




Augmenting photometric redshift estimates using spectroscopic nearest neighbours

F. Tosone¹, M. S. Cagliari¹ , L. Guzzo^{1,2,3} , B. R. Granett² , and A. Crespi¹

¹ Dipartimento di Fisica “Aldo Pontremoli”, Università degli Studi di Milano, Via G. Celoria 16, 20133 Milano, Italy
e-mail: federico.tosone@unimi.it; marina.cagliari@unimi.it

² INAF–Osservatorio Astronomico di Brera, Via Brera 28, 20121 Milano, and Via E. Bianchi 46, 23807 Merate, Italy

³ INFN–Sezione di Milano, Via G. Celoria 16, 20133 Milano, Italy

Received 3 November 2022 / Accepted 2 March 2023

ABSTRACT

As a consequence of galaxy clustering, close galaxies observed on the plane of the sky should be spatially correlated with a probability that is inversely proportional to their angular separation. In principle, this information can be used to improve photometric redshift estimates when spectroscopic redshifts are available for some of the neighbouring objects. Depending on the depth of the survey, however, this angular correlation is reduced by chance projections. In this work, we implement a deep-learning model to distinguish between apparent and real angular neighbours by solving a classification task. We adopted a graph neural network architecture to tie together photometry, spectroscopy, and the spatial information between neighbouring galaxies. We trained and validated the algorithm on the data of the VIPERS galaxy survey, for which photometric redshifts based on spectral energy distribution are also available. The model yields a confidence level for a pair of galaxies to be real angular neighbours, enabling us to disentangle chance superpositions in a probabilistic way. When objects for which no physical companion can be identified are excluded, all photometric redshift quality metrics improve significantly, confirming that their estimates were of lower quality. For our typical test configuration, the algorithm identifies a subset containing $\sim 75\%$ high-quality photometric redshifts, for which the dispersion is reduced by as much as 50% (from 0.08 to 0.04), while the fraction of outliers reduces from 3% to 0.8%. Moreover, we show that the spectroscopic redshift of the angular neighbour with the highest detection probability provides an excellent estimate of the redshift of the target galaxy, comparable to or even better than the corresponding template-fitting estimate.

Key words. galaxies: distances and redshifts – methods: statistical – methods: data analysis

1. Introduction

Knowledge of galaxy distances is of the utmost importance for cosmology to reconstruct the underlying 3D dark matter distribution that encapsulates key information about the evolution and matter content of the Universe. On cosmological scales, the most efficient method for estimating distances is through their cosmological redshift, which directly connects to the standard definitions of distance. Sufficiently precise redshift measurements allow us to test the world model through the redshift-distance relation, coupled with standard rulers and standard candles (e.g. Riess et al. 1998; Perlmutter et al. 1998).

Over the past 25 yr, galaxy clustering measurements from large redshift surveys have been able to quantify the universal expansion and growth histories, pinpointing the value of cosmological parameters to high precision (e.g. Tegmark et al. 2006; Colless et al. 2003; Blake et al. 2011; de la Torre et al. 2017; Alam et al. 2017; Pezzotta et al. 2017; Bautista et al. 2021). Even larger redshift surveys are now ongoing (DESI; DESI Collaboration 2016) or are scheduled to start soon (Euclid; Laureijs et al. 2011), with the goal of further refining these measurements to exquisite precision and find clues for the poorly understood ingredients of the remarkably successful standard model of cosmology.

The redshift is measured from the shift in the position of emission and absorption features identified in galaxy spectra,

typically through cross-correlation techniques with reference templates, which capture the full available information (e.g. Tonry & Davis 1979). Despite the considerable advances of multi-object spectrographs over the past 40 yr, collecting spectra for large samples of galaxies remains an expensive task. A cheaper, lower-precision alternative is offered by photometric estimates, that is, by measurements based on multi-band imaging, in which integrated low-resolution spectral information is collected at once for large numbers of objects over large areas. The price to be paid is that of larger measurement errors, together with a number of catastrophic failures, which limit the scientific usage of such photometric redshifts (photo-*z*s hereafter) to specific applications (e.g. Newman & Gruen 2022). Still, when a sufficient number of photometric bands is available (Benitez et al. 2014; Laigle et al. 2016; Alarcon et al. 2021) or when even information about the ensemble mean spectrum can be obtained (Cagliari et al. 2022), these samples become highly valuable in many respects. Photo-*z*s are traditionally estimated by fitting template spectral energy distributions (SED) to the measured photometric fluxes (see e.g. Bolzonella et al. 2000; Arnouts et al. 2002; Maraston 2005; Ilbert et al. 2006). Detailed reviews can be found in Salvato et al. (2019), Brescia et al. (2021), and Newman & Gruen (2022).

Since the pioneering work of Collister & Lahav (2004; see also Lahav 1994), who first used artificial neural networks (ANN) to obtain photo-*z* estimates, machine-learning (ML) algorithms have seen many further applications in this




Fig. 1. Correlation between the galaxy redshift and that of its n th nearest angular neighbour ($n = \{1, 2, 3, 4\}$, left to right), as seen in the VIPERS redshift survey data, which cover the range $0.5 < z < 1.2$. Clearly, while a tight correlation exists for a number of objects, many other angular pairs just correspond to chance superpositions.

context. These include random forests (Carliles et al. 2010), self-organizing maps (SOM; Masters et al. 2015), and advanced ANNs (Sadeh et al. 2016). A notable recent application uses the full images of galaxies through convolutional neural networks (CNN; Pasquet et al. 2019; Henghes et al. 2022). All these methods provide photo- z estimates by using information that is strictly local, that is, the flux of each object measured in a number of photometric bands, independently of correlations with the other galaxies in the sample.

In the specific case when a photometric survey includes spectroscopic redshifts for a representative sub-sample spread over the same area, these represent additional information, which can be exploited to obtain improved estimates of the missing redshifts. Since galaxies are spatially clustered, angular neighbours on the sky preserve a degree of redshift correlation, depending on the depth of the catalogue. The deeper the catalogue, the weaker the correlation because the projection is made over a deeper baseline. Still, an angular correlation remains, as can be seen explicitly in Fig. 1, in the data of the VIMOS Public Extragalactic Redshift Survey (VIPERS; Guzzo et al. 2014).

This correlation was exploited, for example, to improve our knowledge of the overall sample redshift distribution (Newman 2008), which is a fundamental quantity for many cosmological investigations such as weak-lensing tomography. With VIPERS, instead, it was used to estimate the galaxy density field to fill the gaps due to missing redshifts (Cucciati et al. 2014). Even more finely, Aragon-Calvo et al. (2015) used the fact that galaxies are typically confined within cosmic web structures to obtain a dramatic improvement in the estimate of photo- z s for ~ 200 million Sloan Digital Sky Survey galaxies, starting from only about one million spectroscopically measured redshifts.

Our goal with the work presented here has been to optimally retrieve this non-local information from the neighbouring objects of a given galaxy building upon a specific class of ML architectures, graph neural networks (GNN). The key property of this class is the ability to combine information from unstructured data based on our priors of the task at hand (Bronstein et al. 2017). The end goal is to obtain an improved estimate of the galaxy redshift.

As shown by Fig. 1, the existing correlation between angular neighbours is strongly diluted by the sea of chance superpositions along the line of sight. Thus, the problem can be more appropriately recast into quantifying the probability that a given angular neighbour (with known redshift) is a physical companion for a given galaxy and thus is closely correlated in redshift as well. Our GNN model, dubbed NezNet, combines the intrinsic

features of a target galaxy and a neighbour, that is, their multi-band fluxes, the spectroscopic redshift of the neighbour, and their relative angular distance, to output the probability for the two galaxies to be spatially correlated. We trained and tested NezNet using the spectroscopic sample of VIPERS. We show that discarding targets for which no real physical neighbour is identified with significant probability improves the quality of the associated photo- z catalogue obtained through classic SED fitting, increasing precision and accuracy and reducing the fraction of catastrophic outliers. Moreover, when real neighbours are identified, the redshift of the highest-probability neighbour represents an estimate of the target redshift that is typically more precise than that obtained through the classical SED fitting.

The idea of using GNNs to draw additional redshift information from neighbouring galaxies is not new. Beck & Sadowski (2019) presented preliminary results of an approach based on using only the photometry of a neighbourhood of galaxies, obtaining a 10% improvement on the median absolute deviation of the photo- z s estimated via a single object-based ML algorithm. The main shortcoming of methods that are based on apparent neighbours lies in the large fraction of chance superpositions, as evident in Fig. 1. Here, we reformulated the problem as a detection task that identifies the physical neighbours of the surrounding spectroscopic objects, also including the neighbour's spectroscopic information. In this way, we obtain a significant improvement.

The paper is organised as follows. In Sect. 2, we give a brief description of how GNNs work and specify the architecture of our model. In Sect. 3, we describe the properties of VIPERS data and the way we prepared the training set, in particular, how we defined real or apparent neighbouring objects. Section 4 describes how the model is applied to the data and the metrics we used to quantify the performance of the results. Finally, in Sect. 5 we present and discuss our results, and we conclude in Sect. 6.

2. Model

A neural network model can be summarised as a set of non-linear functions applied to a set of inputs that undergo a linear mapping. Each mapping has many parameters that are optimised through a training process that allows the network model to approximate a wide variety of almost arbitrary functions (LeCun et al. 2015). In its simplest form, a neural network model corresponds to a multi-layer perceptron (MLP), also known as dense neural network (Murtagh 1991). For images, neural architectures such




Fig. 2. Schematic architecture of NezNet. The input features are first processed by a dense network. Message passing between the two layers through Eq. (1) is then applied to take the relative differences and global values of the features into account. Before the final dense layer, the features are summed and then reprocessed with an MLP to output the score probability of two galaxies being actual neighbours.

as CNN are more suited because they take our a priori knowledge about the data structure into account (O’Shea & Nash 2015).

This reasoning can be pushed further by introducing neural networks for graph representations (Zhou et al. 2018). In this work, we make use of one key aspect of GNN, that is, message passing (Gilmer et al. 2017). To fix ideas, the problem we wish to address is the following: we need to find the spectroscopic galaxies with the highest probability of being close to a galaxy for which only photometric information is available. This can be recast as a classification task for each pair of galaxies, in which our aim is to distinguish between apparent and real neighbours when projected on the plane of the sky.

Intuitively, a model that distinguishes between apparent and real neighbours should be based on the relative difference between galaxy features. A neural network like this can be designed by including a layer of the form

$$\mathbf{x}'_i = \sum_{j \in \mathcal{N}(i)} h(\mathbf{x}_i, \mathbf{x}_i - \mathbf{x}_j), \quad (1)$$

where \mathbf{x}_i refers to the array of input features of the node i , $\mathcal{N}(i)$ is the neighbourhood of the same node, \sum is the aggregation function that sums the outcomes from each pair of nodes. The function h is an MLP that explicitly combines the value of the input feature at the node and the relative difference of that feature with respect to the neighbour. It is worth noting that this GNN is both permutation equivariant and permutation invariant, so that it is not affected by a change in the order of the nodes, that is, the input galaxies.

The complete architecture of our model is illustrated in Fig. 2. Each node is a galaxy, whose inputs (e.g. the photometric measurements) were pre-processed through an MLP before undergoing the message passing of Eq. (1). We restricted ourselves to the case of galaxy pairs, so that the neighbourhood $\mathcal{N}(j)$ includes only one galaxy, and the aggregation function simply sums the features $\mathbf{x}'_1 + \mathbf{x}'_2$. This model can be seen as a trivial version of EdgeConv (Wang et al. 2018), where the adjacency matrix is a 2×2 matrix, with 0 entries for diagonal elements and 1 for the off-diagonal elements. Finally, the summed features undergo a last dense layer with a scalar output. All the activation functions are rectified linear units, with the exception of the final layer, where we used a sigmoid, to represent a probability for our classification task.

We call this classification model Nearest- z Network (NezNet). NezNet provides the probability for a pair of galaxies to be real neighbours. The loss function adopted to train NezNet is a standard binary cross entropy,

$$\mathcal{L} = \frac{1}{n} \sum_i^n [y_i \log p_i + (1 - y_i) \log (1 - p_i)], \quad (2)$$

where p_i is the output probability of NezNet for each galaxy pair, while $y_i = 0, 1$ is the corresponding training label, and the sum is averaged over the mini-batch. To design our model, we made use of the Spektral library¹ (Grattarola & Alippi 2020), where the EdgeConv layer is conveniently already implemented.

3. Data

We trained and tested our approach on the final data release of VIPERS (Guzzo et al. 2014; Scodreggio et al. 2018), for which the redshift correlation between angular neighbours is shown in Fig. 1. The survey used the VIMOS multi-object spectrograph at the ESO Very Large Telescope to target galaxies brighter than $i_{AB} = 22.5$ in the Canada-France-Hawaii Telescope Legacy Survey Wide (CFHTLS-Wide) catalogue, with an additional $(r - i)$ vs. $(u - g)$ colour pre-selection to remove objects at $z < 0.5$. The resulting sample covers the redshift range $0.5 \lesssim z \lesssim 1.2$, with an effective sky coverage of 16.3 deg^2 , split over the W1 and W4 fields of CFHTLS-Wide. We used only galaxies with secure redshift measurements, as identified by their quality flag, corresponding to a 96.1% confidence level (see Scodreggio et al. 2018).

For each galaxy in the catalogue, the following information was considered: the spectroscopic redshift measurement z_{spec} , the six magnitudes u, g, r, i, z (not to be confused with redshift) and K_s , the right ascension α (RA), in radians, and the declination δ (Dec), in radians.

The angular separation on the sky between two objects with RA α_1 and α_2 and Dec δ_1 and δ_2 is given by the haversine formula,

$$\Delta\Theta = \arccos(\sin \delta_1 \sin \delta_2 + \cos \delta_1 \cos \delta_2 \cos(\alpha_1 - \alpha_2)). \quad (3)$$

We selected the parent photometric sample by applying the same VIPERS colour and magnitude cuts defined above, so as to be fully coherent with the spectroscopic data.

4. Application

We set up a training set from the VIPERS W1 galaxy catalogue. We randomly selected about 3×10^4 target galaxies, whose spectroscopic redshift during training was ignored. For each of them, we identified the first n_{NN} angular nearest neighbours as defined by Eq. (3), which we called spectroscopic galaxies because their spectroscopic redshift information was used in our model. Each of these spectroscopic neighbours was associated with the same target galaxy, but the pairs can be considered as independent from one another in our model. Each angular pair was

¹ <https://graphneural.network>

assigned label 1 when it was a real physical pair, otherwise, it was assigned a 0. The training set was thus made of galaxy pairs.

A target galaxy of a pair can also be the nearest neighbour of another target galaxy in another pair. We made this choice in order to maximise the number of training examples available in W1. Our final tests on the W4 catalogue show that this does not lead to any over-fitting of VIPERS data, as the model generalises well. We note that this setting assumes a ratio of spectroscopic to photometric objects of 1 : 1. In the Conclusions section (Sect. 6), we also confirm these results in the more realistic case in which the number of spectroscopic redshifts used for training are a fraction of the number of photometric objects.

The definition of a real neighbour is arbitrary; it is reasonable to consider that two angular neighbours form a physical pair when their spectroscopic separation is smaller than a given threshold,

$$\Delta z(1 + z_{\text{spec}}). \quad (4)$$

This means that in setting up the training data, there are two hyper-parameters, the number of nearest neighbours n_{NN} to be considered, and the spectroscopic separation Δz . As we show below, these two hyper-parameters can affect the results significantly, and it is thus relevant to set them up wisely, depending on the specific survey.

For each galaxy in the pairs, the input features of the nodes in NezNet are the photometry, the spectroscopy, and the angular position, as listed in Sect. 3. For the target galaxy, we always set $z_{\text{spec}} = 0$, so that the model considered it as a missing feature, while providing its value for the neighbouring galaxy. Magnitudes were normalised to the range [0, 1], as computed over the whole VIPERS dataset. The angular inputs were provided in terms of relative distance with respect to the target galaxy, so that $\Delta\Theta = 0$ for the latter, while for the neighbour, it corresponded to Eq. (3). By adopting this choice, we guaranteed that the model has translational invariance.

Another tested option (see Sect. 6) is to use the relative distance in the two sky coordinates RA and Dec as input variables instead of the angular separation of the two galaxies. This choice arises because the surface distribution of the sample is not rotationally invariant on the sky because of the technical set-up of the slits in the VIMOS focal plane, with the spectral dispersion oriented along the declination direction. As spectra must not overlap on the detector, targets need to be separated in Dec much more than in RA. As a result, the minimum separation is ~ 1.9 arcmin in Dec and 5 arcsec in RA. More details can be found in Bottini et al. (2005) and Pezzotta et al. (2017, see their Sect. 4.1). Our experiments show that providing the model with the angular separation $\Delta\Theta$ introduces a bias in the redshift metrics, which is not observed when the relative separations along RA and Dec are given. In general, however, we find that the separation information does not significantly improve the classifier, and for this reason, we did not use it in our final model. Spatial information instead comes only from the number of nearest neighbours considered.

The other hyperparameters of the model, that is, the batch size, number of neurons, and learning rate, have a far weaker impact than Δz and n_{NN} , and were set to fiducial values: a batch size of 32, a learning rate of 0.001, and a total number of parameters of the order of a few thousands. We find little difference in the output metrics of the redshift estimates when the complexity of the model is increased, or when the batch size and the learning rate are changed around these fiducial values.

NezNet gives as output the probability for two galaxies to be real neighbours. As each target galaxy corresponds to n_{NN}

independent pairs, we can select the neighbour with the highest probability among them. If this probability is below the classification threshold set to define a positive case, we conclude that there is no physical neighbour for that target galaxy in the catalogue. This implies that the probability for the latter is too high to be an outlier in terms of its properties when compared to its neighbours. Removing these objects from the final catalogue significantly improves the metrics when comparing photo- z and spectroscopic measurements. In particular, the reduction in the number of catastrophic redshifts confirms our assumption. Finding a true neighbour instead reinforces the confidence in the photo- z . At the same time, the spectroscopic redshift of the neighbour in this case is typically an even better estimate of the target redshift than the SED-estimated photo- z . These tests are discussed in the following section.

The quantitative comparison between NezNet results, spectroscopic measurements $z_{\text{spec}}^{(i)}$, and SED-fitting estimated photo- z s was performed using the metrics defined in Salvato et al. (2019). These are the precision (i.e. the dispersion of the estimated values),

$$\sigma = \sqrt{\frac{1}{N} \sum_i \left(\frac{z_{\text{spec}}^{(i)} - z^{(i)}}{1 + z_{\text{spec}}^{(i)}} \right)^2}, \quad (5)$$

the bias

$$b = \frac{1}{N} \sum_i (z_{\text{spec}}^{(i)} - z^{(i)}), \quad (6)$$

and the absolute bias

$$|b| = \frac{1}{N} \sum_i |z_{\text{spec}}^{(i)} - z^{(i)}|, \quad (7)$$

quantifying systematic deviations. Finally, the outliers are defined as objects for which

$$|z_{\text{spec}}^{(i)} - z^{(i)}| \geq 0.15(1 + z_{\text{spec}}^{(i)}). \quad (8)$$

All the results presented in the following section were obtained by applying the trained NezNet to a test catalogue built in a similar fashion to W1, randomly selecting about 2×10^4 galaxies from the twin W4 field of VIPERS.

Finally, in the following discussion about our classifier, we use the notion of the true positive rate (TPR), which is the fraction of correctly predicted positive examples with respect to all the real positive examples. It is defined as

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (9)$$

where TP stands for true positives and FN stands for false negatives. Similarly, we can define the false positive rate (FPR), which is the fraction of negative examples classified as positives with respect to all the real negative examples, which reads

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}}, \quad (10)$$

where FP stands for false positives and TN stands for true negatives.




Fig. 3. Central galaxy spectroscopic redshift versus its photometric redshift measured with and without NezNet. The left panel shows the distribution of photometric vs. spectroscopic estimates in the original data. In the middle panel, we show the same distribution after removing the galaxies with low score probability from the catalogue (fr stands for the fraction of retained data). Finally, the right panel shows redshift estimates by assigning the spectroscopic redshift of the neighbour with the highest detection probability to the target galaxy. The model was trained with $n_{\text{NN}} = 30$ and $\Delta z = 0.08$.




Fig. 4. Same as Fig. 3, but the model was trained with the higher $\Delta z = 0.15$, while $n_{\text{NN}} = 30$ is the same as before.

5. Results

As explained in the previous section, NezNet can be used to simply clean a photo- z sample by discarding low-probability neighbours or to provide an alternative redshift estimate derived from the highest-probability neighbour. This is demonstrated on the test catalogue in Fig. 3 for a model trained using the hyperparameters $\Delta z = 0.08$ and $n_{\text{NN}} = 30$. In addition to the VIPERS spectroscopic redshifts, this comparison also includes the original photo- z s estimated by Moutard et al. (2016) using standard SED fitting. For these and all following results, angular information (i.e. the separation of the two objects on the sky) was not used as an input variable. The reason for this was already mentioned in the previous section, and is discussed again in more detail below.

Figure 3 shows that by simply dismissing the outliers as identified by NezNet, all the metrics improve significantly (central panel). Moreover, when the best neighbour redshifts are adopted for the target galaxies (right panel), we obtain metrics that are comparable to or even better than those of the cleaned photo- z sample. It is worth noting that in this case, the plot shows a characteristic checkerboard pattern because the spectroscopic redshift stripping is reflected, as spectroscopic redshifts are now assigned to target photometric objects.

Figure 3 also shows the limits of the method. Comparing the left panel with the other two, we can note that NezNet tends

to cut off the high-redshift tail of the distribution. This is easily understood considering the magnitude-limited ($i_{\text{AB}} < 22.5$) character of the sample used here, which becomes very sparse at $z \gtrsim 1$, where only rare luminous galaxies are present. This means that the model becomes intrinsically less efficient because fewer real physical neighbours are available both for the training and for inference, as is also evident from the density of points at high redshift in Fig. 1. Devising a different loss function to up-weight the few physical pairs in this regime might improve the classification task, but an intrinsic limit to the method clearly exists when the density of the sample decreases.

Figure 4 shows the same set of plots, but using a higher value for the spectroscopic separation in the training, that is, $\Delta z = 0.15$. As expected, allowing for a larger separation in the definition of real angular neighbours discards fewer data. Conversely, there is in general a lower precision and a small increase in the fraction of outliers.

In principle, using a stricter Δz could remove even more outliers, retaining only pairs that are closer in redshift and leading to a smaller, but more precise sub-sample. We explore this dependence in Fig. 5. Overall, this method is always able to clean poor estimates from the sample, but at the price of discarding many data points. The minor improvement in precision probably does not justify the use of $\Delta z < 0.08$ in the case of VIPERS, because more than half of the sample is excluded.




Fig. 5. Redshift estimates derived from the best nearest neighbour for various Δz at fixed $n_{\text{NN}} = 30$. Increasing the spectroscopic separation to define physical neighbours while diminishing the quality of the metrics increases the fraction of data that are not dismissed from the catalogue.




Fig. 6. ROC curve for a varying redshift threshold Δz at fixed $n_{\text{NN}} = 30$. The performance of our classifier (AUC) improves as we use a less strict definition of what we define as a true neighbour. The probability that an angular neighbour is a physical neighbour increases at larger Δz , which is also reflected by the high detection threshold (thr).

It is apparent that the hyper-parameter Δz is very relevant for the quality of the classifier. This is made clear by the receiver operator characteristic (ROC) curve in Fig. 6, which shows the TPR (Eq. (9)) against the FPR (Eq. (10)), and has been computed from the target galaxies in the test catalogue by considering their neighbour with the highest probability. In general, the area under the curve (AUC) is higher for the better classifier. Increasing Δz increases the AUC, which would tend to unity for very high values of this parameter, as all galaxies would then be considered real neighbours. However, our ultimate goal is not to increase the performance of the classifier per se, but to improve the metrics of our redshift estimates. These show that $\Delta z \gtrsim 0.08$ represents the best choice for VIPERS.

The other hyper-parameter of NezNet, that is, n_{NN} , the number of nearest neighbours considered in the training, has a weaker impact on the classifier. We show this in Fig. 7, where each ROC curve corresponds to a model trained with a different n_{NN} , but all with the same Δz . A drastic change in n_{NN} does not correspond to comparable changes in the AUC. However, n_{NN} has a large impact on the redshift estimates, as Fig. 8 shows. A larger number of angular neighbours increases the probability of finding a physical pair, as is shown by the metrics in Fig. 7. We also experimented with a higher value of n_{NN} up to 50, but found no further gain with respect to using $n_{\text{NN}} = 30$. The redshift metrics start to saturate to the optimal values already above $n_{\text{NN}} = 10$.




Fig. 7. ROC curve for a varying number of nearest neighbours $n_{\text{NN}} = 30$ at fixed $\Delta z = 0.08$. Increasing the number of neighbours that are given in input to the training seems to make the training more difficult. However, this test of the classifier does not reflect the quality of the final redshift estimate, as Fig. 8 shows.

As a further test, we also computed the gradients of the predictions with respect to their input variables to detect the most relevant ones, as shown in Fig. 9. It is interesting to see that the neighbour redshift is a relevant input, as expected, and some of the photometric bands are even more relevant. This confirms the intuition that the photometric information of the neighbours does indeed provide additional information about the relative distance from the target. In this plot, we also show results for the case when the angular separation is considered as one of the input variables. These results show that the angular separation $\Delta\Theta$ between the target and the neighbour does affect the predictions. This manifests itself as a bias in the redshift estimates, as visible in Fig. 10: in this case, NezNet systematically favours neighbours that are closer to us than the target, increasing the value of the bias b (Eq. (6)). We also tested what happens when the angular separation information is rather given in terms of the relative difference in the angular coordinates RA and Dec of the two galaxies. In this case, the bias disappears and the results are comparable to the standard case in which no angle information is provided. However, in this case, the two parameters clearly have smaller gradients than when $\Delta\Theta$ alone is considered, which suggests that they do in fact not contribute to the predicting power of the model. For these reasons, the angular separation is not considered as input variable in our final results.




Fig. 8. Redshift estimates based on the best nearest neighbour for various n_{NN} at fixed $\Delta z = 0.08$. Increasing the number of nearest neighbours for each target improves the performance of NezNet in estimating redshifts, as it increases the probability that physical pairs are considered.

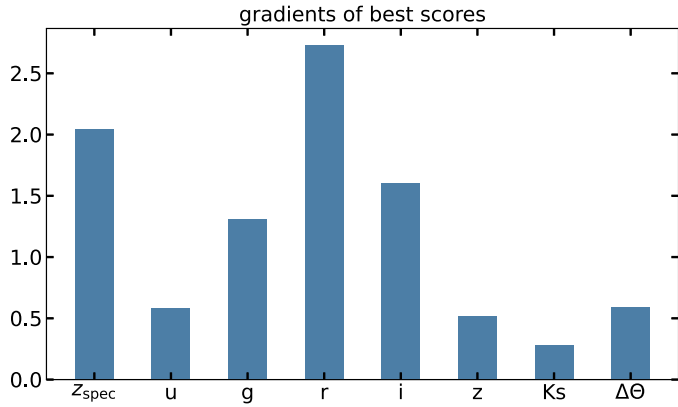


Fig. 9. Average absolute values of the gradients of NezNet with respect to the input features of the neighbours. For each target, we only considered the neighbour with the highest probability.




Fig. 11. Comparison of the redshift distribution for the predictions of NezNet, and a simpler graph model without message passing. While the latter performs reasonably well in general, it tends to cut the tail of the distribution.




Fig. 10. Results of redshift estimates for the target galaxies, in the case where the angular separation Eq. (3) is an explicit input of the model. Many galaxies have slightly lower values than the real spectroscopic value, resulting in a large bias b . Currently, we do not have an explanation of this observed effect.

One of the novelties of NezNet is the message passing between node features. This is where GNNs differ from a standard ANN, where all input variables of both galaxies would be provided directly to dense layers. We also experimented with a simpler graph model, closely resembling the architecture of NezNet, but without message passing. The input features were processed independently by MLP layers for each node (we tried using either just one or several layers). The new architecture is as in Fig. 2, with the exception of \mathbf{h} function blocks, which are now substituted with new MLP blocks, without applying any message passing. The \mathbf{x}_i features are summed by the aggregation function, and the summed features are mapped to the output probability through final dense layers with sigmoid activation output, just like in the model with message passing. This kind of model, which maintains the permutation invariance property of a graph, is often referred to as a deep set (Zaheer et al. 2017). We find that this simple model still works remarkably well and is comparable to NezNet in general. However, it systematically cuts off the high-redshift tail of the catalogue (Fig. 11), even though the overall metrics remain good.

6. Conclusions

We have presented a new ML model, dubbed NezNet, which for a pair of galaxies takes as input their measured fluxes in a number of bands together with the redshift of one of the two




Fig. 12. Redshift estimates based on the best nearest neighbour, obtained by uniformly subsampling the W1 catalogue, at fixed $n_{\text{NN}} = 30$ and $\Delta z = 0.08$. The titles of the panels refer to the surface density of spectroscopic objects of W1 used for training, with Σ referring to the complete W1 sample. Except for minor fluctuations in the redshift statistics, NezNet maintains a performance similar to the case without subsampling. The only noticeable trend is the fraction of central galaxies for which a physical pair is found, which decreases for lower densities. This could be due to the decreasing number of available training data. The percentage of real physical neighbours for a central galaxy, which decreases only slightly from Σ to $\Sigma/8$, remains around 40% and explains why NezNet is still effective.

galaxies. NezNet is capable of probabilistically learning whether their redshift distance is below a given threshold Δz , which is set as a hyper-parameter of the model. The angular separation between the galaxies is implicit in the training set, as for every target galaxy we select its first n_{NN} angular neighbours (another hyper-parameter), but it can be an explicit input variable of the model. The backbone of the model is a GNN, a class of neural networks based on message passing and the aggregation of features (Fig. 2). This message passing is explicitly performed as a relative difference between features (Eq. (1)).

NezNet outputs the score probability for a galaxy pair to be real neighbours. This information that can be used in two ways. On the one hand, if none of the n_{NN} nearest neighbours is identified as a physical neighbour, the target galaxy can be considered an outlier in terms of its properties. This may suggest that it is an interloper, that is, a foreground or background object with respect to the volume sampled by the spectroscopic sample we used for the comparison. It should therefore be discarded from any sample that aims to cover the same redshift range as the spectroscopic catalogue, for instance, via photometrically estimated redshifts. We have proved this to be true using the VIPERS catalogue. On the other hand, if a physical neighbour is identified, the target galaxy can be assigned the spectroscopic redshift of the highest scoring galaxy among the n_{NN} angular neighbours, providing an independent estimate of its redshift in this way.

These results are summarised in Figs. 3 and 4: when outliers as detected by NezNet are discarded, all the metrics of the sample improve considerably. Moreover, the NezNet redshift estimates are comparable to or superior in precision to SED-based photometric redshifts, depending on the values chosen for the hyper-parameters. Increasing Δz increases the goodness of the classifier (Fig. 6), as well as the fraction of retained data (Fig. 5). Changing n_{NN} has a smaller impact on the classifier (Fig. 7), although it significantly affects the redshift quality metrics because a large enough n_{NN} improves the probability of detecting a real neighbour; a value $n_{\text{NN}} \sim 30$ is optimal in the case of VIPERS (Fig. 8).

It is often the case that the fraction of the parent photometric sample without a spectroscopic measurement has a higher density than the spectroscopic sample. VIPERS indeed has a spectroscopic surface density of $\Sigma \sim 6 \times 10^3/\text{deg}^2$, to compare against the photometric surface density $\Sigma_{\text{ph}} \sim 45 \times 10^3/\text{deg}^2$. For

this reason, we tested NezNet by varying the surface density of the spectroscopic sample used during training. We achieved this by repeating the training procedure on a uniformly subsampled catalogue extracted from W1. The test was performed on W4 without any subsampling, so that we tested for the effectiveness of NezNet trained on a lower-density catalogue. Figure 12 shows that NezNet keeps its effectiveness even when using a subsample of one-eighth of the original spectroscopic density Σ , similar to the VIPERS ratio of spectroscopic to photometric objects.

This suggests that NezNet could have an interesting potential also in the context of future experiments, such as *Euclid* or the NASA *Nancy Grace Roman* mission (Akeson et al. 2019). These slitless spectroscopic surveys will indeed naturally deliver overlapping photometric and spectroscopic data, which can be combined using NezNet to improve photometric redshift estimates.

It is worth stressing that some details of the results presented here depend on the specific features of VIPERS and its parent CFHTLS photometric sample. Some of them may have been advantageous, but others could have penalised the success of the method. For example, the slit-placement constraints in VIPERS limits the ability to target close galaxy pairs, which introduces a shadow in the layout of a VIMOS pointing (see Fig. 6 of Guzzo et al. 2014), and forces a lower limit in the separation of observable galaxy pairs (see Sect. 4). This means that the training sample of NezNet was not ideal in our analysis because surely many of the missed angular pairs were also physical pairs. This increases our confidence in the obtained results because it shows that for samples that are characterised by small-scale incompleteness, as is typical of surveys built using fibre or multi-slit spectrographs, the method still also delivers very useful results. In the case of the VIPERS data, an interesting exercise in this respect would be to use the data from the VLT-VIMOS Deep Survey (VVDS; Le Fèvre et al. 2005) as training sample, which used the same spectrograph, but with repeated passes over the same area of 0.5 deg^2 that substantially mitigate the proximity bias. We leave this exercise for a future work.

Acknowledgements. We thank Davide Bianchi for useful suggestions during the development of this work. FT and MSC are thankful to Daniele Grattarola for insightful discussions on GNNs and the use of the Spektral library. We thank the anonymous referee for his comments and suggestions. FT and LG acknowledge financial support by grant MUR PRIN 2017 ‘From Darklight to Dark Matter’.

grant no. 20179P3PKJ. LG and MSC acknowledge financial support from the Italian Space Agency, ASI agreement no. I/023/12/0.

References

- Akeson, R., Armus, L., Bachelet, E., et al. 2019, ArXiv e-prints [arXiv:1902.05569]
- Alam, S., Ata, M., Bailey, S., et al. 2017, *MNRAS*, **470**, 2617
- Alarcon, A., Gaztanaga, E., Eriksen, M., et al. 2021, *MNRAS*, **501**, 6103
- Aragon-Calvo, M. A., van de Weygaert, R., Jones, B. J. T., & Mobasher, B. 2015, *MNRAS*, **454**, 463
- Arnouts, S., Moscardini, L., Vanzella, E., et al. 2002, *MNRAS*, **329**, 355
- Bautista, J. E., Paviot, R., Vargas Magaña, M., et al. 2021, *MNRAS*, **500**, 736
- Beck, R., Sadowski, P., Glaser, Y., & Szapudi, I. 2019, Refined Redshift Regression in Cosmology with Graph Convolution Networks, https://ml4physicalsciences.github.io/2019/files/NeurIPS_ML4PS_2019_80.pdf
- Benitez, N., Dupke, R., Moles, M., et al. 2014, ArXiv e-prints [arXiv:1403.5237]
- Blake, C., Brough, S., Colless, M., et al. 2011, *MNRAS*, **415**, 2876
- Bolzonella, M., Miralles, J. M., & Pelló, R. 2000, *A&A*, **363**, 476
- Bottini, D., Garilli, B., Maccagni, D., et al. 2005, *PASP*, **117**, 996
- Brescia, M., Cavuoti, S., Razim, O., et al. 2021, *Front. Astron. Space Sci.*, **8**
- Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., & Vandergheynst, P. 2017, *IEEE Signal Process. Mag.*, **34**, 18
- Cagliari, M. S., Granett, B. R., Guzzo, L., et al. 2022, *A&A*, **660**, A9
- Carliles, S., Budavári, T., Heinis, S., Priebe, C., & Szalay, A. S. 2010, *ApJ*, **712**, 511
- Colless, M., Peterson, B. A., Jackson, C., et al. 2003, ArXiv e-prints [arXiv:astro-ph/0306581]
- Collister, A. A., & Lahav, O. 2004, *PASP*, **116**, 345
- Cucciati, O., Granett, B. R., Branchini, E., et al. 2014, *A&A*, **565**, A67
- de la Torre, S., Jullo, E., Giocoli, C., et al. 2017, *A&A*, **608**, A44
- DESI Collaboration (Aghamousa, A. et al.) 2016, ArXiv e-prints [arXiv:1611.00036]
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. 2017, ArXiv e-prints [arXiv:1704.01212]
- Grattarola, D., & Alippi, C. 2020, ArXiv e-prints [arXiv:2006.12138]
- Guzzo, L., Scodeggio, M., Garilli, B., et al. 2014, *A&A*, **566**, A108
- Henghes, B., Thyagalingam, J., Pettitt, C., Hey, T., & Lahav, O. 2022, *MNRAS*, **512**, 1696
- Ilbert, O., Arnouts, S., McCracken, H. J., et al. 2006, *A&A*, **457**, 841
- Lahav, O. 1994, *Vistas Astron.*, **38**, 251
- Laigle, C., McCracken, H. J., Ilbert, O., et al. 2016, *ApJS*, **224**, 24
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, ArXiv e-prints [arXiv:1110.3193]
- Le Fèvre, O., Vettolani, G., Garilli, B., et al. 2005, *A&A*, **439**, 845
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Nature*, **521**, 436
- Maraston, C. 2005, *MNRAS*, **362**, 799
- Masters, D., Capak, P., Stern, D., et al. 2015, *ApJ*, **813**, 53
- Moutard, T., Arnouts, S., Ilbert, O., et al. 2016, *A&A*, **590**, A102
- Murtagh, F. 1991, *Neurocomputing*, **2**, 183
- Newman, J. A. 2008, *ApJ*, **684**, 88
- Newman, J. A., & Gruen, D. 2022, *ARA&A*, **60**, 363
- O'Shea, K., & Nash, R. 2015, ArXiv e-prints [arXiv:1511.08458]
- Pasquet, J., Bertin, E., Treyer, M., Arnouts, S., & Fouchez, D. 2019, *A&A*, **621**, A26
- Perlmutter, S., Aldering, G., della Valle, M., et al. 1998, *Nature*, **391**, 51
- Pezzotta, A., de la Torre, S., Bel, J., et al. 2017, *A&A*, **604**, A33
- Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, *AJ*, **116**, 1009
- Sadeh, I., Abdalla, F. B., & Lahav, O. 2016, *PASP*, **128**, 104502
- Salvato, M., Ilbert, O., & Hoyle, B. 2019, *Nat. Astron.*, **3**, 212
- Scodeggio, M., Guzzo, L., Garilli, B., et al. 2018, *A&A*, **609**, A84
- Tegmark, M., Eisenstein, D. J., Strauss, M. A., et al. 2006, *Phys. Rev. D*, **74**, 123507
- Tonry, J., & Davis, M. 1979, *AJ*, **84**, 1511
- Wang, Y., Sun, Y., Liu, Z., et al. 2018, ArXiv e-prints [arXiv:1801.07829]
- Zaheer, M., Kottur, S., Ravanbakhsh, S., et al. 2017, ArXiv e-prints [arXiv:1703.06114]
- Zhou, J., Cui, G., Hu, S., et al. 2018, ArXiv e-prints [arXiv:1812.08434]