A catalogue of photometric redshifts for the SDSS-DR9 galaxies[*]
(Research Note)

M. Brescia[1], S. Cavuoti[1], G. Longo[2], and V. De Stefano[2]

[1] INAF Astronomical Observatory of Capodimonte, via Moiariello 16, 80131 Napoli, Italy
    e-mail: brescia@oacn.inaf.it
[2] University Federico II, Department of Physics, via Cinthia, 80126 Napoli, Italy

ABSTRACT

Context. Accurate photometric redshifts for large samples of galaxies are among the main products of modern multiband digital surveys. Over the last decade, the Sloan Digital Sky Survey (SDSS) has become a sort of benchmark against which to test the various methods.
Aims. We present an application of a new method to the estimation of photometric redshifts for the galaxies in the SDSS Data Release 9 (SDSS-DR9). Photometric redshifts for more than 143 million galaxies were produced.
Methods. The Multi Layer Perceptron with Quasi Newton Algorithm (MLPQNA) model, provided within the framework of the DAta Mining and Exploration Web Application REsource (DAMEWARE), is an interpolative method derived from machine learning models.
Results. The obtained redshifts have an overall uncertainty of $\sigma = 0.023$ with a very small average bias of $\sim 3 \times 10^{-5}$, and a fraction of catastrophic outliers ($|\Delta z| > 2\sigma$) of $\sim 5\%$. This result is slightly better than what was already available in the literature in terms of the smaller fraction of catastrophic outliers as well.

Key words. techniques: photometric – galaxies: distances and redshifts – galaxies: photometry – methods: data analysis – catalogs

1. Introduction

In the last few years, photometric redshifts (photo-z) for large samples of normal or active galaxies have become crucial for a variety of cosmological applications (Scranton et al. 2005; Myers et al. 2006; Hennawi et al. 2006; Giannantonio et al. 2008) and many different methods of evaluating them have been presented and extensively discussed in the literature (cf. Hildebrandt et al. 2010). The problem of deriving accurate photometric redshifts has become even more cogent because of the huge amount of data produced by most ongoing and planned photometric surveys (cf. PANNSTARS: Kaiser 2004, KIDS[1], EUCLID: Laureijs et al. 2011) aimed at explaining weak lensing to prove the dark components of the universe.

Without entering into details that can be found elsewhere, it is worth a reminder that, broadly speaking, all photo-z methods are based on the interpolation of some a priori knowledge represented by sets of templates, and differ only in one or both of the following aspects: (i) the way in which the a priori knowledge base (KB) is constructed (higher accuracy spectroscopic redshifts or, rather, empirically or theoretically derived spectral energy distributions or SEDs); and (ii) the interpolation/fitting algorithm employed.

In all methods, the main source of uncertainty is the fact that the function mapping the colour space into the spectroscopic redshift space is just an oversimplified approximation of the complex and otherwise unknown relation existing between colours and the redshift (as an example, see Csabai et al. 2003). Among the various interpolative methods, we shall just quote a few: i) polynomial fitting (Connolly et al. 1995); ii) nearest neighbours (Csabai et al. 2003); iii) neural networks (D'Abrusco et al. 2007; Yéche et al. 2010 and references therein); iv) support vector machines (Wadadekar 2005); v) regression trees (Carliles et al. 2010); vi) Gaussian processes (Way & Srivastava 2006; Bonfield et al. 2010); and vii) diffusion maps (Freeman et al. 2009).

In this paper, we focus on the application to the galaxies contained in the SDSS Data Release 9 (DR9, Paris et al. 2012), of the Multi Layer Perceptron with Quasi Newton Algorithm (MLPQNA) method already described in detail elsewhere (Brescia et al. 2012, 2013), hence we refer to these papers for all the mathematical and technical details. In the framework of the PHAT1 contest (Hildebrandt et al. 2010), which blindly compared most existing methods for photo-z evaluation, the MLPQNA method proved to be among the two best empirical methods to date (Cavuoti et al. 2012). This in spite of the very limited base of knowledge available for the contest (~ 500 objects only).

The MLPQNA is just one among the many data mining methods publicly available under the DAta Mining and Exploration Web Application REsource infrastructure (DAMEWARE; Brescia et al. 2014).

In the next section (Sect. 2), we describe the data set used as the knowledge base, while in Sect. 3 we describe the experiments and discuss their outcome. In Sect. 4, we describe the resulting catalogue of photometric redshifts and a short summary of the work is given in Sect. 5.

---

[*] The produced catalogue, composed by 58 tables is only available at the CDS via anonymous ftp to
cdsarc.u-strasbg.fr (130.79.128.5) or via
http://cdsarc.u-strasbg.fr/viz-bin/qcat?J/A+A/568/A126
[1] http://www.astro-wise.org/projects/KIDS/
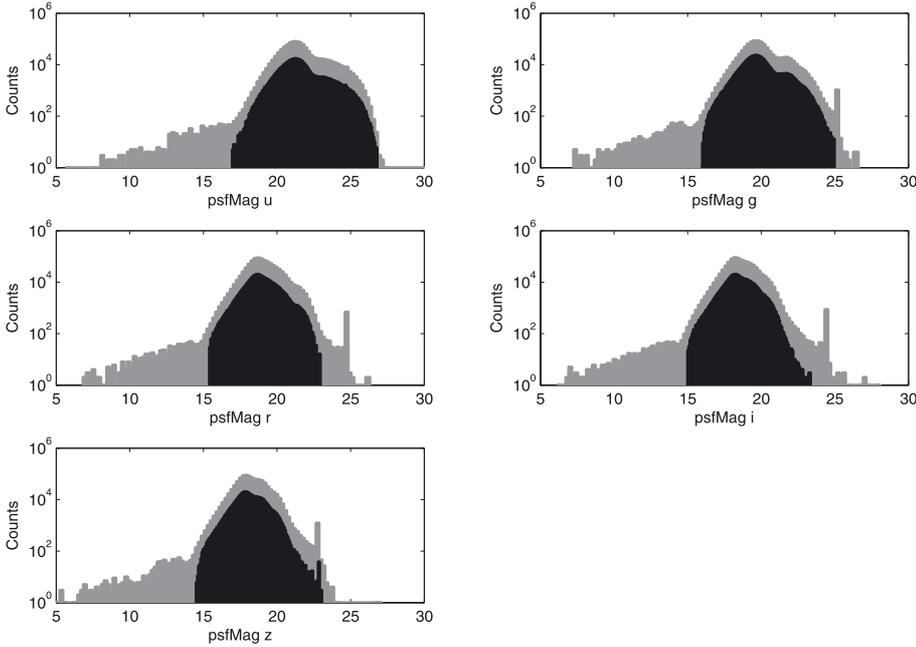
2. The data

The Sloan Digital Sky Survey (SDSS, York et al. 2000), is the forerunner of modern wide-field surveys. It combines multi-band photometry and fiber-based spectroscopy, thus providing both photometric data for a very large number of objects and spectroscopic information for a smaller, but still significant, sub-sample of the same population. Hence, it provides all information needed to constrain the fit of an interpolating function mapping the photometric features into the spectroscopic redshift space. This is the main reason why most, if not all, photometric redshifts methods have been tested on the various data releases of the SDSS which, over the years, has become a sort of benchmark data set against which to test old and new methods.

To form our KB we extracted from the spectroscopic sub-sample of the SDSS-DR9 all objects with specClass galaxy together with their photometry. In particular, we used the ($psfMag$) magnitudes and related colours, rejecting all objects with missing or non-detected information in any of the SDSS photometric bands.

We obtain the cuts in the magnitude by considering the limits within which the photometric parameter space sampled by the spectroscopic objects is significantly covered. Within these limits, the neural algorithm during the training phase is exposed, in every region of the cleaned parameter space, to a number of examples sufficiently large to facilitate learning. Obviously, the less populated the region of the parameter space, the smaller the accuracy of the final result. An additional implication is that the less populated a region of the parameter space, the smaller the capability to correctly learn the rule for peculiar or rare objects. The resulting psf magnitude limits are listed in Table 1, while Fig. 1 shows the psf magnitude distributions in the knowledge base. As also described in Oyaizu et al. 2008, we trained our model on the spectroscopic sample up to the magnitude limit of $psfMag\_r < 23.0$. By considering a photometric limit of $r < 22.2$, the resulting fainter limit in the training set covers the complete photometric region of interest without introducing boundary effects for photometric redshifts of galaxies having magnitudes near the $psfMag\_r$ limit. All this is taken

Table 1. The $psfMag$ type magnitude limits derived in each band during the knowledge base definition.

Band	Lower limit	Upper limit
u	16.97	26.79
g	15.91	24.97
r	15.41	22.94
i	15.00	23.34
z	14.17	23.04

into account by the complete spectroscopic KB, which consists of 497, 339 objects.

3. Experiments and discussion

In machine learning supervised methods it is common practice to use the available KB to build at least three disjoint subsets for every experiment: one (training set) for training purposes, i.e. to train the method in order to acquire the hidden correlation among the input features necessary to perform the regression; the second one (validation set) to check the training, in particular, against a loss of generalization capabilities (a phenomenon also known as overfitting); and the third one (test set) to evaluate the overall performances of the model (Brescia et al. 2013).

In this work, the validation was performed during training by applying the standard leave-one-out k-fold cross validation mechanism (Geisser 1975). We would like to stress that none of the objects included in the training (and validation) sample were included in the test sample and only the test data were used to generate the statistics. In other words, the test was blind, i.e. based only on objects never submitted to the network.

We decided to populate the training and the test set with 30% and 70% of the objects in the KB, namely with 149, 997 and 347, 342 objects, respectively. This decision, which might seem a little anomalous since it is common practice for machine learning methods to operate with data sets of reversed proportion, was dictated by the large number of examples present in the
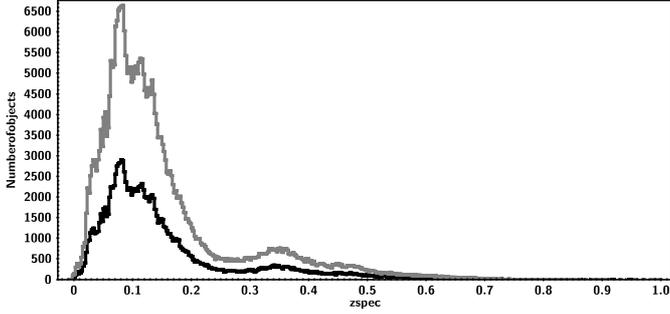
Fig. 2. Spectroscopic redshift distribution of objects included in the training set (black line) and test set (grey line).
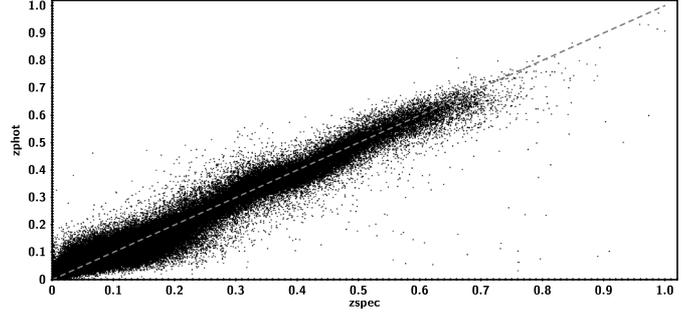
Fig. 3. Spectroscopic versus photometric redshifts for SDSS *DR*9 galaxies of the blind test set.

knowledge base and by the specificity of the MLPQNA method, which can *overfit* (with a loss of generalization capability) the data when exposed to a very large number of examples. The histogram in Fig. 2 shows the distribution of the objects in the KB as a function of the zspec in both training and test sets.

In order to ensure that the KB provided a proper coverage of the parameter space, the data were split into the two data sets by random extraction. In other words, by randomly shuffling and splitting the original dataset, we replicated several times the extraction sequence, and evaluated the average of their output. This mechanism prevents possible biases induced by fluctuations in the coverage of the parameter space, namely small differences in the redshift distribution of training and test samples used in the experiments.

Once the data sets were produced, we checked which types of flux combinations were more effective, in terms of magnitudes or related colours, and therefore we performed and compared two experiments with two different sets of features: (i) MAG; using the five SDSS *psfMag* magnitudes; and (ii) MIXED, replacing the magnitudes with the derived colours and leaving only the *r psfMag* as pivot magnitude. The best combination turned out to be the MIXED type. From the data mining point of view this is rather surprising since the amount of information should not change when applying linear combinations between features. From the physical point of view, however, the better performances of the MIXED experiment can be easily understood by noticing that even though colours are derived as a subtraction of magnitudes, the content of information is quite different, since an ordering relationship is implicitly assumed, thus increasing the amount of information in the final output (i.e., flux gradients instead of fluxes). The additional pivot magnitude used in the experiment serves to remove the degeneracy in the luminosity class for a specific galaxy type.

Individual experiments, as well as their comparison with results provided by others, were evaluated in a consistent and objective manner using a homogeneous and standard set of statistical indicators:

- the bias, defined as the mean value of the residuals $\Delta z = z_{spec} - z_{phot}$;
- the standard deviation (σ) of the residuals;
- the normalized median absolute deviation or NMAD of the residuals, defined as $\text{NMAD}(\Delta z) = 1.48 \times \text{Median}(|\Delta z|)$ and;
- all the above quantities also calculated on the normalized residuals, i.e. $\Delta z_{norm} = \frac{z_{spec} - z_{phot}}{1 + z_{spec}}$.

Furthermore, as an overall estimate of the accuracy of the final results we use the prescription in CLSI 2006, deriving the overall uncertainty (OU) defined as $\text{bias}_{norm} \pm \sigma_{norm}$.

In Table 2, we list our results, which we also compare with those obtained by Laurino et al. (2011). These authors, on the SDSS Data Release 7 objects, used a machine learning model with a slightly more complex architecture, called weak gated experts (WGE) and which, to the best of our knowledge has achieved the highest accuracy to date.

In the second half of the table, we list the fraction of outliers, i.e. the fraction of objects for which the photometric redshift estimate deviates more than 0.15 in absolute value, or deviating more than 1, or 2σ from the spectroscopic value.

Summarizing, MLPQNA achieves the very small bias of $\sim 3 \times 10^{-5}$, and a normalized standard deviation of 0.023. Moreover, our method leads to a very small fraction of outliers, i.e. less than 0.04% and $\sim 5.4\%$ using the $|\Delta z_{norm}| > 0.15$ and the 2σ criteria, respectively.

In Fig. 3, we plot the photometric redshift estimates versus the spectroscopic redshift values for all objects in the test set. After the rejection of catastrophic outliers, as defined by the $|\Delta z_{norm}| > 2\sigma(\Delta z_{norm})$, we obtain a σ_{norm} of ~ 0.0174, which is larger than NMAD_{norm}. This is exactly what is to be expected according to Mobasher et al. (2007). In fact, in the case where photo-*z* are empirical, it is always useful to analyse the direct correlation between the NMAD_{norm} and the standard deviation σ_{norm} calculated on data that are not catastrophic outliers. In these cases, a correct photo-*z* prediction occurs whenever the quantity NMAD_{norm} is lower than the σ_{norm} for the *cleaned* sample.

For empirical methods, the $\sigma(\Delta z_{norm})$ overestimates the theoretical Gaussian σ, mainly due to catastrophic outliers and to the intrinsic training error.

In order to better characterize the performances of the experiment, we also computed the statistics on subsets of the test data binned according to either redshift or magnitude range (see Figs. 4 to 7).

As far as redshifts are concerned, as shown in Table 2, we built the first subset using objects in the redshift range starting from $z = 0$ up to the redshift $z_1 = 0.11549$, which includes 50% of the objects in the test set; the second subset uses the range from $z_1 = 0.115$ to $z_2 = 0.177$ (corresponding to an additional 25% of the sample); the third subset uses the range $z_2 = 0.177$ to $z_3 = 0.345$ (corresponding to an additional 15%) and, finally, the fourth subset includes all remaining objects (redshift $> z_3 = 0.345$). We also derived statistics in the redshift range [0.05, 0.6] corresponding to the same range covered in Laurino et al. (2011) (in order to allow for a fair comparison).

The behaviour of the residuals as a function of the magnitude in the SDSS *r* band was instead studied in the three bins listed in Table 3.

Table 2. Results for the MIXED experiment (see text).

Ref.	|bias|	σ	NMAD	rms	$\|bias_{norm}\|$	σ_{norm}	$NMAD_{norm}$	RMS_{norm}
Overall	0.000030	0.028	0.0196	0.028	0.000574	0.023	0.0172	0.023
$[0, z_1]$	0.008259	0.022	0.0179	0.023	0.007895	0.020	0.0167	0.022
$]z_1, z_2]$	0.008975	0.023	0.0195	0.025	0.007828	0.020	0.0171	0.022
$]z_2, z_3]$	0.006664	0.035	0.0234	0.035	0.005640	0.028	0.0188	0.029
$>z_3$	0.008568	0.037	0.0255	0.038	0.005705	0.024	0.0177	0.025
$[0.05, 0.6]$	0.002224	0.026	0.0188	0.026	0.001643	0.022	0.0163	0.022
Laurino et al.	0.015	0.015	0.016	0.021	0.014	0.013	0.013	0.019
	$\|\Delta z\|$ >0.15	$\|\Delta z\|$ >1σ	$\|\Delta z\|$ >2σ	$\|\Delta z_{norm}\|$ >0.15	$\|\Delta z_{norm}\|$ >1σ	$\|\Delta z_{norm}\|$ >2σ	Skewness	Kurtosis
Overall	0.12%	22.61%	5.16%	0.04%	23.57%	5.43%	1.37×10^{-16}	1.8
$[0, z_1]$	0.03%	27.30%	6.56%	0.02%	27.02%	6.62%	9.73×10^{-17}	1.8
$]z_1, z_2]$	0.02%	28.13%	7.16%	0.01%	28.28%	7.20%	1.23×10^{-15}	1.8
$]z_2, z_3]$	0.36%	22.51%	6.57%	0.08%	22.54%	6.67%	8.03×10^{-16}	1.8
$>z_3$	0.50%	20.56%	4.17%	0.16%	22.02%	4.34%	-1.13×10^{-15}	1.8
$[0.05, 0.6]$	0.11%	22.84%	5.44%	0.03%	23.62%	5.51%	2.24×10^{-15}	1.8

Notes. The table is split in two parts. In the upper half we provide standard statistical indicators (see text for an explanation) used to evaluate the performances of photo-z methods. We also include the same indicators for Laurino et al. (2011). In the lower half of the table, we report the fraction (in percentage) of outliers computed using a fixed threshold of 0.15, the more meaningful 1, and 2σ clipping thresholds and the values of skewness and kurtosis of the σ(Δz_{norm}) distributions. The redshift range starting from $z = 0$ up to the redshift $z_1 = 0.115$ includes 50% of the objects in the test set; the second redshift bin, using the range from $z_1 = 0.115$ to $z_2 = 0.177$, corresponds to an additional 25% of the sample; the third bin, using the range $z_2 = 0.177$ to $z_3 = 0.345$, corresponds to an additional 15% and, finally, the fourth bin includes all remaining objects (redshift $>z_3 = 0.345$). The last row reports the statistics in the redshift range [0.05, 0.6] corresponding to the same range covered in Laurino et al. (2011).

Table 3. Test set data distribution with r mag binning.

Mag bin (r)	Test objects	$\|bias_{norm}\|$	σ_{norm}	Quality flag	Skewness	Kurtosis
All	347, 342	0.000574	0.023	–	1.37×10^{-16}	1.8
[15.4, 20.4]	321, 514	0.000539	0.022	1	-1.92×10^{-17}	1.8
]20.4, 22.2]	25, 375	0.001067	0.032	2	6.34×10^{-18}	1.8
]22.2, 23.0[	453	0.001745	0.028	3	4.37×10^{-16}	1.8

Notes. The table is ordered on the r mag intervals (first column). The second column reports the number of objects within the related bin. Column 3 lists the assigned photo-z quality flag. Finally, Cols. 4 and 5 list the skewness and kurtosis of the distribution of the residuals, respectively.

Using these information, we assigned a photo-z quality flag (from 1 as best quality to 3 as worst quality) to all objects in all r mag bins, following both the r mag completeness limit and the σ_{norm} trend as criteria. The results are summarized in Table 3.

As expected, the σ_{norm} error also still remains acceptable slightly outside the r magnitude completeness limit ($r < 22.2$). In this region, however, the number of training points is rather small, and for the reasons stated above, the predicted redshifts need to be taken with some caution since, given the selection criteria applied to select the targets for the spectroscopic survey, it is very likely that not all galaxy types are present in the knowledge base and that the much wider population of objects with photometric observations only is not well represented in the training set.

In the already mentioned PHAT1 contest, however, MLPQNA obtained very good results using a KB of size (i.e. ~500 objects) similar to that used for training in the last magnitude bin.

In order to better characterize the distribution of the residuals in terms of Gaussianity of the distributions, we fitted a Gaussianity to the residuals in the various quality bins obtaining the kurtosis and skewness listed in Table 3. The distributions

of residuals appear to be quite symmetric even though they are slightly leptokurtic.

4. The photometric catalogue

To produce the final catalogue, we downloaded from the SDSS DR9 data server[2] all objects falling within the declination range [$-30°$, $+85°$] and detected in all SDSS bands.

We underline that all empirical photo-z methods suffer from a poor capability to extrapolate outside the range of distributions imposed by the training. In other words, outside the limits of magnitudes and zspec used in the training set, these methods do not ensure optimal performances. hence, to remain in a safe condition, we performed a selection of objects in the final catalogue according to the same selection done on the training and test sample limits.

Furthermore, the SDSS DR9 hosts objects that are spectroscopically recognized as galaxies, but whose photometric class is different. In most cases, such objects are photometrically classified as stars. From the spectral point of view indeed, there is

---

[2] http://skyserver.sdss3.org/CasJobs/
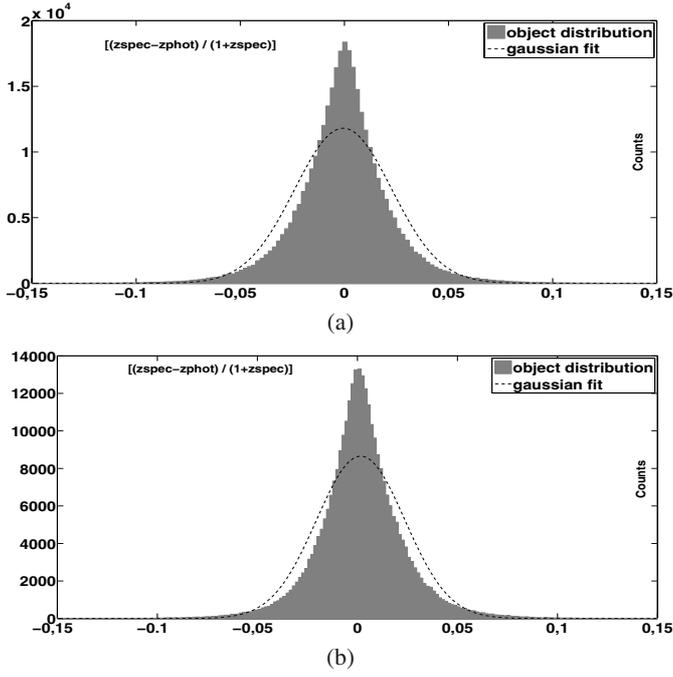
Fig. 4. Distribution of the residuals in two redshift bins. All objects **a)**, zspec within the range [0.05, 0.6] **b)**. The plotted range is $-0.15 \leq |\Delta z_{\mathrm{norm}}| \leq 0.15$.

Fig. 6. Distribution of the residuals in zspec bins, respectively, $[0, z_1]$ **a)**, $]z_1, z_2]$ **b)**. The plotted range is $-0.15 \leq |\Delta z| \leq 0.15$. The limits are $z_1 = 0.115$, $z_2 = 0.177$, respectively.

Fig. 5. Distribution of the residuals in three $psfMag\_r$ magnitude bins, respectively, [15.4, 20.4] **a)**,]20.4, 22.2] **b)** and]22.2, 23.0[**c)**. The plotted range is $-0.15 \leq |\Delta z_{\mathrm{norm}}| \leq 0.15$.
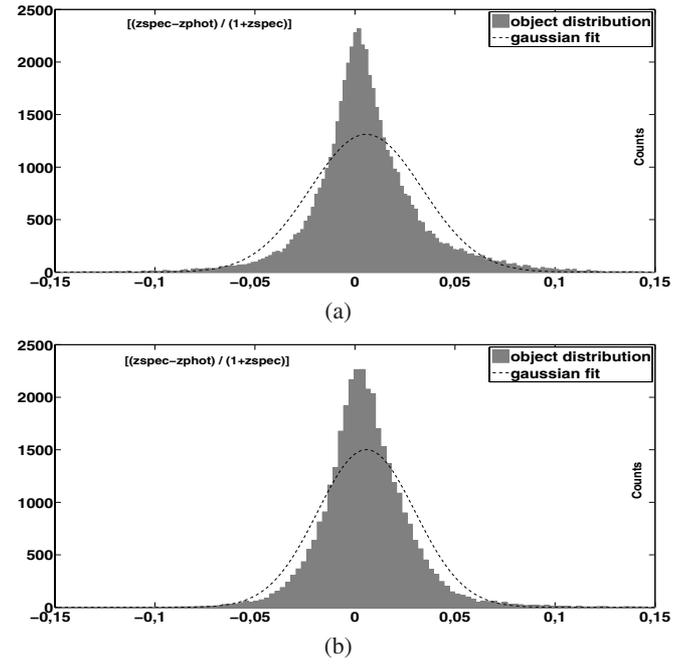
Fig. 7. Distribution of the residuals in zspec bins, respectively, $]z_2, z_3]$ **a)**, $>z_3$ **b)**. The plotted range is $-0.15 \leq |\Delta z| \leq 0.15$. The limits are $z_2 = 0.177$, $z_3 = 0.345$, respectively.

a zspec value assigned to most of such objects, although they are lost from any photometric search based on galaxy type. Hence, for reasons of completeness, we added such objects to the photo-z photometric catalogue, retrieving them through a special SQL query (see Appendix C).

For convenience, the whole catalogue was split in 59 files, containing a total of 143 584 848 objects with the estimated photo-z. Among them, the file with suffix *specialObjects*

includes the photo-*z* for special objects with a mismatch between photometric and spectroscopic class assignment. The other 58 files forming the final catalogue correspond to different declination ranges and each is structured in 24 columns containing:

- Col. 1: the SDSS-DR9 object identification;
- Cols. 2 and 3: right ascension and declination;
- Cols. 4–8: the *u, g, r, i*, and *z* PSF magnitudes;
- Cols. 9–13: the *psfMag_err* error for all magnitudes;
- Cols. 14–18: the extinction for each magnitude;
- Cols. 19–22: the colours derived from *ugriz psfMag* type magnitudes;
- Col. 23: the estimated photo-*z*; and
- Col. 24: quality flag of the photo-*z* obtained from the information gathered during the analysis of the test set. The value 1 stands for the best photo-*z* accuracy, the value 2 for photo-*z* with lower accuracy, and the value 3 for the photo-*z* related to objects outside the *psfMag_r* completeness limit.

5. Conclusions

The MLPQNA neural network was applied to the SDSS-DR9 photometric galaxy data, using a knowledge base derived from the SDSS-DR9 spectroscopic sub-sample.

After a set of experiments, best results were obtained with a two hidden layer network, using a combination of the four SDSS colours (obtained from the SDSS *psfMag*) plus the pivot magnitude *psfMag* in the *r* band. This yields a normalized overall uncertainty of $\sigma = 0.023$ with a very small average bias of $\sim 3 \times 10^{-5}$, a low NMAD, and a low fraction of outliers ($\sim 5\%$ at 2σ and $\sim 0.1\%$ at 0.15). After the rejection of catastrophic outliers, the residual uncertainty is $\sigma = 0.0174$.

The trained network was then used to process the galaxies in the SDSS-DR9, matching the above outlined selection criteria, and to produce the complete photometric catalogue. This catalogue consists of photo-*z* estimates for more than 143 million SDSS-DR9 galaxies.

Appendix A: Spectroscopic query

The following SQL code has been used to obtain the spectroscopic KB to train and test the model.

```
SELECT
    p.objid, s.specObjID, p.ra, p.dec,
    p.psfMag_u, p.psfMag_g, p.psfMag_r, p.psfMag_i,
    p.psfMag_z, p.psfmagerr_u, p.psfmagerr_g,
    p.psfmagerr_r, p.psfmagerr_i, p.psfmagerr_z,
    p.fiberMag_u, p.fiberMag_g, p.fiberMag_r,
    p.fiberMag_i, p.fiberMag_z, p.fibermagerr_u,
    p.fibermagerr_g, p.fibermagerr_r,
    p.fibermagerr_i, p.fibermagerr_z,
    p.petroMag_u, p.petroMag_g, p.petroMag_r,
    p.petroMag_i, p.petroMag_z, p.petromagerr_u,
    p.petromagerr_g, p.petromagerr_r,
    p.petromagerr_i, p.petromagerr_z,
```

```
    p.modelMag_u, p.modelMag_g, p.modelMag_r,
    p.modelMag_i, p.modelMag_z,
    p.modelmagerr_u, p.modelmagerr_g,
    p.modelmagerr_r, p.modelmagerr_i,
    p.modelmagerr_z,
    p.extinction_u, p.extinction_g,
    p.extinction_r, p.extinction_i,
    p.extinction_z, s.z as zspec,
    s.zErr as zspec_err, s.zWarning,
    s.class, s.subclass, s.primTarget
INTO
    mydb.galaxies_spec
FROM
    PhotoObjAll as p,
    SpecObj as s
WHERE
    s.class = 'GALAXY' AND s.zWarning = 0 AND
    p.mode = 1 AND p.SpecObjID = s.SpecObjID AND
    dbo.fPhotoFlags('PEAKCENTER') != 0 AND
    dbo.fPhotoFlags('NOTCHECKED') != 0 AND
    dbo.fPhotoFlags('DEBLEND_NOPEAK') != 0 AND
    dbo.fPhotoFlags('PSF_FLUX_INTERP') != 0 AND
    dbo.fPhotoFlags('BAD_COUNTS_ERROR') != 0 AND
    dbo.fPhotoFlags('INTERP_CENTER') != 0
```

Appendix B: Photometric query

The produced photometric catalogue with the estimated photo-*z* has been taken from SDSS DR9 service, by applying the following SQL query. The reported code here is referred to a DEC range between 60 and 65 deg, as an example.

```
SELECT
    p.objid, p.ra, p.dec,
    p.psfMag_u, p.psfMag_g, p.psfMag_r,
    p.psfMag_i, p.psfMag_z,
    p.psfmagerr_u,  p.psfmagerr_g,
    p.psfmagerr_r,  p.psfmagerr_i,
    p.psfmagerr_z,
    p.extinction_u, p.extinction_g,
    p.extinction_r, p.extinction_i,
    p.extinction_z
INTO
    mydb.p60_p65
FROM
    Galaxy as p
WHERE
    p.mode = 1 AND
    dbo.fPhotoFlags('PEAKCENTER') != 0 AND
    dbo.fPhotoFlags('NOTCHECKED') != 0 AND
    dbo.fPhotoFlags('DEBLEND_NOPEAK') != 0 AND
    dbo.fPhotoFlags('PSF_FLUX_INTERP') != 0 AND
    dbo.fPhotoFlags('BAD_COUNTS_ERROR') != 0 AND
    dbo.fPhotoFlags('INTERP_CENTER') != 0 AND
    p.dec >= 60 AND p.dec <65
```

Appendix C: Special query

Below we report the SQL code used for the query needed to integrate the photo-*z* catalogue with objects spectroscopically recognized as galaxies, but photometrically assigned to different classes within the SDSS DR9.

```
SELECT
    p.objid, p.ra, p.dec,
    p.psfMag_u, p.psfMag_g, p.psfMag_r,
    p.psfMag_i, p.psfMag_z,
    p.psfmagerr_u, p.psfmagerr_g,
    p.psfmagerr_r,  p.psfmagerr_i,
    p.psfmagerr_z, p.extinction_u,
    p.extinction_g, p.extinction_r,
    p.extinction_i, p.extinction_z
INTO
    mydb.photoerror
FROM
    PhotoObjAll as p,
    SpecObj as s
WHERE
    s.class = 'GALAXY' AND p.type != 3 AND
p.mode = 1 AND
    dbo.fPhotoFlags('PEAKCENTER') != 0 AND
    dbo.fPhotoFlags('NOTCHECKED') != 0 AND
    dbo.fPhotoFlags('DEBLEND_NOPEAK') != 0 AND
    dbo.fPhotoFlags('PSF_FLUX_INTERP') != 0 AND
    dbo.fPhotoFlags('BAD_COUNTS_ERROR') != 0 AND
    dbo.fPhotoFlags('INTERP_CENTER') != 0 AND
    p.SpecObjID = s.SpecObjID
```

References

Bonfield, D. G., Sun, Y., Davey, N., et al. 2010, MNRAS, 485, 987

Brescia, M., Cavuoti, S., Paolillo, M., Longo, G., & Puzia, T. 2012, MNRAS, 421, 1155

Brescia, M., Cavuoti, S., D'Abrusco, R., Longo, G., & Mercurio, A. 2013, ApJ, 772, 140

Brescia, M., Cavuoti, S., Longo, G., et al. 2014, PASP, in press [arXiv:1406.3538]

Carliles, S., Budavári, T., Heinis, S., Priebe, C., & Szalay, A. S. 2010, ApJ, 712, 511

Cavuoti, S., Brescia, M., Longo, G., & Mercurio, A. 2012, A&A, 546, A13

Connolly, A. J., Csabai, I., Szalay, A. S., et al. 1995, AJ, 110, 2655

CLSI (Clinical and Laboratory Standard Institute) 2006, User Verification of performance for precision and trueness, 25, No. 17

Csabai, I., Budavári, T., Connolly, A. J., et al. 2003, AJ, 125, 580

D'Abrusco, R., Longo, G., & Walton, N. 2007, ApJ, 663, 752

Freeman, P. E., Newman, J. A., Lee, A. B., Richards, J. W., & Schafer, C. M. 2009, MNRAS, 398, 2012

Geisser, S. 1975, J. Am. Stat. Assoc., 70, 320

Giannantonio, T., Scranton, R., Crittenden, R. G., et al. 2008. Phys. Rev. D, 77, 3520

Hennawi, J. F., Strauss, M. A., Oguri, M., et al. 2006, AJ, 131, 1

Hildebrandt, H., Arnouts, S., Capak, P., et al. 2010, A&A, 523, A31

Kaiser, N. 2004, Proc. SPIE, 5489, 11

Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, ESA/SRE(2011)12, Issue 1.1 [arXiv:1110.3193]

Laurino, O., D'Abrusco, R., Longo, G., & Riccio, G. 2011, MNRAS, 418, 2165

Myers, A. D., Brunner, R. J., Richards, G. T., et al. 2006, ApJ, 638, 622

Mobasher, B., Capak, P., Scoville, N. Z., et al. 2007, ApJSS, 172, 117

Oyaizu, H., Lima, M., Cunha, C. E., et al. 2008, ApJ, 674, 2, 768

Paris, I., Petitjean, P., Aubourg, É., et al. 2012, A&A, 548, A66

Scranton, R., Menard, B., Richards, G. T., et al. 2005, ApJ, 633, 589

Wadadekar, Y. 2005, PASP, 117, 79

Way, M. J., & Srivastava, A. N. 2006, ApJ, 647, 102

Yéche, C., Petitjean, P., Rich, J., et al. 2010, A&A, 523, A14

York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, AJ, 120, 1579