in pink. These regions devoid of spectroscopy also have very poor SNR. They could be the locus of bona fide galaxy populations that were systematically missed as spectroscopic targets due to their low optical SNR but it seems more likely that their colours are wrong and their PDFs useless due to poor image quality. In any case, whether the CNN infers in uncharted territories where machine learning techniques are unable to perform or whether the input data




**Figure 21.** The  $(u - g)/(g - r)$  colour distributions of the ‘clean = 1’ photometric sample at  $r < 17.8$  and  $r > 17.8$  in the top left and right panels, respectively. The distribution are colour-coded by the mean PDF width in the middle panels and by the mean  $r$ -band SNR in the bottom panels. A bright magnitude, the star sequence stands out with the poorest PDFs and the highest SNR. The PDFs are predictably very inconclusive outside of the training contours (shown in pink), where the SNRs are also very poor.

are flawed, the PDF widths clearly signal worthless predictions. We use the threshold introduced in the previous section (equation 1) to discard them.

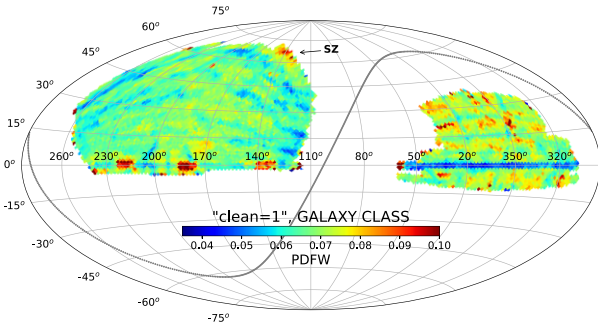
We also build a galaxy/star/QSO classifier. This is a CNN similar to the redshift classifier, with the same type of input data, into which we insert, at the output of the last convolution, two successive layers of 96 neurons and a final 3 neuron layer for the triple classification. We train it with 80 000 sources in each class, randomly extracted from the SDSS spectroscopic catalogue (type 3 and 6 objects were included for stars and QSOs). The results are cross-validated 5 times, with 80 per cent of the data used for training and 20 per cent for validation. Predicted classes assigned according to the highest probability yield completeness and purity scores<sup>2</sup> above 98 per cent for galaxies at all magnitudes (save for a glitch at  $r \sim 12.7$ ) as shown in Fig. 22. The two scores are actually above 99 per cent for red galaxies, and 97.6 per cent and 96 per cent respectively for blue galaxies. Only type 3 sources are considered here, as in the photometric catalogue, which leaves only 6194 stars and 7775 QSOs. (The scores are higher for type 6 stars and QSOs, except for the small fraction of bright,  $r < 18$  QSOs whose purity remains low).

Fig. 23 shows the PDF width-coloured sky map of the ‘clean’ photometric sources classified as galaxies. Redshift prediction quality is unevenly distributed. The ‘suspect zone’ (marked ‘sz’) identified by P19 in the SDSS at  $r < 17.8$  is visible in the northern region. Other

<sup>2</sup>Completeness is defined as the fraction of galaxies (stars or QSOs) that are correctly classified, and purity as the fraction of classified galaxies (stars or QSOs) that really are galaxies (stars or QSOs).



**Figure 22.** Completeness and purity scores of the galaxy/star/QSO classifier as a function of magnitude in the concatenated validation samples. Only type 3 sources are considered.

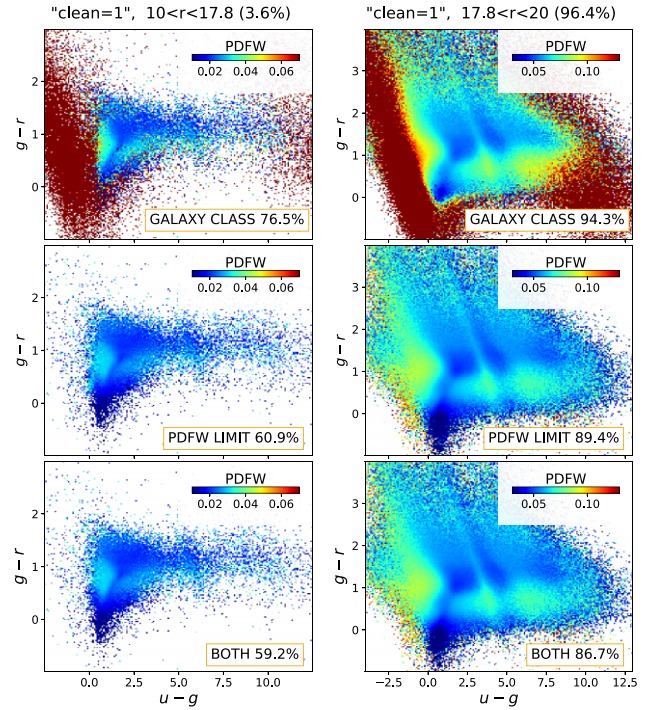


**Figure 23.** PDF width-coloured map of the sources classified as galaxies in the ‘clean = 1’ photometric samples. The ‘suspect zone’ identified by P19 in the SDSS at  $r < 17.8$  in the northern region is marked ‘sz’. Several similarly degraded patches are also shown in the southern region. The larger than average PDF widths in the equatorial GAMA regions are due to their deeper than average magnitude. The lower than average PDF widths in Stripe 82 (equatorial blue stripe) are due to better image quality.


similarly degraded patches also show in the southern region (many more are visible in the ‘clean = 0’ subsample shown in Fig. C2). On the other hand, the larger than average PDF widths in the GAMA regions are due to their deeper than average magnitude (nearly all GAMA sources with  $r < 19.8$  being part of the spectroscopic sample). The lower than average PDF widths in Stripe 82 (blue equatorial stripe in the South) are due to better image quality (P19).

Fig. 24 shows the  $(u - g)/(g - r)$  colour distributions of the sample colour-coded by the mean PDF width and split into the bright and faint magnitude intervals, following three cleansing procedures: the top panels are restricted to sources classified as galaxies (76.5 per cent at  $r < 17.8$ , 94.3 per cent at  $r > 17.8$ ), the middle panels to sources with PDF widths below the threshold (60.9 per cent at  $r < 17.8$ , 89.4 per cent at  $r > 17.8$ ), and finally the bottom panels to classified galaxies with PDF widths below the threshold (59.2 per cent at  $r < 17.8$ , 86.7 per cent at  $r > 17.8$ ). This final procedure rejects 14.3 per cent of the initial data (6.4 per cent classified as stars or QSOs, 11.6 per cent with PDF widths above the threshold).

The PDF width limit efficiently screens both the offending colour regions and most of the classified stars and QSOs. It also clears the red, ‘suspect zone’-like blotches in the sky maps (Figs 23 and C2). Expectedly, the waste is much greater than in the test sample, and much worse still for the ‘clean = 0’ sources, but the procedure greatly and homogeneously improves the quality of the photometric redshifts in the two photometric subsamples, which can thus be combined.




**Figure 24.** The PDF width-coloured  $(u - g)/(g - r)$  distribution of the ‘clean = 1’ photometric sample at  $r < 17.8$  (left panels) and  $r > 17.8$  (right panels), following three cleansing procedures: the top panels restrict the sample to sources classified as galaxies, the middle panels apply the PDF width threshold (equation 1), and the bottom panels use both constraints. The PDF width limit alone efficiently screens both the faulty colour regions and most of the classified stars and QSOs. The equivalent distributions are shown in Fig. C3 for the ‘clean = 0’ sample.



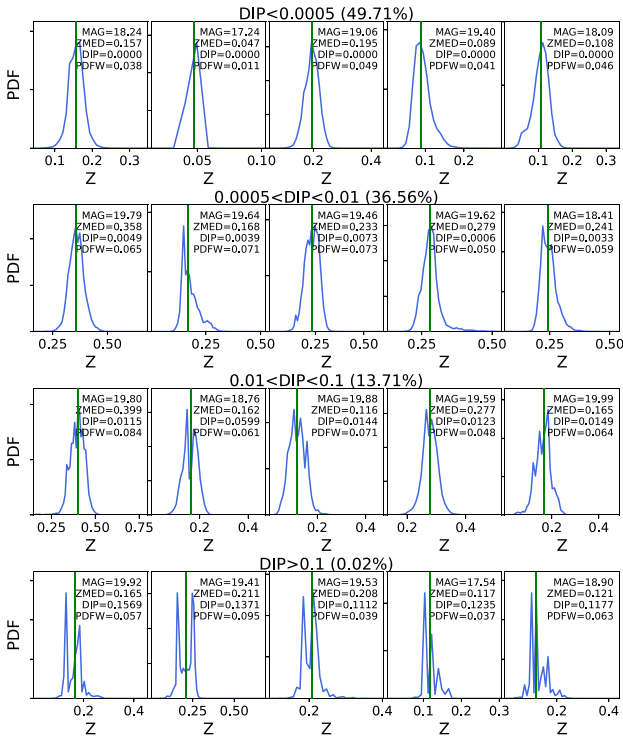
**Figure 25.** PDF width-coloured map of sources classified as galaxies with PDF width below the threshold in the full photometric samples ( $\sim 10.75$  million sources). The colour code is the same as in Figs 23 and C2 to highlight the effect of the applied threshold.

## 6.2 The full photometric sample

The combined ‘clean = 0’ + ‘clean = 1’ photometric sample contains  $\sim 11$ M sources classified as galaxies with PDF widths below the threshold. We simply refer to them as ‘galaxies’ in the rest of this work. Fig. 25 shows the PDF width-coloured sky distribution of these galaxies, with the same colour code as in Figs 23 and C2 to highlight the effect of the applied threshold. The cleansing procedure rejects 20 per cent of the initial data ( $\sim 10$  per cent are classified as stars or QSOs,  $\sim 17$  per cent have PDF widths above the threshold).



**Figure 26.** The DIP score distribution of galaxies in the photometric sample (pink histogram) compared with that of the test sample (grey histogram, accordingly scaled). The green and orange histograms are the DIP score distributions of sources classified as stars or QSOs and of sources with PDF width above the threshold, respectively.



**Figure 27.** Randomly selected PDFs with  $DIP < 0.0005$  (49.7 per cent),  $0.0005 < DIP < 0.01$  (36.5 per cent),  $0.01 < DIP < 0.1$  (13.7 per cent), and  $DIP > 0.1$  (0.02 per cent), from top to bottom.

However, it is not drastic enough to wash away the stripy disparities due to the SDSS observing conditions.

Fig. 26 shows the DIP score distribution of the photometric galaxies compared to the scaled test sample. Also shown for comparison are the DIP score distributions of sources classified as stars or QSOs and of sources with PDF width above the threshold. Fig. 27 shows samples of PDFs of photometric galaxies with DIP scores in the intervals:  $DIP < 0.0005$  (49.7 per cent),  $0.0005 < DIP < 0.01$  (36.5 per cent),  $0.01 < DIP < 0.1$  (13.7 per cent) and  $DIP > 0.1$  (0.02 per cent). The green vertical lines mark the median point estimates. PDFs with  $DIP < 0.01$  (86 per cent) look very close to uni-modal. Of the remaining 14 per cent, only a very small minority with  $DIP > 0.1$  look compellingly multi-modal with probabilities dropping very low between peaks.

### 6.3 Redshift distributions

We compare the shapes of the CNN redshift distributions to those derived from the clustering redshift technique (Ménard et al. 2013; Rahman et al. 2015) using the online platform ‘Tomographer’<sup>3</sup>. This tool is perfectly suited to our case since it relies on the SDSS-BOSS spectroscopic population (Main Galaxy Sample, LRG and quasar samples), covering the north and south galactic caps as well as the redshift range of our bright,  $r \leq 20$  sample (Chiang & Ménard 2019; Chiang, Ménard & Schiminovich 2019). The technique consists in spatially cross-correlating the spectroscopic population in bins of redshift with the sky positions of a test sample. The clustering amplitude is directly related to the redshift distribution of the test population,  $dN/dz$ , scaled by its bias  $b(z)$  with respect to the underlying dark matter density field. In a narrow interval of magnitude, the redshift evolution of the bias may be neglected and the reconstructed distribution compared to the test distribution with a constant scaling factor.

We first run Tomographer on the spectroscopic sample in 12 intervals of magnitude from  $r = 17.6$  to 20 ( $\Delta mag = 0.2$ ) to test the level of accuracy of the reconstruction. The results are shown in Appendix D. A noisy high-redshift tail to  $z \sim 3$  is present in all the magnitude bins, which we choose to ignore in the normalization to allow for a satisfactory, though far from perfect agreement between the observed and reconstructed redshift distributions in the redshift ranges of interest. This comparison gauges the accuracy we may expect for unknown distributions.


In Fig. 28, we compare the redshift distributions of the photometric galaxies with the outputs of Tomographer in the 12 intervals of magnitude. The high-redshift tails are not shown for clarity but as in the spectroscopic case, they need to be excluded in the normalization. Here we neglect them by fitting the outputs with the following ad-hoc, 4 free parameter function  $F(z) = az^b e^{-(z/c)^d}$ . This operation does away with the oscillations around zero and the necessity to cherry-pick the last bin of interest to normalize each distribution and compute Kullback-Leibler divergences ( $dN_{\text{CNN}}/dz || dN_{\text{tom}}/dz$ ). These are reported in each panel, except in the first panel where no fit could be found. The  $z_{\text{med}}$  KL ( $KL_{\text{MED}}$ ) range from 0.006 to 0.039 from the brightest to the faintest bin. The KL of the stacked PDFs tend to be smaller. The B16 distributions are significantly more discrepant. They are plagued with a growing feature at  $z \sim 0.35$  as magnitude increases, presumably related to the red galaxy degeneracy (Section 5.1). The mean redshift values, indicated for the three distributions with vertical lines of the appropriate colour, all agree within  $\sim 5$  per cent.

### 6.4 Cluster membership

Another test of the CNN redshifts is provided by the redMaPPer cluster catalogue (Rykoff et al. 2014).<sup>4</sup> The redMaPPer Cluster Finder is a red sequence cluster finder combining a calibration of the red sequence with spectroscopic redshifts and a matched-filter technique to find the clusters. For each cluster candidate, a redshift and a richness are assigned as well as a list of member galaxies at  $r \lesssim 22$  with their membership probability. The catalogue contains  $\sim 26k$  cluster candidates at  $z \lesssim 0.6$  over the two SDSS galactic caps. Based on the central galaxies with spectroscopic redshifts, the redshift uncertainty of the clusters is estimated to be lower than  $\sigma =$

<sup>3</sup><https://tomographer.org/>

<sup>4</sup>We use the 6.3 version of the redMaPPer catalogues based on the SDSS DR8 and available at: <https://cdsarc.cds.unistra.fr/ftp/J/ApJ/785/104/>

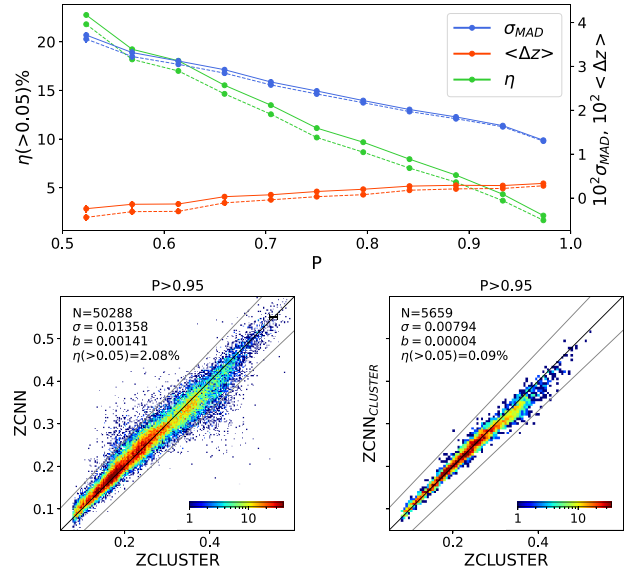


**Figure 28.** Normalized CNN redshift distributions in intervals of magnitude compared to the Tomographer-derived distributions (red dots, fitted by dotted lines). The PDF sums (in green) are shown in their native binning while  $z_{\text{med}}$  (grey shaded histogram) and  $z_{B16}$  (dashed black histogram) use the Tomographer binning. Noted on each subplot are the KL divergences between the three photometric redshift distributions and Tomographer. The vertical segments mark the mean redshift of each distribution in their respective colour.

0.01 at  $z \leq 0.3$  and  $\sigma = 0.015$  at  $z \leq 0.5$ , with a systematic offset  $\Delta z_{\text{norm}} < 0.003$  over the whole redshift range.

Among the  $\sim 1.7\text{M}$  galaxies in the redMaPPer catalogue of probable cluster members,  $\sim 377\text{k}$  belong to our photometric sample of sources classified as galaxies. We assign them the redshift of their associated cluster,  $z_{\text{CL}}$  (the velocity dispersion within such a system rarely exceeding  $\sigma_v = 1000 \text{ km s}^{-1}$  (Clerc et al. 2016), corresponding to an individual redshift uncertainty  $\sigma_z < 0.004$ ) and define  $\Delta z = (z_{\text{CNN}} - z_{\text{CL}})/(1 + z_{\text{CL}})$  using  $z_{\text{med}}$  as CNN point estimate. The top panel of Fig. 29 shows how the deviation, bias and rate of catastrophic failures ( $|\Delta z| > 0.05$ ) evolve with the membership probability. The dashed lines correspond to galaxies with PDF width below our quality threshold ( $\sim 360\text{k}$ ). The accuracy gradually improves from  $\sigma_{\text{MAD}} \sim 0.035$  for galaxies with  $0.50 < P < 0.55$  to  $\sigma_{\text{MAD}} \sim 0.013$  for galaxies with  $0.95 < P < 1.0$ , while the catastrophic fraction decreases from 20 per cent to 2 per cent.

The bottom left panel compares the redMaPPer redshift of the clusters to the CNN redshifts of the galaxies attributed to them with a probability  $P > 0.95$ , and with PDF width below the threshold



**Figure 29.** Comparison between the photometric redshifts of rich redMaPPer galaxy clusters and the CNN  $z_{\text{med}}$  of probable cluster members in the photometric sample. Top panel:  $\sigma_{\text{MAD}}$ , bias and fraction of catastrophic failures as a function of galaxy membership probability  $P$ . The dashed lines correspond to galaxies with PDF width below the threshold, which we assume in the bottom panels. Bottom left: the redMaPPer redshifts of the clusters versus the CNN redshifts of the galaxies assigned to them with a membership probability  $P > 0.95$ . The straight lines are the catastrophic failure borders and the identity line. Bottom right: the redMaPPer redshifts of the clusters compared to the mean value of the PDF product of their members. Only clusters with at least three members are retained.

( $N = 50,288$ ). The  $\sigma_{\text{MAD}}$ , mean bias and rate of catastrophic failures are comparable to those measured in the test sample (Table 3). In the bottom right panel, we compare the redMaPPer redshifts of the clusters to the weighted mean value of the PDF product of their members, restricting the sample to clusters with at least 3 members with  $P > 0.95$  and PDF width below the threshold ( $N = 5659$ , with only 4 very rich clusters not surviving the PDF product). The catastrophic fraction is negligible and  $\sigma_{\text{MAD}}$  drops to 0.00794. It reaches 0.00628 if we further restrict the sample to clusters with at least 5 members ( $N = 3274$ , with the same 4 clusters having inconsistent members). This very good agreement mutually confirms the completely independent redshift quality of the CNN and of the red sequence clusters, also supported by the SPIDERS cluster follow-up (Clerc et al. 2016).

## 7 CONCLUSION

Inferring from a CNN trained with *ugriz* stamp images of galaxies from the SDSS, GAMA and BOSS spectroscopic surveys, we estimated redshifts for the  $\sim 14$  million sources at  $r \leq 20$  without spectroscopic redshift in the  $11\,500 \text{ deg}^2$  of the SDSS north and south footprints. The redshifts extend to  $\sim 0.7$ . To limit biases, particularly those resulting from the overpopulation of Luminous Red Galaxies in the BOSS data, we extracted from the full spectroscopic catalogue several training samples approximately representative of the general galaxy population at  $r < 20$ , using GAMA as a model. The CNN was built to classify redshifts into narrow, contiguous bins. The classification outputs offer all the benefits of well-behaved bona fide PDFs, passing several statistical tests and efficiently flagging, via their widths, unreliable estimates due to poor photometry or stellar

sources. Based on a representative spectroscopic test sample, the point estimates are more than twice as accurate as the photometric redshifts currently published for the SDSS at  $r < 20$ .

We designed a magnitude-dependent PDF width threshold and a galaxy/star/QSO classifier to clean the inference sample, leaving  $\sim 11$ M sources whose CNN redshift quality we deem reliable and homogeneous over the whole footprint. These redshifts are in very good agreement with the independently derived photometric redshifts of the redMaPPer galaxy clusters for the probable cluster members among them. Their distributions in narrow bins of magnitudes also match the results of tomographic analyses satisfactorily.

Pending the release of spectroscopic redshifts by the Dark Energy Spectroscopic Instrument (DESI) for their Bright Galaxy Sample in a few years these photometric redshifts are of interest for a variety of statistical analyses. They are accurate enough to improve cluster membership and reveal the cosmic web in thin redshift slices (e.g. Laigle et al. 2018), allowing us to extend the spectroscopic analyses probing its impact on galaxy properties (Kraljic et al. 2018; Malavasi et al. 2017). They may also improve the cosmological information retrieved from cross-correlating sparse spectroscopic samples with photometric data, e.g. by reducing the shot noise at the baryon acoustic oscillation scale (Patej & Eisenstein 2018). They can be used to measure, e.g.: the connectivity of groups and clusters to study the properties of their member galaxies as a function of group mass and assembly history (Darragh Ford et al. 2019); the evolution of the mean connectivity with redshift, which depends on the cosmological model and on the nature of the dark energy (Codis, Pogosyan & Pichon 2018); the one-point distribution of the cosmic density field in cylinder at a given radius in a way complementary to standard power spectrum analysis (albeit with different biases and sensitivity to cosmology, Uhlemann et al. 2018).

However, the range of the present CNN redshifts is limited and their precision degrade as they and magnitude increase. Going down in magnitude is a big challenge, especially as spectroscopic data become very sparse. More complex deep learning techniques are necessary for deeper on-going surveys such as HSC-CLAUDS (Sawicki et al. 2019) and future ones such as LSST (Ivezić et al. 2019). Work is underway that promises to reach high photometric redshift accuracy to  $z \sim 1.5$  (Ait Ouahmed et al. in preparation), allowing us to extend the scope of the above cosmological investigations. The challenge of satisfying the LSST science requirements to  $z = 3$  is yet to be met.

## ACKNOWLEDGEMENTS

This work was granted access to the HPC resources of IDRIS under the allocation 2023-AD010414147 made by GENCI. It was carried out with the support of the French National Research Agency (ANR) as part of the DEEPDIP project (ANR-19-CE31-0023, <http://deepdip.net>), and of the Programme National Cosmologie et Galaxies (PNCG) of INSU,CNRS with INP and IN2P3, co-funded by CEA and CNES. This publication makes use of Sloan Digital Sky Survey (SDSS) data. Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org/>. SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofísica

de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

## DATA AVAILABILITY

The CNN input (five-band image cutouts and associated spectroscopic and photometric catalogues) and output (PDFs and derived quantities) for the  $\sim 14$ M sources in the SDSS DR16 at  $r < 20$  are available at: <https://deepdip.iap.fr/treyer+2023>.

## REFERENCES

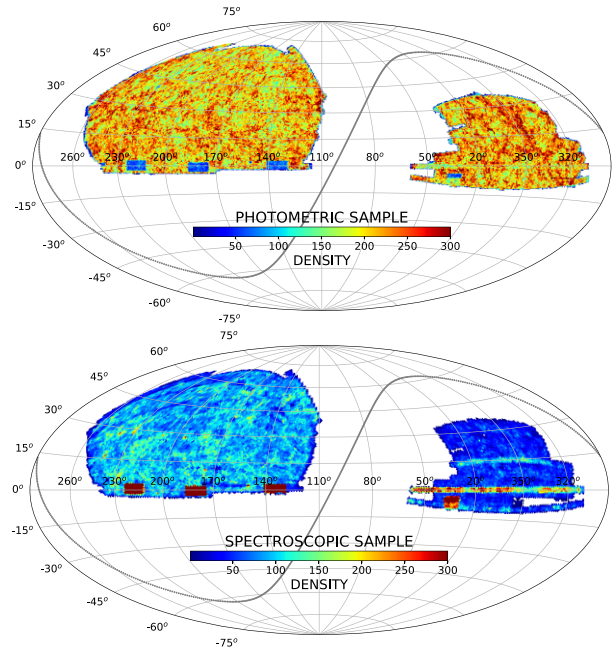
- Ahumada R. et al., 2020, *ApJS*, 249, 3  
 Ait Ouahmed R., Arnouts S., Pasquet J., Treyer M., Bertin E., 2023, preprint (arXiv:2310.02185)  
 Alam S. et al., 2015, *ApJS*, 219, 12  
 Amaro V. et al., 2019, *MNRAS*, 482, 3116  
 Arnouts S., Cristiani S., Moscardini L., Matarrese S., Lucchin F., Fontana A., Giallongo E., 1999, *MNRAS*, 310, 540  
 Baum E. B., Wilczek F., 1987, in Proc. 1987 Int. Conf. on Neural Information Processing Systems. NIPS'87. MIT Press, Cambridge, MA, p. 52, available at: <http://dl.acm.org/citation.cfm?id=2969644.2969650>  
 Beck R., Dobos L., Budavári T., Szalay A. S., Csabai I., 2016, *MNRAS*, 460, 1371  
 Bertin E., Mellier Y., Radovich M., Missonnier G., Didelon P., Morin B., 2002, in Bohlender D. A., Durand D., Handley T. H., eds, ASP Conf. Ser. Vol. 281, Astronomical Data Analysis Software and Systems XI. Astron. Soc. Pac., San Francisco, p. 228  
 Bradshaw E. J. et al., 2013, *MNRAS*, 433, 194  
 Brammer G. B., van Dokkum P. G., Coppi P., 2008, *ApJ*, 686, 1503  
 Brescia M., Cavuoti S., Razim O., Amaro V., Riccio G., Longo G., 2021, *Front. Astron. Space Sci.*, 8, 70  
 Bridle J. S., 1990, in Soulié F. F., Héroult J., eds, Neurocomputing. Springer, Berlin, Heidelberg, p. 227  
 Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, 351, 1151  
 Calabretta M. R., Greisen E. W., 2002, *A&A*, 395, 1077  
 Chiang Y.-K., Ménard B., 2019, *ApJ*, 870, 120  
 Chiang Y.-K., Ménard B., Schiminovich D., 2019, *ApJ*, 877, 150  
 Clerc N. et al., 2016, *MNRAS*, 463, 4490  
 Codis S., Pogosyan D., Pichon C., 2018, *MNRAS*, 479, 973  
 Collister A. A., Lahav O., 2004, *PASP*, 116, 345  
 Csabai I., Dobos L., Trencsényi M., Herczegh G., Józsa P., Purger N., Budavári T., Szalay A. S., 2007, *Astron. Nachr.*, 328, 852  
 D'Isanto A., Polsterer K. L., 2018, *A&A*, 609, A111  
 Dark Energy Survey Collaboration, 2016, *MNRAS*, 460, 1270  
 Darragh Ford E. et al., 2019, *MNRAS*, 489, 5695  
 Dawid A. P., 1984, *J. R. Stat. Soc. A (General)*, 147, 278  
 Dawson K. S. et al., 2013, *AJ*, 145, 10  
 de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2013, *Exp. Astron.*, 35, 25  
 Dey B., Andrews B. H., Newman J. A., Mao Y.-Y., Rau M. M., Zhou R., 2022, *MNRAS*, 515, 5285  
 Driver S. P. et al., 2009, *Astron. Geophys.*, 50, 5  
 Driver S. P. et al., 2011, *MNRAS*, 413, 971  
 Driver S. P. et al., 2022, *MNRAS*, 513, 439  
 Firth A. E., Lahav O., Somerville R. S., 2003, *MNRAS*, 339, 1195  
 Fukushima K., 1969, *IEEE Transact. Syst. Sci. Cybernetics*, 5, 322  
 Hartigan J. A., Hartigan P. M., 1985, *Ann. Stat.*, 13, 70



- Hayat M. A., Stein G., Harrington P., Lukić Z., Mustafa M., 2021, *ApJ*, 911, L33
- Henghes B., Thiyaalingam J., Pettitt C., Hey T., Lahav O., 2022, *MNRAS*, 512, 1696
- Hersbach H., 2000, *Weather Forecast.*, 15, 559
- Hoyle B., 2016, *Astron. Comput.*, 16, 34
- Ilbert O. et al., 2006, *A&A*, 457, 841
- Ivezić Ž. et al., 2019, *ApJ*, 873, 111
- Jones E., Do T., Boscoe B., Singal J., Wan Y., Nguyen Z., 2023, preprint (arXiv:2306.13179)
- Kingma D. P., Ba J., 2015, in Bengio Y., LeCun Y., eds, 3rd Int. Conf. Learning Representations, ICLR 2015
- Kraljic K. et al., 2018, *MNRAS*, 474, 547
- Laigle C. et al., 2018, *MNRAS*, 474, 5437
- Laureijs R. et al., 2011, preprint (arXiv:1110.3193)
- Le Fèvre O. et al., 2013, *A&A*, 559, A14
- LeCun Y., Bottou L., Bengio Y., Haffner P., 1998, *Proc. IEEE*, 86, 2278
- Lilly S. J. et al., 2007, *ApJS*, 172, 70
- Liske J. et al., 2015, *MNRAS*, 452, 2087
- Malvasi N. et al., 2017, *MNRAS*, 465, 3817
- McLure R. J. et al., 2013, *MNRAS*, 428, 1088
- Ménard B., Scranton R., Schmidt S., Morrison C., Jeong D., Budavari T., Rahman M., 2013, preprint (arXiv:1303.4722)
- Newman J. A. et al., 2013, *ApJS*, 208, 5
- Pasquet J., Bertin E., Treyer M., Arnouts S., Fouchez D., 2019, *A&A*, 621, A26
- Patej A., Eisenstein D. J., 2018, *MNRAS*, 477, 5090
- Polsterer K. L., D’Isanto A., Lerch S., 2022, in Ruiz J. E., Pierfederici F., Teuben P., eds, ASP Conf. Ser. Vol. 532, Astronomical Data Analysis Software and Systems XXX. Astron. Soc. Pac., San Francisco, p. 173
- Rahman M., Ménard B., Scranton R., Schmidt S. J., Morrison C. B., 2015, *MNRAS*, 447, 3500
- Richard M., Lippmann R., 1991, *Neural Comput.*, 3, 461
- Rojas R., 1996, *Neural Comput.*, 8, 41
- Rykoff E. S. et al., 2014, *ApJ*, 785, 104
- Sadeh I., 2014, in Heavens A., Starck J.-L., Krone-Martins A., eds, IAU Symp. 306, Statistical Challenges in 21st Century Cosmology. p. 316
- Sadeh I., Abdalla F. B., Lahav O., 2016, *PASP*, 128, 104502
- Sawicki M. et al., 2019, *MNRAS*, 489, 5202
- Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, 500, 525
- Schuldts S., Suyu S. H., Cañameras R., Taubenberger S., Meinhardt T., Leal-Taixé L., Hsieh B. C., 2021, *A&A*, 651, A55
- Scodreggio M. et al., 2018, *A&A*, 609, A84
- Solla S., Levin E., Fleisher M., 1988, *Complex Syst.*, 2, 625
- Tagliaferri R., Longo G., Andreon S., Capozziello S., Donalek C., Giordano G., 2003, in, Vol. 2859, Lecture Notes in Computer Science. p. 226
- Treyer M. et al., 2018, *MNRAS*, 477, 2684
- Uhlemann C., Pichon C., Codis S., L’Huillier B., Kim J., Bernardeau F., Park C., Prunet S., 2018, *MNRAS*, 477, 2772
- Wittman D., Bhaskar R., Tobin R., 2016, *MNRAS*, 457, 4005
- York D. G. et al., 2000, *AJ*, 120, 1579
- Zamo M., Naveau P., 2018, *Math. Geosci.*, 50, 209

## APPENDIX A: DATA SKY DISTRIBUTION

Fig. A1 shows the sky distribution of the photometric and spectroscopic samples down to the dereddened petrosian magnitude of  $r = 20$  in the two main regions of the SDSS.



**Figure A1.** Top: Mollweide projection of the SDSS photometric sample at  $r \leq 20$  without spectroscopy ( $\sim 13.8\text{M}$  sources). The colour code indicates the source density in HEALPix cells with  $n_{\text{side}} = 140$  ( $0.18 \text{ deg}^2 \text{ pixel}^{-1}$ ). The sky coverage is  $\sim 11\,529 \text{ deg}^2$ . Bottom: Mollweide projection of the spectroscopic sample at  $r \leq 20$  ( $\sim 1.5\text{M}$  galaxies). The colour code indicates the source density in HEALPix cells with  $n_{\text{side}} = 64$  ( $0.84 \text{ deg}^2 \text{ pixel}^{-1}$ ). The sky coverage is  $\sim 11\,029 \text{ deg}^2$ . The four high-density regions in red are the GAMA fields.

## APPENDIX B: GALAXY TYPE CLASSIFICATION


To classify galaxies as either blue/star-forming or red/passive, we use the sSFR (the star formation rate per unit stellar mass) derived for the spectroscopic sample in the SDSS DR12 by Brinchmann et al. (2004). Fig. B1 shows this sSFR (mean value per pixel) in the observed  $(u - g)/(g - r)$  plane in several bins of redshifts. The narrow greenish demarcations between the blue and the red zones highlight the bimodal distribution of the sSFR at a given redshift that prompts the distinction between blue and red galaxies. The black dashed lines running through these demarcations are modelled as followed:

$$\begin{aligned} y_{\text{lim}}[(u - g) < 3] &= y_b + 0.05((u - g) - 2)^2 \\ y_{\text{lim}}[(u - g) > 3] &= y_b + 0.05 - 0.02((u - g) - 3), \end{aligned} \quad (\text{B1})$$

where

$$\begin{aligned} y_b[z < 0.32] &= 0.65 + 1.59z + 1.19z^2 \\ y_b[z > 0.32] &= 1.28. \end{aligned} \quad (\text{B2})$$

We compare this classification to that derived for GAMA at  $z < 0.3$  by SED fitting the rich multi-band (UV to IR) photometry (Kraljic et al. 2018; Treyer et al. 2018). The limit between blue and passive




**Figure B1.** The  $(u - g)/(g - r)$  distribution of the SDSS spectroscopic sample at  $r < 20$ , coloured with the mean specific star formation rate ( $\text{dex yr}^{-1}$ ) per pixel, in different intervals of redshift. The narrow demarcations between the blue and the red zones highlighting the bimodal distribution of the sSFR are used to model boundaries between blue and red galaxies as a function of redshift.

galaxies was set at  $\text{sSFR} \approx -10.5 \text{ dex yr}^{-1}$ . The red completeness, defined as the number of galaxies classified as passive according to both criteria over the number of passive galaxies according to the GAMA sSFR criterion, is 87 per cent. The red purity, defined as the number of galaxies classified as passive according to both criteria over the number of passive galaxies according to the  $(u - g)/(g - r)$ /redshift criterion, is 75 per cent. Likewise, the blue completeness and purity are 87 per cent and 93 per cent respectively. The present recipe tends to overestimate red galaxies compared to the GAMA sSFR limit but the two classifications are in reasonably good agreement. It is however inadequate for the LRGs in the highest redshift bin, many of which would be classified as blue. BOSS galaxies are considered red regardless. Rough as it is, we will use this prescription to uncover statistical differences in CNN performance, if any, between the two types. Optimizing it is beyond the scope of this paper.


### APPENDIX C: THE ‘DIRTY’ PHOTOMETRIC SAMPLE

Fig. C1 shows the  $(u - g)/(g - r)$  colour distribution of the ‘clean = 0’ photometric sources at  $r < 17.8$  and  $17.8 < r < 20$  (left and right panels respectively). The distributions are colour-coded by the mean PDF width in the middle panels and by the mean  $r$ -band SNR in the bottom panels. We use the same colour code as in Fig. 21 to emphasize the differences between the two samples. The SNR and PDF quality are very degraded nearly everywhere compared to the ‘clean’ sample, included within the training contours. The striking star sequence at  $r < 17.8$  in Fig. 21 is drowned in sources with equally poor PDFs.

Fig. C2 shows the PDF width coloured sky map of the sources classified as galaxies, to be compared with Fig. 23. The difference between the two samples is all the more striking that the present data are brighter ( $\langle r \rangle = 18.81$ ) than the ‘clean’ data ( $\langle r \rangle = 19.24$ ). The ‘suspect zone’ identified by P19 is much more prominent in this data set, which also contains many other similarly degraded areas, especially in the southern region. The enhanced image quality in Stripe 82 (greenish equatorial stripe in the South) remains visible.

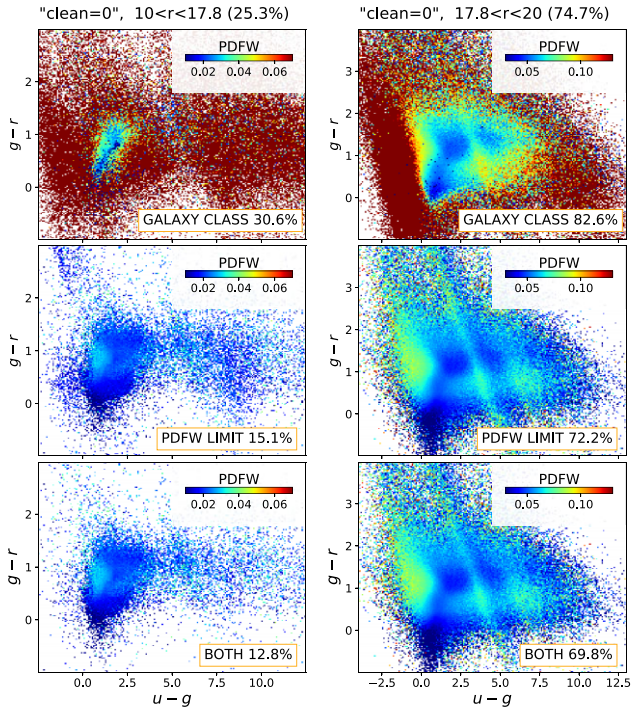


**Figure C1.** The  $(u - g)/(g - r)$  colour distribution of the ‘clean = 0’ photometric sample at  $r < 17.8$  (left panels) and  $r > 17.8$  (right panels), colour-coded by the density (top), the mean PDF width (middle) and the  $r$ -band SNR (bottom). The quality of the data is significantly inferior to the ‘clean = 1’ sample shown in Fig. 21.



**Figure C2.** PDF width coloured map of sources classified as galaxies in the ‘clean = 0’ photometric sample, to be compared with Fig. 23. The difference between the two samples is all the more striking that the clean data are deeper. The larger than average PDF widths in the GAMA regions are due to their deeper average magnitude. The higher image quality in Stripe 82 (greenish equatorial stripe in the South) remains visible.

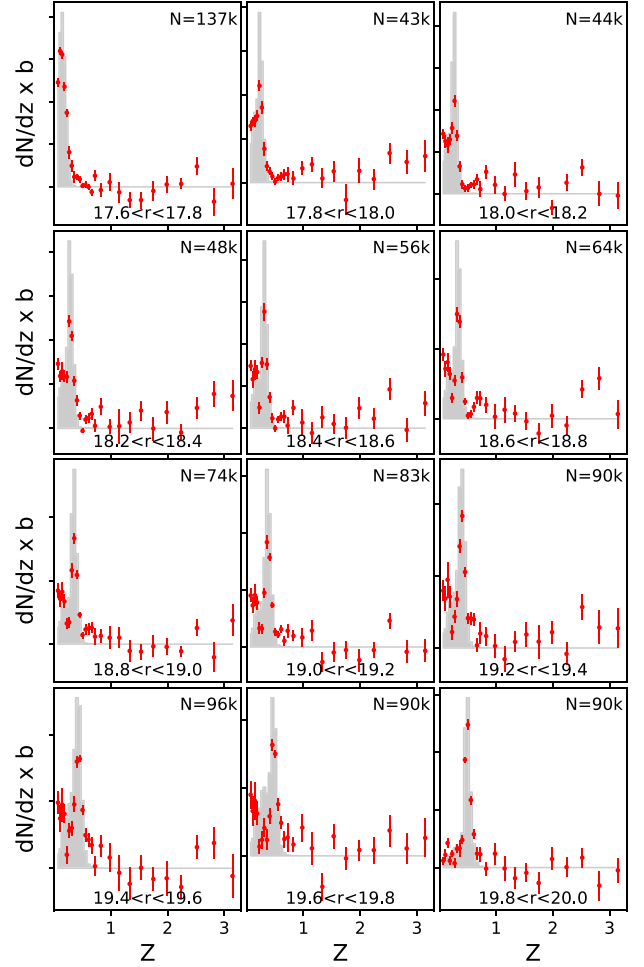
Fig. C3 shows the  $(u - g)/(g - r)$  colour distributions of the sample colour-coded by the mean PDF width and split into the bright and faint magnitude intervals, following the three cleansing procedures: the top panels restrict the samples to sources classified as galaxies (30.6 per cent at  $r < 17.8$ , 82.6 per cent at  $r > 17.8$ ), the middle panels to sources with PDF widths below the threshold (15.1 per cent at  $r < 17.8$ , 72.2 per cent at  $r > 17.8$ ), and finally the bottom panels to classified galaxies with PDF widths below the threshold (12.8 per cent at  $r < 17.8$ , 69.8 per cent at  $r > 17.8$ ). This final procedure rejects 44.5 per cent of the initial data (30.5 per cent classified as stars or QSOs, 41.8 per cent with PDF widths above the threshold).



**Figure C3.** The  $(u-g)/(g-r)$  colour distribution of the ‘clean = 0’ sample colour-coded by the mean PDF width at  $r < 17.8$  in the left panels and  $r > 17.8$  in the right panels. The top panels restrict the sample to sources classified as galaxies, the middle panels apply the PDF width threshold, and the bottom panels use both constraints.

#### APPENDIX D: TOMOGRAPHER

We run Tomographer on the spectroscopic sample in order to test the level of accuracy of the output distributions. The results are shown as red dots in Fig. D1 for 12 narrow intervals of magnitude ( $\Delta\text{mag} = 0.2$ ), with the spectroscopic redshift distributions as shaded histograms. These need to be normalized to match the Tomographer outputs. Given the many negative and unrealistically high data points in the high-redshift tail, we choose to ignore everything at  $z > 1$  and to normalize the spectroscopic redshift distributions by the Tomographer counts at  $z < 1$ , which allows for a much better, though far from perfect agreement in the redshift ranges of interest. This comparison gauges the accuracy we may expect for unknown distributions.




**Figure D1.** The normalized redshift distributions of the spectroscopic sample in 12 bins of magnitude (grey shaded histograms) compared to the distributions derived from Tomographer (red dots). The normalization ignores the tail at  $z > 1$ , which allows for a better agreement at  $z < 1$ .

#### APPENDIX E: ALTERNATIVE EXPERIMENTS


The following experiments are conducted under the final conditions of this work: we average the outputs of 15 networks trained with a redshift bin width  $\delta z = 0.009$ .

### E1 Classification versus regression


Table E1 shows the CNN performance on the test sample for different training strategies: classification + regression with a RMSE loss function, classification + regression with a MAE loss function, classification alone without regression, regressions alone with either a RMSE or a MAE loss function. The classification alone provides better statistics than both regressions, especially in terms of  $\sigma_{\text{MAD}}$ . Adding a regression to the classification has a minor positive impact on the classification but a significant one on the  $\sigma_{\text{MAD}}$  of the regression. The preferred strategy, used in this work, is highlighted in bold.




**Figure E1.** The mean bias,  $\sigma_{\text{MAD}}$  and rate of catastrophic failures of the 4 point estimates in the test sample for the 15 networks trained under two strategies: 15 samples trained once (blue points), 1 sample trained 15 times (pink points). The metric of the joint network is marked by a yellow line. The black circles and vertical lines show the mean and standard deviations. The right most points in each panel result from averaging the 15 outputs.



**Figure E2.** The  $z_{\text{spec}}$  and  $z_{\text{med}}$  distributions of the additional leftover galaxies at  $r < 17.8$  (left) and  $17.8 < r < 20$  (right), in grey and pink respectively, with the corresponding metrics. The hatched histograms are the  $z_{\text{spec}}$  distributions of the leftover galaxies common to both training strategies (shown in Fig. 20).



**Figure E3.** The  $z_{\text{med}}$  distributions of the leftover galaxies at  $r < 17.8$  and  $17.8 < r < 20$  before and after adding 50k LRGs to the training samples, in blue and pink respectively, with the corresponding metrics in the direction of the arrow. The grey shaded histograms are the  $z_{\text{spec}}$  distributions. The bright sample is unchanged, the bias in the faint interval is reduced by 35 per cent. The deviation and rate of catastrophic failure are also improved.



**Figure E4.** The  $z_{\text{med}}$  distributions of the blue and red galaxies in the test sample before and after adding 50k LRGs to the training samples, in blue and pink respectively, with the corresponding metrics in the direction of the arrow. The grey shaded histograms are the  $z_{\text{spec}}$  distributions. The blue sample is unchanged, the red galaxy bias is increased by 37 per cent. The deviation and rate of catastrophic failure are also degraded.

### E2 Randomness strategy

Instead of training several samples, with the goal of feeding the CNN the largest variety of galaxies from the spectroscopic sample, we test the alternative strategy of training a single sample several times. This leaves twice as many galaxies for testing. We pick one of the 15 trained samples at random and retrain it 14 times.

Fig. E1 shows the different metrics for the different point estimates for the 15 networks trained under the two strategies: 15 samples trained once (blue points), 1 sample trained 15 times (pink points). The yellow lines mark the metrics of the joint network. The scattered points to the left of each panel ( $N = 1$ ) show that the second strategy generates as much variation as the original strategy for  $\sigma_{\text{MAD}}$  and the bias, but less for the catastrophic failures (except for  $z_{\text{peak}}$ ). The general similarity may not be surprising given that the 15 training samples are designed to contain similar galaxies, if not the same. But the overlap between 2 samples is in many cases less than 50 per cent so the impact of randomly initializing the weights at the start of training is as large as replacing half of the training set with different but similar sources. The right most points in each panel ( $N = 15$ ) show that the metrics resulting from averaging the 15 outputs are slightly poorer in the second scenario. Even if the reverse could presumably have happened, it seems that randomizing 15 samples had a higher chance of reaching lower metrics than randomizing the initial training weights of a unique sample.

**Table E1.** The CNN performance on the test sample (Section 4.2) for different training strategies: classification + regression with RMSE loss (used in the present work), classification + regression with MAE loss, classification alone without regression, regressions alone with either an RMSE or MAE loss. The best statistics are highlighted in bold.

	PDF + REG(RMSE)			PDF + REG(MAE)			PDF		REG(RMSE)	REG(MAE)
	$z_{\text{mean}}$	$z_{\text{med}}$	$z_{\text{reg}}$	$z_{\text{mean}}$	$z_{\text{med}}$	$z_{\text{reg}}$	$z_{\text{mean}}$	$z_{\text{med}}$	$z_{\text{reg}}$	$z_{\text{reg}}$
$10^5 \sigma_{\text{MAD}}$	1466	<b>1421</b>	1481	1464	1431	1461	1470	1444	1586	1517
$10^5 \langle \Delta z \rangle$	205	<b>153</b>	228	209	154	196	218	159	240	189
$\eta(> 0.05)(\%)$	4.1	<b>3.99</b>	4.22	4.05	4.06	4.26	4.18	4.16	4.2	4.15

As expected from the above results, the metrics found for the leftover galaxies common to both training scenarios (the sample shown in Fig. 20) are also slightly degraded. Fig. E2 shows the normalized spectroscopic redshift and  $z_{\text{med}}$  distributions of the additional  $\sim 542\text{k}$  galaxies left over from the second scenario. The hatched histograms are the spectroscopic redshift distributions of the leftover galaxies in common, for comparison. The bright fraction of this additional sample is smoother than the hatched one, very slightly improving the predictions. On the contrary the faint fraction contains galaxies in the strong peaks of the GAMA redshift distribution on the left flank and a larger fraction of red galaxies in the region of colour degeneracy than of LRGs. The bias is reduced, compared to the common sample shown in Fig. 20, but the deviation and rate of catastrophic failure are degraded. While it may be interesting to see the effect of mixing heterogeneous populations, the twice as large number of leftover galaxies does not allow for more informative testing. It confirms that the performance of the network is poor for red galaxies at 'normal' redshifts due to the colour degeneracy, very poor for LRGs that are deliberately relegated to the leftover sample, and that the predictions are not able to capture strong redshift structures, also deliberately smoothed in the training samples.

### E3 Luminous Red Galaxies

Fig. E3 shows the  $z_{\text{med}}$  distributions of the leftover sample (shown in Fig. 20) at  $r < 17.8$  and  $17.8 < r < 20$  before and after the addition of 50k LRGs subtracted from it in the training samples. The corresponding metrics are indicated in the direction of the arrow. The bright sample is unchanged, while the bias in the faint interval is reduced by 35 per cent. The deviation and rate of catastrophic failure are also improved. However, this improvement is at the cost of degrading 'normal' galaxies. Fig. E4 shows the  $z_{\text{med}}$  distributions of the blue and red galaxies in the test sample before and after the LRG addition. The blue sample is unchanged, while the red galaxy bias is increased

by 37 per cent. The deviation and rate of catastrophic failure are also degraded.

This demonstrates the utmost importance of matching the training and test samples and that more does not necessarily mean better. Although the spectroscopic sample contains 1.5M galaxies, it cannot be used as is for lack of representativeness. Randomly splitting it for training/validation, e.g. 80 per cent/20 per cent, as is usually done to evaluate the performance of a network, would yield very misleading results. We choose 50 per cent/50 per cent (i.e. 750k randomly selected galaxies for training) to illustrate the point. This is the strategy employed by B16. It more than doubles computing resources (memory and time) compared to the smaller, more representative training samples we adopt. It also doubles the bias for the red population in the test sample ( $b = 0096$ ) and further degrades the MAD and rate of catastrophic failure. Meanwhile the metrics on the validation sample are significantly better than on the test sample ( $b = -0.00015$ ,  $\sigma_{\text{MAD}} = 0.011$ ,  $\eta = 2$  per cent) as it matches the training sample by design, in particular the LRG population. The same applies to the performance reported by B16 on their validation sample, which is significantly lower than this however, save for their even smaller bias. We note that LRGs could be used to train an independent network that would improve their redshift estimates in the inference sample, provided we were able to identify them.

### APPENDIX F: CNN ARCHITECTURE

Fig. F1 lists all the CNN layers with their type, shape, number of parameters and the layer(s) they are connected to upstream.  $N$  is the number of galaxies in a batch (32 for training),  $NZ$  in the 'PDF' output layer is the number of redshift classes.

In the case of the SDSS at  $r < 17.8$  (Section 4.1), the CNN is trained for 45 epochs, with a learning rate of  $10^{-4}$  from epoch 1 to 30, decreasing by a factor of 10 at epoch 30 and 40. At  $r < 20$  (Section 4.2), the network is trained for 50 epochs, with a learning rate of  $10^{-4}$  from epoch 1 to 35, decreasing by a factor of 10 at epoch 35 and 45.

Layer (type)	Output Shape	Param.	Connected to	— continued from previous column	
INPUT (InputLayer)	(N, 64, 64, 5)	0			
conv2d#0 (Conv2D)	(N, 64, 64, 96)	12096	INPUT	avp#5 (Av. Pooling)	(N, 16, 16, 101) 0 conv2d#24
conv2d#1 (Conv2D)	(N, 64, 64, 96)	83040	conv2d#0	conv2d#25 (Conv2D)	(N, 16, 16, 109) 57007 avp#4
avp2d#0 (Av. Pooling)	(N, 32, 32, 96)	0	conv2d#1	concat#3 (Concat.)	(N, 16, 16, 522) 0 conv2d#21
conv2d#2 (Conv2D)	(N, 32, 32, 101)	9797	avp2d#0		conv2d#23
conv2d#4 (Conv2D)	(N, 32, 32, 101)	9797	avp2d#0		avp#5
conv2d#6 (Conv2D)	(N, 32, 32, 101)	9797	avp2d#0		conv2d#25
conv2d#3 (Conv2D)	(N, 32, 32, 156)	394056	conv2d#2	conv2d#26 (Conv2D)	(N, 16, 16, 101) 52823 concat#3
conv2d#5 (Conv2D)	(N, 32, 32, 156)	141960	conv2d#4	conv2d#28 (Conv2D)	(N, 16, 16, 101) 52823 concat#3
avp2d#1 (Av. Pooling)	(N, 32, 32, 101)	0	conv2d#6	conv2d#30 (Conv2D)	(N, 16, 16, 101) 52823 concat#3
conv2d#7 (Conv2D)	(N, 32, 32, 109)	10573	avp#0	conv2d#27 (Conv2D)	(N, 16, 16, 156) 394056 conv2d#26
concat#0 (Concat.)	(N, 32, 32, 522)	0	conv2d#3	conv2d#29 (Conv2D)	(N, 16, 16, 156) 141960 conv2d#28
			conv2d#5	avp#6 (Av. Pooling)	(N, 16, 16, 101) 0 conv2d#30
			avp#1	conv2d#31 (Conv2D)	(N, 16, 16, 109) 57007 concat#3
			conv2d#7	concat#4 (Concat.)	(N, 16, 16, 522) 0 conv2d#27
conv2d#8 (Conv2D)	(N, 32, 32, 101)	52823	concat#0		conv2d#29
conv2d#10 (Conv2D)	(N, 32, 32, 101)	52823	concat#0		avp#6
conv2d#12 (Conv2D)	(N, 32, 32, 101)	52823	concat#0		conv2d#31
conv2d#9 (Conv2D)	(N, 32, 32, 156)	394056	conv2d#8	avp#7 (Av. Pooling)	(N, 8, 8, 522) 0 concat#4
conv2d#11 (Conv2D)	(N, 32, 32, 156)	141960	conv2d#10	conv2d#32 (Conv2D)	(N, 8, 8, 101) 52823 avp#7
avp#2 (Av. Pooling)	(N, 32, 32, 101)	0	conv2d#12	conv2d#34 (Conv2D)	(N, 8, 8, 101) 52823 avp#7
conv2d#13 (Conv2D)	(N, 32, 32, 109)	57007	concat#0	conv2d#36 (Conv2D)	(N, 8, 8, 101) 52823 avp#7
concat#1 (Concat.)	(N, 32, 32, 522)	0	conv2d#9	conv2d#33 (Conv2D)	(N, 8, 8, 156) 394056 conv2d#32
			conv2d#11	conv2d#35 (Conv2D)	(N, 8, 8, 156) 141960 conv2d#34
			avp#2	avp#8 (Av. Pooling)	(N, 8, 8, 101) 0 conv2d#36
			conv2d#13	conv2d#37 (Conv2D)	(N, 8, 8, 109) 57007 avp#7
conv2d#14 (Conv2D)	(N, 32, 32, 101)	52823	concat#1	concat#5 (Concat.)	(N, 8, 8, 522) 0 conv2d#33
conv2d#16 (Conv2D)	(N, 32, 32, 101)	52823	concat#1		conv2d#35
conv2d#18 (Conv2D)	(N, 32, 32, 101)	52823	concat#1		avp#8
conv2d#15 (Conv2D)	(N, 32, 32, 156)	394056	conv2d#14		conv2d#37
conv2d#17 (Conv2D)	(N, 32, 32, 156)	141960	conv2d#16	conv2d#38 (Conv2D)	(N, 6, 6, 96) 451104 concat#5
avp#3 (Av. Pooling)	(N, 32, 32, 101)	0	conv2d#18	conv6 (Conv2D)	(N, 4, 4, 96) 83040 conv2d#38
conv2d#19 (Conv2D)	(N, 32, 32, 109)	57007	concat#1	conv2d#39 (Conv2D)	(N, 2, 2, 96) 83040 conv6
concat#2 (Concat.)	(N, 32, 32, 522)	0	conv2d#15	avp#9 (Av. Pooling)	(N, 1, 1, 96) 0 conv2d#39
			conv2d#17	reshape (Reshape)	(N, 96) 0 avp#9
			avp#3	EBV (InputLayer)	[(N, 1)] 0
			conv2d#19	concat#6 (Concat.)	(N, 97) 0 reshape
avp#4 (Av. Pooling)	(N, 16, 16, 522)	0	concat#2		EBV
conv2d#20 (Conv2D)	(N, 16, 16, 101)	52823	avp#4	dense#0 (Dense)	(N, 1024) 100352 concat#6
conv2d#22 (Conv2D)	(N, 16, 16, 101)	52823	avp#4	dense#1 (Dense)	(N, 1024) 1049600 dense#0
conv2d#24 (Conv2D)	(N, 16, 16, 101)	52823	avp#4	dense#2 (Dense)	(N, 512) 524800 dense#0
conv2d#21 (Conv2D)	(N, 16, 16, 156)	394056	conv2d#20	PDF (Dense)	(N,NZ) NZ×1025 dense#1
conv2d#23 (Conv2D)	(N, 16, 16, 156)	141960	conv2d#22	REG (Dense)	(N, 1) 513 dense#2

**Figure F1.** Successive CNN layers with their output dimension, number of trainable parameters and the layer(s) that connect(s) to them. The CNN is diagrammed in Figs 3 and 4.

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