






Preliminary Results of Using k -nearest-neighbor Regression to Estimate the Redshift of Radio-selected Data Sets*

Kieran J. Luken^{1,2} , Ray P. Norris^{1,2} , and Laurence A. F. Park¹ 

¹ Western Sydney University Locked Bag 1797 Penrith, NSW 2751, Australia; k.luken@westernsydney.edu.au

² CSIRO Astronomy and Space Sciences Australia Telescope National Facility P.O. Box 76, Epping, NSW 1710, Australia
Received 2018 May 31; accepted 2018 October 22; published 2019 August 30

Abstract

In the near future, all-sky radio surveys are set to produce catalogues of tens of millions of sources with limited multiwavelength photometry. Spectroscopic redshifts will only be possible for a small fraction of these new-found sources. In this paper, we provide the first in-depth investigation into the use of k -nearest-neighbor (kNN) regression for the estimation of redshift of these sources. We use Australia Telescope Large Area Survey (ATLAS) radio data, combined with *Spitzer* Wide-Area Infrared Extragalactic Survey infrared, Dark Energy Survey optical, and Australian Dark Energy Survey spectroscopic survey data. We then reduce the depth of photometry to match what is expected from the upcoming Evolutionary Map of the Universe survey, testing against both data sets. To examine the generalization of our methods, we test one of the subfields of ATLAS against the other. We achieve an outlier rate of $\sim 10\%$ across all tests, showing that the kNN regression algorithm is an acceptable method of estimating redshift, and would perform better given a sample training set with uniform redshift coverage.

Key words: methods: analytical – methods: statistical – galaxies: distances and redshifts – galaxies: statistics – (cosmology:) distance scale

Online material: color figures

1. Introduction

Large-scale radio surveys are becoming more common, resulting in catalogues of millions of radio sources with limited multiwavelength data (Norris 2017). Knowledge of their redshift is important to achieve most scientific goals (Norris et al. 2011). While spectroscopic redshifts remain the gold standard, only a few million spectroscopic redshifts will be available in this decade, with the Sloan Digital Sky Survey (SDSS) having measured around three million over the northern sky (Abolfathi et al. 2018); the Taipan Galaxy Survey is expecting to provide two million spectroscopic redshifts out to $z = 0.4$ (da Cunha et al. 2017), and the Wide Area VISTA Extra-galactic Survey is expecting to measure 2.5 million redshifts across the souther sky out to $z = 1.5$ (Driver et al. 2016). Alternatively, redshift can also be measured photometrically, by comparing the magnitudes at different wavelengths to templates (Baum 1962; Butchins 1981; Loh & Spillar 1986). A photometric redshift—or photo- z —is measured using template fitting, and can be highly accurate, estimating redshift to an accuracy of $\sigma_{\Delta z/(1+z_{\text{spec}})} \sim 0.015$ (Salvato et al. 2011). However, this requires high-quality photometry in at least 15 different filter bands, and up to 31 different bands for the high-accuracy results. Unfortunately, this level of photometry will

not be available for large-scale sky surveys. Additionally, photometric template-fitting methods tend to fail catastrophically when attempted on active galactic nuclei (AGNs), particularly radio-selected AGNs (Duncan et al. 2018; Salvato et al. 2018).

Rather than measuring the redshift directly in the form of spectroscopy or indirectly by fitting templates, it has been shown that a photo- z can be estimated empirically using the knowledge of previously measured redshifts from similar astronomical objects.

Machine learning has been applied to this problem in the past in the form of neural networks (Tagliaferri et al. 2003; Firth et al. 2003; Collister & Lahav 2004; Cavuoti et al. 2012, 2014, 2015, 2017; Breschia et al. 2013; Sadeh et al. 2016; Pasquet-Itam & Pasquet 2018), random forests (Carliles et al. 2010; Carrasco et al. 2015; Mountrichas et al. 2017), the combination of template-fitting methods with Bayesian statistics (Duncan et al. 2018), and the stacking of a self-organised map and a decision tree (Zitlau et al. 2016). For the most part, these methods have mainly been concerned with maximizing the accuracy of the measured redshift by testing with optically selected galaxy samples, and have been able to achieve a similar accuracy to template-fitting methods given a large enough training set. The data sets used have been derived mainly from optical surveys like the SDSS, limiting the number

* Released on 2018 June 1.

of possible radio-loud AGNs, which also create issues for photometric template fitting.

Norris et al. (2018) have addressed the problem of relatively low-quality photometry available from all-sky surveys by comparing the performance of different algorithms when using photometry similar to the upcoming Evolutionary Map of the Universe (EMU) survey (Norris et al. 2011), using radio-selected AGNs. This has given a glimpse of the expected accuracies of various techniques when the depth of observations and the coverage of all frequency bands are not complete.

This paper provides a deeper investigation into the effectiveness of the k -nearest-neighbor (kNN) algorithm (Cover & Hart 1967) for estimating redshift. The kNN regression algorithm has previously been applied to photometric redshifts by Polsterer et al. (2013), Kügler et al. (2015) and Cavuoti et al. (2017); however, Kügler et al. used optical spectra from the SDSS, and Polsterer et al. and Cavuoti et al. used spectroscopic redshifts and optical magnitudes only from the SDSS. kNN regression is a regression model, meaning that it uses a training set of objects and their known redshift values to estimate the redshift of new objects. In particular, kNN regression estimates the redshift of each new object as the mean of the redshifts of the k nearest neighbours from the training set. To perform this task, we must define a feature space (so that we can represent each object as a vector), a metric (to provide the distance between object vectors), and the value of the constant k . The feature space dimensions are chosen as the set of variables that are thought to be predictive of the regression response. Each object is represented as a vector in this space using their measurements of the feature space variables (e.g., in our case each object's vectors contain a set of photometry at different wavelengths from a given object in the training set). It is common to find Euclidean distance being used as the kNN metric, and the number of neighbours k typically set within the range of 2 and 20.

Given that the speed of the kNN algorithm does not scale well with the number of dimensions or sources, this algorithm can be modified to use a k - d tree to find the k most similar sources (Polsterer et al. 2013). Alternatively, the problem can be parallelized and run on a GPU. This paper has implemented the kNN regression algorithm using the publicly available TensorFlow³ on GPU (Abadi et al. 2015), which provides a 3–5 times speed improvement over the equivalent CPU version.

Our initial investigation was to determine the suitability of the kNN algorithm for the problem. We then tested how well the algorithm generalizes by testing one sub-field against the other. We next modified our data set to match the depth expected from the EMU survey, and corresponding sky surveys from other astronomical regimes. For all of these tests, we

compared the use of Euclidean distance and Manhattan Taxicab distance.

2. The Data

The data set used in this paper is primarily based on the Australia Telescope Large Area Survey (ATLAS) Data Release 3 (Norris et al. 2006; Franzen et al. 2015), with cross-identifications to other wavelengths provided by J. Swan et al. (2019, in preparation). Other catalogues used include the *Spitzer* Wide-Area Infrared Extragalactic (SWIRE) infrared data set (Lonsdale et al. 2003), the Dark Energy Survey (DES) optical data set (Abbott et al. 2018), and the Australian Dark Energy Survey (OzDES) spectroscopic redshift data set (Yuan et al. 2015; Childress et al. 2017).

2.1. ATLAS

The ATLAS DR3 data set (Franzen et al. 2015) forms the basis for our total catalogue, providing 1.4 GHz radio flux densities on 4780 sources measured using the Australia Telescope Compact Array. The ATLAS data set covers the European Large Area ISO Survey–South 1 (ELAIS-S1) and extended Chandra Deep Field South (eCDFs) fields, down to an r.m.s noise level of $\sim 15 \mu\text{Jy}$.

2.2. SWIRE

The SWIRE data set provides infr-red flux densities at 3.6, 4.5, 5.8 and 8.0 μm , measured using *Spitzer* (Lonsdale et al. 2003), reaching a 5σ sensitivity of 5, 9, 43, and 40 $\mu\text{Jy}/\text{beam}$ respectively. SWIRE identifications were found for 4328 radio sources at 3.6 μm , providing at least a 3.6 μm flux for $\sim 91\%$ of radio sources. All ATLAS sources were initially cross-matched with SWIRE sources, and then the latter were cross-matched to the DES sources.

2.3. DES

The DES data set provides g , r , i , and z optical magnitudes (to a depth of $g = 24.33$, $r = 24.08$, $i = 23.44$, and $z = 22.69$), measured using the Dark Energy Camera mounted on the 4m Blanco telescope at the Cerro Tololo Inter-American Observatory in Chile. (Abbott et al. 2018). The DES data set provides optical counterparts for 3102 of our radio sources at g band, covering $\sim 65\%$ of sources.

2.4. OzDES

The OzDES data set provides the spectroscopic redshifts required to create any empirical model (Yuan et al. 2015; Childress et al. 2017). The spectroscopic redshift master list compiled as the OzDES data set by Childress et al. (2017) provides spectroscopic redshifts for 2012 radio sources, covering $\sim 42\%$ of sources.

³ <https://www.tensorflow.org/>

2.5. Test Sample

The kNN algorithm works best if the reference data set is shaped such that the feature space is populated homogeneously, i.e., avoiding strong concentrations in a certain region, or sparsely populated regions. In the tests described here, we made no correction for any excess sources in any of the given training samples. Instead, we calculated optical and infrared colours as $c_i = \text{mag}_i - \text{mag}_{i+1}$. This transformation improves the distribution of the photometric data over the parameter space for each band (Norris et al. 2018) and also replaces flux, which is brightness and redshift dependent, by colour that depends only on the spectral energy density (Polsterer et al. 2013). We completed this transformation on both the optical magnitudes and log-transformed infrared fluxes. Note that this operation reduces the effective number of dimensions of the feature space by two.

We compiled a full-sensitivity “DEEP” sample containing those sources that have photometry at 1.4 GHz, infrared 3.6, 4.5, 5.8, and 8.0 μm , optical g , r , i , and z bands, and a spectroscopic redshift. This provides us with 1408 sources with complete photometry and spectroscopy for our tests.

We then selected a “WIDE” sample to match the depth of photometry expected from the EMU survey (Norris et al. 2011), which will use the SkyMapper survey which has r and i limits of ≈ 22 (Wolf et al. 2018) and the AllWISE Infrared Survey (Cutri et al. 2013) of 3.6 $\mu\text{m} > 26 \mu\text{Jy}$ and 4.5 $\mu\text{m} > 56 \mu\text{Jy}$. We removed the 5.8 and 8.0 μm data from our sample, and rejected any sources that were undetected at any band at the above limits. The resulting WIDE sample had 760 sources.

As discussed in Section 2.1, the ATLAS catalogue covers two fields: the ELAIS-S1 field, which makes up 553 of the 1408 sources in the DEEP data set and 281 of the 760 sources in the WIDE data set, and the eCDFs field, which makes up 855 of the 1408 sources in the DEEP data set and 479 of the 760 sources in the WIDE data set.

3. Experiments

In these experiments, we use the kNN regression algorithm to estimate redshifts, using the following steps:

1. Compute a distance matrix between all test sources and training sources.
2. Sort the distance matrix by the distance calculated in step 1, identifying the k closest training sources in feature space to the test sources.
3. Take the mean redshift of the known sources identified from step 2, and assign it to the test source.

To apply this method, we split our data into a training and test set. Depending on the test being run, the training set was either 70% of the data set in the case of the full data set tests

(Tests 1–4 in Table 1), or the entire sub-field in the case of the sub-field tests (Tests 5–12 in Table 1). The remaining 30% or sub-field was set aside as the test set.

To avoid under- or over-fitting the data, 10-fold cross-validation was used with our training sets, excluding our test sets, minimizing the number of incorrect estimates. In all tests, the value of k that minimized the outlier rate varied, and is listed in Table 2.

We computed 95% confidence intervals for each redshift prediction using bootstrapping with 1000 iterations. Bootstrapping is a sampling method that allows us to estimate the variance of sample estimates under the assumption that the population from which the sample is taken is approximately many replications of the sample. The estimated intervals provide the range in which the true redshift is likely to occur, while also providing an indication of the uncertainty of the prediction. These confidence intervals are displayed in the form of error bars in our figures in Section 4.

In our investigation of the accuracy of kNN, we examined the effect of varying the following experimental parameters.

1. **Distance Metric:** we evaluated both the Euclidean distance metric:

$$d(\vec{p}, \vec{q}) = \sqrt{\sum_{i=1}^n (q_i - p_i)^2} \quad (1)$$

and the Manhattan Taxicab distance metric:

$$d(\vec{p}, \vec{q}) = \|\vec{p} - \vec{q}\| = \sum_{i=1}^n |p_i - q_i|, \quad (2)$$

where p and q are vectors containing the features of two sources.

2. **Depth of Photometry:** we used both the full-sensitivity DEEP sample and the reduced-sensitivity WIDE sample.
3. **Generalization:** we randomly selected training and test sets from both ATLAS sub-fields as one test, and used one ATLAS sub-field as the training set and the other as the test set and reverse as additional tests.

As these variations are not independent; each needed to be completed in combination with all others, resulting in the 12 experiments listed in Table 1.

4. Results

We present the results of the experiments in Table 2 and Figures 1 to 3.

In Table 2 and Figures 1 to 3, we calculate the outlier rate η as

$$\eta = \frac{\text{count}(|\Delta z| > 0.15 \times (1 + z_{\text{spec}}))}{\text{Number Of Sources}} \times 100, \quad (3)$$

Table 1
Details of Experiments Completed

Experiment Number	Size of Training Set	Size of Test Set	Distance Metric	Data set Used	Training Sample
1	986	422	Manhattan	DEEP	Random
2	986	422	Euclidean	DEEP	Random
3	532	228	Manhattan	WIDE	Random
4	532	228	Euclidean	WIDE	Random
5	553	855	Manhattan	DEEP	ELAIS-S1
6	553	855	Euclidean	DEEP	ELAIS-S1
7	281	479	Manhattan	WIDE	ELAIS-S1
8	281	479	Euclidean	WIDE	ELAIS-S1
9	855	553	Manhattan	DEEP	eCDFS
10	855	553	Euclidean	DEEP	eCDFS
11	479	281	Manhattan	WIDE	eCDFS
12	479	281	Euclidean	WIDE	eCDFS

Note. Includes the experiment number; training and test set sizes; the distance metric used; the data set used; and where the training sample came from.

Table 2
Summary of the Results from All Tests

Experiment Number	Test Size	Best k	Standard Deviation	NMAD	Outlier Rate
1	422	3	0.1	0.06	7.35%
2	422	14	0.12	0.05	7.35%
3	228	3	0.12	0.06	11.40%
4	228	2	0.12	0.05	10.96%
5	855	13	0.13	0.05	11.11%
6	855	9	0.13	0.05	11.46%
7	479	3	0.12	0.06	13.15%
8	479	4	0.12	0.06	12.11%
9	553	4	0.12	0.06	10.31%
10	553	3	0.12	0.06	9.04%
11	281	2	0.13	0.06	13.88%
12	281	2	0.12	0.06	12.10%

Note. Includes the experiment number, the test size, standard deviation calculated typically, and by normalized absolute deviation and outlier rate.

where $\Delta z = z_{\text{spec}} - z_{\text{photo}}$ and the normalized median absolute deviation σ_{NMAD} as

$$\sigma_{\text{NMAD}} = 1.4826 \times \text{median}(|X_i - \text{median}(X)|) \quad (4)$$

where X is a vector of residuals

In Figures 1, 2, and 3, the top panels show the distribution of z_{spec} vs z_{photo} , and the lower panels show the normalised residuals vs the z_{spec} , where the normalised residuals are calculated by

$$\frac{\Delta z}{z_{\text{spec}} + 1}$$

For every source plotted on the top panels, we have provided error bars representing the 95% confidence interval of each redshift, calculated using the bootstrap method. On each of the plots, we display multiple statistics:

1. N —the number of sources in the test sample;

2. σ —standard deviation of the residual error, calculated typically;
3. NMAD—standard deviation of the residual error, calculated using the normalized absolute deviation (Equation (4));
4. η —outlier rate, calculated using Equation (3).

These results have all been summarized in Table 2.

4.1. Distance Metric

Across all of our experiments, we have found negligible difference between using the Manhattan Taxicab distance metric and the Euclidean distance metric. This indicates that there were few significant outliers when finding the k nearest neighbours, as the Manhattan Taxicab distance will minimize the effect of outliers. In Tests 3 versus 4, 7 versus 8, and 11 versus 12, we find that Euclidean distance provides a slightly lower outlier rate. In Test 5 versus 6 and 9 versus 10, Manhattan Taxicab distance provides the better option with Test 1 versus 2 being equal. In no case, however, does one distance metric have a difference in outlier rate greater than 1.78%.

4.2. Depth of Field

In all cases, we find that the outlier rate is higher in the WIDE data set when compared with the DEEP data set, as shown by the right-hand pair of panels in Figures 1 to 3. For the kNN regression algorithm, this is expected in current catalogues. In the process of modifying the DEEP data set to become the WIDE data set, all sources that we removed are at the fainter end of the data set, which are typically the high-redshift sources. This leaves the WIDE data set with a more heavily positively skewed distribution of z_{spec} , with the majority of sources being below $z = 0.5$. For sources at high redshifts, the kNN algorithm tends to fail because of the




Figure 1. Summary of the results from Tests 1–4, from left to right. All results displayed here have been trained on sources randomly sampled from the entire ATLAS data set. The top panels show the distribution of z_{spec} vs z_{photo} , and the lower panels show the normalised residuals vs the z_{spec} . The dashed red line represents $z_{\text{spec}} = z_{\text{photo}}$, and the dashed blue lines represent the outlier boundary, calculated using Equation (3).

(A color version of this figure is available in the online journal.)

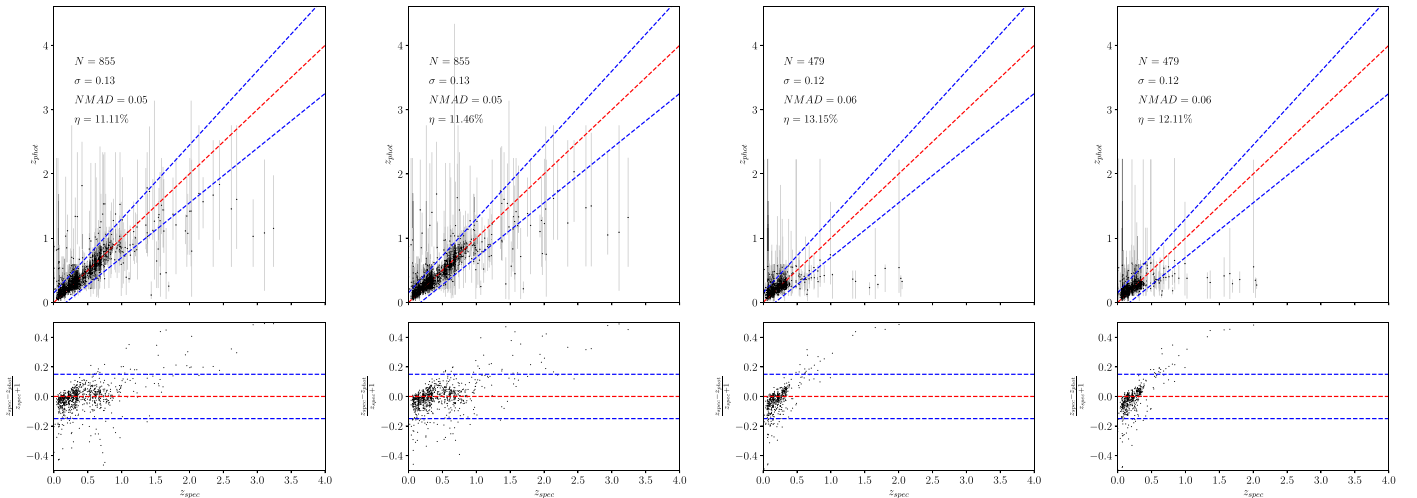


Figure 2. Summary of the results from Tests 5–8, from left to right. All results displayed here have been trained on sources found exclusively in the ELAIS-S1 field. Other details are as in Figure 1.

(A color version of this figure is available in the online journal.)

paucity of these sources, forcing low-redshift sources into the group of nearest neighbours.

A better test would use a larger data set, with a larger population of high-redshift sources, but such a data set is not yet available.

4.3. Generalization

While the best-case DEEP experiments using a random sample as the training set attain the best results, experiments that train on one ATLAS field, and test on a different ATLAS

field, are not much worse. We can attribute the better results in the former case to a more uniform and statistically consistent training sample. This indicates that, with a more consistent training sample, the kNN regression algorithm should perform well across different sections of sky.

5. Implications for Large-scale Radio Surveys

Table 2 shows that redshifts can be measured to an accuracy of about 6% (NMAD) to 12% (standard deviation), with an outlier rate of about 11%, and this result remains true for all




Figure 3. Summary of the results from Tests 9–12, from left to right. All results displayed here have been trained on sources found exclusively in the eCDFs field. Other details are as in Figure 1.

(A color version of this figure is available in the online journal.)

sources for which photometry is available, independently of the depth of the photometry.

If we assume that the DEEP sample has the same radio sensitivity as EMU, and that the WIDE sample has the same optical/infrared photometric depth as that available for EMU sources, then the relative sizes of the DEEP and WIDE samples implies that $\sim 45\%$ of EMU sources will have the required photometry for redshift measurement.

We can therefore conclude, based on these tests, that about 40% of EMU sources will have redshifts available, or a total of about 28 million radio sources.

We plan to extend this work by investigating the effect of (a) using the non-detection information, (b) using a more sophisticated metric that allows missing values and measurement bounds, (c) carefully modelling the sensitivity limits of the available photometric surveys, (d) incorporating other data types (such as radio and X-ray), (e) quantizing redshift to provide a classification problem rather than a regression problem, and (f) obtaining more high-redshift training data from deep surveys in small fields. Future work will continue in this direction.

6. Conclusion

The main result from these preliminary experiments is that, using the kNN algorithm, we can make good estimates of redshifts in large radio surveys given the photometry that is likely to be available, although further work is expected to improve results further. Around 90% of EMU sources with optical/infrared photometry will have a reliable estimated redshift, resulting in redshifts for $\sim 40\%$ of EMU sources. However, we expect that future work will result in an even higher fraction of sources with useful redshifts.

We found no obvious difference in the results provided by Manhattan Taxicab distance and Euclidean distance. However, greater benefits may be obtained from self-learned distance metrics that can weight features based on their relevance, and can deal with missing values.

We found that the DEEP data set reported better results than the WIDE data set, probably because the WIDE survey contains mainly low-redshift sources with the few high-redshift sources having to be estimated from the former. Further work will characterize what fraction of sources have the required photometry at higher redshifts.

We found that there was no obvious difference in success rate if the algorithm used training and test sets from spatially separate fields in the sky. Experiments 7–8 and 11–12 (different field training sets on the WIDE data set) suggest that we can overcome the lack of high-redshift sources in training sets by drawing training sets from small, deep fields and applying them to targets covering the entire sky.

Finally, this paper has demonstrated that, with sufficient redshift coverage in the training set, the kNN algorithm provides acceptable error rates when estimating the redshift of radio galaxies.

We thank Stefano Cavuoti and his colleagues for use of the code to plot our results.

ORCID iDs

Kieran J. Luken  <https://orcid.org/0000-0002-6147-693X>

Ray P. Norris  <https://orcid.org/0000-0002-4597-1906>

Laurence A. F. Park  <https://orcid.org/0000-0003-0201-4409>

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, Software Available from <https://www.tensorflow.org/>
- Abbott, T. M. C., Abdalla, F. B., Allam, S., et al. 2018, arXiv:1801.03181
- Abolfathi, B., Aguado, D. S., Aguilar, G., et al. 2018, *ApJS*, **235**, 42
- Baum, W. A. 1962, in IAU Symp. 15, Problems of Extra-Galactic Research, ed. G. C. McVittie (New York: Macmillan), 390
- Brescia, M., Cavuoti, S., D'Abrusco, R., Longo, G., & Mercurio, A. 2013, *ApJ*, **772**, 140
- Butchins, S. A. 1981, *A&A*, **97**, 407
- Carliles, S., Budavári, T., Heinis, S., Priebe, C., & Szalay, A. S. 2010, *ApJ*, **712**, 511
- Carrasco, D., Barrientos, L. F., Pichara, K., et al. 2015, *A&A*, **584**, A44
- Cavuoti, S., Amaro, V., Brescia, M., et al. 2017, *MNRAS*, **465**, 1959
- Cavuoti, S., Brescia, M., De Stefano, V., & Longo, G. 2015, *ExA*, **39**, 45
- Cavuoti, S., Brescia, M., & Longo, G. 2014, in IAU Symp. 306, Statistical Challenges in 21st Century Cosmology, ed. A. Heavens, J.-L. Starck, & A. Krone-Martins (Cambridge: Cambridge Univ. Press), 307
- Cavuoti, S., Brescia, M., Longo, G., & Mercurio, A. 2012, *A&A*, **546**, A13
- Childress, M. J., Lidman, C., Davis, T. M., et al. 2017, *MNRAS*, **472**, 273
- Collister, A. A., & Lahav, O. 2004, *PASP*, **116**, 345
- Cover, T., & Hart, P. 1967, *IEEE Trans. Inf. Theory*, **13**, 21
- Cutri, R. M., et al. 2013, *yCat*, **2328**, 1
- da Cunha, E., Hopkins, A. M., Colless, M., et al. 2017, *PASA*, **34**, e047
- Driver, S. P., Davies, L. J., Meyer, M., et al. 2016, in The Universe of Digital Sky Surveys, Astrophysics and Space Science Proc., Vol. 42 (Cham: Springer), 205
- Duncan, K. J., Brown, M. J. I., Williams, W. L., et al. 2018, *MNRAS*, **473**, 2655
- Firth, A. E., Lahav, O., & Somerville, R. S. 2003, *MNRAS*, **339**, 1195
- Franzen, T. M. O., Banfield, J. K., Hales, C. A., et al. 2015, *MNRAS*, **453**, 4020
- Kügler, S. D., Polsterer, K., & Hoecker, M. 2015, *A&A*, **576**, A132
- Loh, E. D., & Spillar, E. J. 1986, *ApJ*, **303**, 154
- Lonsdale, C. J., Smith, H. E., Rowan-Robinson, M., et al. 2003, *PASP*, **115**, 897
- Mountrichas, G., Corral, A., Masoura, V. A., et al. 2017, *A&A*, **608**, A39
- Norris, R. P. 2017, *NatAs*, **1**, 671
- Norris, R. P., Afonso, J., Appleton, P. N., et al. 2006, *AJ*, **132**, 2409
- Norris, R. P., Hopkins, A. M., Afonso, J., et al. 2011, *PASA*, **28**, 215
- Norris, R. P., Salvato, M., Longo, G., et al. 2019, *PASP*, **131**, 108004
- Pasquet-Itam, J., & Pasquet, J. 2018, *A&A*, **611**, A97
- Polsterer, K. L., Zinn, P.-C., & Gieseke, F. 2013, *MNRAS*, **428**, 226
- Sadeh, I., Abdalla, F. B., & Lahav, O. 2016, *PASP*, **128**, 104502
- Salvato, M., Ilbert, O., Hasinger, G., et al. 2011, *ApJ*, **742**, 61
- Salvato, M., Ilbert, O., & Hoyle, B. 2019, *NatAs*, **3**, 212
- Tagliaferri, R., Longo, G., Andreon, S., et al. 2003, *LNCS*, **2859**, 226
- Wolf, C., Onken, C. A., Luvaul, L. C., et al. 2018, *PASA*, **35**, e010
- Yuan, F., Lidman, C., Davis, T. M., et al. 2015, *MNRAS*, **452**, 3047
- Zitlau, R., Hoyle, B., Paech, K., et al. 2016, *MNRAS*, **460**, 3152