

Photometric redshift estimation via deep learning

Generalized and pre-classification-less, image based, fully probabilistic redshifts

A. D’Isanto and K. L. Polsterer

Astroinformatics, Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnenweg 35, 69118 Heidelberg, Germany
e-mail: [antonio.disanto;kai.polsterer]@h-its.org

Received 7 June 2017 / Accepted 20 July 2017

ABSTRACT

Context. The need to analyze the available large synoptic multi-band surveys drives the development of new data-analysis methods. Photometric redshift estimation is one field of application where such new methods improved the results, substantially. Up to now, the vast majority of applied redshift estimation methods have utilized photometric features.

Aims. We aim to develop a method to derive probabilistic photometric redshift directly from multi-band imaging data, rendering pre-classification of objects and feature extraction obsolete.

Methods. A modified version of a deep convolutional network was combined with a mixture density network. The estimates are expressed as Gaussian mixture models representing the probability density functions (PDFs) in the redshift space. In addition to the traditional scores, the continuous ranked probability score (CRPS) and the probability integral transform (PIT) were applied as performance criteria. We have adopted a feature based random forest and a plain mixture density network to compare performances on experiments with data from SDSS (DR9).

Results. We show that the proposed method is able to predict redshift PDFs independently from the type of source, for example galaxies, quasars or stars. Thereby the prediction performance is better than both presented reference methods and is comparable to results from the literature.

Conclusions. The presented method is extremely general and allows us to solve of any kind of probabilistic regression problems based on imaging data, for example estimating metallicity or star formation rate of galaxies. This kind of methodology is tremendously important for the next generation of surveys.

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

1. Introduction

In recent years, the availability of large synoptic multi-band surveys increased the need of new and more efficient data analysis methods. The astronomical community is currently experiencing a data deluge. Machine learning and, in particular, deep-learning technologies are increasing in popularity and can deliver a solution to automate complex tasks on large data sets. Similar trends can be observed in the business sector for companies like Google and Facebook. In astronomy, machine learning techniques have been applied to many different uses (Ball & Brunner 2010). Redshift estimation is just one relevant field of application for these statistical methods. Constant improvements in performances had been achieved by adopting and modifying machine learning approaches. The need for precise redshift estimation is increasing due to its importance in cosmology (Blake & Bridle 2005). For example, the *Euclid* mission (Laureijs et al. 2012) is highly dependent on accurate redshift values. Unfortunately, measuring the redshift directly is a time consuming and expensive task as strong spectral features have to be clearly recognized. Therefore, redshifts extracted via photometry based models provide a good alternative (Beck et al. 2016). At the price of a lower accuracy compared to the spectroscopic measurements, photometric redshift estimates enable the processing of huge numbers of sources (Abdalla et al. 2011). Moreover, by combining photometric and spectroscopic techniques, low signal-to-noise spectra of faint objects can be better calibrated and processed (Fernández-Soto et al. 2001).

Photometric redshift estimation methods found in the literature can be divided in two categories: template based spectral energy distribution (SED) fitting (e.g. Bolzonella et al. 2000; Salvato et al. 2009) and statistical and/or machine learning algorithms (e.g. Benítez 2000; Laurino et al. 2011). In this work we will focus on the latter ones and in particular on the application of deep-learning models. Most machine learning approaches demand a large knowledge base to train the model. Once trained, such models allow us to process huge amounts of data and to automatically generate catalogs with millions of sources (as done in Brescia et al. 2014). Instead of generating point estimates only, extracting density distributions that grant access to the uncertainty in the prediction is gaining more focus in the community (Carrasco Kind & Brunner 2013). This scheme is increasingly important for the success of the *Euclid* mission, which depends on highly precise and affordable probabilistic photometric redshifts (Dubath et al. 2017).

Due to the advent of faster and more specialized compute architectures as well as improvements in the design and optimization of very deep and large artificial neural networks, more complex tasks could be solved. In recent years, the so called field of deep-learning was hyped together with big-data-analytics, both in the commercial sector as well as in science. Astronomy has always been a data-intense science but the next generation of survey missions will generate a data-tsunami. Projects such as the Large Synoptic Survey Telescope (LSST) as well as the Square Kilometre Array (SKA) are just some examples that demand

processing and storage of data in the peta- and exabyte regime. Deep-learning models could provide a possible solution to analyze those data-sets, even though those large networks are lacking the ability to be completely interpretable by humans.

In this work, the challenge of deriving redshift values from photometric data is addressed. Besides template fitting approaches, several machine-learning models have been used in the past to deal with this problem (Collister & Lahav 2004; Laurino et al. 2011; Polsterer et al. 2013; Cavuoti et al. 2015; Hoyle 2016). Up to now, the estimation of photometric redshifts was mostly based on features that had been extracted in advance. A Le-Net deep convolutional network (DCN; LeCun et al. 1998) is able to automatically derive a set of features based on imaging data and thereby make the extraction of tailored features obsolete. Most machine-learning based photometric redshift estimation approaches found in the literature just generate single value estimates. Template fitting approaches typically deliver a probability density function (PDF) in the redshift space. Therefore currently the machine learning based approaches are either modified or redesigned to deliver PDFs as well (Sadeh et al. 2016). We have proposed a combined model of a mixture density network (MDN; Bishop 1994) and a DCN. The MDN hereby replaces the fully-connected feed-forward part of the DCN to directly generate density distributions as output. This enables the model to use images directly as input and thereby utilize more information in contrast to using a condensed and restricted feature-based input set. The convolutional part of the DCN automatically extracts useful information from the images which are used as inputs for the MDN. The conceptual idea of combining a DCN with a MDN together with details on the implementation have been presented to the computer science community in D’Isanto & Polsterer (2017) while this publication addresses the challenges of photometric redshift estimation on real world data by using the proposed method. A broad range of network structures have been evaluated with respect to the performed experiments. The layout of the best performing one is presented here.

The performance of the proposed image-based model is compared to two feature-based reference methods: a modified version of the widely used random forest (RF; Breiman 2001) which is able to produce PDFs and a plain MDN. The predicted photometric redshifts are expressed as PDFs using a Gaussian mixture model (GMM). This allows the capture of the uncertainty of the prediction in a compact format. Due to degeneracies in the physical problem of identifying a complex spectral energy distribution with a few broadband filters only, multi-modal results are expected. A single Gaussian is not enough to represent the photometric redshift PDF. When using PDFs instead of point estimates, proper statistical analysis tools must be used, taking into account the probabilistic representation. The continuous ranked probability score (CRPS; Hersbach 2000) is a proper score that is used in the field of weather forecasting and expresses how well a predicted PDF represents the true value. In this work, the CRPS reflects how well the photometrically estimate PDF represents the spectroscopically measured redshift value. The CRPS is calculated as the integral taken over the squared difference between the cumulative distribution function (CDF) and the Heaviside step function of the true redshift value. In contrast to the likelihood, the CRPS is more focused on the location and not on the sharpness of the prediction. We have adopted the probability integral transform (PIT; Gneiting et al. 2005) to check the sharpness and calibration of the predicted PDFs. By plotting a histogram of the CDF values at the true redshift over all predictions, overdispersion, underdispersion and

biases can be visually detected. We demonstrate that the CRPS and the PIT are proper tools with respect to the traditional scores used in astronomy. A detailed description of how to calculate and evaluate the CRPS and the PIT are given in Appendix A.

The experiments were performed using data from the Sloan Digital Sky Survey (SDSS-DR9; Ahn et al. 2012). A combination of the SDSS-DR7 Quasar Catalog (Schneider et al. 2010) and the SDSS-DR9 Quasar Catalog (Pâris et al. 2012) as well as two subsamples of 200 000 galaxies and 200 000 stars from SDSS-DR9 are used in Sect. 3.

In Sect. 2 the machine learning models used in this work are described. The layout of the experiments and the used data are described in Sect. 3. Next, in Sect. 4 the results of the experiments are presented and analyzed. Finally, in Sect. 5 a summary of the whole work is given. In the appendix, an introduction on CRPS and PIT is given, alongside with the discussion of the applied loss function. Furthermore, we motivate the choice of the number of components in the GMM used to describe the predicted PDFs. Finally, the SQL queries used to download the training and test data as well as references to the developed code are listed.

2. Models

In this section the different models used for the experiments are described. The RF is used as a basic reference method while the MDN and the DCN are used to compare the difference between feature based and image based approaches.

2.1. Random forest

The RF is a widely used supervised model in astronomy to solve regression and classification problems. By partitioning the high dimensional features space, predictions can be efficiently generated. Bagging as well as bootstrapping make those ensembles of decision trees statistically very robust. Therefore RF is often used as a reference method to compare performances of new approaches. The RF is intended to be used on a set of input features which could be plain magnitudes or colors in the case of photometric redshift estimation (Carfiles et al. 2010). In its original design, the RF does not generate density distributions. To produce PDFs, the results of every decision tree in the forest are collected and a Gaussian mixture model (GMM) is fitted. The PDF is then presented by a mixture of n components:

$$p(x) = \sum_{j=1}^n \omega_j \mathcal{N}(x|\mu_j, \sigma_j), \quad (1)$$

where $\mathcal{N}(x|\mu_j, \sigma_j)$ is a normal distribution with a given mean μ_j and standard deviation μ_j at a given value x . Each component is weighted by ω_j and all weights sum to one. To calculate the CRPS for a GMM, the equations by Grimit et al. (2006) can be used. For the proposed experiments a model composed of 256 trees is used. As input the five *ugriz* magnitudes and the pairwise color combinations are used, obtaining a feature vector of 15 dimensions per data item.

2.2. Mixture density network

An MDN (Bishop 1994) is a modified version of the widely known multilayer perceptron (MLP; Rosenblatt 1962), producing distributions instead of point estimates. The MLP is a supervised feed-forward neural network, whose main calculation unit

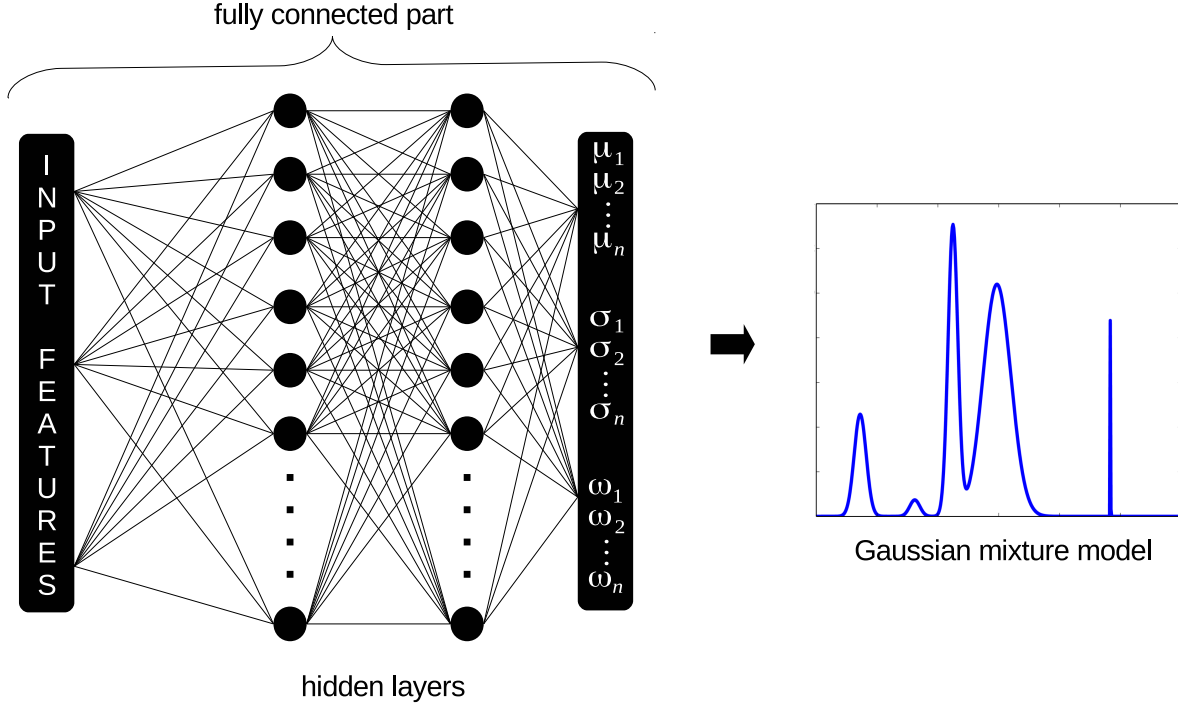


Fig. 1. Architecture of the mixture density network. Next to the input layer, two hidden and fully interconnected layers are depicted. As output vector for each parameter of the GMM is predicted (μ , σ , ω). Based on this compact description, a density distribution can be generated.

is called neuron. The neurons are arranged in layers, an input layer, one or more hidden layers, and an output layer. Several hyperparameters characterize the architecture of an MLP. The activation function is a non-linear function applied to the data as they pass through the neurons. Commonly a sigmoidal function or hyperbolic tangent (\tanh) are utilized. In recent years, the MLP has been widely used in astronomy, for example to estimate redshifts based on photometric features (Brescia et al. 2013).

The MDN interprets the output vector o of the network as the means μ , standard deviations σ and weights ω of a mixture model, using n Gaussians as basis functions:

$$\begin{aligned} \mu_j &= o_j^\mu, \\ \sigma_j &= \exp(o_j^\sigma), \\ \omega_j &= \frac{\exp(o_j^\omega)}{\sum_{i=1}^n \exp(o_i^\omega)} \end{aligned} \quad (2)$$

with $j \in 1..n$ and $o = \{o^\mu, o^\sigma, o^\omega\}$.

Commonly, MDNs are trained using the log-likelihood as loss function. In this application the focus is more on the distribution and shape than on the location of the prediction; hence the CRPS is adopted as loss function. A detailed analysis of the performances of both loss functions is provided in Appendix B. The CRPS increases linearly with the distance from the reference value while the log-likelihood increases with the square of the distance. Like the RF, the MDN is a feature-based model and therefore can use exactly the same input features: plain magnitudes and colors. A generalized architecture of an MDN is shown in Fig. 1.

2.3. Deep convolutional mixture density network

A DCN is a feed-forward neural network in which several convolutional and sub-sampling layers are coupled with a fully-connected part. In some sense, it is a specialization of the fully

interconnected MLP model. By locally limiting the used interconnections in the first part, a spatial sensitivity is realized. Thereby the dimensionality of the inputs is reduced step-wise in every layer. The second part makes use of the so-called feature maps in a flattened representation, using them as input for an ordinary MLP.

This kind of network architecture finds wide application in the fields of image, video and speech recognition due to its capability of performing some kind of dimensionality reduction and automatic feature extraction. A DCN model was chosen because of its ability to deal directly with images, without the need of pre-processing and an explicit features extraction process. The network is trained to capture the important aspects of the input data. By optimizing the condensed representation of the input data in the feature maps, the performance of the fully connected part is improved. Every input data item is a tensor of multi-band images. In the experiments the five *ugriz* filters from SDSS as well as the pixel-wise differences were used. Those image gradients can be compared to the colors in feature based approaches that minimize the effects of object intrinsic variations in luminosity. In the convolutional layers, the input images are spatially convolved with proper filters of fixed dimensions. Hereby “proper” denotes a size that corresponds to the expected spatial extension of distinctive characteristics (filter of the dimension 3×3 and 2×2 have been selected for the application at hand). The filters constitute the receptive field of the model.

After the convolution with the filters, a non-linear function is applied. The \tanh was used in this work as expressed by the following relation:

$$H_k = \tanh(W_k * X + b_k). \quad (3)$$

In general, for every filter k the hidden activation H_k is calculated by using the filter weight matrix W_k and the filter bias b_k . The outputs of the previous layer H are used as inputs X for the succeeding convolutional layer. The first layer is directly fed with the imaging data that should be processed. This filtering system

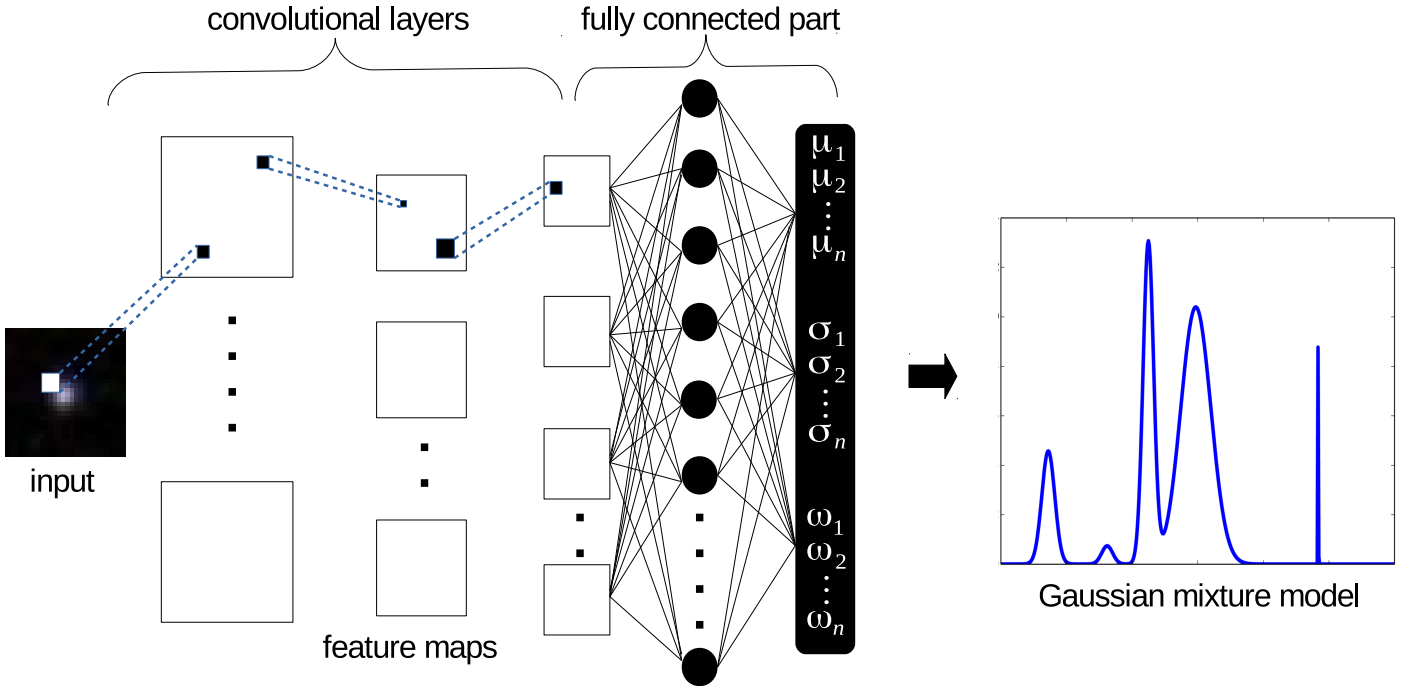


Fig. 2. Architecture of the deep convolutional mixture density network. This architecture directly uses image matrices as input. In this figure, three convolutional layers are drawn, showing the local connection between the different layers. In the fully connected part, a MDN with one hidden layer is used and a vector for each parameter of the GMM is given as result (μ , σ , ω). A sample of a predicted PDF is shown exemplarily.

with local connection make the architecture sensitive to spatial input patterns. The filters are tiled through the entire visual field, with a fixed stride, generating a new representation of the input, that is, the feature maps. A useful tool to pursue the down sampling is to apply max-pooling between the convolutional layers. Those pooling filter typically select only the maximum value in confined region, typically of dimension 2×2 or 3×3 . Convolutional and pooling layers are alternated until the desired complexity in terms of feature maps is obtained. In addition, those layers are alternated with *rectified linear unit (ReLU)* layers, in which the non-linear activation function is substituted by a non saturating function:

$$H_k = \max(0, W_k * X + b_k). \quad (4)$$

This function has many advantages; in particular it increases the efficiency of the gradient descent process, avoiding vanishing gradient problems. Furthermore, using only comparison, addition and multiplication operations it is computationally more efficient. The choice of the activation functions, namely the *tanh* and the *rectified linear unit*, influences the convergence and performance of the neural network. Therefore, the activation function has to be considered as a free parameter when designing the network architecture. In our case, switching to the *rectified linear unit* improved the performance of the predictions notably. Many different combinations have been tested, choosing the best performing one. The feature maps constitute the input for the fully connected part, which has in general the same behavior as a MLP. In the proposed model, we substitute the usual fully connected part with a MDN in order to generate PDFs instead of single point estimates. For this reason, this combined architecture is denoted by us as deep convolutional mixture density network (DCMDN). The structure of the DCMDN is sketched in Fig 2. Furthermore, as for the MDN, the CRPS is used as loss function.

Several hyperparameters influence the layout of the network architecture as well as the training phase. Multiple combinations

have been tested extensively. Due to the immense amount of possible parameter combinations the currently used solution was obtained by clever manual engineering. The most influencing parameters are listed below:

- *global architecture*: this includes the number and types of layers in the local-connected part and the number of hidden layers and neurons characterizing the fully-connected part.
- *activation function*: defines the non-linear function to process the input values of a neuron.
- *number of filters*: influences the number of generated feature maps and therefore the amount of extracted features.
- *filter shape*: characterizes the dimensions of the filters used; it can vary from layer to layer.
- *max-pooling shape*: as for the filters, it specifies the dimension of the area to which the max-pooling is applied.
- *learning rate*: influences the step-size during the gradient descent optimization. This value can decrease with the number of trained epochs.
- *number of epochs*: defines how often the whole training data set is used to optimize the weights and biases.
- *batch size*: as a stochastic gradient descent optimization strategy was chosen, this number defines the amount of training patterns to be used in one training step.

The presented model exhibits many advantages for the application to photometric redshift estimation tasks. It can natively handle images as input and extract feature maps in a fully-automatized way. The DCMDN does not need any sort of pre-processing and pre-classification on the input data and as shown in the experiments section, the performances are far better with respect to the reference models. The reason for an improvement with respect to the estimation performance is the better and extensive use of the available information content of the images. Besides automatically extracting features, their importance with respect to the redshift estimation task is automatically determined too. Feature based approaches make use of the condensed

information that is provided through previously extracted features only. Those features are good for a broad set of applications but not optimized for machine learning methods applied to very specific tasks.

Depending on the size of the network architecture, an extremely high number of weights and biases has to be optimized. This allows the network to adopt to not significant correlations in the data and therefore overfit. While good performances are achieved on the training data, the application to another datasets would exhibit poor results. To limit the effect of overfitting, the dropout technique can be applied; randomly setting a given percentage of the weights in both parts of the network to zero. As deep-learning methods highly benefit from a huge amount of training objects, data augmentation is a common technique. By simply rotating the images by 90° , 180° and 270° and flipping the images, the size of the training data-set can be increased by a factor of four and eight, respectively. This reduces the negative effects of background sources and artifacts on the prediction performance. Moreover, the early stopping technique can be applied to limit the chance of overfitting. As soon as the performance that is evaluated on a separate validation set starts to degrade while the training error is still improving, the training is stopped even before reaching the anticipated number of epochs. The DCMDN is based on the LeNet-5 architecture (LeCun et al. 1998) and realized in Python, making use of the Theano library (Theano Development Team 2016).

The architecture implemented by us is meant to run on graphics processing units (GPUs) as the training on simple central processing units (CPUs) is by far too time consuming. In our experiments a speedup of factor of $\approx 40\times$ between an eight-core CPU (i7 4710MQ 2.50 GHz \times 8) and the GPU hardware was observed. During the experiments a cluster equipped with Nvidia Titan X was intensively used, allowing us to evaluate a larger combination of network architectures and hyperparameters.

3. Experiments

In the following sections the experiments performed with the presented models are described. Those experiments are intended to compare the probabilistic redshift prediction performances of different models on different data-sets. The data-sets used for training the models as well as evaluating the performances are described in the following.

3.1. Data

All data-sets have been compiled using data from SDSS. To generate a set of objects that cover the whole range of redshifts, separate data-sets for quasars and galaxies have been created. The SDSS Quasar Catalog Seventh Data Release (Schneider et al. 2010), containing 105,783 spectroscopically confirmed quasars and the SDSS Quasar Catalog Ninth Data Release (Pâris et al. 2012), containing 87,822 spectroscopically confirmed quasars, are used as basis for the quasar data-set. The two catalogs had to be combined because the former contains confirmed sources from SDSS II only, while the latter is composed of 91% of new objects observed in SDSS III-BOSS (Dawson et al. 2013). In this way a catalog with a much better coverage of the redshift space has been composed (see Fig. 3). Furthermore, two samples composed of 200 000 randomly chosen galaxies and 200 000 randomly picked stars from DR9 (Ahn et al. 2012) have been selected (queries are stated in Appendix F). The objects that have been classified by the spectroscopic pipeline as stars are assigned a redshift of zero. These objects are mandatory to crosscheck

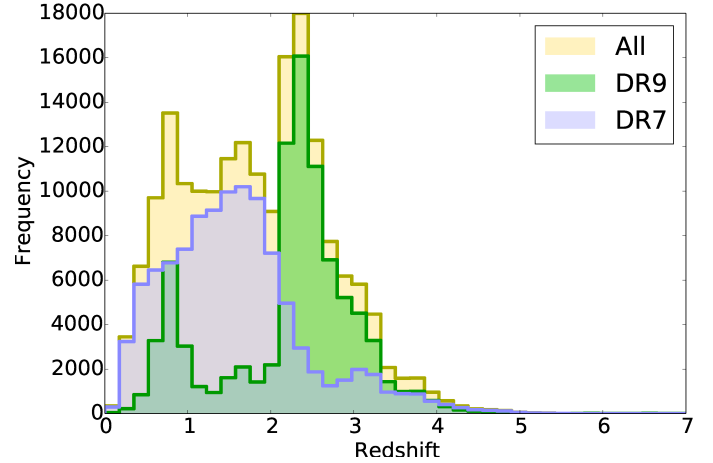


Fig. 3. Redshift distribution of the quasar data from DR7, DR9 and the combined catalog, respectively.

the performance on objects that have not been pre-classified and therefore might be contaminated with stellar sources.

For every catalog the corresponding images from SDSS DR9 in the five *ugriz* filters have been downloaded. The images have been downloaded using the Hierarchical Progressive Surveys (HiPS) data partitioning format (Fernique et al. 2015), in a 28×28 square-pixels format, and subsequently pairwise subtracted, in order to obtain color images corresponding to the colors used in the feature-based experiments. As described above, those pixel wise gradients minimize the effects of object intrinsic variations in luminosity.

3.2. Setup and general configuration

Three different categories of experiments have been performed, each using the three models of Sect. 2 to compare the performances. Before stating the details of the single experiments, a brief description of the general parameters adopted for the models is given. The experiments with the RF and the MDN are feature-based, while the DCMDN is trained on the image representations of the data items. As input features for the two reference models, magnitudes in the *ugriz* filter are used together with all pair-wise color combinations. For the RF models a fixed structure with 256 decision trees was chosen. Bootstrapping was used for creating the individual trees. Neither the depth nor the number of used features have been limited. The MDN architectures use 15 inputs neurons (corresponding to the 15 input features) followed by two hidden layers containing 50 and 100 neurons, respectively. 15 output neurons are used to characterize the parameters of a GMM with $n = 5$. Photometrically estimated redshift distributions are of complex and multimodal nature (Kügler et al. 2016). The choice of 5 Gaussian components is based on experiments where the Bayesian information criterion (BIC) was used as a metric. Depending on the redshift region calculated for, the BIC is indicating values between one and five components. The results presented in Appendix C indicate that five components are, on average, a good choice. For the DCMDN, many different architectures have been tested, comprising less deep and compact convolutional parts. The architecture that gave the best performances was finally created via a clever manual and empirical engineering. The learning rate influences the size of the steps when applying gradient descent during the training with the backpropagation algorithm. For the experiments presented in this work a \tanh was chosen with the learning

Table 1. Layout of the DCMDN architecture for the experiments with the galaxies catalog.

#	Type	Size	# Maps	Activation
1	input	$15 \times 28 \times 28$	/	/
2	convolutional	3×3	256	\tanh
3	pooling	2×2	256	\tanh
4	convolutional	2×2	512	\tanh
5	pooling	2×2	512	\tanh
6	convolutional	3×3	512	$ReLU$
7	convolutional	2×2	1024	$ReLU$
8	MDN	500	/	\tanh
9	MDN	100	/	\tanh
10	output	15	/	Eq. (2)

Notes. Stacks of 15 input images of the size 28×28 square-pixels are used.

rate degrading over the number of epochs. To prevent overfitting, a common technique is to stop training as soon as the training and evaluation error starts to diverge (early stopping). When training both different network architectures an early stopping rule was applied, varying the learning rate from a maximum of 0.01 to a minimum of 0.001 and changing the mini-batch size from 1000 to 500. In the DCMDN, dropout with a ratio of 60% is applied to limit overfitting.

The evaluation of the performances of the models in the different experiments is done by using commonly used scores. The root mean square error (RMSE), the median absolute deviation (MAD), the bias and the σ^2 , are calculated based on the mean of the predicted redshift probability distributions and on the first (most significant) mode. In addition, a normalized version is calculated by weighting the individual prediction errors with the true redshifts before calculating the final scores. Moreover, in the experiments labeled as “selected”, the objects that show a complex and/or multi-modal behavior based on the predicted PDFs are excluded from the score calculation. Those complex objects do not allow a single value as prediction result and therefore do not support a meaningful evaluation through the standard scores. The results without the objects showing ambiguous predictions are presented in addition to the scores obtained on all objects with the previously mentioned performance measures. To be more precise, only objects that fulfill

$$\begin{aligned}
|\mu_1 - \mu_2| &< (\sigma_1 + \sigma_2) \\
\text{or} \\
\omega_1 &> 0.9545,
\end{aligned}
\tag{5}$$

have been selected. Here μ_1, μ_2, σ_1 and σ_2 denote the means and the variances of the first two most significant modes of the predicted PDFs. The modes are chosen based on their weights ω_j , with ω_1 being the weight of the strongest mode. This selection criterion just picks objects where the first two modes are close to each other with respect to their widths or those objects with an extremely dominant first mode. The value of ω_1 was chosen a priori to ensure that 2σ of the joined distribution are represented by the dominant peak.

As PDFs need proper tools and scores for evaluation, the CRPS and PIT for every experiment are reported as well. Those measures are much better able to report the performance of the estimations with respect to the spectroscopic redshifts. In fact, they are capable of taking the location and the shape of the density distribution into account; important characteristics which the

Table 2. Layout of the DCMDN architecture for the experiments with the quasar and the mixed catalog.

#	Type	Size	# Maps	Activation
1	input	$15 \times 16 \times 16$	/	/
2	convolutional	3×3	256	\tanh
3	pooling	2×2	256	\tanh
4	convolutional	2×2	512	\tanh
5	pooling	2×2	512	\tanh
6	convolutional	2×2	1024	$ReLU$
7	MDN	500	/	\tanh
8	MDN	100	/	\tanh
9	output	15	/	Eq. (2)

Notes. Stacks of 15 input images of the size 16×16 square-pixels are used.

commonly used scores can not capture. Typically a k -fold cross-validation is performed to prevent overfitting when training and evaluating machine learning models. Due to the extreme computational demands when training the DCMDN, only a simple hold out strategy was used to evaluate the performance. This is reasonable, as the same training, test and evaluation data-sets are used for all models. The reported performances therefore allow a fair and qualitative comparison between the individual models, even though the presented absolute performances might slightly vary depending on the used training and testing data. To be able to present an absolute performance which can be quantitatively compared to those in other publications, adequate reference data-sets have to be defined and published. The architectures used for the DCMDN are stated in Tables 1 and 2. Two different input sizes for images have been used, depending on the experiment. Larger images are used for galaxies while quasars and the mixed experiments use smaller images. For this reason two different architectures are presented.

3.2.1. Experiment 1 – Galaxies

In the first experiment we perform the prediction of redshift PDFs on galaxies only. The 200 000 patterns contained in the galaxies catalog are split in a training and test-set, both containing 100 000 objects each. The images for all experiments have been cut-out with a size of 28×28 square-pixels. Only the galaxies experiment kept the original size. As galaxies are extended objects this enables a better use of the available information. Together with the five *ugriz* images, all color combinations are built and hence a $15 \times 28 \times 28$ dimensional tensor is retrieved as object representation to be used as input for the DCMDN. The architecture of the DCMDN that was used for this experiment is specified in Table 1. This first experiment is intended to test and compare the performances of the three models on objects in the low redshift range, taking into account the spatial extension. As most objects of the galaxies sample are in a redshift range of $z \in [0..1]$, this provides a good testbed for the nearby Universe. Experiment 1 is just based on galaxies and therefore a strong bias is introduced in the training phase. The derived models are just producing usable redshift estimations when applied to images of galaxies. Such a model is limited to objects with a correct and proper pre-classification and selection.

3.2.2. Experiment 2 – Quasars

In the second experiment PDFs are estimated for quasars. The quasar catalog is composed of 185 000 objects; the DR7 and DR9 catalogs are combined selecting all the objects and removing double ones. The quasar experiments have been performed using 100 000 objects in the training-set and 85 000 in the test-set. This makes the size of the data used for training comparable to the previous experiment. In this experiment the size of the input images for the DCMDN is reduced to 16×16 square-pixels in order to save computational resources and speed up the training. As quasars are more compact sources, a smaller cut-out should be sufficient to capture the details of the spatial distribution and still include information from the hosting galaxy. The same color combinations used in the first experiment were created and a $15 \times 16 \times 16$ tensor is used as input for the DCMDN. The architecture of the DCMDN that was used for this experiment is specified in Table 2. The quasar experiment tests the performance for less extended objects that cover a wider range of redshift $z \in [0..6]$. Similarly to the first experiment, the models of the second experiment were heavily dependent on a correct pre-classification of objects.

3.2.3. Experiment 3 – Mixed

Finally, in the third experiment a mixed catalog was used. By combining quasars, galaxies and stars we are able to test and evaluate the performances of the three models independently from the nature of the sources. The step of pre-classifying objects is hereby made obsolete. The stars that have been added can be considered as contamination; as faint cool stars can be easily confused with quasars. This makes the use-case of photometric redshift estimation more realistic to the challenges of processing yet unseen objects with an uncertain classification. To be able to derive a proper PDF for all objects, stars have been assigned a redshift of $z = 0$. As stated above, the whole catalog is composed of 585 000 objects. In this experiment, 300 000 patterns were used for training and 285 000 for testing and the dimension of the input images is reduced to 16×16 square-pixels. The DCMDN has therefore the same architecture as used in the previous experiment (see Table 2). This experiment is intended to evaluate the performances of the models in a realistic use-case. Hence the results of this experiments are the most notable, as no biases through a pre-classification phase are introduced. Such an experiment should be part of every publication, introducing a new data-driven method for photometric redshift estimation.

4. Results

The experiments of Sect. 3 have been performed on a GPU cluster. The detailed results are presented in the following.

4.1. Experiment 1 – Galaxies

The results of the first experiment are depicted in Figs. 4 and 5. In both figures the estimated redshifts are plotted against the spectroscopic redshifts. To be comparable to results from other publications in the field of photometric redshift estimation, in Fig. 4 the complex estimated PDFs are compressed into a single point estimate. This compression is realized by either taking the plain mean, the first and most dominant mode of the mixture model or by taking the mean of objects that do not exhibit an ambiguous PDF (see Eq. (5)). Based on these three simplified representations of the estimated PDFs, the traditional scores have

been calculated. Those values are reported in Table 3 and show a similar performance as other publications (e.g. Laurino et al. 2011) that used a comparable data-set. All three models used for testing show a very similar performance. It is notable, that due to the nature of the used different models, the MDN and the DCMDN show a slightly better generalization performance especially in those regions where the transition of characteristic spectroscopic features through the broadband filters do not allow a distinct separation. In those redshift regions where the degeneracy of the reverse determination of the spectral energy distribution and distinct spectral features through the low spectral resolution image data is dominating, selecting the first mode instead of using the mean value shows a poorer performance. This is especially the case for $z \approx 0.35$ and $z \approx 0.45$. When selecting only those objects for evaluation that do not have an ambiguous behavior, the mentioned redshift regions become underpopulated in the diagnostic plots. For the DCMDN this effect is not as prominent as for the other two reference models due to the ability to make use of a larger base of information. As presented in Fig. 4, compressing the PDFs into single values does not cover the full complexity of the redshift estimations. In particular, the selection of outliers having no unique single dominant redshift component in the PDF demonstrates the multi-modal nature of the photometric redshift estimation task. Therefore, a visible improvement of the performance can be noticed when selecting only the subset of patterns which show no multi-modal behavior. For these reasons, a proper probabilistic evaluation of the PDFs has to be performed. Thus, in Fig. 5 a diagnostic plot is introduced which preserves the overall information of the density distributions. Alongside this probabilistic comparison between the estimated redshift distributions and the spectroscopic values, the PIT and the CRPS are used as proper tools. In the upper plot, the spectroscopic redshift is compared with the generated predictive density of every data item. Hereby the logarithm of the summed probability density for each redshift bin is plotted. In the two lower plots, the histograms for the PIT and the individual CRPS are given, together with the value of the mean CRPS. Analyzing the plots, it is notable that the RF performs slightly better with respect to the PIT. In all cases the PIT shows a more or less well calibrated uniform distribution of the evaluated CDFs at the corresponding spectroscopic redshifts. The MDN and the DCMDN in particular exhibit a more uniform and cleaner alignment toward the diagonal line which indicates the ideal performance.

4.2. Experiment 2 – Quasars

In Fig. 6 the results of the second experiment are shown in the standard plot. All estimated PDFs have been transferred into point estimates as it was done for the results of the first experiment. In Table 4 the corresponding scores are presented. As expected, the performances achieved with the mean values are comparable to point estimates obtained by other methods (Laurino et al. 2011). When using the first modes of the predicted density distributions, the degeneracies caused by the poor spectral resolution of the broadband photometry become visible. Especially in the redshift areas around $z \approx [0.5..0.9]$ and $z \approx [1.5..2.5]$ a strong multi-modal behavior can be observed. In those areas, picking the first mode does not necessary provide the best estimate, as both redshifts are equally likely. The task of uniquely assigning a redshift based on poorly resolved broadband observations is in some regions a degenerated problem. The better description of the probability distributions results therefore in a symmetry along the line of ideal

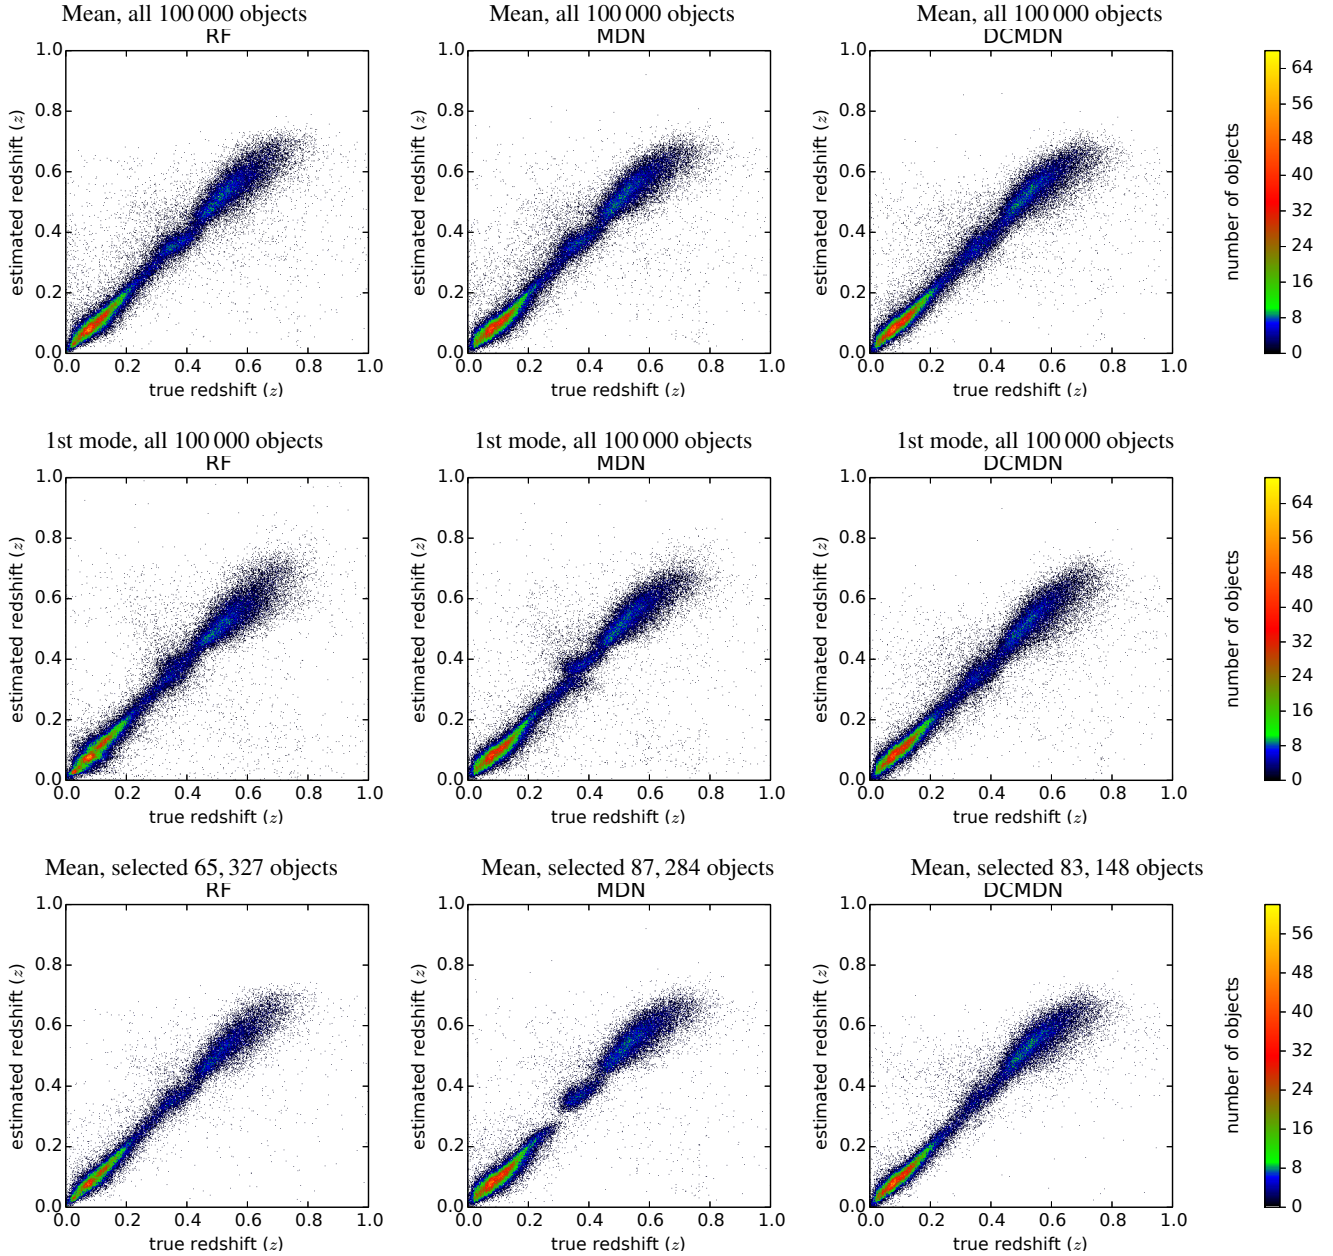


Fig. 4. Results of Experiment 1 (galaxies). The estimated redshifts are plotted against the spectroscopic redshifts, which are considered ground trues. As in Table 3 the plots are done for all three considered models (*left to right*) as well as with the mean, dominant mode and unambiguous objects (*top to bottom*). In addition, the number of unambiguous objects is reported.

Table 3. Results of Experiment 1 (galaxies).

Criterion	Model	Bias(Δz)	$\sigma^2(\Delta z)$	Mad(Δz)	rmse(Δz)	Bias(Δz_{norm})	$\sigma^2(\Delta z_{\text{norm}})$	Mad(Δz_{norm})	rmse(Δz_{norm})
Mean	RF	0.0001	0.0033	0.0164	0.0575	-0.0017	0.0018	0.0133	0.0431
	MDN	0.0016	0.0034	0.0174	0.0589	-0.0003	0.0019	0.0141	0.0442
	DCMDN	0.0018	0.0030	0.0157	0.0548	-0.0003	0.0017	0.0128	0.0409
First mode	RF	0.0031	0.0042	0.0171	0.0652	0.0010	0.0022	0.0139	0.0471
	MDN	0.0029	0.0039	0.0172	0.0628	0.0010	0.0021	0.0140	0.0459
	DCMDN	0.0060	0.0031	0.0167	0.0561	0.0029	0.0016	0.0135	0.0407
Selected	RF	0.0001	0.0023	0.0147	0.0484	-0.0011	0.0013	0.0121	0.0365
	MDN	0.0021	0.0027	0.0165	0.0523	0.0006	0.0015	0.0136	0.0391
	DCMDN	0.0017	0.0023	0.0146	0.0485	-0.0001	0.0013	0.0120	0.0366

Notes. Based on the estimated PDFs the traditional scores have been calculated. This was done by using the mean, the most dominant mode and the mean of the selected unambiguous objects, respectively.

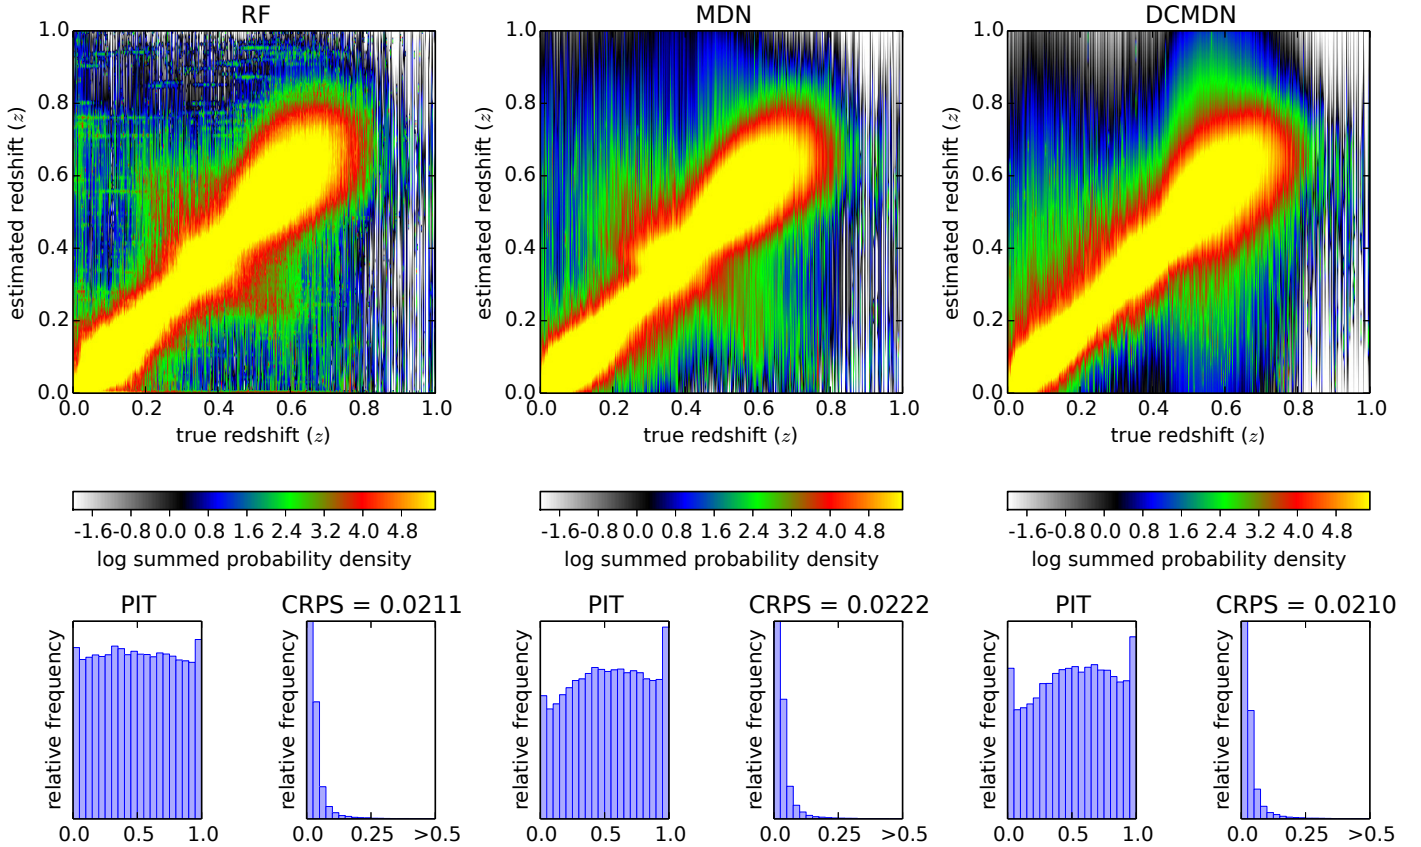


Fig. 5. Results of Experiment 1 (galaxies) with a fully probabilistic evaluation and representation of the estimated PDFs. For each model, three plots are present. In the upper one, the predicted density distribution for each individual object is plotted at the corresponding spectroscopic redshift. The colors hereby indicate the logarithm of the summed probability densities using 500 redshift bins per axis. In the two lower plots, the histogram of the PIT values and the histogram of the CRPS values are shown, respectively. The mean CRPS values are reported, too.

performance. Due to selection effects by observing a limited and cone-shaped volume, the distribution of objects with respect to cosmological scales results in not perfect symmetric behavior. The necessity of a multi-modal description becomes even more obvious when objects with an ambiguous density distribution are not taken into consideration. This way of excluding sources with estimated ambiguous redshift distributions can be considered a correct method to filter possible outliers, hence the selection of objects that show an uni-modal behavior is purely done based on the probabilistic description. In all three different representation of the complex PDFs through a single point estimate, the DCMDN exhibits a better performance. This is also reflected by the scores that are presented in Table 4. For the RF, the partitioning of the high-dimensional feature space orthogonal to the dimension axis does not provide a generalized representation of the regression problem. When covering a wider redshift range the predictions exhibit more differences performance wise. Therefore the performance of the RF drops in this experiment with respect to the other models as the MDN and DCMDN produce much smoother predictions along the ideal diagonal of the used diagnostic plot. With respect to catastrophic outliers, the DCMDN has a superior performance when compared with the other two models. This is consistent with the results from the first experiment. When analyzing the probabilistic representation of the prediction results in Fig. 7, the superior performance of the DCMDN is striking. The CRPS especially indicates the quality of the predictions made by the DCMDN. In fact, by using images the DCMDN can utilize all the contained information and automatically extract the best usable features.

4.3. Experiment 3 – Mixed

The experiment with the mixed catalog is the most challenging one. It tests how well the advantages of a deep convolutional network architecture can be used to render the step of pre-processing and pre-classifying objects obsolete. This experiment is much closer to the real application compared to the previous cases. In Fig. 8 and Table 5 the results of Experiment 3 are reported and evaluated as single point estimates. As in the previous experiments, the PDFs are therefore compressed. When using the mean or the first mode as representation, the DCMDN significantly outperforms the RF and the MDN. In the case of selecting unambiguous PDFs only, the result of the DCMDN is close to the ideal performance. The obtained results confirm the indications of the previous experiments. The fully probabilistic representation of the results of the third experiment are presented in Fig. 9. When using proper tools for evaluation, the DCMDN is the best model with respect to the CRPS. Both, the RF and the MDN exhibit similar CRPS results, as they did in all the three experiments. In the representation of the PITs the differences in calibration can be further analyzed. The RF shows a nearly uniform distribution with an extreme peak in the center. Due to the partitioning of the high-dimensional feature space performed by the RF, most of the stars are perfectly recovered. As stars are fixedly assigned a redshift of $z = 0$, the chosen way of fitting a GMM to the individual decision tree results produces the central peak. This can be seen as an extreme overdispersion of the predictions in relation to the true values. With respect to the estimated PDFs, the CDFs at the truncated true value $z = 0$ is always very close to 0.5. For a stellar object, only a very few

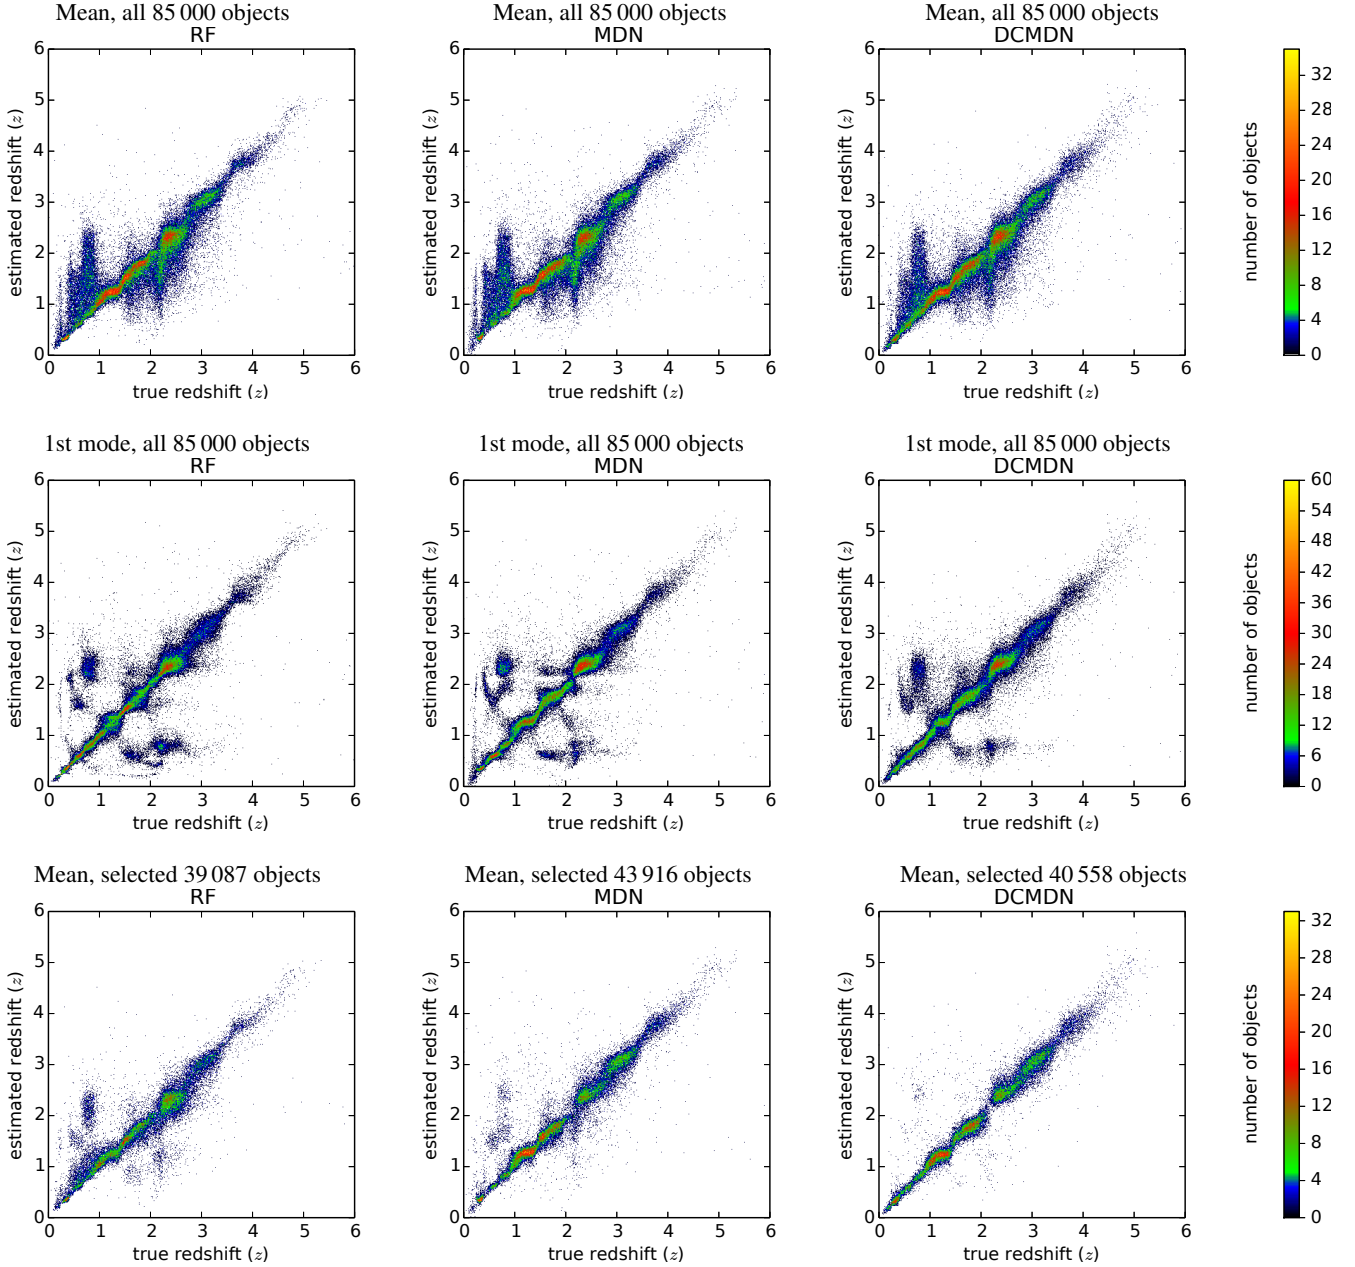


Fig. 6. Results of Experiment 2 (quasars). The estimated redshifts are plotted against the spectroscopic redshifts. As in Fig. 4 the plots are sorted by the models (*left to right*) and the extracted point estimate (*top to bottom*).

Table 4. Results of Experiment 2 (quasars).

Criterion	Model	Bias(Δz)	$\sigma^2(\Delta z)$	Mad(Δz)	rmse(Δz)	Bias(Δz_{norm})	$\sigma^2(\Delta z_{\text{norm}})$	Mad(Δz_{norm})	rmse(Δz_{norm})
Mean	RF	0.007	0.217	0.145	0.466	-0.033	0.048	0.050	0.222
	MDN	-0.002	0.216	0.156	0.465	-0.037	0.048	0.054	0.223
	DCMDN	0.011	0.168	0.128	0.411	-0.023	0.035	0.045	0.189
First mode	RF	0.002	0.319	0.087	0.564	-0.026	0.066	0.031	0.258
	MDN	-0.058	0.282	0.095	0.535	-0.052	0.068	0.034	0.267
	DCMDN	-0.043	0.206	0.095	0.456	-0.038	0.048	0.034	0.222
Selected	RF	0.005	0.162	0.111	0.402	-0.024	0.037	0.039	0.194
	MDN	-0.010	0.098	0.086	0.314	-0.017	0.021	0.030	0.145
	DCMDN	0.004	0.047	0.075	0.217	-0.004	0.009	0.026	0.095

Notes. Similarly to experiment 1, the traditional scores are presented. The score have been calculated with the mean, the most dominant mode and the mean of the selected unambiguous objects.

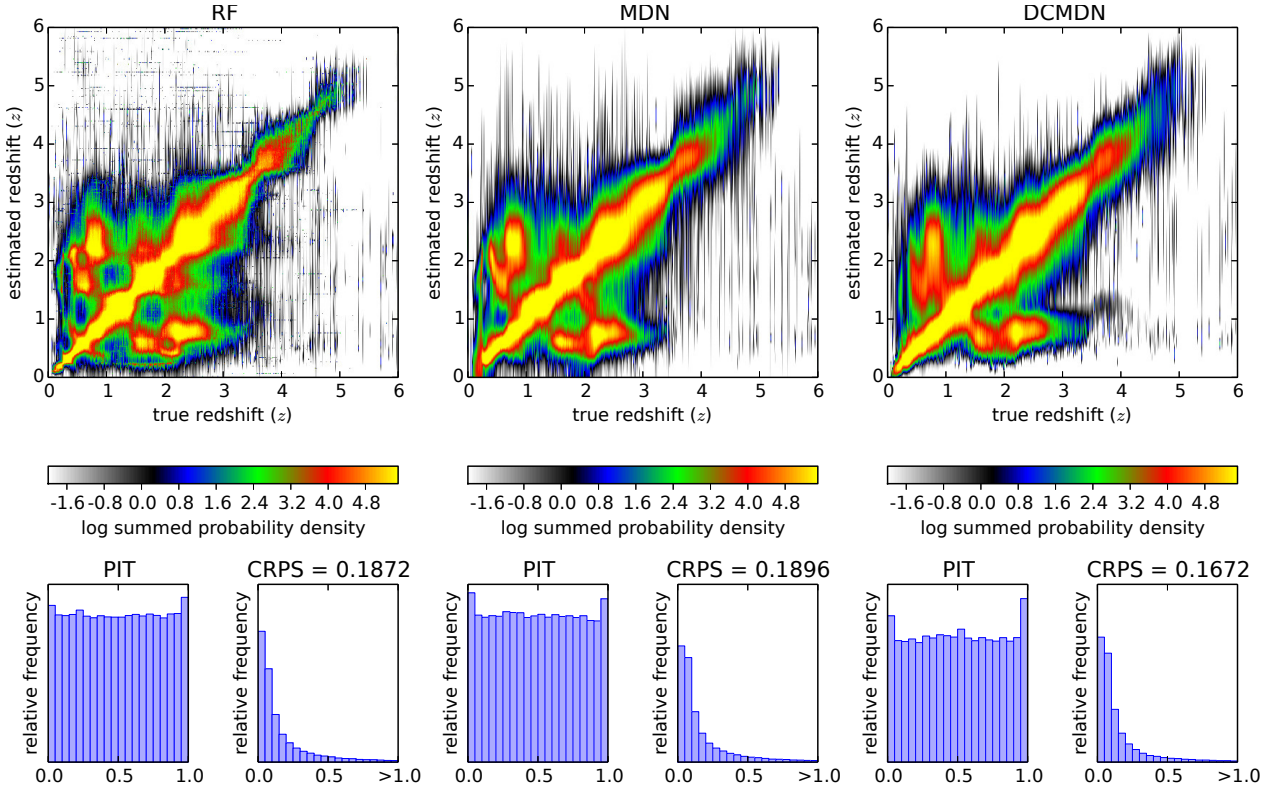


Fig. 7. Results of Experiment 2 (quasars) with a fully probabilistic evaluation. As in Fig. 5 three plots per model are used to visualize the probabilistic performance of the estimated PDFs.

decision trees might return a redshift $z > 0$, given the large number of $\approx 100\,000$ stars used in the training sample. During training the MDN got biased by the large number of stellar components with redshift set to $z = 0$. Therefore the corresponding PIT indicates a tendency to underestimate the redshift. A complex, asymmetric behavior in the shape of the estimated PDFs is not captured when using the classical scores and diagnostic plots only, as the bias on the mean of the PDFs is close to zero for this experiment. For the DCMDN we can observe two shortcomings. Overall, the generated PDFs are overdispersed. This effect is caused by the large fraction of stellar objects with a fixed value and the other objects covering the full redshift range. The model tries to account for this highly unbalanced distribution in the redshift space when being optimized. Generating broader distributions for one part of the objects while creating very narrow estimates for the stellar population is hard to be achieved by a generalizing model. The second effect observed is the presence of outliers where the CDF indicates an extreme value. This is caused by objects in the high redshift range being highly underrepresented in the training sample. The DCMDN architecture has a superior capability of generalization and makes better use of the information contained in the original images, giving a boost in the final performance, together with a good separation of the stars from quasars.

5. Conclusions

The aim of this work is to present and test a new method for photometric redshift estimation. The novelty of the presented approach is to estimate probability density functions for redshifts based on imaging data. The final goal is to make the additional steps of feature-extraction and feature-selection obsolete. To achieve this, a deep convolutional network architecture

was combined with a mixture density network. Essential for a proper training is the use of the CRPS as loss function, taking into account not only the location but also the shape and layout of the estimated density distributions. The new architecture is described in a general and conceptual way that allows using the concept for many other regression tasks in astronomy.

In order to perform a fair evaluation of the performance of the proposed method, three experiments have been performed. Three different catalogs have been utilized for evaluation, containing galaxies, quasars and a mix of the previous catalogs plus a sample of stars. The experiments were chosen to test the performance on different redshift ranges and different sources. The last experiment was designed to test the model in a more realistic scenario, where a contamination with stellar sources and therefore a confusion between no and high redshift objects is synthetically introduced. This experiment is intended to test whether the pre-classification of objects can be omitted, too.

In all three experiments a modified version of the random forest and a mixture density network are used as feature-based reference models. To be comparable to the literature, the traditional scores and diagnostic plots have been used. As we demonstrate, the usual way of expressing the results of the prediction quality through the traditional scores is not able to capture the complexity of often multi-modal and asymmetric distributions that are required to correctly describe the redshift estimates. Therefore, proper scores and tools have been applied in addition to analyzing the results in a probabilistic way, that is, the CRPS and PIT. These indicators can be considered proper tools to estimate the quality of photometric redshift density distributions. As summarized in Table 6, the DCMDN architecture outperforms the two reference models. The same relative performances can be observed with the traditional diagnostics, too. When using the PDFs to find and exclude ambiguous objects, the DCMDN

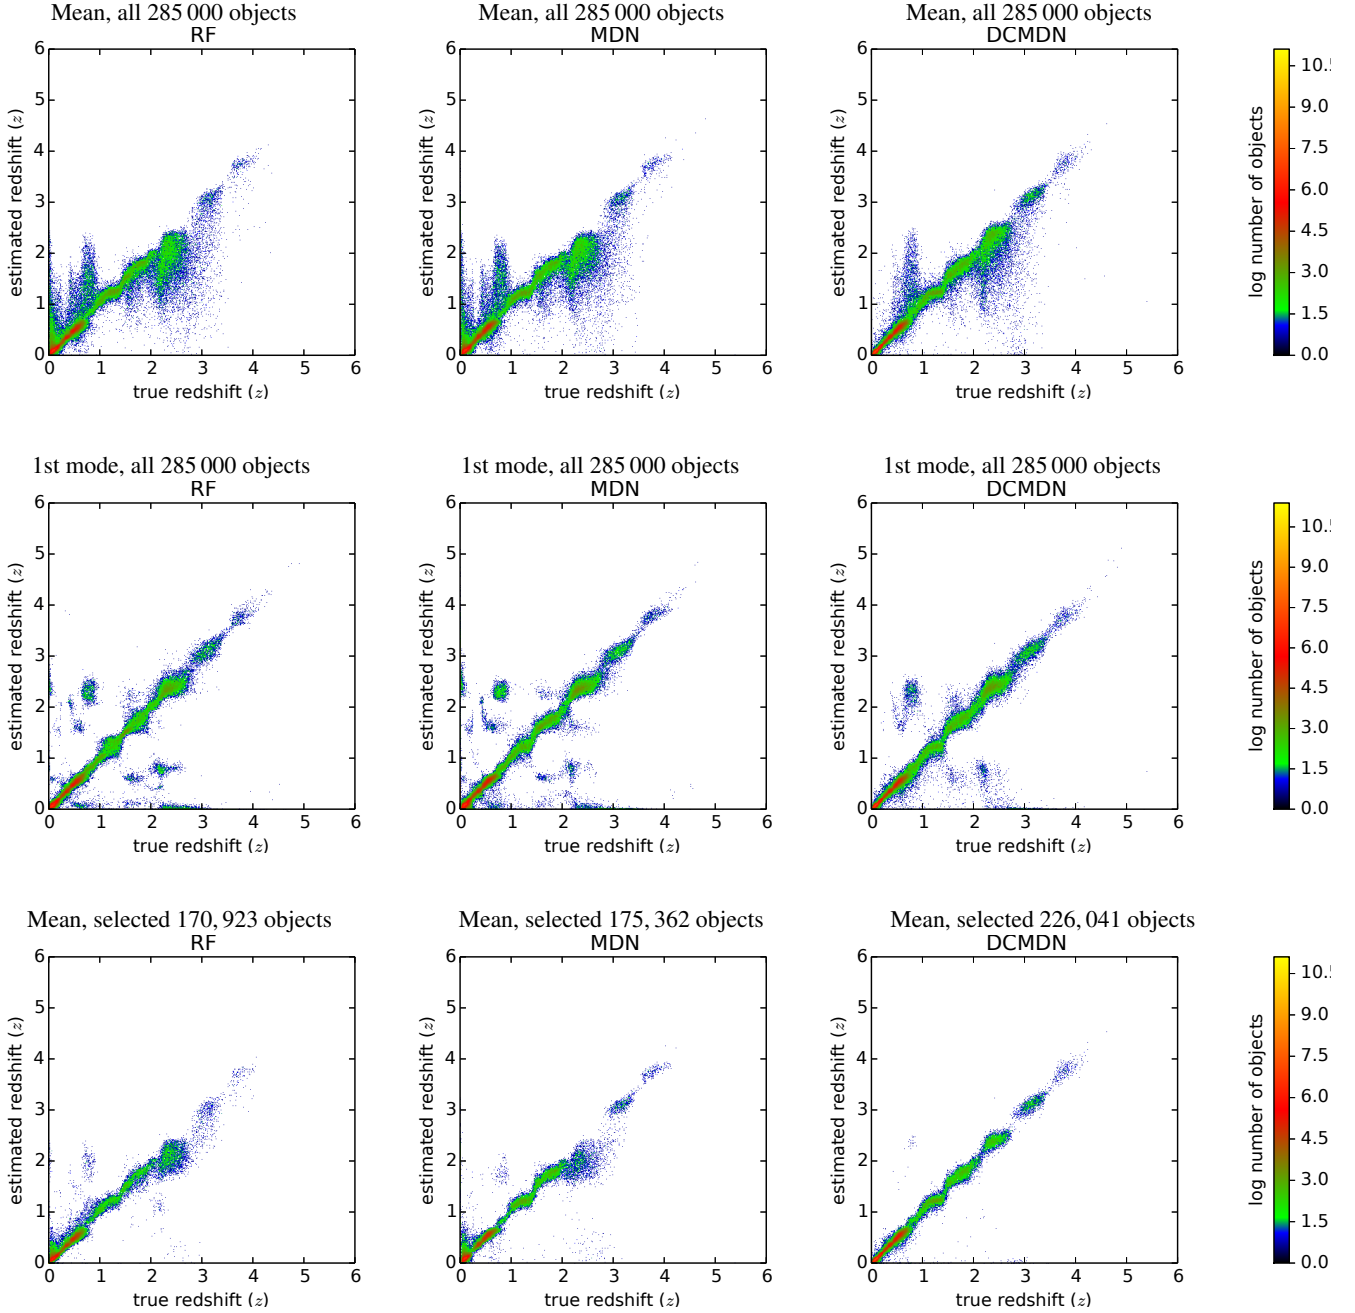


Fig. 8. Results of Experiment 3 (mixed). The estimated redshifts are plotted against the spectroscopic redshifts. As in the figures of the two other experiments, the plots are sorted by the models (*left to right*) and the extracted point estimate (*top to bottom*).

Table 5. Results of Experiment 3 (mixed).

Criterion	Model	Bias(Δz)	$\sigma^2(\Delta z)$	Mad(Δz)	rmse(Δz)	Bias(Δz_{norm})	$\sigma^2(\Delta z_{\text{norm}})$	Mad(Δz_{norm})	rmse(Δz_{norm})
Mean	RF	-0.001	0.288	0.043	0.536	-0.072	0.126	0.028	0.363
	MDN	-0.006	0.279	0.043	0.528	-0.073	0.125	0.031	0.362
	DCMDN	0.007	0.210	0.022	0.458	-0.041	0.089	0.016	0.301
First mode	RF	0.040	0.435	0.020	0.660	-0.029	0.150	0.013	0.388
	MDN	-0.027	0.393	0.024	0.627	-0.067	0.191	0.016	0.442
	DCMDN	-0.001	0.287	0.018	0.536	-0.036	0.124	0.013	0.355
Selected	RF	-0.001	0.118	0.016	0.343	-0.029	0.059	0.012	0.245
	MDN	-0.006	0.114	0.023	0.337	-0.034	0.052	0.017	0.230
	DCMDN	0.002	0.043	0.012	0.206	-0.007	0.014	0.010	0.120

Notes. Similarly to experiment 1 and 2, the traditional scores have been calculated.

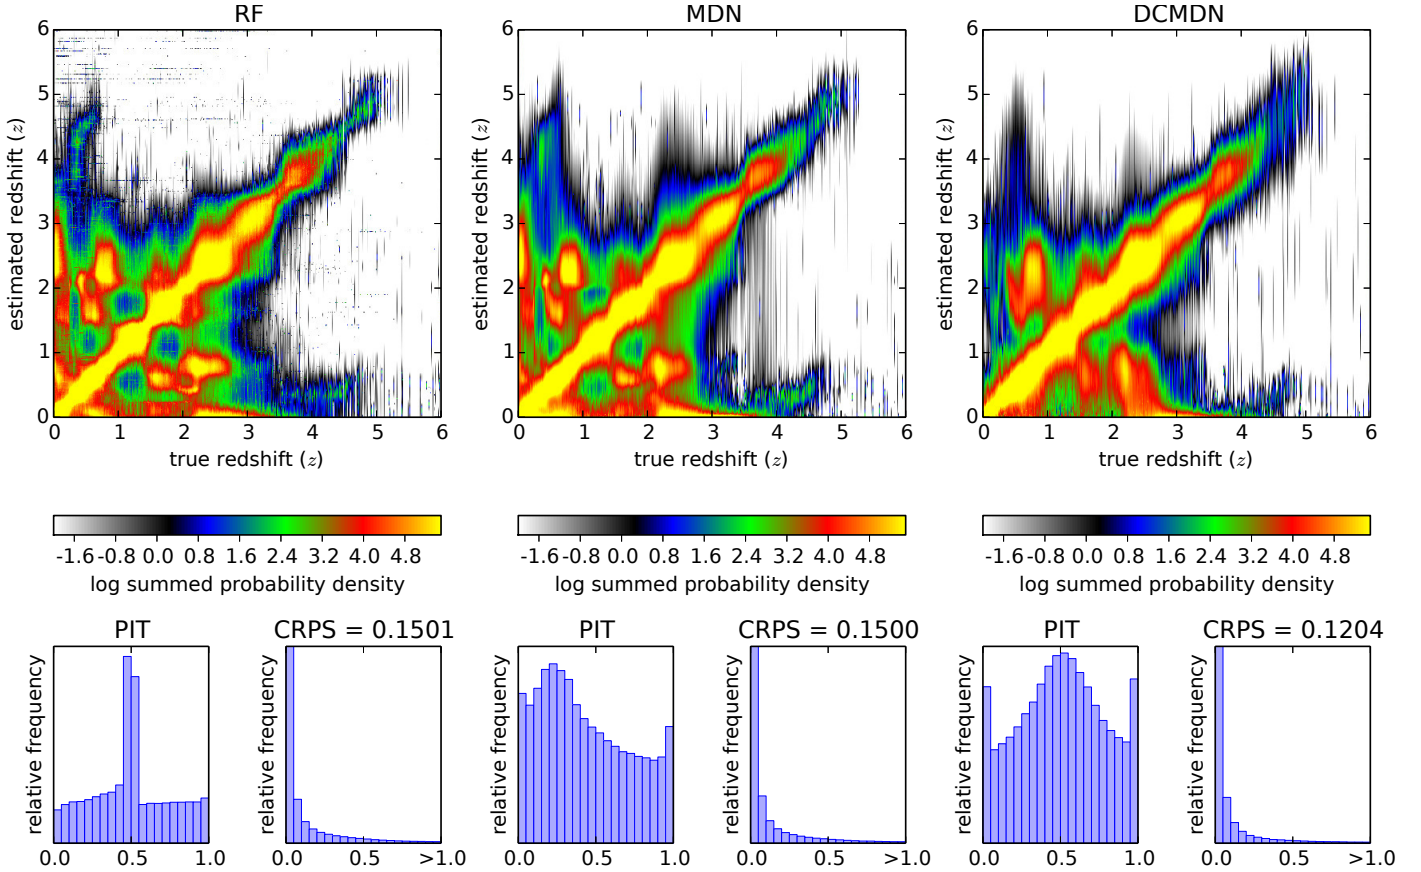


Fig. 9. Results of Experiment 3 (mixed) with a fully probabilistic evaluation. As in Fig. 5 three plots per model are used to visualize the probabilistic performance of the estimated PDFs.

Table 6. Summary table of all the experiments.

Exp.	Model	CRPS	PIT
Galaxies	RF	0.021	well calibrated
	MDN	0.022	biased
	DCMDN	0.021	slightly overdispersed
Quasars	RF	0.187	well calibrated
	MDN	0.190	well calibrated, few outliers
	DCMDN	0.167	well calibrated, few outliers
Mixed	RF	0.150	extremely overdispersed
	MDN	0.150	overdispersed and biased
	DCMDN	0.120	overdispersed, some outliers

always produces the best results. Its ability to generalize from the training data, together with the larger amount of used available information permits a better probabilistic estimate of the redshift distribution and therefore a better selection of spurious predictions.

As shown in the last experiment, the DCMDN performs two very important tasks automatically. During the training of the network, a set of thousands of features is automatically extracted and selected from the imaging data. This minimizes the biases that are introduced when manually extracting and selecting features and increases the amount of utilized information. The second and probably most important capability of the DCMDN is to solve the regression problem without the necessity of pre-classifying the objects. The estimated PDFs reflect in

the distribution of the probabilities the uncertainties of the classification as well as the uncertainties of the redshift estimation. This is extremely important when dealing with data from larger surveys. The errors introduced through a hard classification into distinct classes limits the ability of finding rare but interesting objects, like high redshifted quasars that are easily mistaken with cool faint stars. A wrong initial classification would mark those quasars as stellar components even though the probability of being a high redshifted object is not negligible. A fully probabilistic approach including feature extraction, feature selection and source classification is less affected by selection biases.

The performance of the two feature-based reference methods is in accordance with results from the literature. The boost in performance observed for the DCMDN is related to the better use of information and the automatic selection of the best performing features. As the DCMDN model was trained with data from the SDSS it could in principle be applied to every source in the SDSS database, without concern for the nature of the source, by directly using the images. Only the selection biases that have been introduced when selecting the targets for spectroscopic follow-up observations in SDSS have to be considered, as those are preserved by the trained model. Our approach provides the machinery to deal with the avalanche of data we are facing with the new generation of survey instruments already today. Such fully-automatized and probabilistic approaches, based on deep learning architectures, are therefore necessary.

Acknowledgements. The authors gratefully acknowledge the support of the Klaus Tschira Foundation. We gratefully thank the anonymous referee for the constructive and precise comments, leading to a more detailed discussion

of the loss functions and the number of Gaussian components used. We would like to thank Nikos Gianniotis and Erica Hopkins for proofreading and commenting on this work. This work is based on the usage of Theano (Theano Development Team 2016) to run our experiments. Topcat (Taylor 2005) was used to access and visualize data from the VO and to create the used catalogs. This work is based on data provided by SDSS. 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.

References

- Abdalla, F. B., Banerji, M., Lahav, O., & Rashkov, V. 2011, *MNRAS*, **417**, 1891
- Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2012, *ApJS*, **203**, 21
- Ball, N. M., & Brunner, R. J. 2010, *Int. J. Mod. Phys. D*, **19**, 1049
- Beck, R., Dobos, L., Budavári, T., Szalay, A. S., & Csabai, I. 2016, *MNRAS*, **460**, 1371
- Benítez, N. 2000, *ApJ*, **536**, 571
- Bishop, C. M. 1994, Mixture density networks, Tech. Rep., Aston University
- Blake, C., & Bridle, S. 2005, *MNRAS*, **363**, 1329
- Bolzonella, M., Miralles, J.-M., & Pelló, R. 2000, *A&A*, **363**, 476
- Breiman, L. 2001, *Mach. Learn.*, **45**, 5
- Brescia, M., Cavuoti, S., D'Abrusco, R., Longo, G., & Mercurio, A. 2013, *ApJ*, **772**
- Brescia, M., Cavuoti, S., Longo, G., & De Stefano, V. 2014, *A&A*, **568**, A126
- Carliles, S., Budavári, T., Heinis, S., Priebe, C., & Szalay, A. S. 2010, *ApJ*, **712**, 511
- Carrasco Kind, M., & Brunner, R. J. 2013, *MNRAS*, **432**, 1483
- Cavuoti, S., Brescia, M., Tortora, C., et al. 2015, *MNRAS*, **452**, 3100
- Collister, A. A., & Lahav, O. 2004, *PASP*, **116**, 345
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, **145**, 10
- D'Isanto, A., & Polsterer, K. L. 2017, in ESANN 2017, 25th European Symposium on Artificial Neural Networks, Bruges, Belgium, April 26–28, 2017, Proceedings
- Dubath, P., Apostolakis, N., Bonchi, A., et al. 2017, *IAU Symp.*, **325**, 73
- Fernández-Soto, A., Lanzetta, K. M., Chen, H.-W., Pascarelle, S. M., & Yahata, N. 2001, *ApJS*, **135**, 41
- Fernique, P., Allen, M. G., Boch, T., et al. 2015, *A&A*, **578**, A114
- Gneiting, T., Raftery, A. E., Westveld, A. H., & Goldman, T. 2005, *Mon. Weather Rev.*, **133**, 1098
- Grimit, E. P., Gneiting, T., Berrocal, V. J., & Johnson, N. A. 2006, *Quarterly J. Roy. Meteorol. Soc.*, **132**, 2925
- Hersbach, H. 2000, *Weather and Forecasting*, **15**, 559
- Hoyle, B. 2016, *Astron. Comput.*, **16**, 34
- Kügler, S. D., Gianniotis, N., & Polsterer, K. L. 2016, in 2016 IEEE Symposium Series on Computational Intelligence, SSCI 2016, Athens, Greece, December 6–9, 2016, 1
- Laureijs, R., Gando, P., Duvet, L., et al. 2012, in Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave, *Proc. SPIE*, **8442**, 84420
- Laurino, O., D'Abrusco, R., Longo, G., & Riccio, G. 2011, *MNRAS*, **418**, 2165
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proc. IEEE*, **86**, 2278
- Pâris, I., Petitjean, P., Aubourg, É., et al. 2012, *A&A*, **548**, A66
- Polsterer, K. L., Zinn, P.-C., & Gieseke, F. 2013, *MNRAS*, **428**, 226
- Rosenblatt, F. 1962, Principles of neurodynamics: perceptrons and the theory of brain mechanisms, Report (Cornell Aeronautical Laboratory) (Spartan Books)
- Sadeh, I., Abdalla, F. B., & Lahav, O. 2016, *PASP*, **128**, 104502
- Salvato, M., Hasinger, G., Ilbert, O., et al. 2009, *ApJ*, **690**, 1250
- Schneider, D. P., Richards, G. T., Hall, P. B., et al. 2010, *AJ*, **139**, 2360
- Schwarz, G. 1978, *Ann. Stat.*, **6**, 461
- Taylor, M. B. 2005, in Astronomical Data Analysis Software and Systems XIV, eds. P. Shopbell, M. Britton, & R. Ebert, ASP Conf. Ser., **347**, 29
- Theano Development Team. 2016, ArXiv e-prints [[arXiv:1605.02688](https://arxiv.org/abs/1605.02688)]

Appendix A: CRPS and PIT

The CRPS is defined by the relation:

$$\text{CRPS} = \text{CRPS}(F, x_a) = \int_{-\infty}^{+\infty} [F(x) - F_a(x)]^2 dx \quad (\text{A.1})$$

where $F(x)$ and $F_a(x)$ are the CDFs relative to the predicted PDF $f(t)$ and the observation x_a , respectively. Namely: $F(x) = \int_{-\infty}^x f(t)dt$ and $F_x = H(x - x_a)$, with $H(x)$ being the Heaviside step-function. In case the PDF is given as a GMM, the CRPS can be calculated through the following formula in closed form:

$$\begin{aligned} \text{CRPS} \left(\sum_{m=1}^M \omega_m \mathcal{N}(\mu_m, \sigma_m^2) \right) &= \sum_{m=1}^M \omega_m A(x - \mu_m, \sigma_m^2) \\ &- \frac{1}{2} \sum_{m=1}^M \sum_{n=1}^M \omega_m \omega_n A(\mu_m - \mu_n, \sigma_m^2 + \sigma_n^2) \end{aligned} \quad (\text{A.2})$$

where

$$A(\mu, \sigma^2) = 2\sigma\phi\left(\frac{\mu}{\sigma}\right) + \mu\left(2\Phi\left(\frac{\mu}{\sigma}\right) - 1\right) \quad (\text{A.3})$$

and $\phi\left(\frac{y-\mu}{\sigma}\right)$, $\Phi\left(\frac{y-\mu}{\sigma}\right)$ respectively represent the PDF and the CDF of a normal distribution with a mean of zero and a variance of one evaluated through the normalized prediction error $\frac{y-\mu}{\sigma}$.

The probability integral transform (PIT) is generated with the histogram of the values:

$$p_t = F_t(x_t) \quad (\text{A.4})$$

being F_t the CDF of the predicted PDF evaluated at the observation x_t . In Fig. A.1 some example PITs are given.

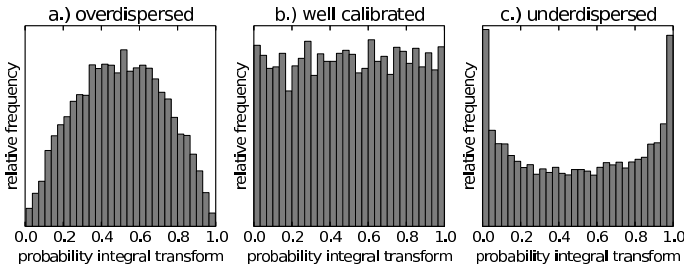


Fig. A.1. Visual guide to the usage of a PIT. In case the estimated PDFs are too broad with respect to the position of the true value, a convex histogram with a peak in the center can be observed (a). As soon as the predicted densities are too narrow, the evaluation of the CDF at the true redshift exhibits in most cases just very low and very high values. Therefore a concave, U-shaped histogram will be produced (c). Only in the case where the widths of the predicted densities is in accordance with the deviations from the true measurements, a uniformly distributed histogram is generated (b). This indicates sharp and well calibrated predictions.

Appendix B: CRPS vs. log-likelihood as loss-function

We compared the performance of the MDN when trained with the log-likelihood and the CRPS as loss functions, in order to investigate the differences in the behavior of the model. The experiment was performed on the data used for the quasar experiment, split in independent subsets for training (100 000 objects)

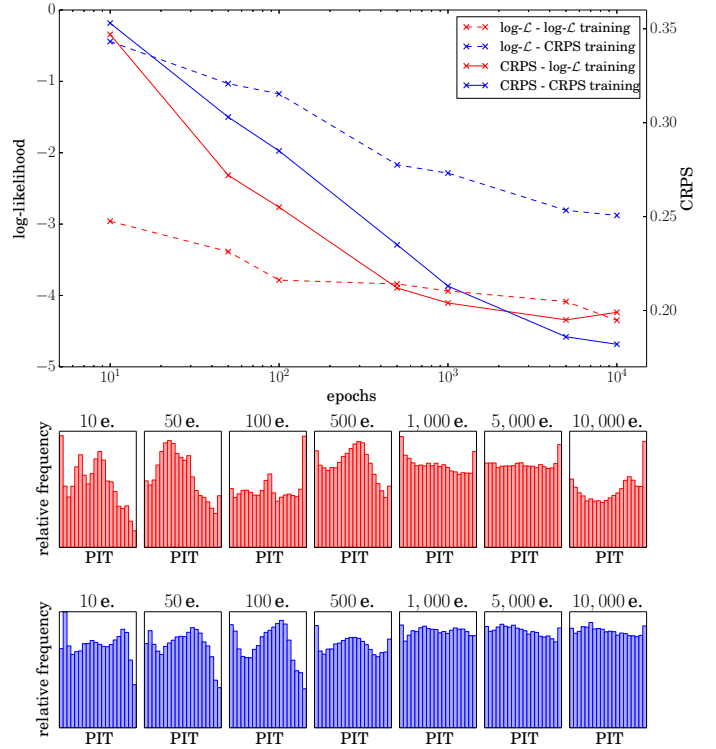


Fig. B.1. Comparison of the performance when using the log-likelihood (red) and the CRPS (blue) as loss-functions for training the MDN. For different epochs the CRPS and log-likelihood scores as well as the PIT histograms are provided.

and testing (85 000 objects). The results are shown in Fig. B.1. In this plot we report the performance expressed by using both the log-likelihood and the CRPS as score functions, at different epochs of the training phase. Moreover, the correspondent PIT histograms for these epochs are shown. Training the network with the log-likelihood improves the performance in terms of the log-likelihood itself, respect to the same architecture trained using the CRPS. As expected, when training with the CRPS, the observed results are opposite to the previous case. The PIT histograms indicate a better performance when the neural network is trained using the CRPS, leading to a well calibrated PIT already after 500 epochs. Moreover, the PIT of the model trained using the log-likelihood starts to degrade again at 10 000 epochs. This does not happen when using the CRPS for training, as the CRPS accounts for calibration and sharpness of the predictions. For this reason, the choice of the CRPS as loss function is reasonable, in order to obtain sharper and better calibrated PDFs.

Appendix C: Number of Gaussian components

In order to choose an appropriate number of Gaussian components, an experiment has been performed using the RF model for the quasar data-set and calculating the BIC following Schwarz (1978). The mean score has been calculated over different redshift bins, for different numbers of Gaussian components. The plot shows that extreme values, like $n = 1$ or $n = 7$ tend to exhibit bad performance in multiple regions. Instead, the results using $n = 3$ and $n = 5$ are comparably good. The BIC score is based on the log-likelihood calculation, therefore we consider a choice of $n = 5$ to be reasonable for this work, keeping in mind that our model is trained using the CRPS.

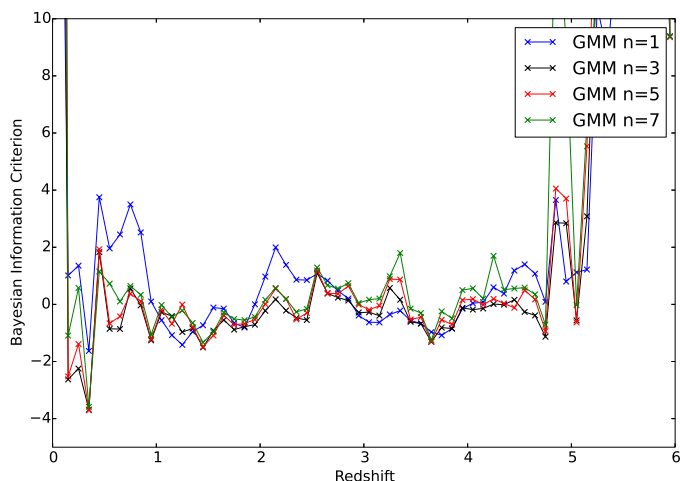


Fig. C.1. Distribution of the mean Bayesian information criterion score, with respect to redshift and number of Gaussian components. The plot depicts that different regions of the redshift range demand a different number of components. On average, the use of $n = 5$ is a reasonable choice. The experiment is performed using the RF and the quasar dataset.

Appendix D: Code

The code used to do this research can be found on the ASCL¹.

Appendix E: Data

The SDSS object IDs of the randomly extracted stars and galaxies are available as supplementary information. In addition the SDSS IDs of the quasars are provided as a plain ASCII file.

The results of the predictions done with the DCMDN architecture for the three experiments are made available as ASCII files, too. Those files contain the spectroscopic redshifts of the test objects followed by 15 outputs that can be used to calculate the GMM parameters (five means, five sigmas, five weights) as described in Eq. (2). This allows reproduction of the performance of our model.

galaxies_id.csv contains the SDSS object IDs of the galaxies used for the experiments.

quasars_id.csv contains the SDSS object IDs of the quasars used for the experiments.

stars_id.csv contains the SDSS object IDs of the stellar objects used for the experiments.

galaxies_output.csv keeps the predictions generated with the DCMDN in the first experiment.

quasars_output.csv keeps the predictions generated with the DCMDN in the second experiment.

mixed_output.csv keeps the predictions generated with the DCMDN in the third experiment.

Appendix F: SQL-queries

The following queries have been used to generate the galaxies and stars catalogs via CasJobs:

“Query used to create the galaxies catalog”

```

1  SELECT TOP 2000000
2     p.objid,p.ra,p.dec ,
3     p.u,p.g,p.r,p.i,p.z ,
4     p.psfMag_u , p.psfMag_g ,
5     p.psfMag_r , p.psfMag_i ,
6     p.psfMag_z , p.modelMag_u ,
7     p.modelMag_g , p.modelMag_r ,
8     p.modelMag_i , p.modelMag_z ,
9     s.specobjid , s.class ,
10    s.z AS redshift
11 INTO mydb.DR9_galaxies_with_modMag
12 FROM PhotoObj AS p
13 JOIN SpecObj AS s ON
14     s.bestobjid = p.objid
15 WHERE s.z BETWEEN 0 AND 6.0
16     AND s.class = 'GALAXY'
17 ORDER BY NEWID()

```

“Query used to extract stellar sources from the SDSS”

```

1  SELECT TOP 2000000
2     p.objid,p.ra,p.dec ,
3     p.u,p.g,p.r,p.i,p.z ,
4     p.psfMag_u , p.psfMag_g ,
5     p.psfMag_r , p.psfMag_i ,
6     p.psfMag_z , p.modelMag_u ,
7     p.modelMag_g , p.modelMag_r ,
8     p.modelMag_i , p.modelMag_z ,
9     s.specobjid , s.class ,
10    s.z AS redshift
11 INTO mydb.DR9_stars
12 FROM PhotoObj AS p
13 JOIN SpecObj AS s ON
14     s.bestobjid = p.objid
15 WHERE s.class = 'STAR'
16 ORDER BY NEWID()

```

¹ <http://www.ascl.net/ascl:1709.006>