A unified framework for constructing, tuning and assessing photometric redshift density estimates in a selection bias setting

P. E. Freeman,¹★ R. Izbicki² and A. B. Lee¹

¹*Department of Statistics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*

²*Department of Statistics, Federal University of São Carlos, 13565-905 São Carlos, Brazil*

Accepted 2017 March 27. Received 2017 March 27; in original form 2016 August 10

ABSTRACT

Photometric redshift estimation is an indispensable tool of precision cosmology. One problem that plagues the use of this tool in the era of large-scale sky surveys is that the bright galaxies that are selected for spectroscopic observation do not have properties that match those of (far more numerous) dimmer galaxies; thus, ill-designed empirical methods that produce accurate and precise redshift estimates for the former generally will not produce good estimates for the latter. In this paper, we provide a principled framework for generating conditional density estimates (i.e. photometric redshift PDFs) that takes into account selection bias and the covariate shift that this bias induces. We base our approach on the assumption that the probability that astronomers label a galaxy (i.e. determine its spectroscopic redshift) depends only on its measured (photometric and perhaps other) properties $\mathbf{x}$ and not on its true redshift. With this assumption, we can explicitly write down risk functions that allow us to both tune and compare methods for estimating importance weights (i.e. the ratio of densities of unlabelled and labelled galaxies for different values of $\mathbf{x}$) and conditional densities. We also provide a method for combining multiple conditional density estimates for the same galaxy into a single estimate with better properties. We apply our risk functions to an analysis of $\approx 10^6$ galaxies, mostly observed by Sloan Digital Sky Survey, and demonstrate through multiple diagnostic tests that our method achieves good conditional density estimates for the unlabelled galaxies.

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

1 INTRODUCTION

Photometric redshift (or photo- z) estimation is an indispensable tool of precision cosmology. The planners of current and future large-scale photometric surveys such as the Dark Energy Survey (Flaugher 2005) and the Large Synoptic Survey Telescope (Ivezić et al. 2008), which combined will observe over one billion galaxies, require accurate and precise redshift estimates in order to fully leverage the constraining power of cosmological probes such as baryon acoustic oscillations and weak gravitational lensing. Numerous estimators currently exist that achieve ‘good’ point estimates of photo- z at low redshifts ($z \lesssim 0.5$), where ‘good’ means that photo- z and spectroscopic (or spec- z) estimates for the *same galaxy* largely match, with only a small percentage of catastrophic outliers. These estimators are conventionally divided into two classes: template fitters, oft-used examples of which include BPZ (Benítez 2000) and EAZY (Brammer, von Dokkum & Coppi 2008), and empirical

methods such as ANNZ (Collister & Lahav 2004).¹ The former utilize sets of galaxy SED templates that are redshifted until a best match with a galaxy’s observed photometry is found, whereas the latter utilize spectroscopically observed galaxies to train machine-learning methods to predict the redshifts of those galaxies that are only observed photometrically.

Less well established within the field of photo- z estimation, however, are methods that (1) produce conditional density estimates (or error estimates) of individual galaxy redshifts *and at the same time* (2) properly take into account the discrepancy between the populations of spectroscopically observed galaxies (roughly closer and brighter) and those observed via photometry only (farther and fainter).

Regarding point (1): the error distributions of photo- z estimates are often asymmetric and/or multimodal, so that single-number summary statistics such as the mean or median are insufficient to

¹ See e.g. Hildebrandt et al. (2010), Dahlen et al. (2013) and Sánchez et al. (2014), who compare and contrast numerous estimators from both classes, and references therein.

★ E-mail: pfreeman@cmu.edu

Figure 1. Estimated distributions of *r*-band model magnitudes and *ugriz* colours for labelled (i.e. spectroscopic; red dashed lines) and unlabelled (i.e. photometric; blue solid lines) galaxies, primarily observed by the SDSS (see Section 7). These distributions indicate the bias for selecting brighter galaxies for spectroscopic follow-up observations, which induces a covariate shift that is here manifested as mismatches in the distributions of colours between labelled and unlabelled data.

describe their shapes. Furthermore, the use of such statistics leads to biased estimation of parameters in downstream cosmological analyses (e.g. Wittman 2009); for instance, Mandelbaum et al. (2008) demonstrate that the use of the conditional density estimate $\hat{f}(z|\mathbf{x})$ [often denoted $p(z)$ in the astronomical literature and often called the probability density estimate, or PDF] reduces systematic error in galaxy–galaxy weak lensing analyses. (Here, $\mathbf{x}$ can represent magnitudes and/or colours and/or other ancillary information measured for a galaxy.) Several other works have touted the use of $\hat{f}(z|\mathbf{x})$ as well, often as a step towards better estimates of ensemble redshift distributions [usually denoted $N(z)$] in tomographic studies (e.g. Cunha et al. 2009; Ball & Brunner 2010; Sheldon et al. 2012; Carrasco Kind & Brunner 2013; Carrasco Kind & Brunner 2014; Bonnett 2015; Rau et al. 2015; De Vicente, Sánchez & Sevilla-Noarbe 2016), and standard methods such as the aforementioned BPZ, EAZY and ANNz provide $\hat{f}(z|\mathbf{x})$ as an available output.

Regarding point (2): it is a well-established truism that in large-scale surveys there is *selection bias*, wherein rare and bright galaxies are preferentially selected for spectroscopic observation. This bias induces a *covariate shift*, since the properties of these bright galaxies do not match those of more numerous dimmer galaxies (see e.g. Fig. 1). This shift affects the accuracy and precision of empirical photo-*z* estimates. One can mitigate covariate shift by estimating importance weights $\beta(\mathbf{x}) = f_U(\mathbf{x})/f_L(\mathbf{x})$, the ratio of the density of galaxies without redshift labels to those observed spectroscopically. For instance, Lima et al. (2008) attempt to directly estimate $N(z)$ in a covariate shift setting with a *k*-nearest-neighbour-based estimator of the importance weights, an estimator since utilized by Cunha et al. (2009), Sheldon et al. (2012) and Sánchez et al. (2014). Rau et al. (2015), who propose a weighted kernel density estimator for $f(z|\mathbf{x})$, offer two other methods for computing the weights (quantile regression forest and ordinal classification PDF). All these weight estimators feature parameters that one must tune for proper performance. One would generally tune estimators by minimizing an estimate of risk using a validation data set, but the authors listed above skirt the issue of tuning by setting the number of nearest neighbours a priori, or, in the case of Rau et al. (2015),

by utilizing a plug-in bandwidth estimate via Scott’s rule (see their equation 24).

In this paper, we describe a principled and unified framework for generating conditional density estimates $\hat{f}(z|\mathbf{x})$ in a selection bias setting: specifically, we provide a suite of appropriate risk estimators, methods for tuning and assessing models, and diagnostic tests that allows one to create accurate density estimates from raw data $\mathbf{x}$.² In Section 2, we define both the problem and our notation. In Sections 3 and 4, we show that if we assume that the probability that a galaxy is labelled depends only on its photometry and not on its true redshift, which is a valid assumption within the redshift regime probed by shallow surveys such as the Sloan Digital Sky Survey (SDSS; York et al. 2000), we can write down risk functions that allow one to properly tune estimators of both $\beta(\mathbf{x})$ and $f(z|\mathbf{x})$. These risk functions also allow us to choose from among competing estimators. In Section 5, we show how one can combine estimators of conditional density to improve upon the results achieved by any one estimator alone. In Section 6, we provide diagnostic tests that one may use to determine the absolute performance of conditional density estimators. In Section 7, we demonstrate our methods by applying them to SDSS data. Finally, in Section 8 we summarize our results. In future works, we will provide methods for variable selection (i.e. the selection of the most informative colours, etc., to retain from a large set of possible covariates) and explore methods in which we relax the galaxy-labelling assumption stated above.

2 PROBLEM STATEMENT: SELECTION BIAS

The data in a conventional photometric redshift estimation problem consist of covariates $\mathbf{x} \in \mathbb{R}^d$ (photometric colours and/or magnitudes, etc.) and redshifts z . We have access to two data samples: an independent and identically distributed (iid) sample $\mathbf{x}_1^U, \dots, \mathbf{x}_{n_U}^U$ consisting of photometric data without associated labels (i.e. redshifts) and an iid labelled sample $(\mathbf{x}_1^L, z_1^L), \dots, (\mathbf{x}_{n_L}^L, z_{n_L}^L)$ constructed from follow-up spectroscopic studies. (For computational efficiency, in our analyses these data sets are samples taken from larger pools of available labelled and unlabelled data.) Our goal is to construct a photo-*z* conditional density estimator, $\hat{f}(z|\mathbf{x})$, that performs well when applied to the unlabelled data (where ‘well’ can be defined by its performance with respect to a number of metrics; see e.g. Section 6 for three examples).

An issue that arises when constructing $\hat{f}(z|\mathbf{x})$ via empirical techniques is that of selection bias. A standard assumption in statistics and machine learning is that labelled and unlabelled data are sampled from similar distributions, which we denote $\mathbb{P}_L$ and $\mathbb{P}_U$, respectively. However, as Fig. 1 demonstrates, these two distributions can differ greatly for sky surveys that mix spectroscopy and photometry; brighter galaxies are more likely to be selected for follow-up spectroscopic observation. To model how selection bias affects learning methods, one needs to invoke additional assumptions about the relationship between $\mathbb{P}_L$ and $\mathbb{P}_U$ (e.g. Gretton et al. 2008; Moreno-Torres et al. 2012). In this work, we assume that the probability that a galaxy is labelled with a spectroscopic redshift depends only upon $\mathbf{x}$ (in accord with Lima et al. 2008; Sheldon et al. 2012); i.e.

$$P(S = 1|\mathbf{x}, z) = P(S = 1|\mathbf{x}), \quad (1)$$

² One may find R functions implementing our framework at github.com/pfreeman/CDESB.

where the random variable S equals 1 if a datum is labelled and 0 otherwise. This assumption implies covariate shift, defined as

$$f_L(\mathbf{x}) \neq f_U(\mathbf{x}), \quad f_L(z|\mathbf{x}) = f_U(z|\mathbf{x}), \quad (2)$$

and thus is, as shown below, critical for establishing the risk function estimators that ultimately allow us to estimate conditional density estimates $\hat{f}(z|\mathbf{x})$. Following the discussion of section 2.3 of Budavári (2009), we point out that assuming $P(S=1|\mathbf{x}, z) = P(S=1|\mathbf{x})$ can be problematic, for instance when only colours are used in analyses in which the training data are selected in limited magnitude regimes. In this work, we apply our framework to galaxies at SDSS depth using colours only; for optimal performance, one should incorporate those covariates that act in concert with z to affect selection, e.g. morphology, size, surface brightness, environment, etc. Nothing in the current framework prevents the incorporation of these covariates.

At first glance, it may seem that covariate shift would not pose a problem for density estimation; if $f(z|\mathbf{x})$ is the same for both labelled and unlabelled samples, one might infer that a good estimator of $f(z|\mathbf{x})$ based on labelled data would also perform well for unlabelled data. However, this is generally untrue. The estimation of $f(z|\mathbf{x})$ depends on the marginal distribution $f(\mathbf{x})$, so an estimator that performs well with respect to $f_L(\mathbf{x})$ may not perform well with respect to $f_U(\mathbf{x})$.

We can mitigate selection bias by pre-processing the labelled data so as to ensure that sufficient labelled data lie where the unlabelled data lie. This allows us to compute expected values with respect to the distribution $\mathbb{P}_U$ using $\mathbb{P}_L$, similar to the idea of importance sampling in Monte Carlo methods. Carrying this mitigation out in practice involves two steps: first, we estimate importance weights as a function of the predictors $\mathbf{x}$:

$$\beta(\mathbf{x}) = f_U(\mathbf{x})/f_L(\mathbf{x}); \quad (3)$$

and secondly, we utilize these weights when estimating conditional densities $f(z|\mathbf{x})$. There are a myriad of estimators both for importance weighting and for conditional density estimation (see e.g. Izbicki, Lee & Schafer 2014; Izbicki & Lee 2016); what we provide here are rigorous procedures for tuning their parameters and for choosing among them. We note here that our overall procedure can be qualitatively summarized by the following dictum: one needs good estimates of importance weights at labelled data points in order to achieve good conditional density estimates at unlabelled data points. (One can observe how this dictum plays out mathematically in equation 11 below: note how the importance weight estimates at labelled points enter into it, in the second term, whereas the importance weight estimates at unlabelled points do not enter into it at all.)

3 IMPORTANCE WEIGHT ESTIMATION

A naive method for computing importance weights $\beta(\mathbf{x})$ would involve estimating f_U and f_L separately and computing the ratio of these densities, but this approach can enhance errors in individual density estimates, particularly in regimes where $f_L \rightarrow 0$ (Sugiyama et al. 2008). Many authors have thus proposed direct estimators of the ratio $\beta(\mathbf{x})$ (e.g. Lima et al. 2008; Sugiyama et al. 2008; Kanamori, Hido & Sugiyama 2009; Kanamori, Suzuki & Sugiyama 2012; Loog 2012; Izbicki et al. 2014; Kremer et al. 2015).

As an example, the estimator of Lima et al. (2008) and Kremer et al. (2015) is

$$\hat{\beta}(\mathbf{x}) = \frac{1}{k} \frac{n_L}{n_U} \sum_{i=1}^{n_U} \mathbb{I}(\mathbf{x}_i^U \in V_k^L), \quad (4)$$

where V_k^L is the region of covariate space containing points that are closer to $\mathbf{x}$ than its k th nearest labelled neighbour, and $\mathbb{I}(\cdot)$ is the indicator function.

To choose between importance weight estimators, one needs first to optimally tune the parameters of each using the training and validation data (model selection), and then assess their performance using the test data (model assessment). We determine optimal values of tuning parameters by minimizing a risk function (equation 5 below). Generating estimates $\hat{\beta}(\mathbf{x})$ [and by extension $\hat{f}(z|\mathbf{x})$, below] implicitly requires smoothing the observed data, with the smoothing bandwidth set such that estimator bias and variance are optimally balanced. (See e.g. Wasserman 2006.) For instance, too much smoothing (e.g. adopting a value of k that is too large in nearest-neighbour-based methods) yields estimates with low variance and high bias, i.e. when applied to independent data sets sampled from the same distribution, the estimates will all look similar (i.e. will have low variance) but will be offset from the truth (i.e. will have high bias). Too little smoothing (e.g. k too small) conversely yields high-variance and low-bias estimates that overfit the training data.

When estimating importance weights, we apply the risk function (Izbicki et al. 2014; Kremer et al. 2015)

$$\begin{aligned} R(\hat{\beta}, \beta) &:= \int (\hat{\beta}(\mathbf{x}) - \beta(\mathbf{x}))^2 dP_L(\mathbf{x}) \\ &= \int \hat{\beta}^2(\mathbf{x}) dP_L(\mathbf{x}) - 2 \int \hat{\beta}(\mathbf{x}) \beta(\mathbf{x}) dP_L(\mathbf{x}) \\ &\quad + \int \beta^2(\mathbf{x}) dP_L(\mathbf{x}) \\ &= \int \hat{\beta}^2(\mathbf{x}) dP_L(\mathbf{x}) - 2 \int \hat{\beta}(\mathbf{x}) dP_U(\mathbf{x}) + C(\beta), \end{aligned} \quad (5)$$

where $dP_U(\mathbf{x}) = \beta(\mathbf{x}) dP_L(\mathbf{x})$ and $C(\beta)$ is a term that does not depend on the estimate $\hat{\beta}$. (We note that here the calculation of risk is with respect to the labelled data set distribution $\mathbb{P}_L$; this is in accord with the dictum stated above that we need good estimates of importance weights for the labelled data in order to achieve good estimates of conditional densities for the unlabelled data. See equation 11 below.) While we utilize an L^2 -loss function above in equation (5) and below in equation (9), one could in principle substitute other functions based on F-divergences, log-densities or notions of L^1 loss. However, functions based on F-divergences and log-densities are overly sensitive to distribution tails and are generally not appropriate for density estimation (see e.g. Hall 1987; Wasserman 2006), while estimating a risk based on L^1 loss requires knowledge of the true $f(z|\mathbf{x})$. Since in model selection and assessment we can ignore $C(\beta)$, we rewrite the above as

$$J(\hat{\beta}) = \int \hat{\beta}^2(\mathbf{x}) dP_L(\mathbf{x}) - 2 \int \hat{\beta}(\mathbf{x}) dP_U(\mathbf{x}), \quad (6)$$

which we estimate as

$$\hat{J}(\hat{\beta}) = \frac{1}{\tilde{n}_L} \sum_{k=1}^{\tilde{n}_L} \hat{\beta}^2(\tilde{\mathbf{x}}_k^L) - \frac{2}{\tilde{n}_U} \sum_{k=1}^{\tilde{n}_U} \hat{\beta}(\tilde{\mathbf{x}}_k^U), \quad (7)$$

where the tildes indicate that the risk is evaluated using either validation data (during model selection) or test data (during model assessment). [Here, we use J to represent a risk function in which the constant term $C(\beta)$ is ignored.] Among multiple estimators of

$\hat{\beta}(\mathbf{x})$, we choose the one that yields the smallest value of $\hat{J}$ when applied to test data.

4 CONDITIONAL DENSITY ESTIMATION

Given an estimate $\hat{\beta}(\mathbf{x})$ of the importance weight, the ratio of densities of the unlabelled and labelled data at the point $\mathbf{x}$, our next step is to compute estimates of the conditional density $\hat{f}(z|\mathbf{x})$. Conditional density estimators include those of Cunha et al. (2009) and Izibicki & Lee (2016); see e.g. Izibicki et al. (2016) for more details. To build intuition here, we write down the estimator of Cunha et al. (2009), as it is particularly simple:

$$\hat{f}(z|\mathbf{x}) \propto \sum_{i \in \mathcal{N}_N(\mathbf{x})} \hat{\beta}(\mathbf{x}_i^L) \mathbb{I}(z_i^L \in b(z)), \quad (8)$$

where $\mathcal{N}_N(\mathbf{x})$ denotes the N neighbours nearest to $\mathbf{x}$ among labelled data, and $b(z)$ denotes the a priori defined bin to which z belongs. This estimator (up)down-weights labelled data in regions where $f_U(\mathbf{x})$ is (larger) smaller than $f_L(\mathbf{x})$.

In a selection bias setting where $\mathbb{P}_L \neq \mathbb{P}_U$, the goal of conditional density estimation is to minimize

$$R(\hat{f}, f) := \iint (\hat{f}(z|\mathbf{x}) - f(z|\mathbf{x}))^2 dP_U(\mathbf{x}) dz, \quad (9)$$

i.e. the risk with respect to the unlabelled data. Under the covariate shift assumption $f_U(z|\mathbf{x}) = f_L(z|\mathbf{x})$, one can rewrite the modified risk (9) up to a constant as

$$\begin{aligned} R(\hat{f}, f) &= \iint \hat{f}^2(z|\mathbf{x}) dP_U(\mathbf{x}) dz \\ &\quad - 2 \iint \hat{f}(z|\mathbf{x}) f(z|\mathbf{x}) dP_U(\mathbf{x}) dz + \iint f^2(z|\mathbf{x}) dP_U(\mathbf{x}) dz \\ &= \iint \hat{f}^2(z|\mathbf{x}) dP_U(\mathbf{x}) dz \\ &\quad - 2 \iint \hat{f}(z|\mathbf{x}) \beta(\mathbf{x}) dP_L(z, \mathbf{x}) + C(f), \end{aligned} \quad (10)$$

where the second equality follows from $f_U(z|\mathbf{x}) dP_U(\mathbf{x}) dz = f_L(z|\mathbf{x}) \beta(\mathbf{x}) dP_L(z, \mathbf{x}) = \beta(\mathbf{x}) dP_L(z, \mathbf{x})$. Again, this risk depends upon unknown quantities; we ignore $C(f)$ and estimate the other terms via the equation

$$\begin{aligned} \hat{J}(\hat{f}) &= \frac{1}{\tilde{n}_U} \sum_{k=1}^{\tilde{n}_U} \left[\int \hat{f}^2(z|\tilde{\mathbf{x}}_k^U) dz \right] \\ &\quad - \frac{2}{\tilde{n}_L} \sum_{k=1}^{\tilde{n}_L} \hat{f}(\tilde{z}_k^L|\tilde{\mathbf{x}}_k^L) \hat{\beta}(\tilde{\mathbf{x}}_k^L), \end{aligned} \quad (11)$$

where again the tildes indicate use of validation data in model selection and test data in model assessment.

5 COMBINING ESTIMATORS

Typical photometric redshift estimation methods utilize one method for computing $\hat{z}|\mathbf{x}$ or $\hat{f}(z|\mathbf{x})$. However, one can improve upon the prediction performances of individual estimators by combining them. Suppose that $\hat{f}_1(z|\mathbf{x}), \dots, \hat{f}_p(z|\mathbf{x})$ are p separate estimators of $f(z|\mathbf{x})$. We define the weighted average to be

$$\hat{f}^\alpha(z|\mathbf{x}) = \sum_{k=1}^p \alpha_k \hat{f}_k(z|\mathbf{x}), \quad (12)$$

where the weights minimize the empirical risk $\hat{R}(\hat{f}^\alpha)$ under the constraints $\alpha_i \geq 0$ and $\sum_{i=1}^p \alpha_i = 1$. One can determine the solution $\alpha = [\alpha_i]_{i=1}^p$ by solving a standard quadratic programming problem:

$$\arg \min_{\alpha: \alpha_i \geq 0, \sum_{i=1}^p \alpha_i = 1} \alpha' \mathbb{B} \alpha - 2\alpha' b, \quad (13)$$

where $\mathbb{B}$ is the $p \times p$ matrix

$$\left[\frac{1}{\tilde{n}_U} \sum_{k=1}^{\tilde{n}_U} \int \hat{f}_i(z|\tilde{\mathbf{x}}_k^U) \hat{f}_j(z|\tilde{\mathbf{x}}_k^U) dz \right]_{i,j=1}^p \quad (14)$$

and b is the vector

$$\left[\frac{1}{\tilde{n}_L} \sum_{k=1}^{\tilde{n}_L} \hat{f}_i(\tilde{z}_k^L|\tilde{\mathbf{x}}_k^L) \hat{\beta}(\tilde{\mathbf{x}}_k^L) \right]_{i=1}^p; \quad (15)$$

the tildes here indicate use of the validation data.

6 DIAGNOSTIC TESTS FOR ESTIMATORS

Risk estimates, such as those given in equations (7) and (11), allow us to tune estimators and to choose between estimators, but they do not ultimately convey how well the estimator performs in an absolute sense. Below, we describe diagnostic tests that one can use to more closely assess the quality of different models. Similar tests can be found in the time series literature (see e.g. Corradi & Swanson 2006).

6.1 Assessing uniformity using empirical CDFs

Let $\hat{F}(z|\mathbf{x}_i)$ denote the estimated conditional cumulative distribution function (CDF), i.e. let

$$U_i = \hat{F}(z_i|\mathbf{x}_i) = \int_0^{z_i} \hat{f}(y|\mathbf{x}_i) dy. \quad (16)$$

If the chosen estimator performs well, then the empirical CDF of the values $U_1, \dots, U_n$ will be consistent with the CDF of the uniform distribution. We can test this hypothesis via, e.g. the Cramér-von Mises (CvM), Anderson–Darling (AD) and Kolmogorov–Smirnov (KS) tests. If the p -value output by the test is > 0.05 , then we fail to reject the null hypothesis that the data $U_1, \dots, U_{n_{\hat{\rho}_0}}$ are sampled from a uniform distribution.

6.2 Assessing uniformity using quantiles

We can use the values $U_1, \dots, U_n$ defined above to build a quantile–quantile (or QQ) plot, by determining the number of data in bins of width ΔU . Let J be the number of bins, each of which has mid-point c_j and the fraction of data $\hat{c}_j$. The QQ plot is that of the values $\hat{c}_j$ against c_j ; if the chosen estimators perform well, the points in this plot will approximately lie on the line $\hat{c} = c$. We assess consistency with uniformity via the chi-square goodness of fit (GoF) test, which utilizes the chi-square statistic

$$\chi_{\text{obs}}^2 = \sum_{j=1}^J \frac{(n\hat{c}_j - n/J)^2}{n/J}. \quad (17)$$

We conclude that the difference data are consistent with constancy if the p -value, the fraction of the time that a value of χ^2 would be observed that is greater than χ_{obs}^2 if the null hypothesis of constancy is true, is > 0.05 . Note that the off-the-shelf GoF test, which allows one to compute the p -value by taking the tail integral of an appropriate chi-square distribution, requires that the number of expected

Algorithm 1 Conditional Density Estimation for Photometric Redshifts

Input: Pools of labeled (i.e. spectroscopic) and unlabeled (i.e. photometric-only) data, which we denote L and U .

- Step 1. Randomly sample n_L data from L and n_U data from U . (Here, $n_L = n_U = 15000$.)
- Step 2. Given number of nearest labeled neighbors k and minimum number of unlabeled data in neighborhood u , repopulate the labeled sample so as to increase the effective sample size, i.e., to decrease the number of labeled data with estimated importance weight zero. Such repopulation will result in photometric covariates distributions that more closely approximates that of the unlabeled sample. (See equation 4 and Algorithm 2.)
- Step 3. Split L and U into training, validation, and test datasets. (Here, $n_{\text{train}} = 7000$, $n_{\text{val}} = 3000$, and $n_{\text{test}} = 5000$ for both samples.)
- Step 4. Apply training and validation data to generate importance weight estimates $\hat{\beta}(\mathbf{x})$ (via equation 4), tuning the estimator so as to minimize the estimated risk given in equation 7.
- Step 5. Apply importance weights $\hat{\beta}(\mathbf{x})$ to the estimation of conditional densities $\hat{f}(z|\mathbf{x})$ (e.g. equation 8), tuning the estimator so as to minimize the risk function given in equation 11. (Here we apply the conditional density estimators dubbed MN-CS, `kerNN`-CS, `Series`-CS, and `Combined` by Izbicki, Lee & Freeman 2016, where CS stands for “covariate shift” and `Combined` stands for the combination of the `kerNN`-CS and `Series`-CS estimators. See Section 5 for details on combining estimator output.)
- Step 6. Use the test data to compute the estimated risk for each conditional density estimator (equation 11). Choose the estimator with the minimum risk.

Output: Estimated conditional densities $\{\hat{f}(z|\mathbf{x}_1^L), \dots, \hat{f}(z|\mathbf{x}_{n_L}^L)\}$ and $\{\hat{f}(z|\mathbf{x}_1^U), \dots, \hat{f}(z|\mathbf{x}_{n_U}^U)\}$.

Algorithm 2 Preprocessing Labeled (i.e. Spectroscopic) Data

Input: Number of labeled and unlabeled data, n_L and n_U .

Labeled dataset $L = \{\mathbf{x}_1^L, \dots, \mathbf{x}_{n_L}^L\}$ and unlabeled dataset $U = \{\mathbf{x}_1^U, \dots, \mathbf{x}_{n_U}^U\}$.

Number of nearest labeled neighbors, k .

Minimum number of unlabeled data required for selection of new labeled datum, u .

- Step 1. Randomly permute all labeled data indices (i.e. those from the larger pool of all available labeled data).
- Step 2. Loop over the permuted set of indices i_p :
 - estimate $\hat{\beta}(\mathbf{x}_{i_p})$ within the volume containing the k nearest neighbors among L (via e.g. equation 4)
 - if $\hat{\beta}(\mathbf{x}_{i_p}) \geq u/k$, accept $\mathbf{x}_{i_p}$ as a new labeled datum
 - if the number of new labeled data equals n_L , terminate loop.

Output: New labeled dataset $\{\mathbf{x}_1^L, \dots, \mathbf{x}_{n_L}^L\}$.

The importance weights for these data are re-estimated and are applied to conditional density estimation (e.g. equation 8, and equation 11).

counts in each bin be $\gtrsim 5$. When that condition is violated, one should use simulations to estimate p -values.

6.3 Assessing uniformity in interval coverage

For every α_j in a grid of values on $[0, 1]$ of length J , and for every observation i in the labelled test sample, we determine the smallest interval $A_{ij} = [z_i^{\text{lo}}, z_i^{\text{hi}}]$ such that

$$\int_{A_{ij}} \hat{f}(z|\mathbf{x}_i) dz = \alpha_j. \quad (18)$$

Then, for every α_j , we determine the proportion of redshifts lie within A_{ij} , i.e. we compute

$$\hat{\alpha}_j = \sum_{i \in S} \mathbb{I}(z_i^L \in A_{ij}). \quad (19)$$

If the chosen estimators perform well, then $\hat{\alpha} \approx \alpha$. We plot the values of $\hat{\alpha}$ against corresponding values of α and assess how close the plotted points are to the line $\hat{\alpha} = \alpha$; we can test for consistency with that line using the chi-square GoF test, as described above.

Note that the construction of coverage plots is also proposed by Wittman, Bhaskar & Tobin (2016), who conclude that the $\hat{f}(z|\mathbf{x})$ produced by the template-based BPZ and EAZY codes are consistently too narrow and approximately correct, respectively.

7 APPLICATION TO SDSS DATA**7.1 Data**

To demonstrate the efficacy of our conditional density estimation method, we apply it to $\approx 10^6$ galaxies that are mostly from Data Release 8 of the SDSS (Aihara et al. 2011).

To build our unlabelled (i.e. photometric) data set, we initially extract model magnitudes for 540 235 objects in the sky patch $\text{RA} \in [168^\circ, 192^\circ]$ and $\delta \in [-1^\circ.5, 1^\circ.5]$. After filtering these data in the manner of Sheldon et al. (2012), namely limiting our selection to those data whose *ugriz* magnitudes were *all* between 15 and 29, and then further limiting ourselves to data for which

$$u < 21 \text{ or } g < 22 \text{ or } r < 22 \text{ or } i < 20.5 \text{ or } z < 20.1,$$

we obtain a sample of 538 974 objects.

We use the labelled (i.e. spectroscopic) data set of Sheldon et al. (2012, Sheldon, private communication). This data set includes 435 875 objects from SDSS DR8 and 31 835 objects from eight other sources, or 467 710 objects in all. As noted by Sheldon et al. (2012), this data set contains a small number of stars. We remove these by excluding all data with spectroscopic redshift $z_s = 0$; after this, we are left with 465 790 objects.

The steps of our analysis are given in Algorithms 1 and 2. As noted in Section 2, the labelled and unlabelled data that we analyse are samples from the larger pools of available data described immediately above. (In this work, we set $n_L = n_U = 15\,000$.) This

is for computational efficiency, both from a standpoint of time and memory; for instance, if we utilize matrices for storing distances between data points, we are currently limited to samples of size $\sim 10^4$ when utilizing typical desktop computers.

7.2 Data pre-processing: construction of the labelled sample

As shown in equation (8), the estimation of conditional densities is partially a function of $\hat{\beta}(\mathbf{x}^L)$, so there is a distinction to be drawn between the stated labelled sample size (e.g. $n_L = 15\,000$, drawn from a pool of size 465 790) and the effective size (the number of data that contribute to estimation, i.e. the number for which $\hat{\beta}(\mathbf{x}^L) > 0$). Thus, one important step of our method involves pre-processing the labelled data to increase their effective size.

The pre-processing of the labelled data requires the specification of a threshold importance weight $\hat{\beta}_{\text{thr}}$. Given n_L and n_U labelled and unlabelled data, respectively, and having specified a minimum number of unlabelled data u that would have to lie closer to a random labelled point $\mathbf{x}^L$ (drawn from the larger pool) than the k th nearest labelled neighbour to that point, we keep $\mathbf{x}^L$ as part of our *new* labelled data set if

$$\hat{\beta}(\mathbf{x}^L) \geq \hat{\beta}_{\text{thr}} = u/k. \quad (20)$$

The value $\hat{\beta}_{\text{thr}}$ is not tunable, per se, as different thresholds yield different labelled data sets, leading to estimated risks that are not directly comparable. One might conjecture that larger values of $\hat{\beta}_{\text{thr}}$ are better, in that the distribution of the labelled data will more closely resemble that of the unlabelled data (see Figs 1 and 2). However, as we demonstrate below, our results are not highly sensitive to the choice of $\hat{\beta}_{\text{thr}}$, so long as $\hat{\beta}_{\text{thr}} > 0$.

Once pre-processing is complete, we repeat the estimation of the importance weights $\hat{\beta}(\mathbf{x}^L)$ and apply these values when estimating conditional densities (e.g. equations 8 and 11).

7.3 Results

Once we generate our new labelled data set, we split the labelled and unlabelled data into training ($n_{\text{train}} = 7000$), validation ($n_{\text{valid}} = 3000$) and test ($n_{\text{test}} = 5000$) sets. We forego model assessment for importance weight estimators in this work; Izbicki et al. (2016) demonstrate that the estimator given in equation (4) consistently performs better than five other competing methods over a number of different levels of covariate shift. We apply equations (4) and (7) to the training and validation data to determine the optimal number of nearest neighbours k and importance weights $\hat{\beta}(\mathbf{x}^L)$ given k . We then apply these importance weights and the estimated risk in equation (11) to the training and validation data within the context of three conditional density estimators detailed in Izbicki et al. (2016): NN-CS, the estimator of Cunha et al. (2009); kerNN-CS, a kernelized variation of NN-CS and Series-CS, the spectral series estimator of Izbicki & Lee (2016).³ In Fig. 3, we demonstrate the importance of pre-processing to generate the labelled data set: without it, the construction of conditional density estimates for the unlabelled data

Figure 2. Same as Fig. 1, except that here we construct the labelled data with the pre-processing scheme outlined in Algorithm 2 so as to increase the effective sample size (i.e. the number of labelled data with non-zero importance weights). The panels exhibit pre-processing with $\hat{\beta}_{\text{thr}} = 0.1, 0.2$ and 0.3 , respectively, from top to bottom (see equation 20). (For larger values of $\hat{\beta}_{\text{thr}}$, we cannot fully populate a labelled subsample of size $n_L = 15\,000$ from the pool of labelled data at our disposal.) In all three cases, the distributions for the labelled data (red dashed lines) more closely resembles those for the unlabelled data (blue solid lines), relative to the labelled distributions exhibited in Fig. 1. Subsequent to pre-processing, we re-estimate the importance weights $\hat{\beta}(\mathbf{x}^L)$ and apply these values to the estimation of conditional densities (e.g. equations 8 and 11).

³ To expand upon the description of the bias-variance trade-off in Section 3, we note that any estimates $\hat{\beta}(\mathbf{x})$ and $\hat{f}(z|\mathbf{x})$ made near (potentially sharp) parameter-space boundaries will suffer from some amount of additional ‘boundary bias’. Mitigating boundary biases in photo- z estimation is an important topic that we will pursue in a future work.

Figure 3. Estimated probability density functions for $\hat{\beta}(\mathbf{x}^L)$ before (red dashed line) and after (blue solid line) application of the labelled data pre-processing scheme described in Section 7.2 and Algorithm 2. Without pre-processing, there are effectively no data with which to learn a statistical model linking the covariates and redshift in the regime $\hat{\beta}(\mathbf{x}^L) \gtrsim 0.5$, a regime that contains >90 per cent of the unlabelled data.

Figure 4. Estimated risk (and standard errors) resulting from the application of the NN-CS, kerNN-CS and Series-CS conditional density estimators, and a combination of the latter two (see Section 5). The four points in each group represent different values for the importance weight threshold $\hat{\beta}_{thr}$ in our scheme for pre-processing labelled data; from left to right, $\hat{\beta}_{thr} = 0.0, 0.1, 0.2$ and 0.3 (black, red, green and blue, respectively). The estimated risks for each $\hat{\beta}_{thr}$ were adjusted to have mean -2.2 . In all four cases, we observe that combining the kerNN-CS and Series-CS estimators leads to better conditional density estimates than either yield themselves.

would be effectively limited to the regime $\hat{\beta}(\mathbf{x}^U) \lesssim 0.5$. The fraction of unlabelled test-set data with $\hat{\beta}(\mathbf{x}^U) \leq 0.5$ is approximately 7.5 per cent; in contrast, the fraction for $\hat{\beta}(\mathbf{x}^U) \leq 3$ is 88.5 per cent. In Fig. 4, we demonstrate that for our particular SDSS data, the kerNN-CS estimator on average outperforms the other two. (Note that the risk can be negative because we ignore the positive addi-

Figure 5. 25 estimates of $\hat{f}(z|\mathbf{x})$ made for a single labelled test-set galaxy via bootstrapping of the labelled training data set. The solid red vertical line indicates the true redshift. As described in the text, we make a preliminary assessment of the noise properties of $\hat{f}(z|\mathbf{x})$ by computing the true redshift quantile for each curve (i.e. the integral of the curve between zero and the true redshift), computing the sample standard deviation of these quantiles, and determining the overall mean standard deviation. Here, the standard deviation of the quantiles is ≈ 0.13 ; on average the standard deviation is ≈ 0.06 .

tive constant $C(f)$; see equation 10.) We use the method outlined in Section 5 to optimally combine the conditional density estimates from kerNN-CS and Series-CS and we determine that combining these estimators, on average, indeed yields better estimates of the conditional densities $\hat{f}(z|\mathbf{x})$.

We make a preliminary assessment of the noise properties of the estimates $\hat{f}(z|\mathbf{x})$ as follows. We create n bootstrap samples of the labelled training data and use them to generate n $\hat{f}(z|\mathbf{x})$ curves for each labelled test datum (see Fig. 5). Then, we compute the quantile of the true redshift given each curve, so that for each labelled test datum we have n quantile values. We use the mean of the standard deviations of each set of quantile values as a metric of uncertainty. For our particular SDSS data, we determine this mean uncertainty to be ≈ 0.065 . We will examine the noise properties of the estimates $\hat{f}(z|\mathbf{x})$ at greater depth in a future work.

In Fig. 6, we demonstrate the trade-off that is inherent when mitigating covariate shift, via the incorporation of importance weights into estimation, using QQ plots. If we do not mitigate covariate shift, we achieve good conditional density estimates within those portions of covariate space in which $\hat{\beta}(\mathbf{x}^L) \lesssim 0.5$ (orange dash-dotted line); however, $\lesssim 8$ per cent of the unlabelled data reside within these portions. Mitigation of covariate shift leads to a worsening of the conditional density estimates within these portions of covariate space (black dashed line), but allows one to make good estimates throughout the remaining space (blue solid line). In Figs 7–9, we show the results of applying hypothesis tests based on QQ plots (Section 6.2), coverage plots (Section 6.3) and the assumption of uniformity (Section 6.1). In the first two figures, we show the results of using the GoF test to determine the consistency of expected and observed quantiles at each unique value of $\hat{\beta}(\mathbf{x}^L)$, in the manner outlined in Section 6. These results are generated using the labelled test-set data, assuming a pre-processing threshold $\hat{\beta}_{thr} = 0.3$. (The

Figure 6. Quantile–quantile results for all labelled test data (green dotted line), labelled test data for which $\hat{\beta}(x^L) < 0.5$ and ≥ 0.5 (black dashed and blue solid lines, respectively), and labelled test data for the case where we do not mitigate covariate shift by incorporating importance weights into estimation (orange dash–dotted line, largely following the blue solid line), which effectively limits modelling to the regime $\hat{\beta}(x^L) \lesssim 0.5$ (see Fig. 3). This plot demonstrates that mitigating covariate shift leads to worse conditional density estimates in the portion of covariate space where $\hat{\beta}(x^L) \lesssim 0.5$ and where $\lesssim 8$ per cent of the unlabelled data are located (orange line versus black line), but yields good estimates throughout the remaining covariate space (blue line versus orange line).

optimal number of nearest neighbours given this threshold is 20, hence the unique values of $\hat{\beta}(x^L)$ are 0, 0.05, 0.1, etc.) In the middle and bottom panels, we observe that for $\hat{\beta}(x^L) \lesssim 0.3$, the chi-square values are much larger, and the p -values much smaller, than what we would expect if $\hat{c} = c$: we do not achieve good behaviour in this regime. (This is consistent with the behaviour of the QQ plots shown in Fig. 6.) For $\hat{\beta}(x^L) \gtrsim 0.3$, on the other hand, the p -values are generally > 0.05 . We thus conclude that our method generates useful conditional density estimates in the regime $\hat{\beta}(x^L) \gtrsim 0.3$. (We note that we come to similar conclusions if we use the pre-processing thresholds $\hat{\beta}_{\text{thr}} = 0.1$ or 0.2 instead.)

Fig. 9 shows the results of applying the CvM, AD and KS tests (from top to bottom, respectively) to the data U generated from the CDFs of the conditional density estimates $\hat{f}(z|x)$ of the labelled test-set data, again assuming $\hat{\beta}_{\text{thr}} = 0.3$. We observe similar behaviour for the p -values here as observed in Fig. 7, with two differences: (1) the p -values are generally above 0.05 for $\hat{\beta}(x^L) \gtrsim 0.5$ as opposed to 0.3 and (2) there are more numerous deviations from uniformity in the regime $\hat{\beta}(x^L) \gtrsim 0.5$ than seen in the bottom panel of Fig. 7, particularly in the case of the AD test (middle panel, Fig. 9).

8 SUMMARY

In this paper, we provide a principled method for generating conditional density estimates $\hat{f}(z|x)$ [elsewhere commonly denoted $p(z)$ and dubbed the ‘photometric redshift PDF’] that takes into account selection bias and the covariate shift that this bias induces. (See Fig. 1 for an example of both: a bias towards brighter galaxies leads to shifted distributions of colours between spectroscopic and photometric data samples. See also Algorithms 1 and 2.) If not mitigated,

Figure 7. Results of testing whether the observed quantile value differences ($\hat{c}_1 - \hat{c}_0, \dots, \hat{c}_B - \hat{c}_{B-1}$) are constant across quantile bins, using the labelled test-set data at each unique value of $\hat{\beta}(x^L)$. Our estimates of $\hat{f}(z|x)$ are generated via the Combined estimator. In all panels, the blue dots/black crosses represent results generated with/without covariate shift mitigation (i.e. with/without the incorporation of importance weights into estimation), with the blue/black triangles specifying data points that fall outside the plot panels. [There are insufficient data above $\hat{\beta}(x^L) \geq 0.5$ to generate results without mitigating covariate shift.] The top panel shows the fraction of data for each unique value, with the dashed red lines indicating 0.1, 1 and 10 per cent, respectively. The middle panel shows the chi-square statistic resulting from each test, with the dashed red line indicating the number of bins ($B = 100$), the value to which chi square should converge if our estimator performs well in regions of covariate space with appreciable densities of unlabelled data. The bottom panel shows the p -value, estimated via simulation since the usual assumption of the chi-square GoF test that there are at least five expected counts in each bin is usually violated; the red dashed line indicates $p = 0.05$, with points above that line indicating that the observed quantile value differences are consistent with constancy across quantile bins. The middle and bottom panels demonstrate that our method achieves desired QQ behaviour at $\hat{\beta}(x^L) \gtrsim 0.3$, while also demonstrating the importance of assuming covariate shift with these data: without this assumption, we can only generate reasonable predictions for unlabelled data in the regime $\hat{\beta}(x^L) \lesssim 0.5$. We assume $\hat{\beta}_{\text{thr}} = 0.3$; analogous plots for $\hat{\beta}_{\text{thr}} = 0.1$ and 0.2 , not shown, indicate similar results.

covariate shift leads to situations where models fit to labelled (i.e. spectroscopic) data will not produce scientifically useful fits to the far more numerous unlabelled (i.e. photometric-only) data, lessening the impact of photo- z estimation within the context of precision cosmology. Here, we mitigate covariate shift by first estimating importance weights, the ratio $\beta(x)$ between the densities of unlabelled and labelled data at point x (Section 3), and then applying these weights to conditional density estimates $\hat{f}(z|x)$ (Section 4).

In order for our two-step procedure to succeed, ultimately, we require good estimates of $\beta(x)$ at labelled data points in order for it to achieve good estimates $\hat{f}(z|x)$ at unlabelled ones. We thus need both rigorously defined risk functions that allow us to tune the free parameters of our importance weight and conditional density estimators, and diagnostic tests that allow us to

Figure 8. Same as Fig. 7, except instead of QQ behaviour we examine the differences between observed and expected coverage as a function of $\hat{\beta}(x^L)$, and only for the case where we mitigate covariate shift. The top and bottom panels demonstrate that our method achieves good coverage for $\hat{\beta}(x^L) \gtrsim 0.5$.

Figure 9. Results of testing whether the observed random variables $U_i = \hat{F}(z_i|x_i)$ are uniformly distributed, using the labelled test-set data at each unique value of $\hat{\beta}(x^L)$. Our estimates of $\hat{F}(z|x)$ are generated via the Combined estimator. The panels show p -values generated via the CvM, AD and KS tests, from top to bottom, respectively. The KS test statistic is the maximum deviation of the empirical CDF for U from the uniform CDF, while the CvM and AD statistics are based on (unweighted and weighted) integrated distances between the empirical and uniform CDFs. These panels demonstrate that our method generally achieves good behaviour with respect to uniformity at $\hat{\beta}(x^L) \gtrsim 0.5$ (cf. 0.3 and 0.5 for tests based on QQ and coverage plots). This plot shows results for the pre-processing threshold $\hat{\beta}_{\text{thr}} = 0.3$; analogous plots for $\hat{\beta}_{\text{thr}} = 0.1$ and 0.2 , not shown, indicate similar results.

determine the quality of the estimates $\hat{f}(z|x)$. Our method is based on the assumption that the probability that astronomers label a galaxy (i.e. determine its spectroscopic redshift) depends only on its (photometric and perhaps other) properties x and not on its true redshift, an assumption currently valid for redshifts $\lesssim 0.5$. This is equivalent to assuming that the conditional densities for labelled

and unlabelled data match [$f_L(z|x) = f_U(z|x)$], even if the marginal distributions differ [$f_L(x) \neq f_U(x)$], which allows us ultimately to substitute out the true unknown quantities $\beta(x)$ and $f(z|x)$ in specifications of risk functions (equations 7 and 11). These risk functions, and their estimates (equations 7 and 11), are the backbone of our method: they allow us to tune parameters in a principled manner [e.g. what is the optimal number of nearest labelled neighbours when estimating $\beta(x)$ via equation 4?], as well as choose between competing estimators (e.g. which is better: the NN-CS, kerNN-CS or Series-CS estimators of conditional density?). An important question to answer in future work is whether we can relax our central assumption and still be able to write down estimated risks that lead to useful estimates of conditional densities in higher redshift regimes.

In Section 5, we demonstrate that once we generate p separate conditional density estimates $\hat{f}_k(z|x)$ (e.g. via the kerNN-CS and Series-CS estimators), tuned via the estimated risk given in equation (11), we can combine them to achieve better predictions (i.e. smaller values of risk). The method we propose utilizes a weighted linear combination, with the weights determinable via quadratic programming, but this is not the only possible way to combine estimates; see e.g. Dahlen et al. (2013), who discuss three methods for combining estimates, including one (Method 2) that adds estimates together as we do, except that while we determine optimal coefficients by minimizing estimated risk, they combine estimates so that 68.3 per cent of the spectroscopic redshifts in their sample fall within their final 1σ confidence intervals.

It is not sufficient to generate estimates $\hat{f}(z|x)$ by minimizing risk; one also needs to demonstrate that the estimates are scientifically useful. There is no unique way to demonstrate the quality of conditional density estimates. In Section 6, we provide alternatives that test (1) whether estimated cumulative densities, evaluated at actual redshifts, are distributed uniformly; (2) whether observed quantiles match expectation via QQ plots and the chi-square GoF test and (3) testing uniformity as a function of interval coverage. The jury is still out as to which of these diagnostics will play a central role in future photo- z analyses; for now, we consider it sufficient to demonstrate in any analysis that these diagnostics yield similar qualitative results.

In Section 7, we demonstrate our method using $\approx 500\,000$ galaxies with, and $\approx 500\,000$ without, spectroscopic redshifts, mostly from the SDSS (see Sheldon et al. 2012 for details). For computational efficiency, we sample 15 000 galaxies from both pools of data. While our initial labelled sample is chosen randomly from the larger pool of labelled data, we implement a pre-processing scheme (see Algorithm 2) to generate a new labelled sample with a larger effective size, whose photometry also more closely resembles that of the unlabelled sample (see Fig. 2). The pre-processing scheme requires the specification of an importance weight threshold ($\hat{\beta}_{\text{thr}}$) whose value cannot be optimized via tuning (since different thresholds yield different labelled data sets, and thus yield not-directly comparable estimated risks). We demonstrate that our results are generally insensitive to the choice of threshold within the regime $\hat{\beta}_{\text{thr}} \lesssim 0.3$. Those results include that (1) as expected, the conditional density of Combined estimator, constructed from those of the kerNN-CS and Series-CS estimators, provide the best estimates as quantified via the risk estimate in equation (11), and (2) via our diagnostic tests, we determine that our Combined estimates exhibit good behaviour in the regimes $\hat{\beta}(x) \gtrsim 0.3$, for QQ-based tests, and $\gtrsim 0.5$, for tests of coverage and cumulative densities. Our results thus demonstrate that our method achieves good, i.e.

scientifically useful, conditional density estimates for unlabelled galaxies.

ACKNOWLEDGEMENTS

The authors would like to thank Jeff Newman (University of Pittsburgh) for helpful discussions about photometric redshift estimation. This work was partially supported by NSF DMS-1520786 and the National Institute of Mental Health grant R37MH057881. RI further acknowledges the support of the Fundação de Amparo à Pesquisa do Estado de São Paulo (2014/25302-2).

REFERENCES

- Aihara H. et al., 2011, *ApJS*, 193, 29
 Ball N. M., Brunner R. J., 2010, *Int. J. Mod. Phys. D*, 19, 1049
 Benítez N., 2000, *ApJ*, 536, 571
 Bonnett C., 2015, *MNRAS*, 449, 1043
 Brammer G., von Dokkum P., Coppi P., 2008, *ApJ*, 686, 1503
 Budavári T., 2009, *ApJ*, 695, 747
 Carrasco Kind M., Brunner R., 2013, *MNRAS*, 432, 1483
 Carrasco Kind M., Brunner R., 2014, *MNRAS*, 438, 3409
 Collister A. A., Lahav O., 2004, *PASP*, 116, 345
 Corradi V., Swanson N. R., 2006, in Elliott G., Granger C., Timmermann A., eds, *Handbook of Economic Forecasting*. Elsevier, Amsterdam, p. 197
 Cunha C. E., Lima M., Oyaizu H., Frieman J., Lin H., 2009, *MNRAS*, 396, 2379
 Dahlen T. et al., 2013, *ApJ*, 775, 93
 De Vicente J., Sánchez E., Sevilla-Noarbe I., 2016, *MNRAS*, 459, 3078
 Flaugher B., 2005, *Int. J. Mod. Phys. A*, 20, 3121
 Gretton A., Smola A., Huang J., Schmittfull M., Borgwardt K., Schölkopf B., 2008, in Quionero-Candela J., Sugiyama M., Schwaighofer A., Lawrence N. D., eds, *Dataset Shift in Machine Learning*. MIT Press, Cambridge, p. 131
 Hall P., 1987, *Ann. Stat.*, 15, 1491
 Hildebrandt H. et al., 2010, *A&A*, 523, A31
 Ivezić Ž., 2008, preprint ([arXiv:0805.2366](https://arxiv.org/abs/0805.2366))
 Izbicki R., Lee A. B., 2016, *J. Comput. Graph. Stat*, 25, 1297
 Izbicki R., Lee A., Schafer C. M., 2014, *Proc. Mach. Learn. Res.*, 33, 420
 Izbicki R., Lee A., Freeman P. E., 2016, *Ann. Appl. Stat.*, preprint ([arXiv:1604.01339](https://arxiv.org/abs/1604.01339))
 James G., Witten D., Hastie T., Tibshirani R., 2014, *An Introduction to Statistical Learning: With Applications in R*. Springer, New York
 Kanamori T., Hido S., Sugiyama M., 2009, *J. Mach. Learn. Res.*, 10, 1391
 Kanamori T., Suzuki T., Sugiyama M., 2012, *Machine Learning*, 86, 335
 Kremer J., Gieseke F., Steenstrup Pedersen K., Igel C., 2015, *Astron. Comput.*, 12, 67
 Lima M., Cunha C. E., Oyaizu H., Frieman J., Lin H., Sheldon E., 2008, *MNRAS*, 390, 118
 Loog M., 2012, in Santamaria I., Arenas-Garcia J., Camps-Valls G., Erdogmus D., Perez-Cruz F., Larsen J. eds, *IEEE International Workshop on Machine Learning for Signal Processing*. IEEE, p. 455
 Mandelbaum R. et al., 2008, *MNRAS*, 386, 781
 Moreno-Torres J. G., Raeder T., Aláiz-Rodríguez R., Chawla N. V., Herrera F., 2012, *Pattern Recognit.*, 45, 521
 Rau M. M., Seitz S., Brimiouelle F., Frank E., Friedrich O., Gruen D., Hoyle B., 2015, *MNRAS*, 452, 3710
 Sánchez C. et al., 2014, *MNRAS*, 445, 1482
 Sheldon E. S., Cunha C. E., Mandelbaum R., Brinkmann J., Weaver B. A., 2012, *ApJS*, 201, 32
 Sugiyama M., Suzuki T., Nakajima S., Kashima H., Bünau P., Kawanabe M., 2008, *Ann. Inst. Stat. Math.*, 60, 699
 Wasserman L., 2006, *All of Nonparametric Statistics*. Springer, New York
 Wittman D., 2009, *ApJ*, 700, L174
 Wittman D., Bhaskar R., Tobin R., 2016, *MNRAS*, 457, 4005
 York D. et al., 2000, *AJ*, 120, 1579

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.