The Overlooked Potential of Generalized Linear Models in Astronomy-II: Gamma regression and photometric redshifts

J. Elliott^a, R. S. de Souza^b, A. Krone-Martins^c, E. Cameron^d, E. E. O. Ishida^e, J. Hilbe^{f,g}, for the COIN collaboration

^aMax-Planck-Institut für extraterrestrische Physik, Giessenbachstraße 1, 85748, Garching, Germany

^bMTA Eötvös University, EIRSA “Lendulet” Astrophysics Research Group, Budapest 1117, Hungary

^cSIM, Faculdade de Ciências, Universidade de Lisboa, Ed. C8, Campo Grande, 1749-016, Lisboa, Portugal

^dDepartment of Zoology, University of Oxford, Tinbergen Building, South Parks Road, Oxford, OX1 3PS, United Kingdom

^eMax-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany

^fArizona State University, 873701, Tempe, AZ 85287-3701, USA

^gJet Propulsion Laboratory, 4800 Oak Grove Dr., Pasadena, CA 91109, USA

Abstract

Machine learning techniques offer a precious tool box for use within astronomy to solve problems involving so-called *big data*. They provide a means to make accurate predictions about a particular system without prior knowledge of the underlying physical processes of the data. In this article, and the companion papers of this series, we present the set of Generalized Linear Models (GLMs) as a fast alternative method for tackling general astronomical problems, including the ones related to the machine learning paradigm. To demonstrate the applicability of GLMs to inherently positive and continuous physical observables, we explore their use in estimating the photometric redshifts of galaxies from their multi-wavelength photometry. Using the gamma family with a log link function we predict redshifts from the *PHoto-z Accuracy Testing* simulated catalogue and a subset of the *Sloan Digital Sky Survey* from Data Release 10. We obtain fits that result in catastrophic outlier rates as low as $\sim 1\%$ for simulated and $\sim 2\%$ for real data. Moreover, we can easily obtain such levels of precision within a matter of seconds on a normal desktop computer and with training sets that contain merely thousands of galaxies. Our software is made publicly available as a user-friendly package developed in Python, R and via an interactive web application. This software allows users to apply a set of GLMs to their own photometric catalogues and generates publication quality plots with minimum effort. By facilitating their ease of use to the astronomical community, this paper series aims to make GLMs widely known and to encourage their implementation in future large-scale projects, such as the Large Synoptic Survey Telescope.

Keywords: techniques: photometric – methods: statistical – methods: analytical – galaxies: distances and redshifts

1. Introduction

Generalized Linear Models (GLMs), as introduced by Nelder and Wedderburn (1972), offer a well es-

tablished statistical framework for robust modelling and prediction making. It allows the application of regression analysis when the observed quantity originates from an *exponential family* distribution rather than a Gaussian (or Normal; e.g., Hardin and Hilbe, 2012; Hilbe, 2014). As a result, GLMs offer a readily interpretable and physically-motivated approach (via family distributions) to machine learning (ML) that can be applied to a variety of astronomical data sets.

Email addresses: jonnyelliott@mpe.mpg.de (J. Elliott), rafael.2706@gmail.com (R. S. de Souza), algol@sim.ul.pt (A. Krone-Martins), dr.ewan.cameron@gmail.com (E. Cameron), emille@mpa-garching.mpg.de (E. E. O. Ishida), j.m.hilbe@gmail.com (J. Hilbe)

Despite being widely used across a range of scientific disciplines, such as biology (Brown et al., 1993; Ahrestani et al., 2013), medicine (Lindsey, 1999), and economics (Pindyck and Rubinfeld, 1998; de Jong and Heller, 2008), and its availability within the overwhelming majority of contemporary statistical software packages (e.g., R, R Core Team 2014; SAS, Inc. 2003; and STATA, StataCorp 2009), GLMs remain almost *terra incognita* within the astronomical community (de Souza et al., 2014a).

One particular problem which presents itself as a candidate for the implementation of GLMs is the photometric redshift (photo- z) estimation of galaxies. Although precise redshifts can in principle be directly determined through identification of known absorption or emission lines in the optical and/or near-infrared spectrum of each target galaxy, the observational cost of this procedure can quickly become prohibitive for large scale surveys. The only feasible alternative in such cases is to use available multi-wavelength photometry to infer approximate photo- z s instead, but this is not always a simple task.

There exist a plethora of different spectra emitted from galaxies throughout the Universe. Their characteristic features carry signatures from the galaxy’s morphology, age, metallicity, star formation history, merging history, and a host of other confounding factors in addition to its redshift. Thus, making photo- z estimation a far from trivial task. There exist several techniques which are commonly used to estimate redshifts from photometry and can be divided in to: (i) template fitting techniques (e.g., Benítez, 2000; Bolzonella et al., 2000; Ilbert et al., 2006), and (ii) ML (or empirical) techniques (e.g. Connolly et al., 1995; Collister and Lahav, 2004; Wadadekar, 2005; Miles et al., 2007; O’Mill et al., 2011; Reis et al., 2012; Krone-Martins et al., 2014). In template fitting techniques, a set of synthetic spectra are determined from synthesised stellar population models for a given set of metallicities, star formation histories and initial mass functions, among other properties. The photo- z is calculated by determining the synthetic photometry (and thus spectral template and redshift) which best fits the photometric observations. ML techniques, on the other hand, usually require a data set with spectroscopically measured redshifts to train the chosen method.

Many studies have examined the individual advantages of each photo- z code (for a glimpse on the diversity of existent methods, see Hildebrandt et al., 2010; Abdalla et al., 2011; Zheng and Zhang, 2012; Sánchez et al., 2014, and references therein). Abdalla et al. 2011 investigated the differences between five commonly used template fitting codes and a neural network. The neural network proved to be more reliable in redshift ranges with a higher density of training data, while the template fitting methods depended heavily on the underlying templates. Despite these caveats, the overall performance of all codes were, to first order, consistent and displayed catastrophic errors ranging from 5 – 9%, which is considered good in terms of photo- z estimates (Abdalla et al., 2011). More recently, methods which combine several photo- z techniques in a Bayesian approach, coined ensemble learning, have begun to be implemented with the hope that they can complement each other’s drawbacks (Carrasco Kind and Brunner, 2014).

One of the largest practical difficulties for the current photo- z methods is the time necessary to either fit the templates or train the underlying ML method; on top of that, the required size of the training set is often highly influential for empirical methods (Firth et al., 2003). *Big data* catalogues expected from large sky surveys, like the *Large Synoptic Survey Telescope*¹ (LSST Science Collaboration et al., 2009), *EUCLID*² (Refregier et al., 2010) or the *Wide-Field Survey Infrared Telescope*³ (Green et al., 2012), warrant the need for fast and reliable photo- z methods that are capable of processing such large volumes of data in minutes to days rather than years, thereby facilitating higher level analyses and model refinements for downstream data products.

In this work, we introduce a new technique based on robust principal component analysis (PCA) and GLMs to estimate photo- z s. The method runs in a matter of seconds on a single core computer, even for millions of objects. In addition, we achieved very low levels of catastrophic errors when using training sets of a few thousands objects. The combination of short computational run time, moderate training set size,

¹<http://www.lsst.org/lsst>

²<http://sci.esa.int/euclid>

³<http://wfirst.gsfc.nasa.gov>

and small catastrophic errors makes GLMs a robust and implementable technique for future large scale surveys.

The outline of this article is as follows. In §2, we give a broad overview of GLMs, in §3 we provide a description of the data set utilised. The methodology implemented is outlined in §4. We then present our results and compare with the recent literature in §5 and summarise our conclusions in §6.

2. Overview of Regression Methods

Before we delve into the details of GLMs and the gamma family, we make a brief overview of linear regression, a common tool used within astrophysics. Afterwards, we explicitly outline the details of GLMs with the gamma family and explain how it can be applied to determine photo-zs for a particular data set.

2.1. Overview of Linear Regression

Consider a given data set containing N (distinct objects; e.g., galaxies),

$$\mathcal{D} = \{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\},$$

where the x_i are observations of the independent variable, X , and the y_i are observed values of a dependent random variable (RV), Y , which is a function of X , $Y = f(X)$. Traditionally, X is called the *explanatory variable* and Y the *response variable*. The expected value and variance of Y are denoted by $E(Y)$ and $\text{var}(Y)$, respectively. In this context, a linear model describes the response variable (Y) as a linear function of the explanatory variable (X):

$$Y = \beta_0 + \beta_1 X + \epsilon = \eta + \epsilon, \quad (1)$$

where $\{\beta_0, \beta_1\}$ are scalars called *slope coefficients* or *covariates*, $\eta = \beta_0 + \beta_1 X$ is the linear component (or *predictor*) of this simple model. Finally, ϵ is an error term considered to be independent and identically distributed, $\epsilon \sim N(0, \sigma^2)$.

When a standard linear regression approach is applied, the linear predictor in equation 1 is assumed to fully describe the response variable. The measured values are used to determine the covariates of the linear predictor that uniquely identify a straight line through the chosen data set minimising the error term.

Having the scalar coefficients determined, the model provides a direct relation between X and Y , allowing one to predict the mean value of Y for a given measurement of X .

In order to clarify the procedure described in the next subsections, we invite the reader to approach this simple linear regression problem from an alternative perspective. Consider now each measurement, $\{x_i, y_i\}$, as a realisation of different variables $\{X_i, Y_i\}$ from a common family of probability density functions (PDFs), but with distinct parameters μ_i for each index i . The underlying PDF driving the behaviour of the response variable (Y_i) will be denoted by $f(y_i; \kappa_i)$, where κ_i is the parameter vector of the PDF underlying the i -th measurement. If Y_i follows a Normal PDF with mean μ_i and variance σ_i^2 , then

$$f(y_i; \kappa_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left[-\frac{1}{2} \frac{(y_i - \mu_i)^2}{\sigma_i^2}\right], \quad (2)$$

where $\kappa_i = \{\mu_i, \sigma_i\}$. This is summarised as $Y_i \sim N(\mu_i, \sigma_i)$. For reasons which will be clarified later, we consider σ_i a fixed value and, thus, determining μ_i is enough to completely characterise $f(y_i; \kappa_i)$. In this context, we can relate the measured x_i to the expected value of the corresponding response variable, $y_i = E(Y_i)$, though the slope coefficients⁴:

$$y_i = E(Y_i) = \mu_i = \mathbf{x}_i^T \boldsymbol{\beta}, \quad (3)$$

with $\mathbf{x}_i = \{1, x_i\}$ and $\boldsymbol{\beta} = \{\beta_0, \beta_1\}$. We can now use the chosen PDF family (e.g., equation 2) and the observed data $\mathcal{D}$ in order to find values of $\boldsymbol{\beta}$ which better describes the data. In this case, the model is composed of two main ingredients:

1. a PDF underlying the behaviour of each response variable, $f(y_i; \mu_i)$ and
2. a relation connecting $E(Y_i)$ with measured values of the explanatory variable, also called the *link function*, $g(E(Y_i)) = \mu_i$.

Once the values of parameter $\boldsymbol{\beta}$ are determined, we can use the inverse link, g^{-1} , to calculate the expected

⁴If our model is correct, for each measurement the expected value should correspond on average to the measured one.

value for the response variable given a measured input x .

Note, that in this simple example, although the mathematical expressions in both approaches are the same, their interpretations are different. In the standard analysis, the best-fitted parameters define completely the connection between X and Y . In the second approach, they characterize a linear relationship between measurements of X_i and a parameter which uniquely identifies a PDF underlying each response variable Y_i , κ_i . This approach allows us to extend the same reasoning to situations where a linear relation is not a good description of the process driving the data.

This example can be generalised for the case with more than one explanatory variable, forming a multiple linear model. The *general linear model*⁵ goes one step further, including situations where there are more than one response variable. In such models the errors are considered to be uncorrelated and follow a multivariate normal distribution. Although a very useful tool, general linear models are not suited for situations where there are restrictions on Y (e.g., binary, count or strictly positive data) or when the variance depends on the mean. The GLMs are a generalisation of this framework, capable of handling both scenarios.

2.2. Generalized Linear Models

This subsection presents an overview of a vast subject. For a detailed theoretical review see Dobson (2002) and Hardin and Hilbe (2012).

The example from the previous section clearly shows that the PDF is a key ingredient in the construction of a GLM framework. Moreover, the procedure relies on one main feature of the PDF: within the chosen family, a distribution should be uniquely identified through one single parameter μ (called *location* or *mean*). Determining this parameter is the ultimate goal of the GLM methodology.

In order to fulfil this requirement, GLMs are constructed for any distribution that belongs to the *exponential family* of distributions (Gaussian/normal, gamma, inverse Gaussian, Bernoulli, binomial, Poisson, and negative binomial). The PDF for any mem-

ber of this family can be written as

$$f(y; \theta) = s(y)t(\theta)e^{a(y)b(\theta)}, \quad (4)$$

where s , t , a and b are known functions and θ is the *canonical parameter*. The reader may also find them as,

$$f(y; \theta) = \exp [a(y)b(\theta) + c(\theta) + d(y)], \quad (5)$$

with $c(\theta) = \ln[t(\theta)]$ and $d(y) = \ln[s(y)]$. When $a(y) = y$, the distribution is said to be in the *canonical form*. *Natural exponential distributions* are exponential ones in the canonical form with $b(\theta) = \theta$. If there is an extra parameter ϕ it is considered known⁶. Natural exponential distributions can always take the form

$$f(y; \theta, \phi) = \exp \left\{ \frac{y\theta - A(\theta)}{B(\phi)} + C(y; \phi) \right\} \quad (6)$$

(see Hardin and Hilbe, 2012, section 2.3). The above expression contains $A(\theta)$, which is called the *cumulant*, $B(\phi)$ the *scale parameter*, ϕ the *dispersion parameter*, and $C(\cdot; \cdot)$, the normalisation term that scales the integral to unity. For these distributions,

$$\begin{aligned} E(Y) &= A'(\theta), \\ \text{var}(Y) &= A''(\theta)B(\phi). \end{aligned} \quad (7)$$

In the specific normal distribution example described previously, it is simple to show that $E(y) = \mu$ and $\text{var}(Y) = \sigma^2$.

Two other important ingredients in the structure of a GLM model are the linear component η and the link function $g(\cdot)$. In the simple linear regression case, η was equivalent to our assumption of a linear relation between the explanatory and response variables and g was merely the identity function ($g(\mu) = \mu$). However, when Y_i relates to X_i through a non-linear expression, η will play the important role of linearising the connection between X_i and $E(Y_i)$. In other words, even if we are dealing with non-linear data, we can still define the link function as

$$g(\mu) = \eta = \mathbf{X}^T \boldsymbol{\beta}, \quad (8)$$

⁵Not to be confused with Generalized Linear Models (GLMs).

⁶In the Gaussian example described previously, $\phi = \sigma^2 = \text{var}(Y)$.

given that the PDF belongs to the one parameter exponential family. After the determination of β , the inverse link function is used to determine μ and from this we know that g must be invertible. For PDFs in the canonical form the *canonical link function* is given by $b(\theta)$ (equation 5).

In summary, all GLMs share a similar structure and are characterized by:

- A random response component whose mean μ is to be estimated. The response variable, Y , is assumed to be theoretically derived as a random sample of an underlying single parameter PDF belonging to the GLM family of distributions. The goal of modelling Y is to find an unbiased estimate of the mean parameter which better describes the data.
- A systematic (or linear) component, η , built from the explanatory variables, $\mathbf{X}$ (sometimes called *covariates*), and their associated slope coefficients, β . Their multiplication produces a linear predictor for each observation (equation 8).
- A link function, $g(\cdot)$, which defines how the mean is associated with the explanatory variable. The link function linearises the relationship between the mean response and predictors (X_i). Once the slope coefficients are determined, we are able to use the inverse link and the observed X_i to estimate the mean, i.e.,

$$\mu = g^{-1}(\eta). \quad (9)$$

- Conversion of the PDF to a log-likelihood function for the observed data, which is used as the basis for the determination of the slope coefficients.

2.3. Gamma Family and Regression

The gamma distribution is characterised by the response variable, Y , taking only positive real values. It is optimal when fitting positive-only values with a shape determined by its estimated parameter. As a single parameter model the GLM gamma model is limited to a specific set of values. If a second scale

parameter is employed, the range of shapes allowed by the model are greatly enlarged. This is not a GLM model though. On the other hand, if the GLM gamma model does fit a given data situation, the model is more efficient and easier to interpret.

In this study we predict the photo- z of a galaxy from multi-wavelength photometry and compare it to the measured spectroscopic redshift, z_{spec} . As redshift is always positive and continuous, we can apply a gamma family distribution, which in its exponential family form can be expressed as (Hardin and Hilbe, 2001):

$$\begin{aligned} f(y; \mu, \phi) &= \frac{1}{y\Gamma(1/\phi)} \left(\frac{y}{\mu\phi}\right)^{1/\phi} \exp\left(-\frac{y}{\mu\phi}\right) \\ &= \exp\left\{\frac{y/\mu - (-\ln \mu)}{-\phi} + C(y; \phi)\right\}, \end{aligned} \quad (10)$$

where

$$C(y, \phi) = \frac{1 - \phi}{\phi} \ln y - \frac{\ln \phi}{\phi} - \ln \Gamma\left(\frac{1}{\phi}\right). \quad (11)$$

From equations 6 and 10, we recognise the canonical parameter as $\theta = 1/\mu$, the cumulant $A(\theta) = -\ln \mu$ and the scale parameter $B(\phi) = -\phi$. Consequently, from equation 7, $E(y) = \mu$, and $\text{var}(y) = \sigma^2 = \mu^2 \phi$. It is important to note that the GLM gamma model describes a response variable as a distribution with constant coefficient of variation, σ/μ (standard deviation / mean). Thus, determining μ is enough to univocally identify a distribution.

A useful feature of GLMs is that a link function can be assigned to the estimation algorithm, and it does not only have to be a natural or canonical link. For example, the log link, $g(\mu) = \ln(\mu)$, is commonly associated with the gamma model. When this is employed in place of the natural link function, the associated inverse link, $\mu = \exp(\eta)$, ensures a positive mean for any η . Unless the mean and data are inversely related statisticians typically use the log link for this continuous response model.

This is also the approach chosen for this work. In what follows, we use the log-link function $\mu^T = \log(\beta^T \mathbf{X})$, where $\mathbf{X}$ is a $m \times n$ matrix containing m magnitudes for each of the n galaxies and β is the covariate column vector. The covariates are estimated

by maximising the log-likelihood of the regression model (Eqn. 10) utilising iteratively re-weighted least-squares (see Hardin and Hilbe, 2001, and references therein).

In addition, to avoid numerical instabilities and identifiability issues it is preferable for the predictor variables, $\mathbf{X}$, not to exhibit strong correlations (i.e., *multicollinearity*). This is not necessarily the case for the magnitudes of galaxies that can be strongly correlated across different broadband filters. As a precaution against multicollinearity we carry out *Principal Component Analysis* and adopt the principal components (PCs) of the observed magnitude set as our explanatory variables⁷. Beyond ensuring non-correlated features, using the PCs also optimises the use of computational resources, as the calculation time required increases non-linearly with the number of explanatory variables. As a dimensionality reduction technique, PCA allows a robust way to increase the execution speed of the redshift estimation.

As a final remark, we emphasise that the key to understanding a statistical model is to determine how well the predicted values fit the observed ones. Each type of model, being based on a specific probability distribution, or mixture of distributions, is limited to a specific range of predicted values. That is, the distributional assumptions of a gamma distribution determine the range and shape of predicted values it can have. If the observed response values differ greatly from the possible predicted values, the model cannot fit the data well. In such a case it is important to find another model more suited to fitting the observed values. In the next sections we quantify the ability of the gamma family GLM in predicting photo- z s of galaxies by comparison with simulations and measured spectroscopic redshifts.

3. Data

To compare GLMs with other methods and to better put our results in to context, we adopted a publicly available data set that was previously submitted to

different photo- z codes. The *PHoto- z Accuracy Testing* (PHAT) was an international initiative to identify the most promising photo- z methods and guide future improvements. Two observational photometric catalogues were provided: PHAT0 with simulations, and PHAT1 with real observations. A total of 17 photo- z codes were submitted. As a direct comparison using PHAT1 is not possible, since the answers of the challenge are not openly available, we applied GLMs to PHAT0 and compare its results to those reported by Hildebrandt et al. (2010). PHAT0 has 169,520 simulated galaxies with redshifts ranging from $z = 0.02 - 2.24$, and magnitudes in 11 filters ($u, g, r, i, z, Y, J, H, K, IRAC1$, and $IRAC2$).

In addition, we apply the same technique to a real data set obtained from the *Sloan Digital Sky Survey* (SDSS; York et al., 2000). We obtain a galaxy sample using the same selection criteria and SQL query (Appendix C) outlined in Krone-Martins et al. 2014 from the most recent Data Release 10 (Ahn et al., 2014). This results in a sample of 1,347,640 galaxies with a redshift range of $z = 0 - 1.0$, with magnitudes in 5 filters (u', g', r', i' , and z'). To compare the same data set, but with dereddened magnitudes, we also use the SQL query outlined in Carrasco Kind and Brunner 2014.

4. Methodology

On the basis of making GLM modelling easily accessible to the community, we have developed a set of photo- z packages which can be used on any multi-wavelength data set (de Souza et al., 2014b). The codes are written in both R⁸ (Appendix A), a programming language commonly used in the statistical sciences, and Python⁹ (Appendix B) which is becomingly widely adopted in the astronomical community. In addition, we have implemented a web application using the Shiny¹⁰ (Appendix A.1) platform whereby users can upload their data set and have the photo- z s and diagnostic plots delivered. We adopted the following step-by-step methodology to determine the photo- z s of a sample of galaxies using their multi-wavelength photometry.

⁷For a general review on PCA, see Jolliffe (2002). For examples of PCA use in astronomy, see also e.g. Conselice (2006); Ishida et al. (2011); Ishida and de Souza (2011, 2013); Jeon-Daniel et al. (2011); Krone-Martins and Moitinho (2014).

⁸www.r-project.org/

⁹www.python.org/

¹⁰shiny.rstudio.com/

Figure 1: The 2D probability density of the predicted redshift from the GLM fit vs. the spectroscopic redshift (central plots). The upper and right subplots in each panel depict the redshift distribution along photo- z and z_{spec} , respectively. We note that this is a randomised subsample amounting to 10,000 galaxies. **Top:** Results for the PHAT0 sample (green). **Bottom:** Results from the SDSS sample (purple).

1. The data was randomly split into training and test sets with the training sample holding at least 10% of the number of galaxies (see section 5.3 for a detailed analysis on the influence of this choice in our final results.).
2. Robust principal component analysis (e.g., Candès et al. 2011; de Souza et al. 2014c) was carried out on the complete data set, training and test, to ensure the PCs are not dominated by one of the two samples. We note that this corresponds to a semi-supervised technique since, that data without any measured redshifts can help determine the PCs (see, e.g., Shah et al., 2008). The threshold on cumulative percentage of total variance was set to $\sim 99.5\%$ in order to determine the number of PCs to be used with the GLM.
3. We utilised a gamma family distribution to reflect the fact that measured redshifts are positive and continuous. The relationship between the redshift and the explanatory variables (our linear predictor; equation 8) took the R formula form of

$$z_{\text{phot}} \sim PC_1^2 + PC_1 * PC_2 * \dots * PC_n + C, \quad (12)$$

where n is the number of principal components and C is a constant. R formula of this kind are called *simple factorial model formulae*, where the $*$ is a *crossing* operator, which allows the inclusion of interaction terms. For example, $A * B = A + B + A \cdot B$ (for a full description see, Nelder, 1965). This formula was fit using iteratively re-weighted least squares. We note that tests using a simple formulation without interaction terms are inadequate in capturing the complexity of the data.

4. The predicted photo- z for the test data was calculated using the principal component projections of the test data set and the best-fit GLM using the training sample.
5. To measure how well the photo- z s were estimated, we employed a metric commonly used in the literature, specifically, the catastrophic error (or outlier rate/fraction). We applied two definitions of this metric: one used for the PHAT0 sample by Hildebrandt et al. 2010, and the second used more

commonly in the literature and the PHAT1 sample by Hildebrandt et al. 2010. We define them respectively as:

$$(a) \quad O^{(a)} = \Delta z^{(a)} > 0.10, \quad (13)$$

where $\Delta z^{(a)} = |z_{\text{phot}} - z_{\text{spec}}|$, and

$$(b) \quad O^{(b)} = \Delta z^{(b)} > 0.10, \quad (14)$$

where $\Delta z^{(b)} = \frac{|z_{\text{phot}} - z_{\text{spec}}|}{1 + z_{\text{spec}}}$.

5. Results

5.1. Simulation

The PHAT0 data set was fit using the gamma family with a log link. We split the data into training and testing with a ratio of 1:9, respectively. This criterion resulted in a training size of 33,904 galaxies and a testing set of 135,616 galaxies. Decomposition by PCA showed that at least 6 components needed to be used to ensure that 99.5% of the variance was retained. The best-fit GLM is shown in Fig. 1 (top panel) which presents a catastrophic error rate of $O^{(a)} = 4.4\%$. The time taken to fit amounted to 170 seconds on an AMD Athlon X2 Dual-Core QL-64 processor with 1.7 GB RAM on the Ubuntu 10.04 operating system, which represents an old laptop at today's standards. We note that changing either the training set size or the number of PCs does not dramatically alter the solution. For example, using 8 PCs (99.95% variance) results in $O^{(a)} = 2.5\%$ catastrophic errors determined in 1200 seconds or 10 PCs results in $O^{(a)} = 1.4\%$ in 4948 seconds.

The central plot in Fig. 1 (top panel) shows z_{spec} compared with the photo- z calculated from our GLM gamma model, while the top and right plots represent their individual distributions for the entire redshift range. Numerical diagnostics for such results are displayed in Table 1. Fig. 2 (left panel) details the redshift distribution of photo- z calculated for PHAT0 in bins of $\Delta z_{\text{spec}} = 0.1$. There is a tendency for a larger bias to be present in higher redshifts ranges, reflecting the characteristics of the training set. In these regions, not only does the data quality decrease significantly, but there is also an observational effect that favours brighter galaxies. Thus, the presence of such biases are not completely unexpected. However, investigating if there is a preference for under/over estimation of photo- z within the GLM framework

would require a k -fold cross-validation analysis which is out of the scope of this work. We do highlight though, that such a study is crucial for a potential user that aims to optimise their photo- z results.

We note that our results give a slightly larger catastrophic error than compared with the literature. However, we highlight that the simulated data set was created from template techniques and, as such, a lower catastrophic error would suggest that the technique has results that reflect those to template fitting rather than the absolute performance of the technique (Hildebrandt et al., 2010). Finally, after demonstrating that the GLM gamma model has competitive capabilities in comparison with other techniques, we investigated its absolute ability with observed data sets.

5.2. Real Data

The SDSS data set was randomly split into training and test sets. The size of the training set was selected such that the catastrophic error plateaued in value. Initial tests show that the catastrophic errors become stable for training sets that contain more than 500 galaxies and plateaus at around 4,000 galaxies. We adopt 10,000 galaxies for the training data set. The PCA decomposition requires only 4 PCAs to retain 99.5% of the variance. This combination of training sample and number of principal components results in a catastrophic error of $O^{(a)} = 8.087\%$ within 10 seconds. Including all 5 components and increasing to a training set size of 50,000 results in an improvement to $O^{(a)} = 6.439\%$. We outline a selection of other diagnostics in Table 1.

5.3. Caveats & Improvements

The use of a GLM to predict photo- z s can achieve competitive estimates with catastrophic errors of the order of 1 – 5%, comparable with current techniques. It is important to note that there is a wide range of choices that the user can make for their own data set to meet the needs of its use. The most dominant parameters are the size of the training set and the number of principal components used. As one would expect, the time taken increases when both of these parameters are increased and there will be some variation in the resulting catastrophic errors. Figures 4 and 5 depict the catastrophic error rate as a function of training set size and number of PCAs used, respectively.

Table 1: Diagnostic comparison of samples

Code	Type ^a	bias _z ^b	rms(Δz) ^c	Outlier rate %
PHAT0				
This work	Empirical	0.033	0.025	1.367 (10 PCs) 2.511 (8 PCs) 4.438 (6 PCs)
Le PHARE ^d	Template	0.000	0.010	0.044
Bayesian Photo- z 's ^d	Template	-0.005	0.011	0.026
Easy and Accurate Z_{phot} from Yale ^d	Template	-0.001	0.012	0.000
GALaxy EVolution and GAZELLE ^d	Template	0.000	0.014	0.053
GOODZ ^d	Template	0.000	0.012	0.018
HyperZ ^d	Template	-0.002	0.013	0.185
Low-Resolution Spectral Templates ^d	Template	0.000	0.011	0.026
Purger (template repair) ^d	Template	-0.005	0.011	0.053
Zurich Extragalactic Bayesian Redshift Analyzer ^d	Template	0.000	0.011	0.062
Zurich Extragalactic Bayesian Redshift Analyzer (modified) ^d	Template	-0.005	0.011	0.044
Artificial Neural Network photo- z ^d	Empirical	0.000	0.011	0.018
Boosted Decision Trees ^d	Empirical	-0.004	0.019	0.389
Purger (nearest-neighbour Fit) ^d	Empirical	0.000	0.017	0.053
Polynomial fitting ^d	Empirical	0.001	0.019	1.669
Regression Trees ^d	Empirical	0.000	0.013	0.010
Singal Neural Networks ^d	Empirical	-0.005	0.049	18.202
SDSS				
This work	Empirical	0.038	0.029	6.44 (5 PCs) 8.09 (4 PCs)
Eureqa (Symbolic Regression) ^e	Empirical	0.009	0.045	1 – 10
Eureqa (Symbolic Regression) ^f	Empirical	0.009	0.045	10 – 100
This work [†]	Empirical	0.030	0.034	2.25 (5 PCs) 2.90 (4 PCs)
Trees for Photo- z ^{g,†}	Empirical	0.019	0.014	0.78
Self-Organizing Maps Photo- z ^{g,†}	Empirical	0.020	0.015	0.70
Bayesian Photo- z 's ^{g,†}	Template	0.023	0.016	1.34
Weighted Average ^{g,†}	Hybrid	0.020	0.014	0.82
Weighted Average, oracle weighting scheme ^{g,†}	Hybrid	0.019	0.014	0.67
Weighted Average, shape weighting scheme ^{g,†}	Hybrid	0.019	0.014	0.81
Weighted Average, fit weighting scheme ^{g,†}	Hybrid	0.020	0.014	0.90
Bayesian Model Averaging ^{g,†}	Hybrid	0.018	0.013	0.60
Bayesian Model Combination ^{g,†}	Hybrid	0.018	0.013	0.59
Hierarchical Bayes ^{g,†}	Hybrid	0.020	0.014	0.84

Notes. All codes use the $O^{(a)}$ outlier fraction, unless a dagger ($\dagger$) is placed next to them, for which the $O^{(b)}$ definition is used. ^a Methods classification: *Empirical*: uses a machine learning technique that usually require supervised or unsupervised learning from a training data set, *Template*: uses synthetic galaxy spectra created from stellar population models, and *Hybrid*: uses a combination of Template and Empirical methods, applying different weighting schemes to aggregate their results. ^b Average of Δz without the outliers. ^c Root mean square of Δz without the outliers. ^d Diagnostics presented in Hildebrandt et al. (2010). ^e Krone-Martins et al. (2014) for $z < 0.7$. ^f Krone-Martins et al. (2014) for $z \geq 0.7$. ^g Fit diagnostics obtained from Carrasco Kind and Brunner (2014).

Figure 2: Violin plots that depict the probability density of photo- z errors per redshift bin, $\Delta z_{\text{spec}} = 0.1$, determined from the entire galaxy sample. In each violin the central dashed line is the median and the dotted lines are the 25% and 75% quartiles. **Left:** Results from PHAT0 sample (green). **Right:** Outcomes from SDSS sample (purple).

Figure 3: The probability density of the number of outliers. **Top:** Results from PHAT0 sample (green). **Bottom:** Results from the SDSS sample (purple).

In addition, as with all ML techniques, one must choose the type of formula used so as to ensure the training data is not over fitted. Increasing the order of the polynomial in equation 12 or including more cross terms will decrease the overall catastrophic error, but will also increase the running time. Such caveats are best left for the user, the catalogue in question, and the needs and requirements of the output photometry.

Another viable use of GLMs is to make extrapolated predictions of photo- z s of high-redshift galaxies by learning on low-redshift galaxies. Like other machine learning techniques, GLMs can over-fit the data and make incorrect predictions. However, GLMs have the advantage that they are not (and should not) be treated as black boxes, like many other empirical and template techniques, e.g. neural networks (Werner and Guven, 2007). Instead, the underlying assumptions can be modified, for example: the formula adopted (equation 12), and the underlying effect of the covariates (PCs). As already noted, the relevance of each covariate can be investigated to determine their importance in the model and reveal a physical interpretation. Even though in our analysis we have used PCA to uncorrelate the magnitudes, the resulting influence of each magnitude in the model (the physical interpretation) can still be obtained indirectly from the PCs using techniques such as *projection to latent structures* (Sasdehli et al., 2015). In combination with a cross-validation process, this would allow the user to optimise the GLM for their needs (e.g.

Figure 4: **Top:** Catastrophic error of photo- z vs. the size of the training set, for the PHAT0 data. **Bottom:** Time taken to complete the GLM fitting as a function of number of elements in the training set.

extrapolating predictions) and would highlight the relevance of the covariates in the determination of their wanted values. Furthermore, the extrapolated prediction will still have two inherent gamma features: (i) non-negative values and (ii) heteroscedasticity, consistent with what we expect from the galaxy photometric redshift distributions.

In this study, we have shown that gamma GLMs can speedily and efficiently compute photometric redshifts of galaxies. However, this initial case study could be built upon and improved. This is not within the scope of this paper, but we outline some techniques that could be used to make our methodology better. We emphasize that, although there are a variety of different techniques available to approach the photo- z problem, GLMs provide a straightforward and statistically coherent way of including the mod-

Figure 5: **Top:** Catastrophic error of the fitted photo- z vs. the number of principal components for the PHAT0 data, considering a fixed training set with 8000 galaxies. **Bottom:** Time taken to complete the GLM fitting as a function of the number of PCs used.

elling (prior) within the regression process. This is done through the choice of the family distribution. Traditional regression methods can sometimes yield negative, non-physical values, which we have avoided from the start by using the gamma family distributions. Thus, the enhancements cited below are all within this framework.

To determine how to further develop this methodology, we utilise a standard diagnostic, in statistics, called the Q-Q (quantile-quantile) plot. We compare the residuals of the best-fit gamma model of the PHAT0 data set to a gamma distribution, as seen in Fig. 6. When the two distributions are similar it is expected that a straight line will go through the data set, as is seen in our case (red line). As the best-fit line is shallower than a normal $y = x$ line, this implies that the theoretical gamma distribution is more dispersed than the sample PHAT0 distribution. There is a deviation at the lower quantiles (< 1) that suggests the residuals are deviating from a gamma. This skewing towards lower quantiles is alleviated when selecting galaxies that lie in the redshift range $0.25 < z < 1.0$. One way to address this problem would be to consider more complex models such as generalized additive models for location, scale and shape (GAMLSS; Stasinopoulos and Rigby, 2007). GAMLSS are more lenient and do not require an exponentiated family distribution to be used, but instead a general distribution family that can be highly skewed, kurtotic, and/or continuous and discrete. Investigation of such models is not within the scope of this paper, but are expected to be addressed in future versions of the code.

6. Conclusions

Generalized linear models are widely used through a multitude of academic disciplines, but have been relatively untouched throughout the astronomical community. Their straight forward implementation and possibility of allowing physically relevant predictions make GLMs a great candidate for competing with conventional methods of modelling that are more often turning to techniques involving ML, and more specifically neural networks. For the case of photometric redshifts, an adoption of a gamma family reflects two important characteristics of the data: (i)

Figure 6: A Q-Q plot of the residuals of the best fit gamma model compared to a gamma distribution. For visual reasons we have clipped one outlier at $y = 6.5$. See the text for an interpretation of the plot.

a non-negative and continuous measurement, and (ii) heteroscedasticity, i.e., the variance of the photo- z measurements changes according to the redshift. The gamma GLM intrinsically assumes that higher values of photo- z have a larger intrinsic scatter.

Upcoming wide field sky surveys, such as the LSST, will take the challenge of determining photo- z s to an unprecedented scale. To this end, we have outlined the use of GLMs to tackle the problem of estimating photo- z s for large samples of galaxies from their multi-band photometry in a semi-supervised learning manner. We demonstrate that GLMs can be trained on 50,000 galaxies with 5 principle components in 3 seconds and 10 principle components in 10 minutes, using a standard laptop (AMD Athlon X2 Dual-Core QL-64 processor with 1.7 GB RAM on the Ubuntu 10.04 OS), and can reach catastrophic errors of 1 – 5%, comparable with current techniques involving template fitting or ML.

In summary, GLMs offer a simple and efficient way of tackling many problems within astronomy that are usually computationally heavy or require large training samples. To promote their use within the astronomical community we have developed a suite of libraries and a web application to allow GLMs to be used to determine photo- z s from a user’s own galaxy sample with a simple click of a button.

Acknowledgments

We thank V. Busti, E. D. Feigelson, M. Killedar, J. Buchner, and A. Trindade for interesting suggestions and comments. JE, RSS and EEOI thank the SIM Laboratory of the *Universidade de Lisboa* for hospitality during the development of this work. Cosmostatistics Initiative (COIN)¹¹ is a non-profit organisation whose aim is to nourish the synergy between astrophysics, cosmology, statistics and machine learning communities. This work was partially supported by the ESA VA4D project (AO 1-6740/11/F/MOS). AKM thanks the Portuguese agency *Fundação para Ciência e Tecnologia – FCT*, for financial support (SFRH/BPD/74697/2010). EEOI is partially supported by the Brazilian agency CAPES (grant number 9229-13-2). Work on this paper has substantially benefited from using the collaborative website AWOB¹² developed and maintained by the Max-Planck Institute for Astrophysics and the Max-Planck Digital Library. This work was written on the collaborative WriteLatex platform¹³, and made use of the GitHub¹⁴ a web-based hosting service and git version control software. This work made use of the cloud based hosting platform ShinyApps.io¹⁵. This work used the following public scientific Python packages `scikit-learn v0.15`¹⁶ (Pedregosa et al., 2011), `seaborn v0.3.1`¹⁷, and `statsmodels v0.6.0`¹⁸. Funding for SDSS-III¹⁹ has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science.

References

Abdalla, F.B., Banerji, M., Lahav, O., Rashkov, V., 2011. A comparison of six photometric redshift methods applied to 1.5 million luminous red galaxies. *MNRAS* 417, 1891–1903. doi:10.1111/j.1365-2966.2011.19375.x, arXiv:0812.3831.

¹¹<https://asaip.psu.edu/organizations/iaa/iaa-working-group-of-cosmostatistics>

¹²<http://awob.mpg.de>

¹³www.writelatex.com

¹⁴www.github.com

¹⁵<http://www.shinyapps.io>

¹⁶github.com/scikit-learn

¹⁷github.com/mwaskom/seaborn

¹⁸github.com/statsmodels

¹⁹<http://www.sdss.org>

- Ahn, C.P., Alexandroff, R., Allende Prieto, C., Anders, F., Anderson, S.F., Anderton, T., Andrews, B.H., Aubourg, É., Bailey, S., Bastien, F.A., et al., 2014. The Tenth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the SDSS-III Apache Point Observatory Galactic Evolution Experiment. *ApJS* 211, 17. doi:10.1088/0067-0049/211/2/17, arXiv:1307.7735.
- Ahrestani, F.S., Hebblewhite, M., Post, E., 2013. The importance of observation versus process error in analyses of global ungulate populations. *Sci Rep* 3, 3125. URL: <http://www.biomedsearch.com/nih/importance-observation-versus-process-error/24201239.html>.
- Benítez, N., 2000. Bayesian Photometric Redshift Estimation. *ApJ* 536, 571–583. doi:10.1086/308947, arXiv:arXiv:astro-ph/9811189.
- Bolzonella, M., Miralles, J.M., Pelló, R., 2000. Photometric redshifts based on standard SED fitting procedures. *A&A* 363, 476–492. arXiv:arXiv:astro-ph/0003380.
- Brown, D., Rothery, P., et al., 1993. Models in biology: mathematics, statistics and computing. John Wiley & Sons Ltd.
- Candès, E.J., Li, X., Ma, Y., Wright, J., 2011. Robust principal component analysis? *J. ACM* 58, 11:1–11:37. URL: <http://doi.acm.org/10.1145/1970392.1970395>, doi:10.1145/1970392.1970395.
- Carrasco Kind, M., Brunner, R.J., 2014. Exhausting the information: novel Bayesian combination of photometric redshift PDFs. *MNRAS* 442, 3380–3399. doi:10.1093/mnras/stu1098, arXiv:1403.0044.
- Collister, A.A., Lahav, O., 2004. ANNz: Estimating Photometric Redshifts Using Artificial Neural Networks. *PASP* 116, 345–351. doi:10.1086/383254, arXiv:arXiv:astro-ph/0311058.
- Connolly, A.J., Csabai, I., Szalay, A.S., Koo, D.C., Kron, R.G., Munn, J.A., 1995. Slicing Through Multicolor Space: Galaxy Redshifts from Broadband Photometry. *AJ* 110, 2655. doi:10.1086/117720, arXiv:arXiv:astro-ph/9508100.
- Conselice, C.J., 2006. The fundamental properties of galaxies and a new galaxy classification system. *MNRAS* 373, 1389–1408. doi:10.1111/j.1365-2966.2006.11114.x, arXiv:astro-ph/0610016.
- de Souza, R.S., Cameron, E., Killedar, M., Hilbe, J., Vilalta, R., Maio, U., Biffi, V., Ciardi, B., Riggs, J.D., 2014a. The Overlooked Potential of Generalized Linear Models in Astronomy - I: Binomial Regression and Numerical Simulations. arxiv:1409.7696 arXiv:1409.7696.
- de Souza, R.S., Elliott, J., A., K., Ishida, E.E.O., Hilbe, J., Cameron, E., 2014b. CosmoPhotoz: Photometric redshift estimation using generalized linear models. *Astrophysics Source Code Library* arXiv:1408.018.
- de Souza, R.S., Maio, U., Biffi, V., Ciardi, B., 2014c. Robust PCA and MIC statistics of baryons in early minihaloes. *MNRAS* 440, 240–248. doi:10.1093/mnras/stu274, arXiv:1308.6009.
- Dobson, A.J., 2002. An Introduction to Generalized Linear Models. Second ed., Chapman and Hall/CRC. URL: <http://>

- [//www.worldcat.org/isbn/1584881658](http://www.worldcat.org/isbn/1584881658).
- Firth, A.E., Lahav, O., Somerville, R.S., 2003. Estimating photometric redshifts with artificial neural networks. *MNRAS* 339, 1195–1202. doi:10.1046/j.1365-8711.2003.06271.x, arXiv:astro-ph/0203250.
- Green, J., Schechter, P., Baltay, C., Bean, R., Bennett, D., Brown, R., Conselice, C., Donahue, M., et al., 2012. Wide-Field InfraRed Survey Telescope (WFIRST) Final Report. arxiv:1208.4012 arXiv:1208.4012.
- Hardin, J.W., Hilbe, J., 2001. Generalized Linear Models and Extensions. College Station, Texas: Stata Press.
- Hardin, J.W., Hilbe, J.M., 2012. Generalized Linear Models and Extensions, 3rd Edition. 3rd ed., StataCorp LP.
- Hilbe, J., 2014. Modeling Count Data. Cambridge University Press.
- Hildebrandt, H., Arnouts, S., Capak, P., Moustakas, L.A., Wolf, C., Abdalla, F.B., Assef, R.J., Banerji, M., et al., 2010. PHAT: PHoto-z Accuracy Testing. *A&A* 523, A31. doi:10.1051/0004-6361/201014885, arXiv:1008.0658.
- Ilbert, O., Arnouts, S., McCracken, H.J., Bolzonella, M., Bertin, E., Le Fèvre, O., Mellier, Y., Zamorani, G., et al., 2006. Accurate photometric redshifts for the CFHT legacy survey calibrated using the VIMOS VLT deep survey. *A&A* 457, 841–856. doi:10.1051/0004-6361:20065138, arXiv:arXiv:astro-ph/0603217.
- Inc., S.I., 2003. SAS/STAT Software, Version 9.1. Cary, NC. URL: <http://www.sas.com/>.
- Ishida, E.E.O., de Souza, R.S., 2011. Hubble parameter reconstruction from a principal component analysis: minimizing the bias. *A&A* 527, A49. doi:10.1051/0004-6361/201015281, arXiv:1012.5335.
- Ishida, E.E.O., de Souza, R.S., 2013. Kernel PCA for Type Ia supernovae photometric classification. *MNRAS* 430, 509–532. doi:10.1093/mnras/sts650, arXiv:1201.6676.
- Ishida, E.E.O., de Souza, R.S., Ferrara, A., 2011. Probing cosmic star formation up to $z=9.4$ with gamma-ray bursts. *MNRAS* 418, 500–504. doi:10.1111/j.1365-2966.2011.19501.x, arXiv:1106.1745.
- Jeerson-Daniel, A., Dalla Vecchia, C., Haas, M.R., Schaye, J., 2011. The correlation structure of dark matter halo properties. *MNRAS* 415, L69–L73. doi:10.1111/j.1745-3933.2011.01081.x, arXiv:1103.5467.
- Jolliffe, I.T., 2002. Principal Component Analysis. Springer-Verlag. URL: <http://www.amazon.com/Principal-Component-Analysis-I-T-Jolliffe/dp/0387954422>.
- de Jong, P., Heller, G.Z., 2008. Generalized Linear Models for Insurance Data. Cambridge University Press. URL: <http://dx.doi.org/10.1017/CB09780511755408>. cambridge Books Online.
- Krone-Martins, A., Ishida, E.E.O., de Souza, R.S., 2014. The first analytical expression to estimate photometric redshifts suggested by a machine. *MNRAS* 443, L34–L38. doi:10.1093/mnrasl/slu067, arXiv:1308.4145.
- Krone-Martins, A., Moitinho, A., 2014. UPMASK: unsupervised photometric membership assignment in stellar clusters. *A&A* 561, A57. doi:10.1051/0004-6361/201321143, arXiv:1309.4471.
- Lindsey, J.K., 1999. A review of some extensions to generalized linear models. *Statistics in medicine* 18, 2223–2236. URL: <http://view.ncbi.nlm.nih.gov/pubmed/10474135>.
- LSST Science Collaboration, Abell, P.A., Allison, J., Anderson, S.F., Andrew, J.R., Angel, J.R.P., Armus, L., Arnett, D., Asztalos, S.J., Axelrod, T.S., et al., 2009. LSST Science Book, Version 2.0. arxiv:0912.0201 arXiv:0912.0201.
- Miles, N., Freitas, A., Serjeant, S., 2007. Estimating Photometric Redshifts Using Genetic Algorithms, in: Ellis, R., Allen, T., Tuson, A. (Eds.), Applications and Innovations in Intelligent Systems XIV. Springer London, pp. 75–87. URL: <http://dx.doi.org/10.1007/978-1-84628-666-76>, doi:10.1007/978-1-84628-666-76.
- Nelder, J.A., 1965. Analysis of randomized experiments with orthogonal block structure .I. Block structure and null analysis of variance. *Proceedings of the Royal Society of London Series A-Mathematical and Physical Sciences* 283, 147. doi:10.1098/rspa.1965.0012.
- Nelder, J.A., Wedderburn, R.W.M., 1972. Generalized linear models. *Journal of the Royal Statistical Society, Series A, General* 135, 370–384.
- O’Mill, A.L., Duplancic, F., García Lambas, D., Sodré, Jr., L., 2011. Photometric redshifts and k-corrections for the Sloan Digital Sky Survey Data Release 7. *MNRAS* 413, 1395–1408. doi:10.1111/j.1365-2966.2011.18222.x, arXiv:1012.3752.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- Pindyck, R.S., Rubinfeld, D.L., 1998. Econometric models and economic forecasts. volume 4. Irwin/McGraw-Hill Boston.
- R Core Team, 2014. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing. Vienna, Austria. URL: <http://www.R-project.org/>.
- Refregier, A., Amara, A., Kitching, T.D., Rassat, A., Scaramella, R., Weller, J., Euclid Imaging Consortium, f.t., 2010. Euclid Imaging Consortium Science Book. arxiv:1001.0061 arXiv:1001.0061.
- Reis, R.R.R., Soares-Santos, M., Annis, J., Dodelson, S., Hao, J., Johnston, D., Kubo, J., Lin, H., Seo, H.J., Simet, M., 2012. The Sloan Digital Sky Survey Co-add: A Galaxy Photometric Redshift Catalog. *ApJ* 747, 59. doi:10.1088/0004-637X/747/1/59, arXiv:1111.6620.
- Sánchez, C., Carrasco Kind, M., Lin, H., Miquel, R., Abdalla, F.B., Amara, A., Banerji, M., Bonnett, C., Brunner, R., Capozzi, D., Carnero, A., Castander, F.J., da Costa, L.A.N., Cunha, C., Fausti, A., Gerdes, D., Greisel, N., Gschwend, J., Hartley, W., Jouvel, S., Lahav, O., Lima, M., Maia, M.A.G., Martí, P., Ogando, R.L.C., Ostrovski, F., Pellegrini, P., Rau, M.M., Sadeh, I., Seitz, S., Sevilla-Noarbe, I., Sypniewski, A., de Vicente, J., Abbot, T., Allam, S.S., Atlee, D., Bernstein,

G., Bernstein, J.P., Buckley-Geer, E., Burke, D., Childress, M.J., Davis, T., DePoy, D.L., Dey, A., Desai, S., Diehl, H.T., Doel, P., Estrada, J., Evrard, A., Fernández, E., Finley, D., Flaugh, B., Frieman, J., Gaztanaga, E., Glazebrook, K., Honscheid, K., Kim, A., Kuehn, K., Kuropatkin, N., Lidman, C., Makler, M., Marshall, J.L., Nichol, R.C., Roodman, A., Sánchez, E., Santiago, B.X., Sako, M., Scalzo, R., Smith, R.C., Swanson, M.E.C., Tarle, G., Thomas, D., Tucker, D.L., Uddin, S.A., Valdés, F., Walker, A., Yuan, F., Zuntz, J., 2014. Photometric redshift analysis in the Dark Energy Survey Science Verification data. *MNRAS* 445, 1482–1506. doi:10.1093/mnras/stu1836, arXiv:1406.4407.

Sasdeselli, M., Hillebrandt, W., Aldering, G., Antilogus, P., Aragon, C., Bailey, S., Baltay, C., Benitez-Herrera, S., Bongard, S., Buton, C., Canto, A., Cellier-Holzem, F., Chen, J., Childress, M., Chotard, N., Copin, Y., Fakhouri, H.K., Feindt, U., Fink, M., Fleury, M., Fouchez, D., Gangler, E., Guy, J., Ishida, E.E.O., Kim, A.G., Kowalski, M., Kromer, M., Lombardo, S., Mazzali, P.A., Nordin, J., Pain, R., Pécontal, E., Pereira, R., Perlmutter, S., Rabinowitz, D., Rigault, M., Runge, K., Saunders, C., Scalzo, R., Smadja, G., Suzuki, N., Tao, C., Taubenberger, S., Thomas, R.C., Tilquin, A., Weaver, B.A., 2015. A metric space for Type Ia supernova spectra. *MNRAS* 447, 1247–1266. doi:10.1093/mnras/stu2416, arXiv:1411.4424.

Shah, A.R., Oehmen, C.S., Webb-Robertson, B.J., 2008. Svm-hustle—an iterative semi-supervised machine learning approach for pairwise protein remote homology detection. *Bioinformatics* 24, 783–790.

Stasinopoulos, D.M., Rigby, R.A., 2007. Generalized additive models for location scale and shape (GAMLSS) in R. *Journal of Statistical Software* 23.

StataCorp, 2009. Stata Statistical Software: Release 11. StataCorp LP. College Station, TX. URL: <http://www.stata.com>.

Wadadekar, Y., 2005. Estimating Photometric Redshifts Using Support Vector Machines. *PASP* 117, 79–85. doi:10.1086/427710, arXiv:arXiv:astro-ph/0412005.

Werner, G., Guven, S., 2007. Glm basic modeling: Avoiding common pitfalls, in: *Casualty Actuarial Society Forum*, pp. 257–272.

York, D.G., Adelman, J., Anderson, Jr., J.E., Anderson, S.F., Annis, J., Bahcall, N.A., Bakken, J.A., Barkhouser, R., et al., SDSS Collaboration, 2000. The Sloan Digital Sky Survey: Technical Summary. *AJ* 120, 1579–1587. doi:10.1086/301513, arXiv:astro-ph/0006396.

Zheng, H., Zhang, Y., 2012. Review of techniques for photometric redshift estimation, in: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 10. doi:10.1117/12.925314.

Appendix A. R package

The R package is publicly available. It can be obtained either through CRAN (Comprehensive R Archive Network) using the package name *CosmoPhotoz*,

or through the COIN GitHub repository²⁰. The stable CRAN release can be easily installed from within R using the standard function `install.packages()`, while the GitHub version can be installed via the function `install_github()` from the package *devtools*.

There are two ways to perform the photometric redshift estimation using the R package. The simplest, but less flexible, way is to perform a direct call to the `CosmoPhotoZestimator()` function using two data frames: one containing the data adopted for training, and another containing the photometric data for photometric redshift estimation. The code below shows how to perform the redshift estimate using the PHAT0 data included in the package.

```
library(CosmoPhotoz)
data(PHAT0train)#Data for training
data(PHAT0test)#Data for estimation
# Run the analysis
photoZest <- CosmoPhotoZestimator(←
  PHAT0train,
  PHAT0test, numberOfPcs=6)
# Not using robust PCA is considerably←
  faster,
# but the results are worse
photoZestN <- CosmoPhotoZestimator(←
  PHAT0train,
  PHAT0test, numberOfPcs=6,
  robust=FALSE)
# Create a boxplot showing the results
plotDiagPhotoZ(photoz = photoZest,
  specz = PHAT0test$redshift,
  type = "box")
```

The most flexible way to use the package, however, is to perform a step by step analysis using the individual functions provided. The following code exemplifies how it is possible to perform the redshift estimate using some of such functions.

```
library(CosmoPhotoz)
data(PHAT0train)# Data for training
data(PHAT0test)# Data for estimation

# Combine the training and test data ←
  and
# calculate the principal components
PC_comb <- computeCombPCA(
  subset(trainData,select=c(-redshift)),
```

²⁰<https://github.com/COINtoolbox/COSMOPhotoz/>


```

9 subset(testData, select=c(-redshift)),
10 robust=TRUE)
11 Trainpc <- cbind(PC_comb$x,
12 redshift=trainData$redshift)
13 Testpc <- PC_comb$y
14
15 # Formula based on the PCs
16 formM <- redshift~poly(Comp.1,2)*
17 poly(Comp.2,2)*Comp.3*Comp.4*
18 Comp.5*Comp.6
19
20 # GLM fitting
21 Fit <- glmTrainPhotoZ(Trainpc, formula←
22 =formM,
23 method="Bayesian", family="gamma")
24
25 # Photo-z estimation
26 photoZtest <- glmPredictPhotoZ(Fit$←
27 glmfit, newdata=Testpc,
28 type="response")
29
30 # Print Photo-z estimation
31 print(photoZtest$photoz)
32
33 # Estimate confidence intervals errors
34 photoz_temp <- predict(Fit$glmfit, ←
35 newdata=Testpc, type="link", se.←
36 fit = TRUE)
37 photoz <- photoz_temp$fit
38
39 critval <- 1.96 ## approx 95% ←
40 Confidence Interval
41 upr <- photoz_temp$fit + (critval * ←
42 photoz_temp$se.fit)
43 lwr <- photoz_temp$fit - (critval * ←
44 photoz_temp$se.fit)
45 fit <- photoz_temp$fit
46
47 fit2 <- Fit$glmfit$family$linkinv(fit)
48 upr2 <- Fit$glmfit$family$linkinv(upr)←
49 #upper limit
50 lwr2 <- Fit$glmfit$family$linkinv(lwr)←
51 #lower limit
52
53 print(upr2)# upper limit
54 print(lwr2)# lower limit
55
56 # Create a boxplot showing the results
57 plotDiagPhotoZ(photoz = photoZtest$←
58 photoz,
59 specz = PHAT0test$redshift,
60 type = "box")

```

The functions provided by the package and that are visible to the user are:

1. `computeCombPCA()` Computes PCA projections of the combined data sets.
2. `computeDiagPhotoZ()` Computes a list of summary statistics of the redshift estimation (mean, sd, median, mad, outliers).
3. `CosmoPhotoZestimator()` Computes redshift estimates from photometric data and a training data set with photometry and spectroscopy. The estimation is based on GLMs.
4. `glmPredictPhotoZ()` Predicts photometric redshifts for a given a GLM fit object.
5. `glmTrainPhotoZ()` Fits a GLM for photometric redshift estimation. A Bayesian fit or a normal fit may be adopted, using the link functions gamma and inverse Gaussian.
6. `plotDiagPhotoZ()` Plots diagnostics for redshift estimations. The following types of plot are available: 1D kernel density estimation of the errors (errordist), observed versus predicted 2D density plot (predobs), a violin and a box plot with errors at redshift bins (errorviolins or box).

The packages additionally includes two dataframes, PHAT0train and PHAT0test, containing 161042 and 8478 objects, and comprising 12 variables (11 bands plus the redshift).

A detailed reference manual of the package can be found at the documentation in CRAN webpage or inside the package.

Appendix A.1. Shiny Package

The R package is also accompanied by a shiny²¹ application that can be hosted locally or deployed by the user at a webserver. This application allows the user to run the photometric redshift estimation, to configure many parameters of the code visually and experiment with the results. It also allows the user to either use the PHAT0 data, or to upload data files (the expected format can be found at the application's help tab). A screenshot of this application can be seen in Fig. A.7.

To run the interface, after installing the package it is necessary to use the following command:

²¹<http://shiny.rstudio.com/>


```
1 runApp(paste(find.package("CosmoPhotoz"),
2             ), "/CosmoPhotoz_shiny", sep=""))
```

Finally, this application can also be used via the web. This option requires no local installation, but the actual processing may be slower. This web interface is hosted by the shinyapps.io platform²², and can be accessed directly at <https://cosmostatisticsinitiat.shinyapps.io/CosmoPhotoz>.

Appendix B. Python Package

The Python package is publicly available. The source code can be accessed directly at the COIN GitHub repository. The package can also be easily installed via the Python Package Index²³ using `easy_install` or `pip`. Complete documentation is fully accessible from the *Read-the-docs*²⁴ platform. The package can be used in two ways. The first is a binary file that can be run on the command line with little interaction required:

```
1 run_glm.py --dataset sample.csv
2
3 run_glm.py --dataset train.csv test.csv
4
5 run_glm.py --dataset sample.csv
6             --num_components 3
7             --training_size 10000
8             --family Gamma
9             --link log
```

The library can also be imported in the standard way by the user. There is a `PhotoSample` class that can be fully manipulated. We give a few examples of how the class can be utilised for personal use with photometric catalogues. Firstly, just executing the analysis like the binary.

```
1 from CosmoPhotoz.photoz import
2     PhotoSample
3 # Use PHATO catalogue supplied
4 # with the software
5 UserCatalogue = PhotoSample(filename="PHATO", \
```

```
5 family="Gamma", \link="log")
6 UserCatalogue.run_full()
```

Secondly, selecting a log link and only making the violin plot.

```
1 # Import the PhotoSample class
2 from CosmoPhotoz.photoz import
3     PhotoSample
4 # Instantiate an object of PhotoSample
5 class
6 UserCatalogue = PhotoSample(filename="PHATO", \
7 family="Gamma", \
8 link="log")
9 # Select the link
10 UserCatalogue.link = "log"
11 # Carry out PCA
12 UserCatalogue.do_PCA()
13
14 # Split into training and test data
15 UserCatalogue.split_sample(random=True)
16
17 # GLM fitting and photo-z estimation
18 UserCatalogue.do_GLM()
19
20 # Create a violin plot showing the
21 results
22 UserCatalogue.make_violin()
```

Finally, to determine the number galaxies required to achieve at least a catastrophic error of 5.937%.

```
1 # Import the PhotoSample class
2 from CosmoPhotoz.photoz import
3     PhotoSample
4 import numpy as np
5 # Instantiate the class
6 UserCatalogue = PhotoSample(filename="PHATO", \
7 family="Gamma", \link="log")
8
9 # Make a training size array to loop
10 through
11 train_size_arr = np.arange(
12     (500, 10000, 500))
13 catastrophic_error = []
14
15 # Select your number of components
16 UserCatalogue.num_components = 4
```

²²<http://www.shinyapps.io>

²³<https://pypi.python.org/pypi/CosmoPhotoz/0.1>

²⁴<http://cosmophotoz.readthedocs.org>

CosmoPhotoz - GLM PhotoZ estimation

Figure A.7: A screenshot of the Shiny web application running on the shinyapps.io cloud. In this screenshot the application was estimating redshifts for a subset of the PHAT0 data set. This application is publicly available at <https://cosmostatisticsinitiative.shinyapps.io/CosmoPhotoz>.

```

16 for i in range(len(train_size_arr)):
17     UserCatalogue.do_PCA()
18     UserCatalogue.test_size = train_size_arr[i]
19     UserCatalogue.split_sample(random=True)
20     UserCatalogue.do_GLM()
21     catastrophic_error.append(\
22         UserCatalogue.catastrophic_error)
23
24 min_indx = np.array(catastrophic_error) < 5.937
25 optimum_train_size = train_size_arr[min_indx]
26
27 # Print the output to the user
28 print(optimum_train_size)

```

The main methods of the PhotoSample class are the following:

1. `__init__()` Constructor of the class. This is used to define the public attributes of the class, e.g., number of PCA components, size of the training sample, file name, etc.
2. `do_PCA()` Carries out PCA on the data set. Using the variance per principal component, it deter-

mines the optimal number of principal components ensuring that 99.95% of the variance is retained.

3. `split_sample()` Randomly splits the sample if the user gives a single file, for which they wish to use a random sample to train and test the GLM. The method also reconstructs a final Pandas DataFrame object to be used in the GLM fitting.
4. `do_GLM()` Instantiates the GLM class and selects the family and link type to be used. If the user does not interfere, the gamma family and inverse power link will be used, which may not be fully optimal in all situations.
5. `make_kde_1d()` Creates a plot that shows the probability density function of the outlier fraction.
6. `make_kde_2d()` Creates a plot that shows the probability density function of the photometric and spectroscopic redshifts on the upper and right most axis. In the centre a 2D probability density plot is shown.

7. `make_violin()` Creates a violin plot. A violin plot shows the probability density of the outlier fraction for a given redshift interval.
8. `write_to_file()` Writes the GLM photo-z predictions to a file named *glmPhotoZresults.csv* when the user supplies a *filename_test* and *filename_train*.

We note that there are slightly differences to how the R and Python codes solve the SVD matrices in the PCA step and also how they minimise the GLM. We find that the predicted redshifts from both packages have a Pearson-R value of 1, and have a spread that is of the order and smaller than the intrinsic spread of the predicted vs. measured redshifts, and are thus negligible.

Appendix C. SDSS SQL Photometry Query

The following SQL code can be run on the SDSS CasJobs service ²⁵ to obtain the photometry of the galaxy sample investigated within this article.

```

1
2  SELECT s.specObjID, g.dered_u, g.dered_
      _g, g.dered_r, g.dered_i, g.dered_
      z, s.z AS redshift INTO mydb.
      specObjAllz_cleanphoto
3  FROM SpecObj As s JOIN Galaxy as g ON
      s.specobjid = g.specobjid,
      PhotoObj
4  WHERE class = "GALAXY"
5      AND zWarning = 0
6      AND g.objID = PhotoObj.objID
7      AND PhotoObj.CLEAN=1
8      AND s.z > 0
9      AND g.dered_u > 0
10     AND g.dered_g > 0
11     AND g.dered_r > 0
12     AND g.dered_i > 0
13     AND g.dered_z > 0

```

²⁵<http://skyserver.sdss3.org/CasJobs/>