

A GALAXY PHOTOMETRIC REDSHIFT CATALOG FOR THE SLOAN DIGITAL SKY SURVEY DATA RELEASE 6

HIROAKI OYAIZU,^{1,2} MARCOS LIMA,^{2,3} CARLOS E. CUNHA,^{1,2} HUAN LIN,⁴
JOSHUA FRIEMAN,^{1,2,4} AND ERIN S. SHELDON⁵
Received 2007 July 30; accepted 2007 September 12

ABSTRACT

We present and describe a catalog of galaxy photometric redshifts (photo- z 's) for the Sloan Digital Sky Survey (SDSS) Data Release 6 (DR6). We use the neural network (NN) technique to calculate photo- z 's and the nearest neighbor error (NNE) method to estimate photo- z errors for ~ 77 million objects classified as galaxies in DR6 with $r < 22$. The photo- z and photo- z error estimators are trained and validated on a sample of $\sim 640,000$ galaxies that have SDSS photometry and spectroscopic redshifts measured by SDSS, the Two Degree Field, the SDSS Luminous Red Galaxy and Quasi-stellar Object Survey (2SLAQ), the Canada-France Redshift Survey (CFRS), the Canadian Network for Observational Cosmology Field Galaxy Survey (CNOC2), the Team Keck Redshift Survey (TKRS), the Deep Extragalactic Evolutionary Probe (DEEP), and DEEP2. For the two best NN methods we have tried, we find that 68% of the galaxies in the validation set have a photo- z error smaller than $\sigma_{68} = 0.021$ or 0.024. After presenting our results and quality tests, we provide a short guide for users accessing the public data.

Subject headings: catalogs — distance scale — galaxies: distances and redshifts — large-scale structure of universe

Online material: color figures

1. INTRODUCTION

While spectroscopic redshifts have now been measured for over one million galaxies, in recent years digital sky surveys have obtained multiband imaging for on the order of a hundred million galaxies. Deep, wide-area surveys planned for the next decade will increase the number of galaxies with multiband photometry to a few billion. Due to technological and financial constraints, obtaining spectroscopic redshifts for more than a small fraction of these galaxies will remain impractical for the foreseeable future. As a result, over the last decade substantial effort has gone into developing photometric redshift (photo- z) techniques, which use multiband photometry to estimate approximate galaxy redshifts. For many applications in extragalactic astronomy and cosmology, the resulting photometric redshift precision is sufficient for the science goals at hand, provided one can accurately characterize the uncertainties in the photo- z estimates.

Two broad categories of photo- z estimators are in wide use: template-fitting and training set methods. In template fitting, one assigns a redshift to a galaxy by finding the redshifted spectral energy distribution (SED), selected from a library of templates, that best reproduces the observed fluxes in the broadband filters. By contrast, in the training set approach, one uses a training set of galaxies with spectroscopic redshifts and photometry to derive an empirical relation between photometric observables (e.g., magnitudes, colors, and morphological indicators) and redshift. Examples of empirical methods include polynomial fitting (Connolly

et al. 1995), the nearest neighbor method (Csabai et al. 2003), the nearest neighbor polynomial (NNP) technique (C. Cunha et al. 2008, in preparation), neural networks (NN; Collister & Lahav 2004; Vanzella et al. 2004; d'Abrusco et al. 2007), and support vector machines (Wadadekar 2005). When a large spectroscopic training set that is representative of the photometric data set to be analyzed is available, training set techniques typically outperform template-fitting methods, in the sense that the photo- z estimates have smaller scatter and bias with respect to the true redshifts (C. Cunha et al. 2008, in preparation). On the other hand, template fitting can be applied to a photometric sample for which relatively few spectroscopic analogs exist. For a comprehensive review and comparison of photo- z methods, see C. Cunha et al. (2008, in preparation).

In this paper, we present a publicly available galaxy photometric redshift catalog for the sixth data release (DR6) of the Sloan Digital Sky Survey (SDSS) imaging catalog (Blanton et al. 2003; Eisenstein et al. 2001; Gunn et al. 1998; Ivezić et al. 2004; Strauss et al. 2002; York et al. 2000). We use the NN photo- z method, which we have shown to be a superior training set method (C. Cunha et al. 2008, in preparation), and briefly compare the results using different photometric observables. We also compare the NN results with those from NNP, an empirical method which achieves similar performance to the NN method (C. Cunha et al. 2008, in preparation). Since the SDSS photometric catalog covers a large area of sky, a number of deep spectroscopic galaxy samples with SDSS photometry are available to use as training sets, as shown in Figure 1. In combination, these spectroscopic samples cover the full apparent magnitude range of the SDSS photometric sample.

The paper is organized as follows. In § 2, we briefly describe the SDSS DR6 photometric catalog and the selection criteria used to obtain the galaxy photometric sample from the catalog. In § 3, we describe the spectroscopic catalogs used to construct the photo- z training and validation sets. In § 4, we outline the photo- z methods, as well as the photo- z error estimator technique applied to the galaxy sample. Statistical results for photometric

¹ Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637.

² Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637.

³ Department of Physics, University of Chicago, Chicago, IL 60637.

⁴ Center for Particle Astrophysics, Fermi National Accelerator Laboratory, Batavia, IL 60510.

⁵ Center for Cosmology and Particle Physics and Department of Physics, New York University, New York, NY 10003.

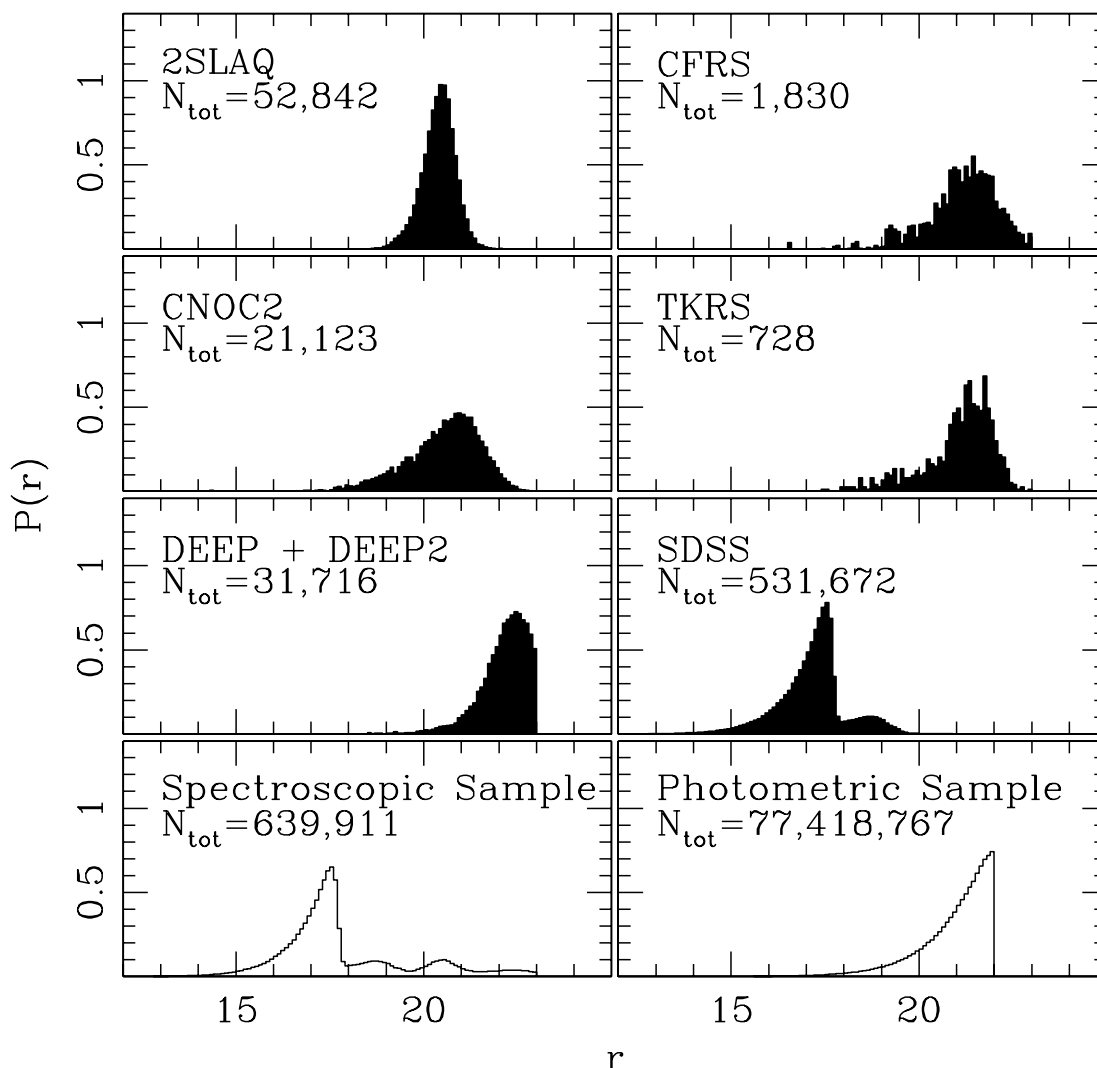


FIG. 1.—Normalized r -magnitude distributions for various catalogs. *Top three rows*: Distributions of the spectroscopic catalogs used for photo- z training and validation for 2SLAQ, CFRS, CNOC2, TKRS, DEEP, and DEEP2, and the SDSS spectroscopic sample. N_{tot} denotes the total number of galaxy measurements used from each catalog; for galaxies in regions with repeat SDSS imaging, each independent photometric measurement is counted separately. *Bottom row*: Distribution of the combined spectroscopic sample (*left*) and distribution for the SDSS photometric galaxy sample (*right*), where objects were classified as galaxies according to the photometric TYPE flag (see text). [See the electronic edition of the Journal for a color version of this figure.]

redshift performance, errors, and redshift distributions are presented in § 5. In § 6, we make recommendations for possible additional cuts on the photo- z catalog based on our own flags and those in the SDSS database. In § 7, we briefly describe how to access the photo- z catalog from the public SDSS data server, and in § 8, we present our conclusions. For completeness, Appendix A provides the database query used to select the photometric sample, Appendix B discusses issues of star-galaxy separation, and Appendix C briefly describes an earlier version of the photo- z algorithm used for SDSS DR5 (Adelman-McCarthy et al. 2007).

2. SDSS PHOTOMETRIC CATALOG AND GALAXY SELECTION

The SDSS comprises a large-area imaging survey of the north Galactic cap, a multiepoch imaging survey of an equatorial stripe in the south Galactic cap, and a spectroscopic survey of roughly 10^6 galaxies and 10^5 quasars (York et al. 2000). The survey uses a dedicated, wide-field, 2.5 m telescope (Gunn et al. 2006) at Apache Point Observatory, New Mexico. Imaging is carried out in drift-scan mode using a 142 megapixel camera (Gunn et al.

2006) that gathers data in five broad bands, $ugriz$, spanning the range from 3000 to 10,000 Å (Fukugita et al. 1996), with an effective exposure time of 54.1 s per band. The images are processed using specialized software (Lupton et al. 2001; Stoughton et al. 2002), and are astrometrically (Pier et al. 2003) and photometrically (Hogg et al. 2001; Tucker et al. 2006) calibrated using observations of a set of primary standard stars (Smith et al. 2002) observed on a neighboring 20 inch (50.8 cm) telescope.

The imaging in the sixth SDSS data release (DR6) covers an essentially contiguous region of the north Galactic cap, with only a few small patches remaining to be observed. In any region where imaging runs overlap, one run is declared primary⁶ and is used for spectroscopic target selection; other runs are declared secondary. The area covered by the DR6 primary imaging survey, including the southern stripes, is 8417 deg², but DR6 includes both the primary and secondary observations of each area and source (Adelman-McCarthy et al. 2008).

⁶ For the precise definition of primary objects, see <http://cas.sdss.org/dr6/en/help/docs/glossary.asp#P>.

TABLE 1
PHOTOMETRIC SAMPLE PROPERTIES

AB Magnitude Limits	
<i>u</i>	22.0
<i>g</i>	22.2
<i>r</i>	22.2
<i>i</i>	21.3
<i>z</i>	20.5
rms Photometric Calibration Errors (%)	
<i>r</i>	2
<i>u</i> - <i>g</i>	3
<i>g</i> - <i>r</i>	2
<i>r</i> - <i>i</i>	2
<i>i</i> - <i>z</i>	3

NOTES.—Magnitude limits are for 95% completeness for point sources in typical seeing; 50% completeness numbers are generally 0.4 mag fainter (Adelman-McCarthy et al. 2007). The median seeing for the SDSS imaging survey is 1.4".

The SDSS database provides a variety of measured magnitudes for each detected object. Throughout this paper, we use dereddened model magnitudes to perform the photometric redshift computations. To determine the model magnitude, the SDSS photometric pipeline fits two models to the image of each galaxy in each passband: a de Vaucouleurs (early-type) and an exponential (late-type) light profile. The models are convolved with the estimated point spread function (PSF), with arbitrary axis ratio and position angle. The best-fit model in the *r* band (which is used to fix the model scale radius) is then applied to the other passbands and convolved with the passband-dependent PSFs to yield the model magnitudes. Model magnitudes provide an unbiased color estimate in the absence of color gradients (Stoughton et al. 2002), and the dereddening procedure removes the effect of Galactic extinction (Schlegel et al. 1998).

To construct the photometric sample of galaxies for which we wish to estimate photo-*z*'s, we obtained a catalog drawn from the SDSS CasJobs Web site.⁷ We checked some of the SDSS photometric flags to ensure that we have obtained a reasonably clean galaxy sample. In particular, we selected all primary objects from DR6 that have the TYPE flag equal to 3 (the type for galaxy), and that do not have any of the flags BRIGHT, SATURATED, or SATUR_CENTER set. For the definitions of these flags, we refer the reader to the PHOTO flags entry at the SDSS Web site⁸ or to Appendix A. We also took into account the nominal SDSS flux limit (see Table 1) by only selecting galaxies with dereddened model magnitude $r < 22.0$. The full database query we used is given in Appendix A.

The photometric galaxy catalog we selected suffers from impurity and incompleteness at some level, since the photometric pipeline cannot separate stars from galaxies with 100% success at faint magnitudes. We describe some of our tests of star-galaxy separation in Appendix B, where we show that the SDSS TYPE flag provides star-galaxy separation performance similar to other methods.

The final photometric sample comprises 77,418,767 galaxies. The *r*-magnitude distribution of this sample is shown in Figure 1

(bottom right); the *g* - *r* and *r* - *i* color distributions are shown in Figure 2 (bottom).

3. SPECTROSCOPIC TRAINING AND VALIDATION SETS

Since our methods to estimate photo-*z*'s and photo-*z* errors are training set based, we would ideally like the spectroscopic training set to be fully representative of the photometric sample to be analyzed, i.e., to have similar statistical properties and magnitude/redshift distributions. Training set methods can be thought of as inherently Bayesian, in the sense that the training set distributions form effective priors for the analysis of the photometric sample; to the extent that the training set distributions reflect those of the photometric sample, we may expect the photo-*z* estimates to be unbiased (or at least they will not be biased by the prior). Given the practical difficulties of carrying out spectroscopy at faint magnitudes and low surface brightness, such an ideal generally cannot be achieved. Realistically, all we can hope for is a training set that (1) is large enough that statistical fluctuations are small and (2) spans the same magnitude, color, and redshift ranges as the photometric sample. Fortunately, our tests indicate that the estimated photo-*z*'s depend only weakly on the shape of the redshift and magnitude distributions of the training set for the SDSS.

We have constructed a spectroscopic sample consisting of 639,911 galaxies that have SDSS photometry measurements (counting repeats; see below) and that have spectroscopic redshifts measured by the SDSS or by other surveys, as described below. We imposed a magnitude limit of $r < 23.0$ on the spectroscopic sample and applied additional cuts on the quality of the spectroscopic redshifts reported by the different surveys. Since we impose a limit of $r < 22.0$ for the SDSS photometric sample, the fainter limit chosen for the spectroscopic training sample accommodates the full photometric range of interest without creating boundary effects for photo-*z*'s of galaxies with magnitudes near the photometric sample limit of $r = 22$. Each survey providing spectroscopic redshifts defines a redshift quality indicator; we refer the reader to the respective publications listed below for their precise definitions. For each survey, we chose a redshift quality cut roughly corresponding to 90% redshift confidence or greater. The SDSS spectroscopic sample provides 531,672 redshifts, principally from the MAIN and luminous red galaxy (LRG) samples, with a confidence level $z_{\text{conf}} > 0.9$. The remaining redshifts are: 21,123 from the Canadian Network for Observational Cosmology Field Galaxy Survey (CNOC2; Yee et al. 2000); 1830 from the Canada-France Redshift Survey (CFRS; Lilly et al. 1995) with $\text{CLASS} > 1$; 31,716 from the Deep Extragalactic Evolutionary Probe (DEEP; Davis et al. 2001) with $q_z = \text{A or B}$ and from DEEP2 (Weiner et al. 2005)⁹ with $z_{\text{quality}} \geq 3$; 728 from the Team Keck Redshift Survey (TKRS; Wirth et al. 2004) with $z_{\text{quality}} > -1$; and 52,842 LRGs from the Two Degree Field, SDSS LRG, and Quasi-stellar Object Survey (2SLAQ; Cannon et al. 2006)¹⁰ with $z_{\text{op}} \geq 3$.

We positionally matched the galaxies with spectroscopic redshifts against photometric data in the SDSS BestRuns Catalog Archive Server (CAS) database, which allowed us to match with photometric measurements in different SDSS imaging runs. The above numbers for galaxies with redshifts count independent photometric measurements of the same objects due to multiple SDSS imaging of the same region; in particular, SDSS Stripe 82 has been imaged a number of times. The numbers of *unique* galaxies used from these surveys are 1435 from CNOC2, 272 from CFRS, 6049

⁷ See <http://casjobs.sdss.org/casjobs/>.

⁸ See <http://cas.sdss.org/dr6/en/help/browser/browser.asp>.

⁹ See <http://deep.berkeley.edu/DR2/>.

¹⁰ See http://lrg.physics.uq.edu.au/New_dataset2/.

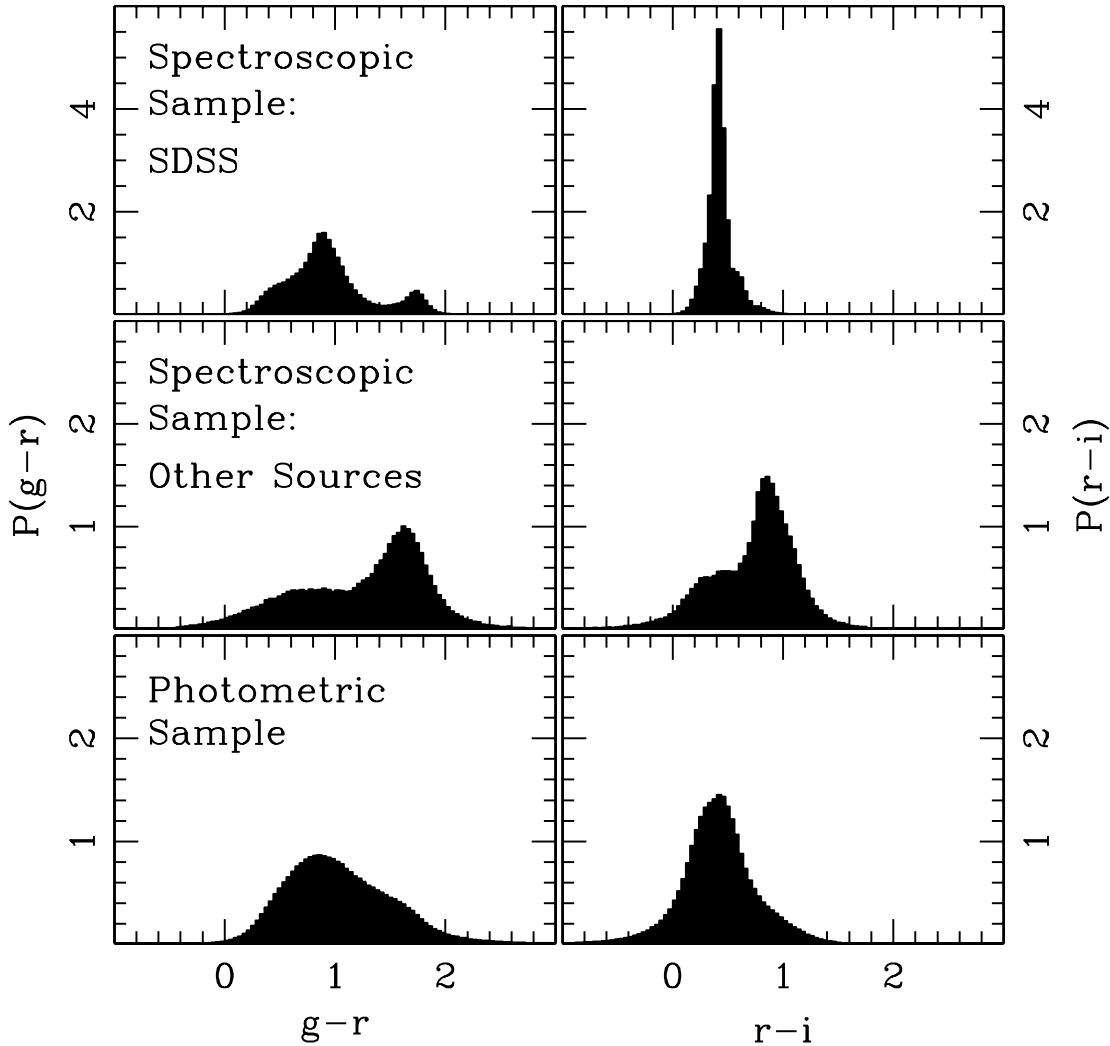


FIG. 2.—Distribution of $g-r$ and $r-i$ colors for different SDSS samples. *Top*: Color distributions for galaxies in the SDSS spectroscopic sample. *Middle*: Color distributions for galaxies in the other (non-SDSS) spectroscopic training samples. *Bottom*: Color distributions for galaxies in the photometric sample. As above, galaxy/star classification used the photometric TYPE flag. [See the electronic edition of the Journal for a color version of this figure.]

from DEEP and DEEP2, 389 from TKRS, and 11,426 from 2SLAQ. The SDSS spectroscopic samples were drawn from the SDSS primary galaxy sample, and therefore are all unique.

The spectroscopic sample obtained by combining all these catalogs, including the repeats, was divided into two catalogs of the same size ($\sim 320,000$ objects each). One of these catalogs was taken to be the *training set* used by the photo- z and error estimators, and the other was used as a *validation set* to carry out tests of photo- z quality (see § 4.1). In our tests, we have found that including the repeat imaging reduces the photo- z scatter at the few percent level compared to photo- z methods trained without the repeats, provided all the photometric measurements for a given object are confined to either the training set or the validation set, and are not mixed. This improvement is due to the fact that repeat measurements allow the training set methods to more fully characterize the effects of magnitude errors on the photometric redshifts. In addition, the repeats help the training methods to avoid overfitting to a particular realization of the objects' photometry. We find that this feature of the repeats is particularly helpful for faint training set galaxies ($r > 21$) whose magnitude errors are large. Furthermore, excluding such multiple images from the spectroscopic sample would result in much smaller training and validation sets; these would be very sparse at faint magnitudes,

leading to much diminished photo- z quality. However, we were careful not to place the repeat images of a single galaxy in both the training and validation sets, so as to not lose the independence of the two sets (see § 4.1).

The r -magnitude and color ($g-r$ and $r-i$) distributions for the spectroscopic samples and for the photometric sample are shown in Figures 1 and 2. While the magnitude and color distributions of the combined spectroscopic sample are not identical to those of the photometric sample, the spectroscopic sample does span the range of apparent magnitude and color of the photometric sample. To test the impact of having a training set that is not fully representative of the photometric sample, we divided the spectroscopic sample into smaller, alternate training and validation sets. For instance, to test the effect of the training-set magnitude distribution on the photo- z estimates, we created a training set with a flat r -magnitude distribution and another with an r -magnitude distribution similar to that of the photometric sample. Our tests indicated that the photo- z quality is not strongly affected by the magnitude distribution of the training set. The changes in the photo- z performance metrics (the rms scatter and the 68% CL region, defined below in § 5) were smaller than 10% when the training-set magnitude distribution was varied between these different choices. Since using the entire spectroscopic sample for the

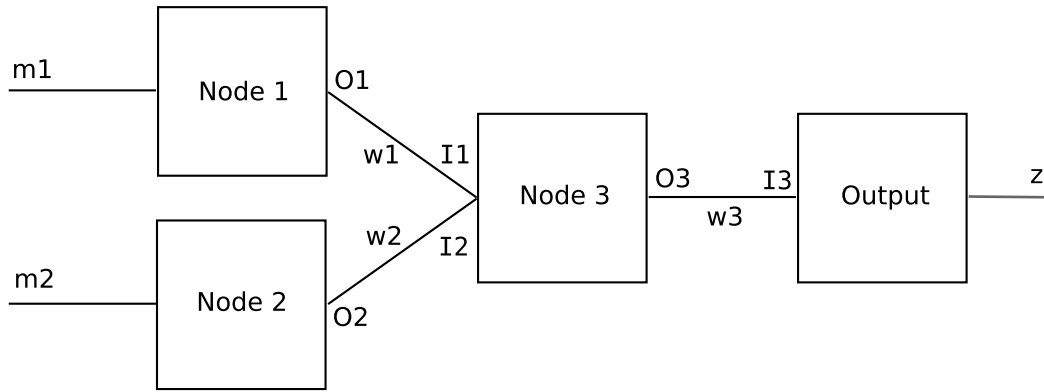


FIG. 3.— Simple FFMP network with three layers and configuration 2:1:1. The inputs are the two magnitudes, m_1 and m_2 . I_x denotes the input from node x , and O_x is the corresponding output of this node. The weights w associated with each connection are found by training the network using training and validation sets (see text).

training and validation sets produced marginally better results than all other cases tested, we have adopted this as our final choice. In addition, we tested the effect of the size of the training set on our photo- z calculations. We found that the photo- z performance metrics defined in § 5.1 are degraded by no more than 10% when the training set is artificially reduced to 10% of its original size. Even when the training set is reduced to $\sim 1\%$ of its original size, the photo- z performance metrics are degraded by less than 25%. This gives us confidence that the spectroscopic training set size used here is sufficient for extracting nearly optimal photo- z estimates.

4. METHODS

4.1. NN and NNP Photometric Redshifts

The NN method that we use to estimate galaxy photo- z 's is a general classification and interpolation tool used successfully in an array of fields such as handwriting recognition, automatic aircraft piloting,¹¹ credit card fraud detection,¹² and the extraction of astronomically interesting sources in a telescope image (Bertin & Arnouts 1996).

We use a particular type of NN called a feed forward multilayer perceptron (FFMP) to map the relationship between photometric observables and redshifts. An FFMP network consists of several input nodes, one or more hidden layers, and several output nodes, all interconnected by weighted connections (see Fig. 3). We follow the notation of Collister & Lahav (2004) and denote a network with N_i input nodes, N_{h_j} nodes in hidden layer j , and N_o output nodes as $N_i:N_{h_1}:N_{h_2}:\dots:N_{h_m}:N_o$. For each input object, the input photometric data (e.g., magnitudes, colors, concentrations, etc.) are fed into the input nodes of the FFMP, which fire signals according to the values of the input data. Each node in a hidden layer receives as a total input a weighted sum of the outputs from the nodes in the previous layer; i.e., node i in a hidden layer receives an input I_i given by

$$I_i = \sum_j w_{ij} O_j, \quad (1)$$

where O_j is the output of the j th node of the previous layer, and w_{ij} is the weight of the connection between node i in the hidden layer and node j in the previous layer. Given the input I_i , the output O_i of node i is a function f of the input,

$$O_i = f(I_i), \quad (2)$$

where f is the activation function. Repeating this process, signals propagate up to the output nodes. The activation function is typically a sigmoid function:

$$f(I_i) = \frac{1}{1 + e^{-I_i}}. \quad (3)$$

However, there are various alternatives, such as step functions and hyperbolic tangents. Vanzella et al. (2004) show that the choice of activation functions makes no significant difference in the result.

We use $X:20:20:20:1$ networks to estimate photo- z 's, where X is the number of input photometric parameters per galaxy. The corresponding number of degrees of freedom (the number of weights) is roughly 1000, depending on the actual value of X . We use hyperbolic tangent functions as the activation function of the hidden layers and a linear activation function for the output layer.

Despite the occasional aura of mystery surrounding neural networks, an FFMP is nothing more than a complex mathematical function; in fact, one can always write down the analytic expression corresponding to a neural network function.

Once the network configuration is specified, it can be trained to output an estimate of redshift given the input photometric observables. The training process involves finding the set of weights w_{ij} that minimize a score function E , chosen here to be

$$E = \frac{1}{2} \sum_i (z_{\text{spec}}^i - z_o^i)^2, \quad (4)$$

where z_{spec} is the measured spectroscopic redshift, z_o is the output redshift of the output node, and the sum is over all galaxies in the training set. Note that the choice of score function is not unique, and different choices will in general lead to different photo- z estimates. The minimization of this score function can be done efficiently, because its derivatives with respect to the weights are available analytically. We use a variable metric method as described in Press et al. (1992) for the minimization.

In machine learning, overfitting refers to the tendency of an algorithm with many adjustable parameters to fit to the noise in the training set data. In order to avoid overfitting, we use the technique of early stopping. The spectroscopic sample is divided into two independent subsets, the *training* and *validation* sets, and the formal minimizations are done using the training set. After each minimization step, the network is evaluated on the validation set, and the set of weights that performs best on the validation set is chosen as the final set. Another issue in machine learning is that

¹¹ See <http://www.nasa.gov/centers/dryden/news/NewsReleases/2003/03-49.html>.

¹² See <http://www.visa.ca/en/about/visabenefits/innovation.cfm>.

minimization procedures that start at different initial choices of weights generally end at different local minima of the score function. To reduce the chance of ending in a less than optimal local minimum, we minimize five networks starting at different positions in the space of weights. Among these, we choose the network that gives the lowest photo- z scatter (cf. eq. [4]) in the validation set. For more details of our implementation of the NN and its performance on mock catalogs and real data, see C. Cunha et al. (2008, in preparation).

The NN photo- z algorithm is very flexible in the sense that it is easy to change the input parameters, the training set, and the network configurations. We tried a variety of combinations of possible input photometric observables to see their effects on photo- z quality. We calculated photo- z 's using galaxy magnitudes, colors, and the concentration indices for some or all of the passbands. The concentration index c_i in passband i is defined as the ratio of PetroR50 and PetroR90, which are the radii that encircle 50% and 90% of the Petrosian flux, respectively. Early-type (E and S0) galaxies, with centrally peaked surface brightness profiles, tend to have low values of the concentration index, while late-type spirals, with quasi-exponential light profiles, typically have higher values of c . Previous studies (Morgan 1958; Shimasaku et al. 2001; Yamauchi et al. 2005; Park & Choi 2005) have shown that the concentration parameter correlates well with galaxy morphological type, and we used it to help break the degeneracy between redshift and galaxy type. We present the photo- z results for different combinations of input parameters in § 5.

For comparison, we also computed photo- z 's for the validation set using another empirical method, the nearest neighbor polynomial (NNP) technique (C. Cunha et al. 2008, in preparation). In NNP, to derive a photo- z for a galaxy in the photometric sample, we look for its training set nearest neighbors in the space of photometric observables (magnitudes, colors, etc.). Suppose we have N_D photometric data entries for each galaxy. The data vector for the galaxy of interest in the photometric sample is denoted by $D^\mu = (D^1, D^2, \dots, D^{N_D})$, while the data vector for the i th galaxy in the training set is $D_i^\mu = (D_i^1, D_i^2, \dots, D_i^{N_D})$. The distance d_i between the photometric object and the i th training set galaxy is defined using a flat metric in data space,

$$d_i^2 = \sum_{\mu=1}^{N_D} (D^\mu - D_i^\mu)^2. \quad (5)$$

The nearest neighbors are the training set objects for which d_i is minimum. Once the nearest neighbors for a given galaxy are identified, they are used to fit the coefficients of a local, low-order polynomial relation between photometric observables and redshift. The galaxy photo- z is then obtained by applying the derived relation to the photometric object.

For the NNP method employed in this work, we take the photometric data D^μ in equation (5) to be the four ‘‘adjacent’’ galaxy colors $u - g$, $g - r$, $r - i$, and $i - z$; we found that this choice produces results marginally better than those obtained using the galaxy magnitudes. We use the nearest 1000 neighbors to fit a quadratic polynomial relation between redshift and the photometric data, here chosen to be the five magnitudes in each passband (*ugriz*) and their corresponding concentration indices. We note that Wang et al. (2007) used a similar technique to estimate photo- z 's for a small sample of SDSS *spectroscopic* galaxies. They applied the kernel regression method of order 0, weighting the training set neighbors and computing photo- z 's by using the weighted average of the neighbors' redshifts. Our NNP method is closer to a kernel regression of order 2, since we perform quadratic fits; however, we

do not apply variable weights to the neighbors, but treat them equally in the fit.

Whereas the NN method provides a single, nonlinear, global fit using the whole training set and applies the derived photo- z relation to all photometric objects, the NNP method yields a separate, linear (in parameters), local fit for each photometric object using its neighbors. If the galaxy magnitude-concentration-redshift hypersurface is a differentiable manifold, i.e., if it can be locally approximated by a hyperplane, even though it is globally curved, then these two photo- z methods should be roughly equivalent. Indeed, as we show in § 5, their performance is very similar.

4.2. Photometric Redshift Errors

We estimated photo- z errors for objects in the photometric catalog using the nearest neighbor error (NNE) estimator (Oyaizu et al. 2007). The NNE method is training set based, with a neighbor selection similar to the NNP photo- z estimator; it associates photo- z errors to photometric objects by considering the errors for objects with similar multiband magnitudes in the validation set. We use the validation set, because the photo- z 's of the training set could be overfit, which would result in NNE underestimating the photo- z errors.

The procedure to calculate the redshift error for a galaxy in the photometric sample is as follows. We find the validation set nearest neighbors to the galaxy of interest. In contrast to NNP, for which the distance in equation (5) is defined in color space, the NNE distance is defined in magnitude space, since photo- z errors correlate strongly with magnitude. Since the selected nearest neighbors are in the spectroscopic sample, we know their photo- z errors, $\delta z = z_{\text{phot}} - z_{\text{spec}}$, where z_{phot} is computed using the NN or the NNP method. We calculated the 68% width of the δz distribution for the neighbors, and assigned that number as the photo- z error estimate for the photometric galaxy. Here we selected the nearest 200 neighbors of each object to estimate its photo- z error. In studies of photo- z error estimators applied to mock and real galaxy catalogs, we found that NNE accurately predicts the photo- z error when the training set is representative of the photometric sample (Oyaizu et al. 2007).

4.3. Estimating the Redshift Distribution

As we shall see in § 5.1, estimates for galaxy photo- z 's suffer from statistical biases that in general cannot be completely removed on an object-by-object basis. However, we can seek an unbiased estimate of the true redshift *distribution* for the photometric sample that is independent of individual galaxy photo- z estimates. For some statistical applications, the redshift distribution of the photometric sample, as opposed to individual galaxy photo- z 's, is all that is required. One way to estimate this distribution is to assign a weight to every galaxy in the spectroscopic sample such that the *weighted* spectroscopic sample has the same distributions of magnitudes and colors as the photometric sample. The z_{spec} distribution of this weighted spectroscopic sample provides an estimate of the true, underlying redshift distribution of the photometric sample.

The weight W^α of the α th spectroscopic galaxy is calculated by comparing the local density around the galaxy in the spectroscopic sample with the density of the corresponding region in the photometric sample. The local density is evaluated by counting the number of nearest neighbors using the distance measured in the space of photometric observables, as in equation (5). We fix the number of spectroscopic neighbors, N_S , which determines the distance d_{max} to the N_S th-nearest spectroscopic neighbor. We then find the number of neighbors N_P in the photometric sample within

the same distance d_{\max} of the spectroscopic galaxy. Up to an arbitrary normalization factor, the weight is defined as

$$W^\alpha \sim \frac{N_P}{N_S}. \quad (6)$$

For our estimates, we chose $N_S = 20$, which provides a good match of the weighted spectroscopic distributions of magnitudes and colors to those of the photometric sample. We note that if additional cuts in magnitude or color are applied to the photometric sample, then this procedure must be repeated for the newly selected photometric sample. More details and tests of this method and comparisons with other methods for estimating the underlying redshift distribution (e.g., deconvolving the error distribution from the z_{phot} distribution) will be presented separately (M. Lima et al. 2008, in preparation).

5. RESULTS

5.1. Photometric Redshifts

The photo- z precision (variance) and accuracy (bias) are limited by a number of factors. There are intrinsic degeneracies in magnitude-redshift space: low-luminosity, intrinsically red galaxies at low redshift can have apparent magnitudes similar to those of high-luminosity, intrinsically blue galaxies at high redshift. This natural degeneracy is amplified by photometric errors, since magnitude uncertainties propagate to photo- z errors. In addition to these observational limitations, which are determined by the photometric precision and the number of passbands of a survey, the photo- z estimator itself may have inherent limitations. For example, for training set methods, the size and representativeness of the training set are important factors, as are the number of parameters or weights in the fitting functions.

To test the quality of the photo- z estimates, we use four photo- z performance metrics. The first two metrics are the photo- z bias, z_{bias} , and the photo- z rms scatter, σ , both averaged over all N objects in the validation set, defined by

$$z_{\text{bias}} = \frac{1}{N} \sum_{i=1}^N (z_{\text{phot}}^i - z_{\text{spec}}^i), \quad (7)$$

$$\sigma^2 = \frac{1}{N} \sum_{i=1}^N (z_{\text{phot}}^i - z_{\text{spec}}^i)^2. \quad (8)$$

The third performance metric, denoted by σ_{68} , is the range containing 68% of the validation set objects in the distribution of $\delta z = z_{\text{phot}} - z_{\text{spec}}$. This metric is useful because the probability distribution function $P(\delta z)$ is in general non-Gaussian and asymmetric (for a Gaussian distribution, σ and σ_{68} coincide). Explicitly, σ_{68} is defined by the value of $|z_{\text{phot}} - z_{\text{spec}}|$ such that 68% of the objects have $|z_{\text{phot}} - z_{\text{spec}}| < \sigma_{68}$. We also use the 95% region σ_{95} , defined similarly. In addition to these global metrics, we also define local versions of them in bins of redshift or magnitude.

To search for an optimal photo- z estimator, we computed photo- z 's using the NN method with different combinations of input photometric observables. Five of these combinations are listed in Table 2. In the first case, dubbed O1, the training and photo- z estimation are carried out using only the five magnitudes *ugriz*. In case C1, we use the five magnitudes and the five concentration indices $c_u, c_g, c_r, c_i,$ and c_z as the input parameters. In case CC1, we use only the four colors $u - g, g - r, r - i,$ and $i - z$. In case CC2, we combine the four colors with the concentration indices $c_g, c_r,$ and c_i in the *gri* filters. Finally, in case D1, we use the *ugriz* magnitudes and the $c_u, c_g, c_r, c_i,$ and c_z con-

TABLE 2
SUMMARY OF NN CASES

Case	Inputs/Description	σ	σ_{68}
O1.....	<i>ugriz</i>	0.0525	0.0229
C1.....	<i>ugriz</i> + $c_u c_g c_r c_i c_z$	0.0519	0.0224
D1.....	<i>ugriz</i> + $c_u c_g c_r c_i c_z$ (split training)	0.0519	0.0209
CC1.....	$u - g, g - r, r - i, i - z$	0.0668	0.0272
CC2.....	$u - g, g - r, r - i, i - z$ + $c_g c_r c_i$	0.0593	0.0245

NOTE.—Photo- z performance metrics σ and σ_{68} for the validation set using different input parameters (magnitudes, colors, and concentration indices) and training procedures.

centration indices, but we split the training set and the photometric sample into five bins of r magnitude and perform separate NN fits in each bin. In all five cases, we use an NN with three hidden layers, and tune the number of hidden nodes to keep the total number of degrees of freedom of the network roughly the same for all cases.

Table 2 provides a summary of the performance results of the different NN cases. We find that using concentration indices in addition to magnitudes (C1 vs. O1) helps break some degeneracies, and reduces the photo- z scatter by a few percent. Using only colors (CC1) degrades the photo- z performance by as much as 20%, mostly because the degeneracy between intrinsically red, nearby galaxies and intrinsically blue, distant galaxies (with red observed colors) cannot be broken. Adding concentration indices to color-only training (CC2) helps break such a degeneracy, because the concentration index correlates with galaxy type and hence intrinsic color. Of the five, case CC2 also yields the most realistic photometric redshift distribution for the photometric sample (see § 5.2). Finally, splitting the training set and photometric sample into magnitude bins (D1) produces results with the best performance metrics (σ and σ_{68}) of all the NN cases we have tested. We choose D1 and CC2 as the best NN cases, and describe their results in more detail below; their outputs for the photometric sample are included in the public DR6 database.

In Figure 4, we plot photometric redshift, z_{phot} , for all objects in the validation set versus true spectroscopic redshift, z_{spec} , for the different photo- z methods and cases, and in different ranges of r magnitude. The top row shows results for NN case D1, the middle row shows the performance of NN case CC2, and the bottom row shows results for the NNP method using magnitudes and concentration indices as the input parameters. In each panel, the values of the corresponding global photo- z performance metrics σ and σ_{68} are shown. The redshift bias z_{bias} is typically much smaller than σ or σ_{68} , since the photo- z methods are designed to minimize it (see Fig. 5). In each panel of Figure 4, the solid line traces $z_{\text{phot}} = z_{\text{spec}}$, i.e., the line for a perfect photo- z estimator. The dashed and dotted lines show the corresponding 68% and 95% regions, defined as above but in z_{spec} bins. Although each photo- z method probes the hypersurface defined by the photometric observables and redshift in a different way, they produce very similar results, suggesting that our results are limited not by the photo- z technique employed, but by the intrinsic degeneracies in magnitude-concentration-redshift space and by the photometric errors.

In Figures 5 and 6, we show the performance metrics $z_{\text{bias}}, \sigma,$ and σ_{68} as functions of r magnitude and z_{spec} for the validation set for the two preferred NN cases. We see that the photo- z precision degrades considerably for objects with $r > 20$. This increased scatter is expected, since the relative photometric errors increase as the nominal detection limit of the SDSS photometry is approached (see Table 1). While the bias for CC2 increases at

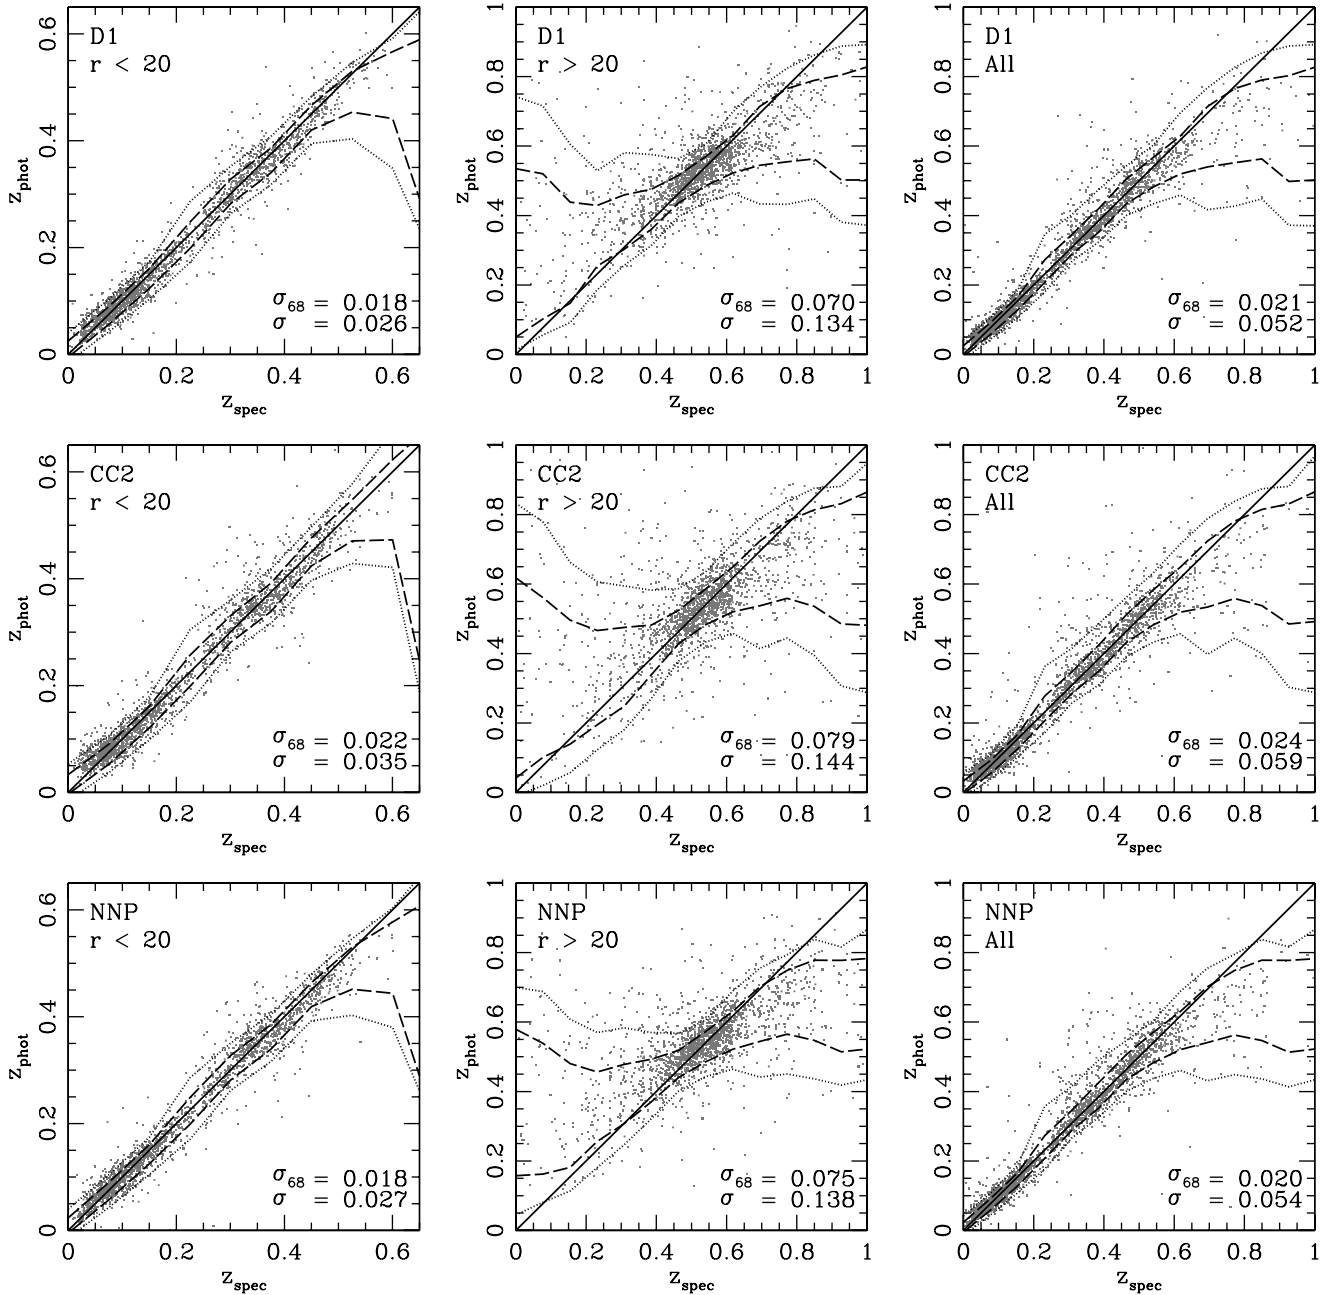


FIG. 4.— z_{phot} vs. z_{spec} for the validation set for different ranges of r magnitude and for different photo- z techniques. *Left column*: Objects with $r < 20$. *Middle column*: Objects with $r > 20$. *Right column*: All objects. *Top row*: NN case D1, for which the input photometric data comprise the five magnitudes ($ugriz$) and the five concentration parameters, and the training is split into five bins of r magnitude. *Middle row*: NN case CC2, for which the input data are the four colors $u - g$, $g - r$, $r - i$, and $i - z$, and the three concentration parameters c_g , c_r , and c_i . *Bottom row*: Results for the NNP method, for which the input data are the five magnitudes and five concentration parameters. In all cases, the photo- z methods used a training set with $\sim 320,000$ objects, and the derived solutions were applied to an independent validation set with $\sim 309,000$ objects and $r < 22$, reflecting the magnitude limit of the photometric sample. The solid line in each panel indicates $z_{\text{phot}} = z_{\text{spec}}$; the dashed and dotted lines show the 68% and 95% confidence regions as a function of z_{spec} . The points display results for a random 10% subset of the validation set in each magnitude range.

$r < 17$, we note that the fraction of objects in the photometric sample that are that bright is very small. As a function of redshift, σ and σ_{68} increase dramatically beyond $z \sim 0.6$ for the validation set. For the $r < 20$ part of the sample, the number of spectroscopic objects with $z > 0.6$ is simply too small to characterize the redshift-magnitude surface, as shown in Figure 7 (*left*). For the faint objects ($r > 20$), the scatter is low for z between 0.4 and 0.6, and increases outside of that range. It is important to note that the photo- z performance metrics were calculated independently of spectral type. Since the neural network and the training set were not optimized for any specific galaxy population (e.g.,

galaxies in clusters) it is possible that certain galaxy types may have photo- z 's with worse (or better!) biases and dispersion.

In Figure 7, we plot $g - r$ color versus spectroscopic redshift for the validation set for both bright ($r < 20$) and faint ($r > 20$) galaxies. The 2SLAQ and DEEP2 galaxies are highlighted by different shades of gray, and the expected color-redshift relations for the four spectral templates from Coleman et al. (1980) (from early to late types) are indicated by the solid lines. We see that for the faint sample, in the range $0.4 < z < 0.6$, the galaxies come mostly from the 2SLAQ survey, which used specific color cuts to select early-type galaxies at $z \sim 0.5$. Because early-type galaxies

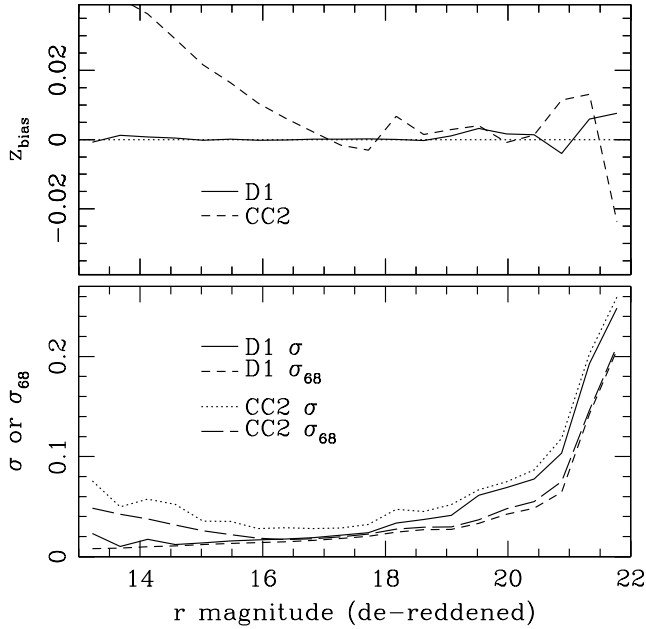


FIG. 5.—Performance metrics z_{bias} , σ , and σ_{68} for the NN D1 and NN CC2 validation sets, shown as a function of r magnitude. CC2 performs relatively poorly for bright objects ($r < 16$), for which the color-redshift relation is contaminated by faint objects with similar colors. In D1, this problem is alleviated by the effective magnitude prior imposed by the training set. At faint magnitudes, the performance degrades as the photometric errors increase.

have a well-defined 4000 Å break feature, their photo- z 's are well determined, and their photo- z scatter is low. Outside the range $0.4 < z < 0.6$, the validation set at faint magnitudes is dominated by bluer galaxies that do not have strong, broad spectral features, resulting in the larger photo- z scatter seen in Figure 6.

Figure 6 shows that the common assumption that the photo- z scatter scales as $(1+z)$ is not consistent with our estimates for

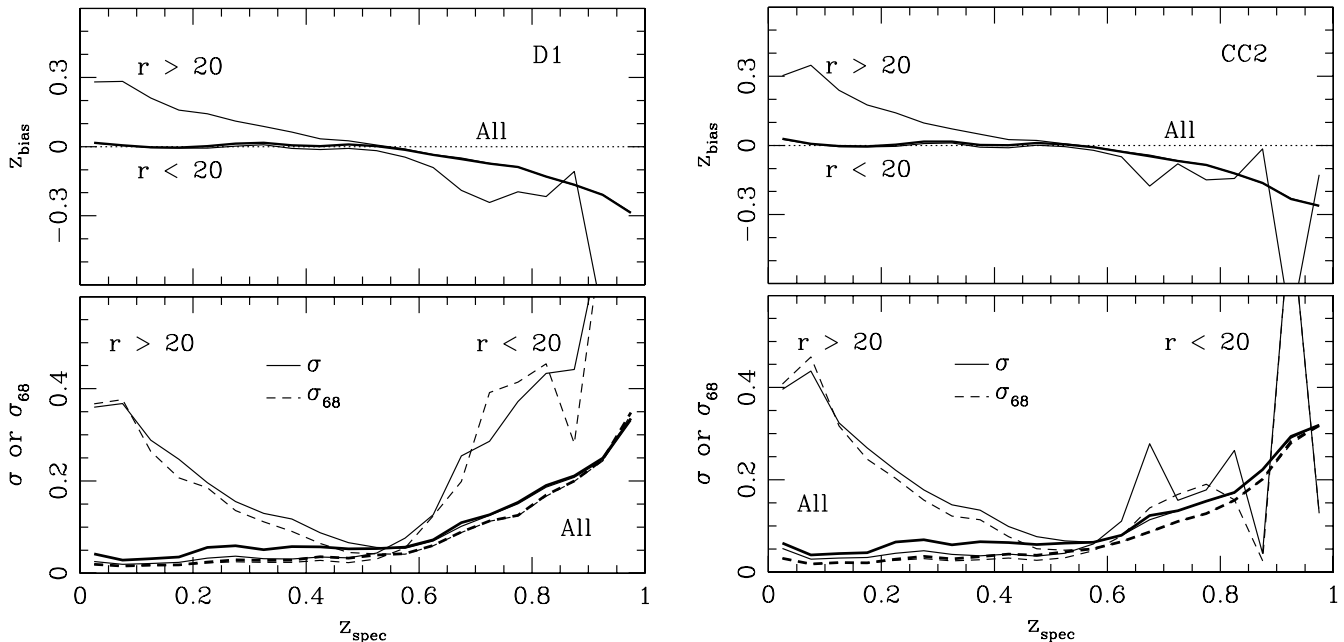


FIG. 6.—Performance metrics z_{bias} , σ , and σ_{68} for the NN D1 and NN CC2 validation sets as functions of z_{spec} for $r < 20$ and $r > 20$. The increased scatter for objects with $z > 0.6$ is due to the 4000 Å break shifting out of the r passband at around $z = 0.7$; beyond that redshift, the estimator effectively relies on only two passbands (i and z) to determine the photo- z 's. Note that faint objects ($r > 20$) have worse scatter at low redshifts for both cases. This is likely due to the fact that the faint, low-redshift objects in the validation set are predominantly blue dwarfs or irregular galaxies that do not have strong 4000 Å breaks; in this case, the photo- z estimator must rely on less pronounced spectral features, resulting in larger photo- z scatter. [See the electronic edition of the Journal for a color version of this figure.]

the SDSS sample. The functional form of the scatter versus redshift depends strongly on the underlying galaxy type distribution.

5.2. Redshift Distributions

So far, we have considered the scatter and bias of photo- z estimates. As discussed in § 4.3, it is also of interest to consider the predicted photo- z distribution as a whole. Different photo- z estimators may achieve similar values for the metrics z_{bias} , σ , and σ_{68} , but predict different forms for the photo- z distribution of the photometric sample. As we shall see, this is the case with the two NN cases D1 and CC2. We therefore define two additional performance metrics to quantify the quality of the predicted photo- z distribution. The first metric, σ_{dist} , measures the rms difference between the binned z_{phot} and z_{spec} distributions of the validation set,

$$\sigma_{\text{dist}}^2 = \frac{1}{N_{\text{bin}}} \sum_{i=1}^{N_{\text{bin}}} (P_{\text{phot}}^i - P_{\text{spec}}^i)^2, \quad (9)$$

where P_{phot}^i is the height of the i th redshift bin of the z_{phot} distribution, P_{spec}^i is the height of the same redshift bin of the z_{spec} distribution, and N_{bin} is the total number of redshift bins used. Here we use $N_{\text{bin}} = 120$ equally spaced redshift bins running from $z = 0$ to $z = 1.2$.

The second redshift distribution metric we employ is the Kolmogorov-Smirnov (K-S) statistic D , the maximum value of the absolute difference between the two (z_{phot} and z_{spec}) cumulative redshift distribution functions. An advantage of the K-S statistic is that it does not require binning the data in redshift. However, our use of the K-S statistic to quantify the difference between the z_{phot} and z_{spec} distributions of the validation set likely does not adhere to formal statistical practice, since it turns out that the probability for the K-S statistic for both cases we consider is very close to zero (Press et al. 1992).

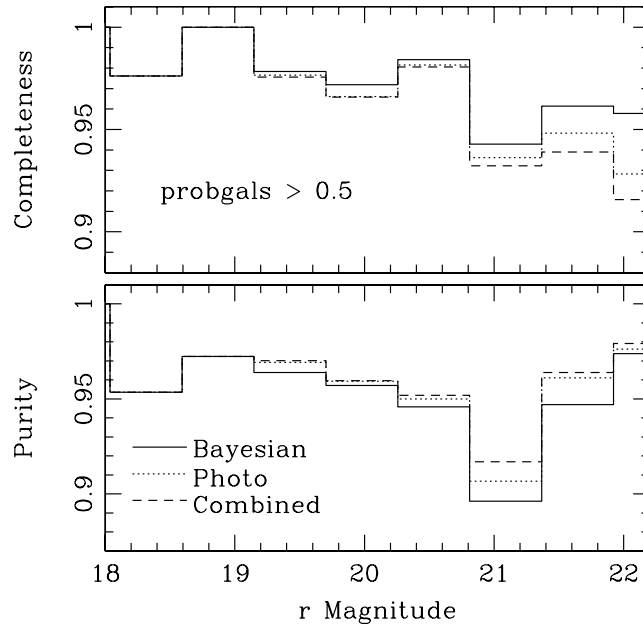


FIG. 12.—Completeness (*top*) and purity (*bottom*) for the Bayesian and PHOTO TYPE galaxy classifications, as well as for a combination of the two, using a sample of galaxies with spectroscopic classification. Results for the Bayesian separator have the `probgals` lower bound set to 0.5. [See the electronic edition of the *Journal* for a color version of this figure.]

SATUR_CENTER indicates that the object center is close to at least one saturated pixel. Note that in selecting PRIMARY objects (using PhotoPrimary), we have implicitly selected objects that either do *not* have the BLENDED flag set, or else have NODEBLEND set or `nchild` equal to zero. In addition, the PRIMARY catalog contains no BRIGHT objects, so the cut on BRIGHT objects in the query above is in fact redundant. BLENDED objects have multiple peaks detected within them, which PHOTO attempts to deblend into several CHILD objects. NODEBLEND objects are BLENDED, but no deblending was attempted on them, because they are too close to an EDGE, they are too large, or one of their children overlaps an edge. A few percent of the objects in our photometric sample have NODEBLEND set; some users may wish to remove them.

We also suggest that users require objects to have the BINNED1 flag set. BINNED1 objects were detected at $\geq 5\sigma$ significance in the original imaging frame.

The SDSS Web site¹⁵ provides further recommendations about flags, which we strongly recommend that users read.

APPENDIX B

TESTS ON STAR-GALAXY SEPARATION

We used the SDSS database TYPE flag to select the galaxy photometric sample for our photo-*z* catalogs. To study the robustness of the TYPE flag in separating galaxies from stars, we also carried out tests using an independent star-galaxy classifier. Here we briefly describe both of these techniques, and show the results obtained on photometric and spectroscopic samples.

The TYPE flag is based on the star-galaxy separator in the SDSS PHOTO pipeline, described in Lupton et al. (2001) and updated in Abazajian et al. (2004). For a given object, the pipeline computes the PSF and `cmodel` magnitudes in each passband,¹⁶ where the `cmodel` magnitude is a measure of the flux using a composite of the best-fit de Vaucouleurs and exponential models of the light profile. If the condition

$$m_{\text{PSF}} - m_{\text{cmodel}} > 0.145 \quad (\text{B1})$$

is satisfied, TYPE is set to GALAXY for that band; otherwise, TYPE is set to STAR. The object's global TYPE is determined by the same criterion, but now applied to the summed PSF and `cmodel` fluxes from all passbands in which the object is detected. Lupton et al. (2001) show that an earlier version of this simple cut works at the 95% confidence level for SDSS objects brighter than $r = 21$.

The second star-galaxy separator we tested is the galaxy probability defined in Scranton et al. (2002). The galaxy probability (hereafter `probgals`) is a Bayesian probability estimate that an object is a galaxy (and not a star), given the object's magnitudes and concentration parameter. Here the concentration parameter is *not* the ratio of Petrosian radii, but is defined as the difference between an object's PSF and exponential-model *r* magnitudes. This concentration parameter is close to zero for stars, is positive for bright galaxies, and approaches zero as galaxies become fainter.

We conducted some simple tests to compare these classification schemes. If we set the Bayesian `probgals` threshold to a value between 0.5 and 0.9, then both methods agree on the classification of more than 90% of the objects for a random 1% subset of the SDSS

¹⁵ See <http://cas.sdss.org/dr5/en/help/docs/algorithm.asp?key=flags>.

¹⁶ See <http://www.sdss.org/dr5/algorithms/photometry.html>.

TABLE 4
DR5 CATALOG Flag

Flag	No. Galaxies	Object Description
.....	86.1 million	All
0.....	12.6 million	Complete and bright
1.....	0.6 million	Incomplete and bright
2.....	59.0 million	Complete and faint
3.....	13.9 million	Incomplete and faint

NOTES.—The flag scheme for the DR5 catalog is based on object detection in some/all passbands and the r magnitude. Incomplete objects are undetected in at least one of the passbands ($ugriz$), and faint objects have $r > 20$.

photometric sample. We also tested the methods on a spectroscopic sample of 29,229 galaxies and stars (counting independent photometric measurements of each object) from the 2SLAQ and DEEP2 catalogs with $r < 22$. Defining stars as objects with $z_{\text{spec}} < 0.01$, the sample contains 24,541 galaxies and 4688 stars. We wish to compare this spectroscopic “truth table” with the photometric classification of the two methods and with a combined method that classifies an object as a galaxy if and only if both separators classify it as a galaxy. For the purposes of this test, we say that the Bayesian scheme classifies an object as a galaxy if `probga1s` > 0.5 . We define galaxy completeness as the ratio of correctly identified galaxies to the total number of galaxies in the spectroscopic sample. Purity is defined as the ratio of correctly identified galaxies to the number of objects identified (correctly or not) as galaxies by the classifier. The purity depends in part on the relative numbers of galaxies and stars in the spectroscopic sample.

Figure 12 shows the completeness and purity of the resulting galaxy catalogs in bins of r magnitude for this spectroscopic sample. Overall, the Bayesian separator and PHOTO TYPE produce similar results for galaxy purity and completeness. Moreover, the agreement between the two classification methods is quite good on an object-by-object basis. The Bayesian separator with `probga1s` ≥ 0.5 achieves slightly higher completeness and slightly lower purity. By varying the `probga1s` boundary, we could improve the purity of the Bayesian galaxy sample at the expense of degrading its completeness. We note that the best value of `probga1s` to use in defining a galaxy photometric sample depends on the scientific applications of the sample, i.e., on whether completeness or purity is the more important feature. In statistical applications, instead of defining a galaxy sample, one can also choose to weight objects by their Bayesian probability (Scranton et al. 2002).

Based on this test, we conclude that the photometric sample for which we have estimated photo- z ’s has better than 90% galaxy purity.

APPENDIX C

PHOTOMETRIC REDSHIFTS FOR SDSS DR5

An earlier version of the photo- z catalog, produced for SDSS Data Release 5 (DR5), is publicly available on the SDSS DR5 Web site (and is also called `photoz2`). The methods used to construct that photo- z catalog were similar to the ones employed here for DR6, but the latter incorporates a number of important improvements. Here we briefly outline the differences between the two. We *strongly* recommend use of the DR6 photo- z catalog instead of the DR5 catalog.

The photometric galaxy sample selection has improved from DR5 to DR6, because we used more stringent cuts in defining the DR6 sample. The DR6 sample selection is described above in Appendix A. The DR5 photometric galaxy sample selection required the `cmodel` and model r magnitudes to lie in the ranges $r_{\text{cmodel}} \in (14.0, 22.0)$ and $r_{\text{model}} \in (13.5, 22.5)$, and also required the value of the smear polarizability (Sheldon et al. 2004) to be $m_r > 0.8$. Also, for DR5, the Bayesian estimator was used to separate stars from galaxies (see Appendix B), with the value `probga1s` > 0.8 , while for DR6 we used PHOTO TYPE. The additional cuts used for the DR6 catalog have produced a cleaner and more reliable galaxy sample.

The DR5 photo- z catalog included a number of flags describing the expected photo- z quality, shown in Table 4. These flags were based on the detection or nondetection of the object in all passbands, and on the value of the r model magnitude. An object was classified as bright (faint) if $r < 20$ ($r > 20$). An object was flagged as “incomplete” if it was not detected in all five SDSS passbands. Table 4 shows the corresponding flag values and the number of objects assigned each flag value. For the DR6 sample, given the stricter sample selection, a very small number of objects would have been classified as incomplete by the definition above, and they have been removed from the sample. As a result, for DR6, we only supply the bright/faint flag, as shown in Table 5.

TABLE 5
DR6 CATALOG Flag

Flag	No. Galaxies	Object Description
.....	77.4 million	All
0.....	11.5 million	Bright
2.....	65.9 million	Faint

NOTE.—The flag scheme for the DR6 catalog is based solely on the r magnitude: faint objects have $r > 20$.

The spectroscopic training set used for the DR6 photo-*z* catalog has important additions compared to the one used for the DR5 catalog. In particular, for DR6 we added the DEEP2 spectroscopic catalog (which became publicly available), which made the training set more complete at faint magnitudes. We also implemented more stringent spectroscopic quality cuts to the training set used for DR6.

Unlike the DR5 training set, the DR6 training set does not contain objects from the SDSS “special” plates, which are extra spectroscopic observations designed to target specific objects for various scientific studies (Adelman-McCarthy et al. 2006). In our tests, we find that the lack of special plates does not result in any degradation of the photo-*z* quality.

The photo-*z* algorithm also changed from DR5 to DR6: we increased the number of hidden-layer nodes in the NN, and we added the concentration indices to the data inputs. Our tests indicated that this leads to improved photo-*z* performance according to our metrics. In addition, the CC2 method further differs from DR5 photo-*z*'s in that CC2 uses only the color information and not the raw magnitudes. For general purpose, full-sample photo-*z*'s, we recommend using CC2 photo-*z*'s over both DR5 and D1 photo-*z*'s. Finally, we carried out more extensive tests of the DR6 photo-*z*'s than were done for DR5, increasing our confidence in the robustness of the photo-*z* estimates.

REFERENCES

- Abazajian, K., et al. 2004, *AJ*, 128, 502
 Adelman-McCarthy, J. K., et al. 2006, *ApJS*, 162, 38
 ———. 2007, *ApJS*, 172, 634
 ———. 2008, *ApJS*, in press (arXiv:0707.3413)
 Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393
 Blanton, M. R., Lin, H., Lupton, R. H., Maley, F. M., Young, N., Zehavi, I., & Loveday, J. 2003, *AJ*, 125, 2276
 Cannon, R., et al. 2006, *MNRAS*, 372, 425
 Coleman, G. D., Wu, C. C., & Weedman, D. W. 1980, *ApJS*, 43, 393
 Collister, A. A., & Lahav, O. 2004, *PASP*, 116, 345
 Connolly, A. J., et al. 1995, *AJ*, 110, 2655
 Csabai, I., et al. 2003, *AJ*, 125, 580
 d’Abrusco, R., et al. 2007, in *Proc. Astron. Data Analysis*, ed. J. L. Starck, in press (astro-ph/0701137)
 Davis, M., Newman, J. A., Faber, S. M., & Phillips, A. C. 2001, in *Proc. of the ESO Workshop, Deep Fields*, ed. S. Cristiani, A. Renzini, & R. E. Williams (Berlin: Springer), 241
 Eisenstein, D. J., et al. 2001, *AJ*, 122, 2267
 Fukugita, M., Ichikawa, T., Gunn, J. E., Doi, M., Shimasaku, K., & Schneider, D. P. 1996, *AJ*, 111, 1748
 Gunn, J. E., et al. 1998, *AJ*, 116, 3040
 ———. 2006, *AJ*, 131, 2332
 Hogg, D. W., Finkbeiner, D. P., Schlegel, D. J., & Gunn, J. E. 2001, *AJ*, 122, 2129
 Ivezić, Ž. et al. 2004, *Astron. Nachr.*, 325, 583
 Lilly, S. J., Le Fevre, O., Crampton, D., Hammer, F., & Tresse, L. 1995, *ApJ*, 455, 50
 Lupton, R., Gunn, J. E., Ivezić, Z., Knapp, G. R., & Kent, S. 2001, in *ASP Conf. Ser. 238, Astronomical Data Analysis Software and Systems X*, ed. F. R. Harnden, Jr., F. A. Primini, & H. E. Payne (San Francisco: ASP), 269
 Morgan, W. W. 1958, *PASP*, 70, 364
 Oyaizu, H., Lima, M., Cunha, C., Lin, H., & Frieman, J. 2007, *ApJ*, submitted (arXiv:0711.0962)
 Park, C., & Choi, Y.-Y. 2005, *ApJ*, 635, L29
 Pier, J. R., et al. 2003, *AJ*, 125, 1559
 Press, W. H., et al. 1992, *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge: Cambridge Univ. Press)
 Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, 500, 525
 Scranton, R., et al. 2002, *ApJ*, 579, 48
 Sheldon, E. S., et al. 2004, *AJ*, 127, 2544
 Shimasaku, K., et al. 2001, *AJ*, 122, 1238
 Smith, J. A., et al. 2002, *AJ*, 123, 2121
 Stoughton, C., et al. 2002, *AJ*, 123, 485
 Strauss, M. A., et al. 2002, *AJ*, 124, 1810
 Tucker, D. L., et al. 2006, *Astron. Nachr.*, 327, 821
 Vanzella, E., et al. 2004, *A&A*, 423, 761
 Wadadekar, Y. 2005, *PASP*, 117, 79
 Wang, D., Zhang, Y. X., Liu, C., & Zhao, Y. H. 2007, preprint (arXiv:0706.2704)
 Weiner, B. J., et al. 2005, *ApJ*, 620, 595
 Wirth, G. D., et al. 2004, *AJ*, 127, 3121
 Yamauchi, C., et al. 2005, *AJ*, 130, 1545
 Yee, H. K. C., et al. 2000, *ApJS*, 129, 475
 York, D. G., et al. 2000, *AJ*, 120, 1579

A GALAXY PHOTOMETRIC REDSHIFT CATALOG FOR THE SLOAN DIGITAL SKY SURVEY DATA RELEASE 6

Hiroaki Oyaizu,^{1,2} Marcos Lima,^{2,3} Carlos E. Cunha,^{1,2} Huan Lin,⁴
Joshua Frieman,^{2,4} and Erin S. Sheldon⁵

Received 2007 July 30; accepted 2007 September 12

ABSTRACT

We present and describe a catalog of galaxy photometric redshifts (z_{photo}) for the Sloan Digital Sky Survey (SDSS) Data Release 6 (DR6). We use the neural network (NN) technique to calculate z_{photo} and the nearest neighbor error (NNE) method to estimate photometric errors for 77 million objects classified as galaxies in DR6 with $r < 22$. The photometric redshifts and errors are available for all objects in the catalog.