



SCUOLA INTERNAZIONALE SUPERIORE DI STUDI AVANZATI

SISSA Digital Library

StratLearn-z: Improved photo-z estimation from spectroscopic data subject to selection effects

*Original*

StratLearn-z: Improved photo-z estimation from spectroscopic data subject to selection effects / Moretti, C.; Autenrieth, M.; Serra, R.; Trotta, R.; Van Dyk, D. A.; Mesinger, A.. - In: OPEN JOURNAL OF ASTROPHYSICS. - ISSN 2565-6120. - 8:(2025), pp. 1-16. [10.33232/001c.137525]

*Availability:*

This version is available at: 20.500.11767/150703 since: 2026-05-07T13:54:50Z

*Publisher:*

*Published*

DOI:10.33232/001c.137525

*Terms of use:*

Testo definito dall'ateneo relativo alle clausole di concessione d'uso

*Publisher copyright*

note finali coverpage

(Article begins on next page)

# STRATLEARN-Z: IMPROVED PHOTO-Z ESTIMATION FROM SPECTROSCOPIC DATA SUBJECT TO SELECTION EFFECTS

CHIARA MORETTI<sup>1,2,3,4</sup> , MAXIMILIAN AUTENRIETH<sup>5,6,7</sup> , RICCARDO SERRA<sup>1</sup>, ROBERTO TROTTA<sup>1,2,3,8</sup>, DAVID A. VAN DYK<sup>5</sup> , AND ANDREI MESINGER<sup>2,9</sup>

<sup>1</sup>SISSA - International School for Advanced Studies, Via Bonomea 265, 34136 Trieste, Italy

<sup>2</sup>Centro Nazionale “High Performance Computer, Big Data and Quantum Computing”

<sup>3</sup>INAF – Osservatorio Astronomico di Trieste, Via Tiepolo 11, I-34143 - Trieste, Italy

<sup>4</sup>INFN sezione di Trieste

<sup>5</sup>Statistics Section, Department of Mathematics, Imperial College London, 180 Queen’s Gate, London SW7 2AZ, UK

<sup>6</sup>Statistical Laboratory, DPMMS, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK

<sup>7</sup>Institute of Astronomy and Kavli Institute for Cosmology, Madingley Road, Cambridge CB3 0HA, UK

<sup>8</sup>Department of Physics, Imperial College London, Blackett Laboratory, Prince Consort Rd, SW7 2AZ London, UK and

<sup>9</sup>Scuola Normale Superiore (SNS), Piazza dei Cavalieri 7, Pisa, PI, 56125, Italy

Version May 1, 2025

## ABSTRACT

A precise measurement of photometric redshifts (photo- $z$ ) is crucial for the success of modern photometric galaxy surveys. Machine learning (ML) methods show great promise in this context, but suffer from covariate shift in training sets due to selection bias where interesting sources, e.g., high redshift objects, are underrepresented, and the corresponding ML models exhibit poor generalisation properties. We present an application of the *StratLearn* method to the estimation of photo- $z$  (*StratLearn-z*), validating against simulations where we enforce the presence of covariate shift to different degrees. *StratLearn* is a statistically principled approach which relies on splitting the combined source and target datasets into strata, based on estimated propensity scores. The latter is the probability for an object in the dataset to be in the source set, given its observed covariates. After stratification, two conditional density estimators are fit separately within each stratum, and then combined via a weighted average. We benchmark our results against the GPz algorithm, quantifying the performance of the two algorithms with a set of metrics. Our results show that the *StratLearn-z* metrics are only marginally affected by the presence of covariate shift, while GPz shows a significant degradation of performance, specifically concerning the photo- $z$  prediction for fainter objects for which there is little training data. In particular, for the strongest covariate shift scenario considered, *StratLearn-z* yields a reduced fraction of catastrophic errors, a factor of 2 improvement for the RMSE as well as one order of magnitude improvement on the bias. We also assess the quality of the predicted conditional redshift estimates using the probability integral transform (PIT) and the continuous rank probability score (CRPS). The PIT for *StratLearn-z* indicates that predictions are well-centered around the true redshift value, if conservative in their variance; the CRPS shows marked improvement at high redshifts when compared with GPz. Our JULIA implementation of the method, *StratLearn-z*, is publicly available at <https://github.com/chiamoretti/StratLearn-z>.

## 1. INTRODUCTION

The main science driver for current and planned cosmological experiments is the exploration of the dark sector. Stage-IV Dark Energy surveys, such as Euclid (Laureijs et al. 2011; Euclid Collaboration et al. 2024), the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST, The LSST Dark Energy Science Collaboration et al. (2018); Ivezić et al. (2019)), and Roman Space Telescope (Akeson et al. 2019), will pursue this goal by mapping the Universe over unprecedented volumes, delivering high precision measurements of cosmological observables. Of key importance for such measurements is our ability to produce accurate estimates of redshift for billions of sources.

The most precise way to estimate redshifts is via spectroscopic observations. These, however, are demanding in terms of observational time, and challenging for faint, high-redshift objects: the sheer number of sources in future surveys will prevent spectroscopic follow-up for the vast majority. A viable and well established alternative is provided by the so-called

photometric redshift (photo- $z$ , see e.g., Salvato et al. (2019); Brescia et al. (2021); Newman & Gruen (2022) for detailed reviews), which are extracted from broadband flux measurements performed at different wavelengths. While spectroscopy allows for highly accurate redshift measurements for individual objects, the sample obtained is often biased and unrepresentative of the full population, whereas photometric redshifts, despite being less precise for individual objects, allow the measurement of several objects simultaneously.

Accuracy and precision in the estimation of photo- $z$  are essential for the success of both current and future imaging surveys. In fact, systematic errors in the determination of redshifts from photometric observations can introduce biases into inferred cosmological parameters. Specifically, the lensing analysis relies on the construction of tomographic redshift bins, with minimal overlap between adjacent bins and an accurate determination of the mean redshift of each bin (Ma et al. 2006; Amara & Réfrégier 2007), as well as knowledge on the  $n(z)$ . Additionally, the clustering signal used in the 3x2pt analysis requires both the mean redshift and the width of the redshift distribution to be known with high accu-

racy (Tutusaus et al. 2020; Porredon et al. 2022). There are two main approaches for estimating photo- $z$ : template-fitting methods and machine learning (ML) based methods. The former are based on matching the observed photometry to catalogs of template galaxy spectra in order to extract the galaxy redshifts; these methods rely on the completeness of such catalogs. Publicly available template-fitting codes include LEPHARE (Arnouts et al. 1999; Ilbert et al. 2006), BPZ (Benítez 2000), HYPERZ (Bolzonella et al. 2000) and EAZY (Brammer et al. 2008). On the other hand, the advent of ML techniques has opened new ways to improve photo- $z$  estimation, taking advantage of artificial neural networks (Firth et al. 2003; Colister & Lahav 2004; Graff et al. 2014) random forests (Carliles et al. 2010; Carrasco Kind & Brunner 2013), advanced ANNs (Sadeh et al. 2016), Gaussian processes (Almosallam et al. 2016a,b), nearest neighbors (Graham et al. 2018), convolutional neural networks (Pasquet et al. 2019; Henghes et al. 2022), graph neural networks (Tosone et al. 2023), and self-organizing maps (Carrasco Kind & Brunner 2014; Masters et al. 2015). Self-organizing maps project high-dimensional data points onto a two-dimensional grid. During training, the grid nodes (or neurons) are iteratively adjusted to capture the underlying structure of the data. After training, photometric target galaxies are assigned to their closest grid nodes based on a distance metric (typically Euclidean distance), and photometric redshift (photo- $z$ ) estimates are derived by averaging the redshifts of the training galaxies associated with each node. This approach can be interpreted as a form of stratification of the data space, where each node effectively defines a stratum. In the original implementations by Carrasco Kind & Brunner (2014) and Masters et al. (2015), the grids consist of approximately 1,000 to 11,000 nodes, making SOMz an extreme example of a stratified approach.

ML methods have successfully been applied to Stage-III surveys: both the Dark Energy Survey (DES, Abbott et al. (2018)) and the Kilo-Degree Survey (KiDS, Heymans et al. (2021)) have used self-organizing maps to estimate photo- $z$  (see Myles et al. (2021) and Hildebrandt et al. (2021) respectively), while the Hyper Suprime-Cam (HSC, Aihara et al. (2018)) adopted conditional density estimators (Rau et al. 2023; Sugiyama et al. 2023). Performance analyses comparing different approaches have also been presented (Hildebrandt et al. 2010; Sánchez et al. 2014). In preparation for Stage-IV surveys, several ML methods have been compared with each other and to template-fitting methods, in particular for Euclid (Euclid Collaboration et al. 2020) and for LSST-DESC (Schmidt et al. 2020).

ML methods rely on source datasets used to train the algorithms where both spectroscopic and photometric observations are available, in order to learn the relationship between the observed photometry and redshift. A significant issue is posed by non-representative training sets, i.e., spectroscopic training samples that are not random samples from the target photometric survey. Indeed, such non-representative samples are the usual situation in astronomy, where the lack of representativeness is a generic and widespread consequence of selection bias. An example is ‘Malmquist bias’, (Malmquist 1922, 1925)), which results in spectroscopic catalogs being biased towards brighter, lower redshift sources. Observational biases can thus result in different distributions for the covariates of the source datasets, used to train the ML algorithms, and the target dataset, for which we wish to estimate photo- $z$ . Under the assumption that the conditional distribution of the outcome (in this context, the redshift) given the covariates (colors/magnitudes) is the same in the source and target

group, the effect is known as covariate shift (Moreno-Torres et al. 2012) in the ML literature. The impact of covariate shift on photo- $z$  estimation has been investigated in Freeman et al. (2017), where the authors propose to mitigate the effect of covariate shift with the use of importance weights, and apply the proposed method to data from DR8 of the Sloan Digital Sky Survey. More recently, Stylianou et al. (2022) investigates the performance of GPz (Almosallam et al. 2016a,b), an algorithm based on Gaussian processes, applying it to simulations that feature imperfections in the spectroscopic training set. Additionally, Toribio San Cipriano et al. (2023) studies the impact of incompleteness of the training set in the context of DES, while Moskowitz et al. (2024) suggests a method to augment the training sample to improve photo- $z$  estimation.

In this work, we focus on *StratLearn* (Autenrieth et al. 2024b; Autenrieth et al. 2024a), a statistically motivated approach that improves the performance of ML-based algorithms in the presence of covariate shift. *StratLearn* splits the data into subgroups, or strata, based on the estimated propensity scores, which in the context of this paper is the probability of a galaxy in the data set being included in the spectroscopic source set, given its observed covariates (magnitudes or colors). Propensity scores are balancing scores, and a pivotal methodology in causal inference (Rosenbaum & Rubin 1983). Conditioning on the propensity scores via stratification approximately balances within strata the covariates of source and target data, which then improves the fitting of ML methods within each stratum (Autenrieth et al. 2024b). In this paper we apply an implementation of the *StratLearn* method, *StratLearn- $z$* , to photo- $z$  estimation, and assess its performance on simulated data with different degrees of covariate shift, with a view to future surveys such as Euclid and LSST. Our goal is to showcase the ability of *StratLearn- $z$*  to provide accurate photo- $z$  estimates even in presence of strong covariate shift, which can reduce systematic biases and therefore be highly beneficial for such experiments.

Throughout, we refer to the implementation of *StratLearn* adopted here as *StratLearn- $z$* .

## 2. STRATLEARN – ADDRESSING COVARIATE SHIFT

In this section we present in more detail the covariate shift problem in the context of photo- $z$  estimation, and summarize the *StratLearn* approach developed in Autenrieth et al. (2024b).

### 2.1. Covariate shift in non-representative spectroscopic data

Starting from a source dataset of labelled observations with spectroscopic redshift  $z^{(i)}$  and covariates  $x^{(i)}$  (photometric magnitudes or colors), we aim to obtain redshift estimations for a much larger target set of unlabelled sources, i.e., for sources with only photometric information/covariates available. The source dataset  $D_S = \{(x_S^{(i)}, z_S^{(i)})\}_{i=1}^{N_S}$  consists of  $N_S$  galaxies sampled from the joint distribution  $p_S(x, z)$ , while the target dataset  $D_T = \{x_T^{(i)}\}_{i=1}^{N_T}$  consists of  $N_T$  galaxies sampled from the joint distribution  $p_T(x, z)$ , with  $N_T \gg N_S$ . Selection effects induce differences between the source and target distributions,  $p_S(x) \neq p_T(x)$ , leading to  $p_S(x, z) \neq p_T(x, z)$ . However, we assume that the conditional distribution of redshifts given the covariates is the same for both source and target,  $p_S(z|x) = p_T(z|x)$ . This situation is known to affect the performances of ML algorithms, because the patterns learned from the (unrepresentative) source set do not gener-

alise properly to the target set. <sup>1</sup>

## 2.2. StratLearn

*StratLearn* is designed to improve the generalisation properties of ML algorithms in presence of covariate shift. While the method is general and applicable to both classification and regression problems, we focus here on photo- $z$  estimation, as discussed in the previous section. In the following we denote all estimated quantities with a hat symbol, e.g.,  $\hat{p}$ .

The approach relies on stratification of data into sub-groups, or strata, based on propensity scores, i.e., the estimated probability for an object  $i$  in the dataset to be in the source set ( $s_i = 1$ ) given the observed covariates  $x_i$ :

$$e(x_i) = P(s_i = 1|x_i), \quad (1)$$

with the assumption that  $0 < e(x_i) < 1$ . For each object we obtain an estimate of the propensity score  $\hat{e}(x_i)$  via binary classification of source and target data. Specifically, we employ a logistic regression model with all photometric magnitudes/colors as covariates, and group the data in  $K$  strata based on the quantiles of the estimated propensity score distribution. We use  $K = 5$  strata, as it has been shown that this number is able to remove at least 90% of the bias for many distributions (Cochran 1968; Rosenbaum & Rubin 1983). The stratification process balances the covariate distribution in the source and target sets (Rosenbaum & Rubin 1983), resulting in  $p_{T_j}(z, x) \simeq p_{S_j}(z, x)$ , for  $j = 1, \dots, K$ , where the subscript  $T_j$  denotes conditioning on assignment to the  $j$ -th target stratum (and analogously,  $S_j$  for source stratum).

We use a weighted average of two supervised full conditional density estimators, trained separately in each stratum, to obtain a non-parametric estimate of the full galaxy photo- $z$  conditional densities of each object in the target set,  $\hat{p}_T(z|x)$ , conditional on its observed covariates. We adopt the *ker-NN* estimator proposed in Freeman et al. (2017) and the *Series* estimator described in Izbicki et al. (2016): the former relies on a kernel smoothed histogram of the redshift of the  $k$  nearest neighbors to  $x$ , within each stratum, to compute the conditional redshift distribution, while the latter adapts a lower-dimensional subspace of the covariates,  $x$ , as the intrinsic dimension of the data, based on data-dependent eigenfunctions of a kernel based estimator. Our choice of conditional density estimators is based on good performance shown in previous work, e.g., Autenrieth et al. (2024b). However, generally, the *StratLearn* framework could be combined with any model.

We fit the two models using the labelled spectroscopic source data as a training set, and obtain an estimate of the redshift probability distribution for each galaxy in the photometric target data. If the source set is representative of the target set, conditional density estimators commonly aim to minimize the generalized risk under the  $L^2$ -loss:

$$\hat{R}_S(\hat{p}) = \frac{1}{N_S} \sum_{k=1}^{N_S} \int \hat{p}^2(z | x_S^{(k)}) dz - 2 \frac{1}{N_S} \sum_{k=1}^{N_S} \hat{p}(z_S^{(k)} | x_S^{(k)}), \quad (2)$$

where the subscript ‘ $S$ ’ indicates the source set,  $z_S^{(k)}$  is the redshift and  $x_S^{(k)}$  refers to the covariates for the  $k$ -th object in the source set. We notice however that the presence of covariate

shift requires minimisation of the *target* risk  $\hat{R}_T(\hat{p})$  in order to obtain accurate target estimates, which in turn would require having access to the target redshifts. The stratification procedure described above allows for minimisation of  $\hat{R}_{T_j}(\hat{p})$  via minimisation of  $\hat{R}_{S_j}(\hat{p})$  within each stratum  $j$  (see Autenrieth et al. (2024a) for a more detailed description). Finally, the two conditional density estimators are combined via

$$\hat{p}(z|x) = (1 - \gamma)\hat{p}_{\text{Series}}(z|x) + \gamma\hat{p}_{\text{ker-NN}}(z|x), \quad (3)$$

where  $0 \leq \gamma \leq 1$  is a weight parameter. This requires a further optimisation to compute  $\gamma$  (within each stratum) via minimisation of the generalised risk:

$$\hat{R}_{S2_j}(\hat{p}) = \frac{1}{N_{T_j}} \sum_{k=1}^{N_{T_j}} \int \hat{p}^2(z | x_T^{(k)}) dz - 2 \frac{1}{N_{S_j}} \sum_{k=1}^{N_{S_j}} \hat{p}(z_S^{(k)} | x_S^{(k)}). \quad (4)$$

## 3. SIMULATED DATASET

To assess the performance of *StratLearn* we make use of a simulated dataset derived from the Buzzard Flock<sup>2</sup> simulation suite (DeRose et al. 2019), initially developed for DES. In particular, we adopt the LSST-DESC DC1 mock catalog used in Stylianou et al. (2022)<sup>3</sup>, consisting of a sample of 100,000 galaxies with redshift range  $0 < z < 2.3$  and complete photometry in the *ugrizy* bands, from which we compute colors to serve as input covariates to our algorithm.

The dataset was constructed from the catalogs described in Schmidt et al. (2020) (specifically, see Sec. 2 of that work and the references therein for a more detailed description of the dataset). Stylianou et al. (2022) uses this dataset to study the impact of systematics in the training dataset on the GPz algorithm (Almosallam et al. 2016a,b). Particular attention is paid to two possible degradations in the training dataset that might affect the performance of the algorithm: redshift incompleteness and emission line confusion. The redshift incompleteness in particular is similar to what we explore here, we therefore use GPz to benchmark the performance of *StratLearn-z*. However, we introduce incompleteness (and therefore, lack of representativeness) in the training data via a different prescription, as detailed in Sec. 3.1.

### 3.1. Introducing covariate shift

We enforce covariate shift in the dataset by performing rejection sampling on the  $r$ -band, to obtain the target data, following Izbicki et al. (2016). This approach is designed to mimic a selection effect that biases the source set towards brighter objects, similar to what happens with real observations. Specifically, we resample the full dataset via rejection sampling, assigning each galaxy to the target set with probability  $p(s = 0 | x) = f_{B(\alpha,\beta)}(x_r) / \max_{x_r} f_{B(\alpha,\beta)}(x_r)$ , where  $f_{B(\alpha,\beta)}$  is the density of a beta random variable with parameters  $\alpha, \beta$  and the  $r$ -band magnitudes,  $x_r$ , are re-scaled to be between 0 and 1. Galaxies not assigned to the target set are assigned to the source data.

We consider four different scenarios with different degrees of covariate shift (CS):

- no covariate shift: data are randomly split between source and target sets;

<sup>2</sup> <https://buzzardflock.github.io/index.html>

<sup>3</sup> Specifically, the catalogue distributed with the *pzflow* package (Crenshaw et al. 2024) at [https://github.com/jfcrenshaw/pzflow/blob/main/pzflow/example\\_files/galaxy-data.pkl](https://github.com/jfcrenshaw/pzflow/blob/main/pzflow/example_files/galaxy-data.pkl)

<sup>1</sup> We note that the assumption  $p_S(z|x) = p_T(z|x)$  may not always hold, e.g. in case of additional selection bias in the identification of a reliable spectroscopic redshift (Hartley et al. 2020) (see also the discussion in Sec. 5). We plan to explore such cases in a future work.




FIG. 1.— Source and target normalised distributions of  $r$ -band magnitudes for the four covariate shift (CS) cases, as stated in each panel. Orange histograms represent the source (labelled) datasets, while in blue we plot the target (unlabelled) datasets. Notice how the rejection sampling procedure introduces a shift – to different degrees, based on the chosen parameters (i.e. scenarios) – between the source and target sets.

- weak covariate shift:  $\alpha = 4, \beta = 5$ ;
- mild covariate shift:  $\alpha = 5, \beta = 6$ ;
- strong covariate shift:  $\alpha = 5, \beta = 7$ .

From this procedure we obtain four source datasets, three of which have  $r$ -band distributions that are shifted with respect to their corresponding target datasets, as shown in Fig. 1 for the various covariate shift cases. Although in a simplified form, the rejection sampling procedure mimics the effect of selection bias in real data, where the limiting magnitude of the  $r$ -band distribution of spectroscopic observations is usually at lower values than for photometric observations (see e.g., Fig. 3 of Izbic et al. 2016), because of higher probability of detecting brighter objects. Such bias is also visible as a shift in Fig. 2: the source set features a peak at lower redshifts, that is more pronounced the stronger the covariate shift. We also plot the distributions for all other photometric bands in Appendix A.

The rejection sampling procedure described above results in source and target datasets of very different sizes among the four covariate shift scenarios. To focus on the impact of covariate shift, and remove dependence of the results on the size of the datasets, we want the source sets for all scenarios to have the same size (and similarly for the target sets). Thus, we randomly resample the sets to construct the source datasets, used to train the conditional density estimators, and target datasets, used to compute performance metrics.

Specifically, we randomly sub-sample 20,000 objects for each training set and 63,000 in each target set, except the target dataset in the weak covariate shift scenario was not sub-sampled because rejection sampling yielded only 62,000 objects.

We choose the target sets to be significantly larger than the training sets to reproduce what happens with real data, where labels (spectroscopic redshifts) are only available for a small fraction of the total photometric measurements. For example, Euclid is expected to observe photometry for  $\sim 1.5$  billion galaxies, and will rely on complementary ground-based




FIG. 2.— Source and target normalised redshift distributions for the four covariate shift (CS) cases, as obtained after the rejection-sampling procedure described in Sec. 3.1. Orange histograms represent the source datasets, while in blue we plot the target datasets. We plot the true redshift distribution for the target data, which in reality is of course unknown.

spectroscopic samples to calibrate the relation between galaxy colours and redshifts, with significant effort put into targeted spectroscopic observations to span the full galaxy colour space (see Euclid Collaboration et al. (2024) for more details).

### 3.2. StratLearn computation

For each scenario, we estimate the propensity scores using a logistic regression model to predict the binary classification into the source and target data. In particular, we use the  $r$ -band data and colors obtained from all other available photometric bands as covariates (i.e.,  $u-g, g-r, r-i, i-z, z-y$ ). We scale all covariates to have mean zero and standard deviation one. We then partition each of the datasets (combining source and target data under each scenario) into five strata, based on the quintiles of the corresponding propensity score distributions. The resulting strata are populated as detailed in Table 1. The no covariate shift case results in approximately equally sized strata:  $\sim 4000$  in the source set and  $\sim 12500$  in the target set, per stratum. For the datasets with covariate shift there is significantly more data in the first source strata (corresponding to a lower mean redshift), as expected.

Fig. 3 plots the distributions of the  $r$ -band magnitudes for the source and target sets within each stratum of the strong covariate shift scenario. Stratification achieves approximately balanced distributions within each strata. Additionally, stratification results in a redshift partition, where in each stratum the mean redshift of the source data approximately matches the mean redshift of the target data – an indication that stratification is approximately removing covariate shift (Autenrieth et al. 2024b) within each stratum. Redshift distribution for source and target sets for the strong covariate shift case are shown in Fig. 16 of Appendix B, showcasing once again the effectiveness of StratLearn to improve balance within each stratum.

We train the ker-NN and Series conditional density estimators separately within each stratum, and combine them using Eq. 3 to obtain an estimate of the conditional distribution of redshift for each object in the target set. We emphasize that only the source redshifts are used to train the models, while the target redshifts are used exclusively to evaluate perfor-

TABLE 1

Balance between strata after propensity score estimation, for the four covariate shift (CS) scenarios. We also report the mean redshift of each stratum for both the source and target datasets. A similar value for the mean redshift within strata indicates that the source and target data are balanced, although we note that this diagnostic is not possible in practice with real source/target data, as  $z$  is not observed in the target dataset.

Stratum	No CS		weak CS		mild CS		strong CS	
	Size	Mean $z$	Size	Mean $z$	Size	Mean $z$	Size	Mean $z$
1 (source)	4113	1.298	9959	0.625	10470	0.617	11038	0.571
1 (target)	12487	1.306	6441	0.708	6130	0.705	5562	0.666
2 (source)	4053	1.313	4479	0.899	4643	0.894	4764	0.850
2 (target)	12547	1.311	11921	0.904	11957	0.887	11836	0.854
3 (source)	3978	1.051	2799	0.979	2644	0.954	2432	0.963
3 (target)	12622	1.058	13601	0.984	13956	0.972	14168	0.962
4 (source)	3929	0.740	1917	1.113	1549	1.130	1253	1.061
4 (target)	12671	0.753	14483	1.121	15051	1.108	15347	1.075
5 (source)	3927	0.644	846	1.498	694	1.469	513	1.406
5 (target)	12673	0.624	15554	1.579	15906	1.585	16087	1.525
total (source)	20000	1.014	20000	0.820	20000	0.795	20000	0.737
total (target)	63000	1.009	62000	1.121	63000	1.117	63000	1.087

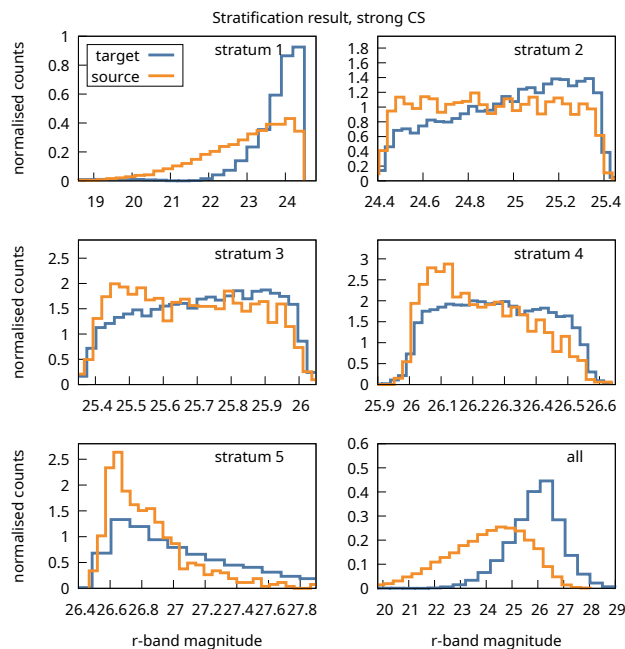


FIG. 3.— Distribution of r-band magnitudes for source (orange histograms) and target (blue histograms) sets for the strong covariate shift (CS) scenario, after performing stratification based on estimated propensity scores. In the lower right panel, where we plot the full source and target sets, the distributions are significantly different, while in the other panels, representing each a different stratum, we are able to achieve improved balance.

mance metrics. The estimated (photometric) redshift  $\hat{z}_{\text{photo}}$ , computed as the mean of the estimated conditional redshift distribution, is compared to the true (spectroscopic) redshift  $z_{\text{spec}}$  for each galaxy in the target dataset, and performance in each covariate shift scenario is assessed using several metrics (see Sec. 4). As discussed in Sec. 3, we benchmark our results against the GPz code run on the same datasets.

#### 4. *StratLearn-z* PERFORMANCE

We evaluate the performance of *StratLearn-z* in terms of both its point estimate  $\hat{z}_{\text{photo}}$  and its estimate of the conditional distribution of redshift via several metrics applied to the conditional densities obtained for the target dataset.

To obtain more robust numerical results for the performance metrics, we run 15 independent realisations for each covariate shift scenario. For each realisation, the datasets are constructed as follows:

- rejection sampling, as described in Sec. 3.1, is applied to the full dataset to generate four distinct source/target pairs corresponding to the different covariate shift scenarios;
- from each of these sets, we randomly sample 20,000 objects for the source set and 63,000 for the target set (62,000 for the weak covariate shift scenario; see Sec. 3.1).

The procedure is repeated 15 times, in order to obtain a total of 60 source/target sets, 15 for each covariate shift scenario. The models are trained separately on each source set, and conditional density estimations constructed for the respective target set. For each realisation, several performance metrics are then computed as described in the next section.

##### 4.1. Point estimate metrics

The quality of the photo- $z$  point estimate  $\hat{z}_{\text{photo}}$  is assessed by comparing it to the true redshift  $z_{\text{spec}}$ . Because we are working with simulated data,  $z_{\text{spec}}$  is available for all objects in the target set, which can therefore be used to compute performance metrics. We focus on the root mean square error (RMSE), the bias, and the fraction retained (FR). The latter provides an estimate of (100 times) the fraction of objects for which  $\hat{z}_{\text{photo}}$  does not exhibit catastrophic errors, defined as those objects for which residual differences between  $\hat{z}_{\text{photo}}$  and  $z_{\text{spec}}$  are above some large threshold, defined relative to

specific survey requirements. FR corresponds therefore to the fraction of redshift estimates that are not catastrophic, i.e., are ‘good’ (see Eq.s 7 and 8 for the thresholds adopted here). The metrics are evaluated as follows:

- Root mean square error (RMSE):

$$\sqrt{\frac{1}{N_T} \sum_{i=1}^{N_T} \left( z_{\text{spec}}^{(i)} - \hat{z}_{\text{photo}}^{(i)} \right)^2}; \quad (5)$$

- mean error, commonly referred to in the literature as ‘bias’:

$$\frac{1}{N_T} \sum_{i=1}^{N_T} \left( z_{\text{spec}}^{(i)} - \hat{z}_{\text{photo}}^{(i)} \right); \quad (6)$$

- FR15:

$$\frac{100}{N_T} \sum_{i=1}^{N_T} \left[ \left| \frac{z_{\text{spec}}^{(i)} - \hat{z}_{\text{photo}}^{(i)}}{1 + z_{\text{spec}}^{(i)}} \right| < 0.15 \right]; \quad \text{and} \quad (7)$$

- FR05:

$$\frac{100}{N_T} \sum_{i=1}^{N_T} \left[ \left| \frac{z_{\text{spec}}^{(i)} - \hat{z}_{\text{photo}}^{(i)}}{1 + z_{\text{spec}}^{(i)}} \right| < 0.05 \right]. \quad (8)$$

The RMSE and mean error used in the context of photo- $z$  estimation are sometimes normalised by a factor  $1 + z_{\text{spec}}$ : this serves to down-weight the deviation from the spectroscopic redshift for high-redshift sources. For example, [Almosallam et al. \(2016a\)](#) and [Euclid Collaboration et al. \(2020\)](#) use the normalised definitions, while [Hartley et al. \(2020\)](#) and [Toribio San Cipriano et al. \(2023\)](#) use the un-normalised ones. We note that the choice of performance metrics in the literature is far from homogeneous. Given our benchmark against GPz, we adopt the metrics used in [Stylianou et al. \(2022\)](#), but report our results in terms of both the normalised and un-normalised metrics.

#### 4.2. Conditional redshift distribution metrics

Measuring the precision of the conditional redshift distribution is more complicated, because there are no ‘true’ conditional redshift distributions available. One possibility is using the distribution of the probability integral transform (PIT, [Dawid \(1984\)](#), see also [Polsterer et al. \(2016\)](#)), defined as

$$\text{PIT} = \int_0^{z_{\text{spec}}} \hat{p}_T(z|x) dz, \quad (9)$$

where  $\hat{p}_T(z|x)$  is the conditional redshift distribution prediction for the target set given the observed photometry  $x$ . Under the assumption that  $\hat{p}_T(z|x)$  exactly represents the conditional distribution of the true redshift,  $z_{\text{spec}}$ , given  $x$ , PIT is uniformly distributed over the unit interval. On the other hand, shifts from uniformity in the PIT distribution suggest systematic deviations in the estimates of the conditional distributions of redshift, while over-estimation (under-estimation) of the spread of the conditional distribution of redshift, i.e., the uncertainty associated with the estimated  $\hat{z}_{\text{photo}}$ , appear as a bump (dip) in the histogram of PIT. The PIT distribution is commonly used to visually evaluate the quality of estimated photo- $z$  distributions (see e.g. [Bordoloi et al. 2010](#); [Polsterer](#)

[et al. 2016](#); [D’Isanto & Polsterer 2018](#); [Tanaka et al. 2018](#); [Euclid Collaboration et al. 2020](#)).

However, we note that the PIT distribution is not an indicator of the information content of the estimated conditional density. Indeed, [Schmidt et al. \(2020\)](#) shows that a method with no predictive power can produce a perfectly uniform PIT distribution. As an example, the TRAINZ estimator described in [Schmidt et al. \(2020\)](#) assigns the marginal redshift distribution of the source data to each galaxy in the target set. This distribution does not depend on the photometric data and produces the same photo- $z$  estimate for every galaxy in the target set, however, TRAINZ is able to achieve a perfect PIT.

Additionally, we compute the continuous rank probability score (CRPS, [Hersbach \(2000\)](#); [Polsterer et al. \(2016\)](#)), defined as the integral of the squared difference between the redshift cumulative distribution function and the Heaviside step function centered at the true redshift value. More precisely, for each galaxy  $i$ , the CRPS is given by

$$\text{CRPS}^{(i)} = \int_{-\infty}^{+\infty} \left[ \text{CDF}^{(i)}(z) - \text{H}(z - z_{\text{spec}}^{(i)}) \right]^2 dz, \quad (10)$$

where  $\text{CDF}^{(i)}(z)$  is the cumulative distribution function of the estimated conditional redshift distribution, and  $\text{H}(z - z_{\text{spec}}^{(i)})$  is the Heaviside step function, i.e.,

$$\text{H}(z - z_{\text{spec}}^{(i)}) = \begin{cases} 0 & \text{if } z < z_{\text{spec}}^{(i)}, \\ 1 & \text{if } z \geq z_{\text{spec}}^{(i)}. \end{cases} \quad (11)$$

The CRPS quantifies the discrepancy between the estimated conditional redshift distribution and a Dirac delta function at  $z_{\text{spec}}^{(i)}$ ; it deviates from zero when the conditional distribution is broad or miscentered from the true redshift  $z_{\text{spec}}^{(i)}$ . The median CRPS across all objects in the target dataset serves as a measure of the overall quality of the estimated redshift distributions.

#### 4.3. Results and comparison to GPz

We compare the performance of *StratLearn-z* to the output of GPz ([Almosallam et al. 2016a,b](#)) for each of the four covariate shift scenarios, by running both codes on the same datasets<sup>4</sup>. GPz is based on sparse Gaussian processes and produces input-dependent (in this context, photometry-dependent) variance estimates which take into account both the density of the source data and noise in the photometric observations. One drawback of GPz is that the predicted conditional redshift distribution is forced to be a unimodal Gaussian. The GPz code was compared to other existing photo- $z$  software in [Schmidt et al. \(2020\)](#) for LSST and in the photometric challenge paper for Euclid ([Euclid Collaboration et al. 2020](#)); as is the case for other ML approaches, its performance is strongly dependent on the quality of the spectroscopic set used for training.

The 2D histograms in Fig. 4, 5, 6, and 7 show the true  $z_{\text{spec}}$  versus the estimated  $\hat{z}_{\text{photo}}$ , obtained as the mean of the predicted conditional redshift distribution, for each covariate shift scenario and for both *StratLearn-z* (left panels) and GPz (right panels). Visual inspection of the 2D histograms indicates that the GPz predictions are strongly degraded for high-redshift

<sup>4</sup> We use the GPz version distributed with [Stylianou et al. \(2022\)](#) and available from [https://github.com/nataliastylianou/photo-\\$z/tree/main](https://github.com/nataliastylianou/photo-$z/tree/main).




FIG. 4.— 2D histograms for the estimated  $\hat{z}_{\text{photo}}$  versus the true  $z_{\text{spec}}$  for the no covariate shift (CS) case,  $\alpha = 1$ ,  $\beta = 1$ . Left: *StratLearn-z*, right: GPz. The colorbar indicates the number of sources in each 2D bin.




FIG. 5.— Same as Fig. 4, but for the weak covariate shift case,  $\alpha = 4$ ,  $\beta = 5$ . Left: *StratLearn-z*, right: GPz.

objects in the presence of covariate shift, where the distribution of  $z_{\text{spec}} - \hat{z}_{\text{photo}}$  exhibits high variance. On the other hand, the *StratLearn-z* predictions are only marginally impacted by covariate shift, even in the worst case scenario shown in Fig. 7. We note that the plots refer to one specific realisation out of the 15 runs we perform for each case, however, a similar pattern is apparent for the remaining realisations.

We perform a more quantitative comparison of the redshift point-estimates by means of the performance metrics described in Sec. 4.1. The metrics are computed separately for

each of the 15 resampled datasets and then averaged. Numerical results for the average performance metrics and standard deviations are collected in Table 2 and illustrated in Fig. 8.

We note that *StratLearn-z* outperforms GPz for all metrics and in all cases considered, even in the no covariate shift case, which can likely be ascribed to better performance of the conditional density estimators. However, the difference is more pronounced with more covariate shift. Both the averaged RMSE and the normalised averaged RMSE, shown in the top panels of Fig. 8, show a similar trend: while for GPz




FIG. 6.— Same as Fig. 4, but for the mild covariate shift case,  $\alpha = 5$ ,  $\beta = 6$ . Left: *StratLearn-z*, right: *GPz*.




FIG. 7.— Same as Fig. 4, but for the strong covariate shift case,  $\alpha = 5$ ,  $\beta = 7$ . Left: *StratLearn-z*, right: *GPz*.

TABLE 2

Average performance metrics and the corresponding standard deviations for *StratLearn-z* and *GPz* for the four covariate shift (CS) scenarios considered.

	no CS		weak CS		mild CS		strong CS	
	<i>StratLearn-z</i>	<i>GPz</i>	<i>StratLearn-z</i>	<i>GPz</i>	<i>StratLearn-z</i>	<i>GPz</i>	<i>StratLearn-z</i>	<i>GPz</i>
RMSE	$0.0583 \pm 0.0036$	$0.0842 \pm 0.0079$	$0.0697 \pm 0.0073$	$0.1159 \pm 0.0149$	$0.0724 \pm 0.0054$	$0.1332 \pm 0.0249$	$0.0817 \pm 0.0010$	$0.1526 \pm 0.0299$
Norm. RMSE	$0.0302 \pm 0.0025$	$0.0363 \pm 0.0020$	$0.0313 \pm 0.0019$	$0.0431 \pm 0.0032$	$0.0323 \pm 0.0019$	$0.0483 \pm 0.0071$	$0.0373 \pm 0.0052$	$0.0547 \pm 0.0086$
Bias	$0.0012 \pm 0.0008$	$0.0043 \pm 0.0027$	$0.0033 \pm 0.0017$	$0.0218 \pm 0.0080$	$0.0038 \pm 0.0017$	$0.0292 \pm 0.0114$	$0.0048 \pm 0.0030$	$0.0376 \pm 0.0155$
Norm. bias	$-0.0005 \pm 0.0003$	$0.0005 \pm 0.0010$	$0.0001 \pm 0.0005$	$0.0068 \pm 0.0029$	$0.0002 \pm 0.0006$	$0.0094 \pm 0.0039$	$0.0003 \pm 0.0009$	$0.0122 \pm 0.0053$
FR15	$99.46 \pm 0.08$	$98.96 \pm 0.15$	$99.24 \pm 0.15$	$98.31 \pm 0.44$	$99.19 \pm 0.12$	$97.43 \pm 0.97$	$98.99 \pm 0.24$	$96.27 \pm 1.53$
FR05	$97.81 \pm 0.27$	$94.48 \pm 0.78$	$96.59 \pm 0.65$	$88.73 \pm 2.30$	$96.14 \pm 0.66$	$86.59 \pm 3.29$	$95.25 \pm 1.35$	$84.99 \pm 3.33$

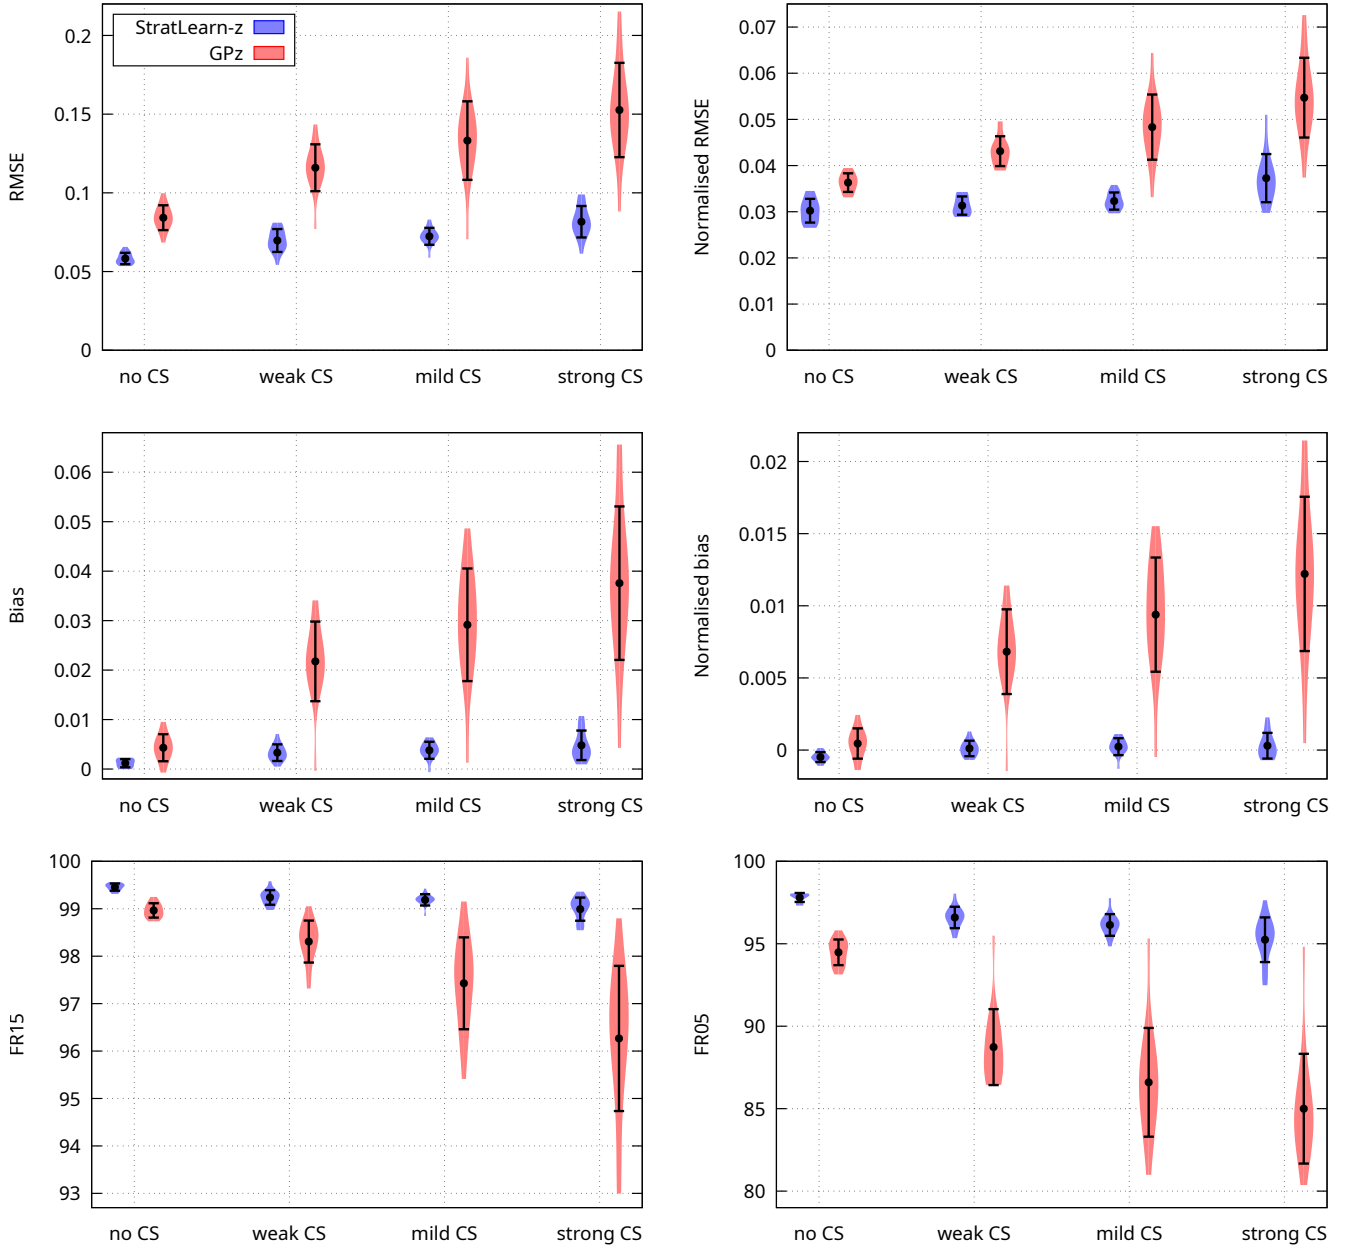


FIG. 8.— Violin plots showing the distribution of the metrics (averaged over 15 data realisations) for the four covariate shift (CS) scenarios. *StratLearn-z* results are shown in blue, while GPz results are shown in red. Black points with errorbars represent the mean and standard deviation for each case. There is a clear trend in the GPz results, which are strongly degraded by increasing covariate shift. On the other hand, the *StratLearn-z* results are only slightly impacted.

they are nearly doubled when going from the case with no covariate shift to the case with strong covariate shift, they are only marginally impacted by the presence of covariate shift in the *StratLearn-z* results. Turning to the mean error, or bias (central left panel of Fig. 8), the impact of covariate shift on *StratLearn-z* leads to a mean error that is four times larger in the strong covariate shift scenario with respect to the no covariate shift case, while for GPz the increase is almost one order of magnitude. Moreover, we note that the averaged normalised error for *StratLearn-z* is always consistent with zero, except for the no covariate shift case where it is slightly lower. Concerning the FR metric we find again that with *StratLearn-z* FR is consistent across covariate shift scenarios, with the FR15 (FR05) being only 0.5% (2.5%) worse in the strong covariate shift case with respect to the no covariate shift case. For GPz on the other hand FR15 decreases from  $\sim 99$  in the no covariate shift case to  $\sim 96$  in the strong covariate shift case. Similarly, the FR05 ranges from  $\sim 98$  to  $\sim 95$  for *StratLearn-z*, and from  $\sim 95.5$  to  $\sim 85$  for GPz.

To evaluate the predicted conditional redshift distributions, Fig. 9 plots the histogram of PIT obtained from *StratLearn-z* and GPz for the four scenarios, focusing on one specific realisation. We checked that the PIT histograms are not significantly different across the 15 realisations. There is an accumulation of mass near 0.5 in the PIT distribution obtained from *StratLearn-z*, indicating an over-estimation of uncertainty in the conditional density estimates. This may prove to be an issue when high-precision estimates are needed, however, we note that the peak is centered at 0.5 and the distributions are symmetric, which suggests that most of the *StratLearn-z* conditional density estimates are well-centered around the true redshift values. Similar considerations hold for all covariate shift cases, highlighting once again the robustness of *StratLearn-z* results in presence of unrepresentative source datasets. On the other hand, the PIT distributions obtained from GPz are significantly flatter, but progressively less symmetric the stronger the covariate shift.

The mode in the PIT distribution of *StratLearn-z* indicated that our conditional redshift distribution estimates are too broad, i.e., too conservative (see also Appendix C for a toy model demonstration). To better understand this result, we inspected the PIT distribution obtained by replacing our conditional density estimators with GPz, i.e., by running GPz on the stratified data. The resulting PIT distributions are very similar to those obtained from GPz run on the unstratified data (i.e., the distribution shown in the right panel of Fig. 9), suggesting that the mode in the PIT obtained from *StratLearn-z* can be attributed to the conditional density estimators rather than the stratification. In future work, a possible improvement of the overall *StratLearn-z* pipeline could be to add GPz to our framework as an additional estimator to potentially flatten the PITs. We note that the covariate shift scenarios considered here are relatively simple and performance as measured by PIT may vary with more realistic covariate shift scenarios.

Tab. 3 reports the median CRPS results for all four scenarios, both for *StratLearn-z* and GPz, averaged over all 15 realizations. GPz slightly outperforms *StratLearn-z* in all cases considered, which is to be expected given that GPz performs better than *StratLearn-z* as measured by PIT. However, we expect that the ability of *StratLearn-z* to better predict the mode of conditional redshift distributions (as reflected by its better performance in terms of the point estimate metrics reported in Table 2 and Fig. 8) should also be reflected in the CRPS, particularly at higher redshifts. To investigate the relative per-

TABLE 3

Average of the median CRPS and its standard deviation over 15 realisations for the four covariate shift (CS) scenarios considered, comparing *StratLearn-z* and GPz (a smaller value indicates better performance).

	<i>StratLearn-z</i>	GPz
no CS	$0.0079 \pm 0.0001$	$0.0069 \pm 0.0005$
weak CS	$0.0101 \pm 0.0006$	$0.0092 \pm 0.0012$
mild CS	$0.0105 \pm 0.0007$	$0.0093 \pm 0.0013$
strong CS	$0.0112 \pm 0.0010$	$0.0092 \pm 0.0011$

formance of the two methods in different redshift ranges, we computed the CRPS in ten redshift bins by taking the median CRPS in each bin (for each realisation), and then averaging over the realisations in each bin. The results are shown in Fig. 10 for the two extreme cases of no covariate shift (left) and strong covariate shift (right). Both cases show similar behavior: in low redshift bins GPz outperforms *StratLearn-z* due to its tighter variance in the predicted redshift distributions. However, at higher redshifts *StratLearn-z* consistently achieves a much lower CRPS, thanks to its improved prediction of the mode. The lower overall (i.e., across all redshifts) median CRPS of GPz is thus explained as the result of there being many more lower redshift objects (where GPz outperforms *StratLearn-z*) than higher redshift ones (where *StratLearn-z* vastly outperforms GPs). We also note that the CRPS obtained from *StratLearn-z* is roughly consistent across all redshifts, and shows a reduced dependence on the specific realisation than GPz, as demonstrated by *StratLearn-z*'s smaller variance across realizations.

## 5. CONCLUSIONS

We assess the performance of *StratLearn-z*, an implementation of the *StratLearn* method proposed in Autenrieth et al. (2024b) aimed at photo-*z* estimation. The method relies on stratification of the source and target datasets based on estimated propensity scores, and is designed to improve the performance of ML algorithms for non-representative training datasets, i.e., in presence of covariate shift. We apply *StratLearn-z* to a simulated dataset of 100,000 galaxies (DeRose et al. 2019), where we enforce the presence of covariate shift by performing rejection sampling on the *r*-band magnitudes, following a similar approach as that of Freeman et al. (2017): this prescription is designed to approximately mimic observational bias in spectroscopic (source) datasets, which are biased towards brighter objects with respect to photometric (target) dataset. We vary the parameters involved in the rejection sampling procedure, thus obtaining four different source-target scenarios, with increasing levels of covariate shift. We then estimate propensity scores using logistic regression, and fit two full conditional density estimators to the resulting five strata, combining them via a weighted average. With our trained models we finally obtain predictions for the conditional redshift distributions of objects in the target datasets, which includes brighter and higher redshift objects with respect to the source set.

To evaluate the performance of *StratLearn-z*, we compute a number of metrics on the target predictions, and benchmark against results obtained from the GPz code (Almosallam et al. 2016a,b), using the version released with Stylianou et al. (2022). To obtain more robust estimates for the metrics and investigate their scatter, we randomly permute the full dataset and perform rejection sampling and model training 15 times, thus obtaining 15 sets of photo-*z* predictions (for each covari-




FIG. 9.— PIT distribution for *StratLearn-z* (left) and GPz (right) for all cases, as detailed in the legend: in blue the no covariate shift (CS) case, in green the weak covariate shift case, in orange the mild covariate shift case and in purple the strong covariate shift case.

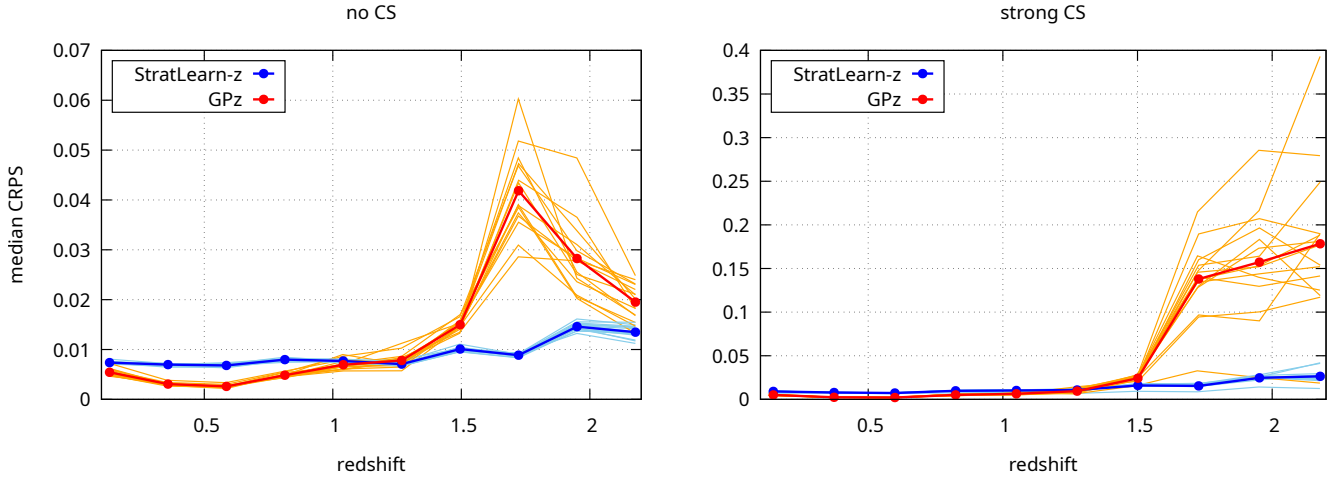


FIG. 10.— Binned CRPS as a function of redshift for no covariate shift (left) and strong covariate shift (right). Thin orange and light blue lines show the results for each realisation for GPz and *StratLearn-z* respectively, while thick lines with points represent the average over all realisations.

ate shift scenario) from which we compute averaged metrics. We find *StratLearn-z* outperforms GPz for all metrics considered and for all degrees of covariate shift. Additionally, the *StratLearn-z* results are only marginally impacted by the unrepresentative training data, while GPz results are significantly degraded.

To evaluate the quality of the predicted conditional redshift distribution we plot the histogram of the PIT, for each covariate shift scenario and for both codes. While *StratLearn-z* overestimates the spread of the conditional redshift distribution, its PIT distribution shows a reduced fraction of catastrophic outliers and symmetry, with a large bump around 0.5, i.e., the *StratLearn-z* conditional densities appear well-centered around the true redshift values. Despite the overestimation of photo- $z$  pdf spread, we note that the PIT distribution for *StratLearn-z* is also consistent across covariate shift scenarios, highlighting again the robustness of the method. The GPz histograms are, however, more uniform than the *StratLearn-z* histograms, which suggests that incorporating GPz or other conditional density estimators within our *StratLearn-z* framework could lead to improved results. Additionally, we assessed the conditional redshift distributions via the CRPS metrics,

which is informative on both on the variance of the redshift pdf and its mode with respect to the true redshift  $z_{\text{spec}}$ . We find that, when comparing median CRPS computed over the whole redshift range, GPz slightly outperforms *StratLearn-z* for all covariate shift scenarios. However, when comparing the CRPS binned as a function of redshift, we showed that *StratLearn-z* consistently achieves lower values in the high redshift range, indicating once again that an optimal combination of conditional density estimators with our propensity score-based stratification procedure can further improve these results.

The fundamental assumption of this work is that the covariate shift condition holds true, i.e., that there are no unmeasured (confounding) covariates that are predictive for redshift, and also associated to the selection process into source/target set. Covariate shift is often assumed when estimating photo- $z$ , e.g. (Izbicki et al. 2017; Freeman et al. 2017), but may not be completely accurate for real data. This is because the probability of obtaining a spectroscopic redshift may in fact depend on additional factors such as the type of object (e.g., passive galaxies are less likely to deliver a secure redshift compared to emission-line galaxies). Moreover, certain redshift ranges

(such as the so-called ‘redshift desert’  $1.4 \leq z \leq 2.5$ ) consistently lack reference redshifts, which occurs when the [O II] emission line shifts out of the optical range at approximately  $0.9 \mu\text{m}$ , leaving a gap until strong UV lines redshift into the optical domain. Furthermore, spectroscopic redshifts measurements can in some cases deviate substantially from the true values, a situation known as ‘redshift failure’ (Hartley et al. 2020). To minimize the impact on subsequent analyses of poor quality redshifts, quality flags based on spectral features (e.g., S/N of emission and absorption lines, and the strength of the  $4000 \text{ \AA}$  break) are used to identify and exclude unreliable redshifts (Hartley et al. 2020). Of course, such spectral features are only observed for spectroscopic confirmed galaxies and usually not available for galaxies in the photometric set, so they cannot be used for stratification. As discussed in Autenrieth et al. (2024a), the covariate shift condition remains valid if the spectroscopic covariates used for quality flagging do not provide additional information about redshift beyond what is already captured by the photometric covariates. Since the detailed procedure of obtaining secure redshift, as well as the definition of quality cuts are survey-specific, the definition and application of a general approach to this question is challenging. We will investigate potential violations of the covariate shift condition as well as potential correction strategies in a dedicated future study, focusing on topical and realistic survey selection procedures, e.g., as applied for Euclid (Laureijs et al. 2011; Euclid Collaboration et al. 2024).

In summary, we found the propensity-score based stratification procedure implemented in *StratLearn-z* provides stable predictions even in the most extreme scenario considered, in particular for high redshift regions where source/training data are lacking. Our analysis shows promising results, specifically concerning applications of the method to current and upcoming galaxy surveys, that will have to face the problem posed by incomplete and/or unrepresentative spectroscopic source sets to maximise the exploitation of photometric data.

**Acknowledgements:** CM’s work is supported by the Fondazione ICSC, Spoke 3 Astrophysics and Cosmos Observa-

tions, National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) Project ID CN\_00000013 “Italian Research Center on High-Performance Computing, Big Data and Quantum Computing” funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di “campioni nazionali di R&S (M4C2-19)” - Next Generation EU (NGEU). RT acknowledges co-funding from Next Generation EU, in the context of the National Recovery and Resilience Plan, Investment PE1 – Project FAIR “Future Artificial Intelligence Research”. This resource was co-financed by the Next Generation EU [DM 1555 del 11.10.22]. RT is partially supported by the Fondazione ICSC, Spoke 3 “Astrophysics and Cosmos Observations”, Piano Nazionale di Ripresa e Resilienza Project ID CN00000013 “Italian Research Center on High-Performance Computing, Big Data and Quantum Computing” funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di “campioni nazionali di R&S (M4C2-19)” - Next Generation EU (NGEU). DvD and MA acknowledge partial support from the UK Engineering and Physical Sciences Research Council [EP/W015080/1, EP/W522673/1] and from the Marie-Skłodowska-Curie RISE [H2020-MSCA-RISE-2019-873089] Grant provided by the European Commission. MA also acknowledges partial support from the European Union’s Horizon 2020 research and innovation programme under European Research Council Grant Agreement No 101002652 (PI K. Mandel) and Marie Skłodowska-Curie Grant Agreement No 873089. RT and AM have been partially supported by the PRO3 Scuole Programme ‘DS4ASTRO’. This work used the Julia language and in particular the *DataFrames.jl*, *CSV.jl*, *Statistics.jl*, *Distances*, *GLM.jl* packages, and the *Gnuplot.jl* package for Julia to generate the plots.

#### DATA AVAILABILITY

The original dataset used in this paper can be found at [https://github.com/jfcrenshaw/pzflow/blob/main/pzflow/example\\_files/galaxy-data.pkl](https://github.com/jfcrenshaw/pzflow/blob/main/pzflow/example_files/galaxy-data.pkl). We publicly release the version of *StratLearn* used for this work with example datasets – one for each covariate shift case considered – at [github.com/chiamoretti/StratLearn-z](https://github.com/chiamoretti/StratLearn-z)

#### REFERENCES

- Abbott T. M. C., et al., 2018, *ApJS*, **239**, 18  
 Aihara H., et al., 2018, *PASJ*, **70**, S8  
 Akeson R., et al., 2019, *arXiv e-prints*, p. [arXiv:1902.05569](https://arxiv.org/abs/1902.05569)  
 Almosallam I. A., Lindsay S. N., Jarvis M. J., Roberts S. J., 2016a, *MNRAS*, **455**, 2387  
 Almosallam I. A., Jarvis M. J., Roberts S. J., 2016b, *MNRAS*, **462**, 726  
 Amara A., Réfrégier A., 2007, *MNRAS*, **381**, 1018  
 Arnouts S., Cristiani S., Moscardini L., Matarrese S., Lucchin F., Fontana A., Giallongo E., 1999, *MNRAS*, **310**, 540  
 Autenrieth M., Wright A. H., Trotta R., van Dyk D. A., Stenning D. C., Joachimi B., 2024a, *arXiv e-prints*, p. [arXiv:2401.04687](https://arxiv.org/abs/2401.04687)  
 Autenrieth M., van Dyk D. A., Trotta R., Stenning D. C., 2024b, *Statistical Analysis and Data Mining: The ASA Data Science Journal*, **17**, e11643  
 Benítez N., 2000, *ApJ*, **536**, 571  
 Bolzonella M., Miralles J. M., Pelló R., 2000, *A&A*, **363**, 476  
 Bordoloi R., Lilly S. J., Amara A., 2010, *MNRAS*, **406**, 881  
 Brammer G. B., van Dokkum P. G., Coppi P., 2008, *ApJ*, **686**, 1503  
 Brescia M., Cavuoti S., Razim O., Amaro V., Riccio G., Longo G., 2021, *Frontiers in Astronomy and Space Sciences*, **8**, 70  
 Carliles S., Budavári T., Heinis S., Priebe C., Szalay A. S., 2010, *ApJ*, **712**, 511  
 Carrasco Kind M., Brunner R. J., 2013, *MNRAS*, **432**, 1483  
 Carrasco Kind M., Brunner R. J., 2014, *MNRAS*, **438**, 3409  
 Cochran W. G., 1968, *Biometrics*, pp 295–313  
 Collister A. A., Lahav O., 2004, *PASP*, **116**, 345  
 Crenshaw J. F., Kalmbach J. B., Gagliano A., Yan Z., Connolly A. J., Malz A. I., Schmidt S. J., The LSST Dark Energy Science Collaboration 2024, *AJ*, **168**, 80  
 D’Isanto A., Polsterer K. L., 2018, *A&A*, **609**, A111  
 Dawid A. P., 1984, *Journal of the Royal Statistical Society: Series A (General)*, **147**, 278  
 DeRose J., et al., 2019, *arXiv e-prints*, p. [arXiv:1901.02401](https://arxiv.org/abs/1901.02401)  
 Euclid Collaboration et al., 2020, *A&A*, **644**, A31  
 Euclid Collaboration et al., 2024, *arXiv e-prints*, p. [arXiv:2405.13491](https://arxiv.org/abs/2405.13491)  
 Firth A. E., Lahav O., Somerville R. S., 2003, *MNRAS*, **339**, 1195  
 Freeman P. E., Izbicki R., Lee A. B., 2017, *MNRAS*, **468**, 4556  
 Graff P., Feroz F., Hobson M. P., Lasenby A., 2014, *MNRAS*, **441**, 1741  
 Graham M. L., Connolly A. J., Ivezić Ž., Schmidt S. J., Jones R. L., Jurić M., Daniel S. F., Yoachim P., 2018, *AJ*, **155**, 1  
 Hartley W. G., et al., 2020, *MNRAS*, **496**, 4769  
 Henghes B., Thiyaalingam J., Pettitt C., Hey T., Lahav O., 2022, *MNRAS*, **512**, 1696  
 Hersbach H., 2000, *Weather and Forecasting*, **15**, 559  
 Heymans C., et al., 2021, *A&A*, **646**, A140  
 Hildebrandt H., et al., 2010, *A&A*, **523**, A31  
 Hildebrandt H., et al., 2021, *A&A*, **647**, A124  
 Ilbert O., et al., 2006, *A&A*, **457**, 841  
 Ivezić Ž., et al., 2019, *ApJ*, **873**, 111  
 Izbicki R., Lee A. B., Freeman P. E., 2016, *arXiv e-prints*, p. [arXiv:1604.01339](https://arxiv.org/abs/1604.01339)  
 Izbicki R., Lee A. B., Freeman P. E., et al., 2017, *The Annals of Applied Statistics*, **11**, 698

Laureijs R., et al., 2011, *arXiv e-prints*, p. [arXiv:1110.3193](#)  
 Ma Z., Hu W., Huterer D., 2006, *ApJ*, **636**, 21  
 Malmquist K. G., 1922, *Meddelanden fran Lunds Astronomiska Observatorium Serie I*, **100**, 1  
 Malmquist K. G., 1925, *Meddelanden fran Lunds Astronomiska Observatorium Serie I*, **106**, 1  
 Masters D., et al., 2015, *ApJ*, **813**, 53  
 Moreno-Torres J. G., Raeder T., Alaiz-Rodríguez R., Chawla N. V., Herrera F., 2012, *Pattern Recognition*, **45**, 521  
 Moskowitz I., Gawiser E., Crenshaw J. F., Andrews B. H., Schmidt S., The LSST Dark Energy Science Collaboration 2024, *arXiv e-prints*, p. [arXiv:2402.15551](#)  
 Myles J., et al., 2021, *MNRAS*, **505**, 4249  
 Newman J. A., Gruen D., 2022, *ARA&A*, **60**, 363  
 Pasquet J., Bertin E., Treyer M., Arnouts S., Fouchez D., 2019, *A&A*, **621**, A26  
 Polsterer K. L., D’Isanto A., Gieseke F., 2016, *arXiv e-prints*, p. [arXiv:1608.08016](#)

Porredon A., et al., 2022, *Phys. Rev. D*, **106**, 103530  
 Rau M. M., et al., 2023, *MNRAS*, **524**, 5109  
 Rosenbaum P. R., Rubin D. B., 1983, *Biometrika*, **70**, 41  
 Sadeh I., Abdalla F. B., Lahav O., 2016, *PASP*, **128**, 104502  
 Salvato M., Ilbert O., Hoyle B., 2019, *Nature Astronomy*, **3**, 212  
 Sánchez C., et al., 2014, *MNRAS*, **445**, 1482  
 Schmidt S. J., et al., 2020, *MNRAS*, **499**, 1587  
 Stylianos N., Malz A. I., Hatfield P., Crenshaw J. F., Gschwend J., 2022, *PASP*, **134**, 044501  
 Sugiyama S., et al., 2023, *Phys. Rev. D*, **108**, 123521  
 Tanaka M., et al., 2018, *PASJ*, **70**, S9  
 The LSST Dark Energy Science Collaboration et al., 2018, *arXiv e-prints*, p. [arXiv:1809.01669](#)  
 Toribio San Cipriano L., et al., 2023, *arXiv e-prints*, p. [arXiv:2312.09721](#)  
 Tosone F., Cagliari M. S., Guzzo L., Granett B. R., Crespi A., 2023, *A&A*, **672**, A150  
 Tutusaus I., et al., 2020, *A&A*, **643**, A70




FIG. 11.— Source and target normalised distributions of  $u$ -band magnitudes for the four covariate shift (CS) cases, as stated in each panel. Orange histograms represent the source dataset, while in blue we plot the target datasets.




FIG. 12.— Same as Fig. 11 but for the  $g$ -band magnitudes.

## APPENDIX

### A. DISTRIBUTIONS FOR ALL PHOTOMETRIC BANDS

As mentioned in Sec. 3.1, the introduction of covariate shift by rejection sampling on the  $r$ -band induces a shift also in the redshift distributions for source and target. Naturally, the effect is not limited to redshift, but also extends to other photometric bands. We plot here the distributions for all photometric bands (except for the  $r$ -band, already shown in Fig. 1), specifically Fig. 11 for the  $u$ -band, Fig. 12 for the  $g$ -band, Fig. 13 for the  $i$ -band, Fig. 14 for the  $z$ -band and Fig. 15 for the  $y$ -band. As in the main text, orange histograms represent the source dataset, while we show in blue the target dataset.

### B. STRATIFICATION RESULTS

We show here the effectiveness of the stratification based on propensity scores in removing covariate shift within each stratum for the redshift distributions. We focus on the worst case scenario, i.e., the strong covariate shift case. Each panel of Fig. 16 refers to one of the five strata, except for the bottom right panel which shows the redshift distributions for the full source and target sets. Blue histograms refer to the source set, while orange histograms refer to the target set.




FIG. 13.— Same as Fig. 11 but for the  $i$ -band magnitudes.




FIG. 14.— Same as Fig. 11 but for the  $z$ -band magnitudes.

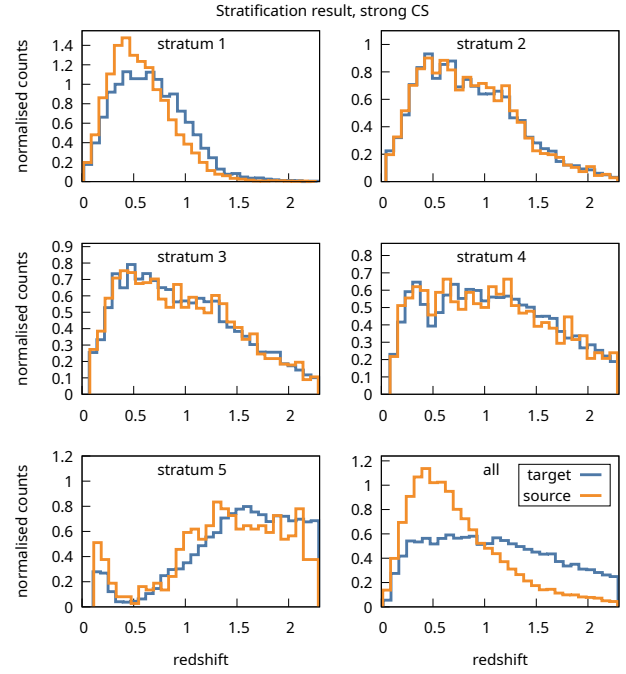


FIG. 16.— Same as Fig. 3 but for the redshift distributions. We note that this diagnostic is not possible in practice with real source/target data, as  $z$  is not observed in the target dataset.




FIG. 15.— Same as Fig. 11 but for the  $y$ -band magnitudes.

### C. TOY MODEL FOR PIT

In this section, we illustrate a simple toy example to demonstrate the effect on the PIT histograms when overestimating (or underestimating) the uncertainties/errorbars of (conditional) density estimates.

More precisely, we simulate data from  $X_i \stackrel{\text{i.i.d.}}{\sim} N(\mu = 0, \sigma = 1)$ , with  $i = 1, \dots, n$ , obtaining  $n = 10,000$  independent samples. In a first step, we calculate the PIT assuming the correct model, i.e.,  $\text{PIT}_i = F_{X_i}(X_i)$ , where  $F_{X_i}(X_i)$  is the CDF of  $X_i$  evaluated at  $X_i$ . As expected, the resulting histogram exhibits a standard uniform distribution, as illustrated in Fig. 17 (left panel).

In a next step, we keep the same data simulating process, however, we assume a model with much larger uncertainties (centered at the correct location). More precisely, we calculate the PITs assuming a much broader Gaussian distribution, i.e., we calculate  $\text{PIT}_i = F_{X'_i}(X_i)$ , where  $F_{X'_i}(X_i)$  is the CDF of a random variable  $X'_i \sim N(\mu = 0, \sigma = 2)$ , evaluated at  $X_i$ . Plotting the histogram of PIT values in Fig. 17 (center panel), we find a large (symmetric) bump around 0.5, with less mass towards the edges zero and one.

Third, we assume a model that underestimates the uncertainties, i.e., we calculate the PITs via  $\text{PIT}_i = F_{X''_i}(X_i)$ , where  $F_{X''_i}(X_i)$  is the CDF of a random variable  $X''_i \sim N(\mu = 0, \sigma = 0.8)$ , evaluated at  $X_i$ . As illustrated in right panel of Fig. 17, this leads to large peaks at the edges (zero and one), indicating a large number of true values (from  $X_i$ ) laying outside the

assumed model distributions. (We note that similar peaks at the edges can be obtained by systematically shifting the center of the assumed model distribution away from the center of the underlying data-generating distributions).

We note that center panel of Fig. 17 exhibits a similar pattern as the PIT histograms obtained for the *StratLearn-z* conditional density estimates in Fig. 9. The symmetric bump in the PIT histogram around 0.5 indicates that a large proportion of the conditional density estimates are well centered around the true values (as also suggested by the low RMSE and low mean error of *StratLearn-z*, as presented in Table 2). However, there is an overestimation of uncertainties/errorbars in the conditional density estimates (overly conservative estimates). On the contrary, underestimating uncertainties leads to large peaks of the PIT histograms at the edges (zero and one) indicating larger fractions of (catastrophic) outliers as it has been observed for the GPz estimates. Finally, while in the optimal case the PIT histograms exhibit a flat uniform distribution (as illustrated in left panel of Fig. 17), we argue that an overly conservative estimate of the uncertainties (with estimates well-centered around the true values) is more desirable than underestimating the uncertainties, which might lead to catastrophic outliers and a disproportionate number of estimates which do not cover the true values.

This paper was built using the Open Journal of Astrophysics L<sup>A</sup>T<sub>E</sub>X template. The OJA is a journal which provides fast and easy peer review for new papers in the `astro-ph` section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at <http://astro.theoj.org>.




FIG. 17.— PIT distributions illustrating the toy examples introduced in Appendix C. Left panel: correct model; center panel: overestimated uncertainties; right panel: underestimated uncertainties.