

The overlooked potential of generalized linear models in astronomy – III. Bayesian negative binomial regression and globular cluster populations

R. S. de Souza,^{1★} J. M. Hilbe,^{2,3★} B. Buelens,⁴ J. D. Riggs,⁵ E. Cameron,⁶
E. E. O. Ishida,⁷ A. L. Chies-Santos^{8,9} and M. Killedar¹⁰

¹*EIRSA ‘Lendulet’ Astrophysics Research Group, MTA Eötvös University, Budapest 1117, Hungary*

²*Arizona State University, 873701 Tempe, AZ 85287-3701, USA*

³*Jet Propulsion Laboratory, 4800 Oak Grove Dr., Pasadena, CA 91109, USA*

⁴*Flemish Astronomical Society, B-3600 Genk, Belgium*

⁵*Northwestern University, Evanston, IL 60208, USA*

⁶*Department of Zoology, University of Oxford, Tinbergen Building, South Parks Road, Oxford OX1 3PS, UK*

⁷*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany*

⁸*Departamento de Astronomia, Instituto de Física, Universidade Federal do Rio Grande do Sul, Porto Alegre, RS 90040-060, Brazil*

⁹*Instituto de Astronomia, Geofísica e Ciências Atmosféricas, Universidade de São Paulo, São Paulo, SP 05508-090, Brazil*

¹⁰*Universitäts-Sternwarte München, Scheinerstrasse 1, D-81679 München, Germany*

Accepted 2015 August 5. Received 2015 August 4; in original form 2015 June 21

ABSTRACT

In this paper, the third in a series illustrating the power of generalized linear models (GLMs) for the astronomical community, we elucidate the potential of the class of GLMs which handles count data. The size of a galaxy’s globular cluster (GC) population (N_{GC}) is a prolonged puzzle in the astronomical literature. It falls in the category of count data analysis, yet it is usually modelled as if it were a continuous response variable. We have developed a Bayesian negative binomial regression model to study the connection between N_{GC} and the following galaxy properties: central black hole mass, dynamical bulge mass, bulge velocity dispersion and absolute visual magnitude. The methodology introduced herein naturally accounts for heteroscedasticity, intrinsic scatter, errors in measurements in both axes (either discrete or continuous) and allows modelling the population of GCs on their natural scale as a non-negative integer variable. Prediction intervals of 99 per cent around the trend for expected N_{GC} comfortably envelope the data, notably including the Milky Way, which has hitherto been considered a problematic outlier. Finally, we demonstrate how random intercept models can incorporate information of each particular galaxy morphological type. Bayesian variable selection methodology allows for automatically identifying galaxy types with different productions of GCs, suggesting that on average S0 galaxies have a GC population 35 per cent smaller than other types with similar brightness.

Key words: methods: data analysis – methods: statistical – globular clusters: general.

1 INTRODUCTION

The current era of astronomy marks the transition from a data-deprived field to a data-driven science, for which statistical methods play a central role. An efficacious data exploration requires astronomers to go beyond the traditional Gaussian-based models which are ubiquitous in the field. Gaussian distributional assumptions fail to hold when the data to be modelled come from *exponential family* distributions other than the Normal/Gaussian¹

(Hardin & Hilbe 2012; Hilbe 2014). For non-Gaussian regression problems there exist powerful solutions already widely used in medical research (e.g. Lindsey 1999), finance (e.g. Jong & Heller 2008), healthcare (e.g. Griswold et al. 2004) and biostatistics (e.g. Marschner & Gillett 2012), but vastly underutilized to date in astronomy. These solutions are known as generalized linear models (GLMs). Despite the ubiquitous implementation of GLMs in general statistical applications, there have been only a handful of astronomical studies applying GLM techniques such as logistic regression (e.g. Raichoor & Andreon 2012, 2014; Lansbury, Lucey & Smith 2014; De Souza et al. 2015), Poisson regression (e.g. Andreon & Hurn 2010), gamma regression (Elliott et al. 2015) and negative binomial (NB) regression (Ata, Kitaura & Müller 2015). The methodology discussed herein focuses on Bayesian count response models (Poisson and NB), suited to handle discrete,

*E-mail: rafaél.2706@gmail.com (RSS); j.m.hilbe@gmail.com (JMH)

¹ The exponential family comprises a set of distributions ranging from both continuous and discrete random variables (e.g. Gaussian, Poisson, Bernoulli, Gamma, etc.).

count-based data sets applied to a catalogue of globular clusters (GCs).

GCs are among the oldest stellar systems in the Universe (formed at $z > 2$; Kruijssen 2014), are pervasive in nearby massive galaxies (Brodie & Strader 2006) and can be found in massive galaxy clusters not necessarily associated with one of its galaxies (e.g. Durrell et al. 2014). Hence, understanding their properties is of utmost importance for drawing a complete picture of galaxy evolution. The past few decades have seen considerable interest in the apparent correlation between the mass of the black hole at the centre of a galaxy, M_{BH} , and the velocity dispersion of the central stellar bulge, σ (e.g. Gebhardt et al. 2000). As part of the process of understanding the nature and origin of the so-called $M_{\text{BH}}-\sigma$ relation, astronomers have investigated links between other properties of the host galaxy. In particular, the correlation between the size of GC populations, N_{GC} , and M_{BH} is tight, possibly more so than the $M_{\text{BH}}-\sigma$ relation, and may reflect an underlying connection to the bulge mass, binding energy, host galaxy stellar mass and total luminosity (Burkert & Tremaine 2010; Harris & Harris 2011; Snyder, Hopkins & Hernquist 2011; Rhode 2012; Harris, Harris & Alessi 2013; Harris, Poole & Harris 2014). This may go some way to explaining the huge range in scales of the regions involved. One notorious outlier is our own Milky Way galaxy, for which there are far too many GCs given the mass of its central supermassive black hole, despite the fact that both are accurately measured. Nevertheless, the otherwise small scatter found in such relations deserves a closer look since it cannot be easily explained by simple scaling relation arguments.

The connection between N_{GC} and the global properties of their host galaxies is an extant astronomical puzzle involving count models, but is treated as a continuous one. Such correlation studies are commonly based on taking pairs of parameters (x, y) in log–log space and searching for solutions in the normal form $y = \alpha + \beta x$, despite the fact that this regression technique assumes continuous variables and a Gaussian error distribution, e.g. χ^2 -minimization (Tremaine et al. 2002).

Our method surpasses the previous χ^2 -minimization approach in several ways. The most obvious being the ability to handle count data without the need of logarithmic transformations of a discrete variable. Hence, we can take into account the cases with zero counts, instead of removing them to accommodate the logarithm transformation, or adding an arbitrary data shift in the form $\log(x + \epsilon)$, with ϵ commonly taken as unity. Our method naturally handles errors in variables in both the x and y axes accommodating the heteroscedasticity of the errors in N_{GC} .² As a further analysis, we introduce one of the most important extensions of GLMs known as generalized linear mixed models (GLMMs). This is done to include in the model information about each galaxy morphological type, allowing discrimination among classes of objects requiring additional adjustments in their regression coefficients.

The outline of this paper is as follows. In Section 2 we provide a brief introduction of GLMs in the context of exponential family distributions. An overview of count data along with Poisson and NB GLMs is presented in Section 3. The data set used in our analysis is summarized in Section 4. In Section 5 we discuss the necessary steps to build our Bayesian model. In Section 6 we discuss GLMMs in the context of random intercepts models. Finally in Section 7, we present our conclusions.

2 GENERALIZED LINEAR MODELS

Classical response-with-covariates models, i.e. general (not generalized) linear models, assume that the response variable and the residual errors, following a normal distribution, are linear in the model parameters and have constant variance. This allows model parameter estimation with ordinary least squares (OLS) methods. As described above, many data sets have response variables that violate one or more of these assumptions. While remedial measures such as transformations on the response variable or the covariates may be applied, these measures may fall short of satisfying the OLS requirements. For data sets for which classical models are ill-suited, the extended class of models, GLMs, is used with model parameters often estimated using maximum likelihood methods (for a brief overview of GLMs in an astronomical context, see e.g. De Souza et al. 2015; Elliott et al. 2015).

Nelder & Wedderburn (1972) introduced a unification of models characterized by being linear on the systematic component (model predictors). For example logistic and probit analysis for binomial variates, contingency tables for multinomial variates and regression for Poisson- and gamma-distributed variates, each a form of the GLM. The random response variable, $Y_i, i = 1, 2, \dots, n$, may be represented as

$$Y_i \sim f(\mu_i, a(\phi)V(\mu_i)),$$

$$g(\mu_i) = \eta_i,$$

$$\eta_i \equiv \mathbf{x}_i^T \boldsymbol{\beta} = \beta_0 + \beta_1 x_1 + \dots + \beta_p x_p. \quad (1)$$

In equation (1), f denotes a response variable distribution from the exponential family (EF), μ_i is the response variable mean, ϕ is the EF dispersion parameter in the dispersion function $a(\cdot)$, $V(\mu_i)$ is the response variable variance function, η_i is the linear predictor, the $\mathbf{x}_i^T = \{x_{i1}, x_{i2}, \dots, x_{ip}\}^T$ is the vector of explanatory variables (covariates or predictors), $\boldsymbol{\beta} = \{\beta_1, \beta_2, \dots, \beta_p\}$ is the vector of covariates coefficients and $g(\cdot)$ is the link function, which connects the mean to the predictor. If $V(\mu_i)$ is a constant for all μ_i , then the mean and variance of the response are independent, which allows using a Gaussian response variable. If the response is Gaussian, then $g(\mu) = \mu$. The general form of the GLM thus allows Gaussian family, \mathcal{N} , linear regression as a subset, taking the form

$$Y_i \sim \mathcal{N}(\mu_i, \sigma^2),$$

$$\mu_i = \beta_0 + \beta_1 x_1 + \dots + \beta_p x_p. \quad (2)$$

The subset of GLMs for count data is the Poisson regression models and the several incarnations of the NB regressions. Poisson regression models assume the count response variable follows a Poisson probability distribution function. Similarly, the NB regression models assume the count response variable follows a NB probability distribution function. Descriptions of the Poisson and NB models follow.

3 MODELLING COUNT DATA

Astronomical quantities can be measured on different scales: nominal (e.g. classes of objects: Type Ia/II supernovae, elliptical/spiral galaxies); ordinal (e.g. ordering planets according to their size or distance to the star) and metric (e.g. galaxy mass, stellar temperature). Observations that have only right-skewed, non-negative integer values belong to a subclass of the metric scale known as count data. Distances between counts are meaningful, hence the counts are metric, but they are not continuous and must be treated as such. Astronomical count data are often log-transformed to satisfy

² Heteroscedastic error structures may remain even after transformation, thus violating the Gaussian assumption of homogeneity of error variance.

Gaussian parametric test assumptions rather than modelled on the basis of a count distribution. Despite the fact that GLMs are better suited to describe count data, a log-transformation of counts has the additional problem of dealing with zeros as observations. With just one observation with value zero, the entire data set needs to be shifted by adding an arbitrary value before transformation. It is well known that such transformations perform poorly, leading to bias in the estimated parameters (O’Hara & Kotze 2010).

We begin our discussion of regression models for count data with the subset of GLMs known as Poisson regression. A common condition accompanying count data is overdispersion, it occurs when the variance exceeds the mean. This condition in Poisson regression suggests that remedial measures, such as the use of NB regression, may be appropriate.

3.1 Poisson regression

Poisson regression was the first model specifically used to deal with count data and still stands as basis for many types of analyses. It assumes a discrete response described by a single parameter distribution which represents the mean or rate, μ , i.e. the expected number of times an event occurs within a fixed time interval. Another important feature is the assumption of equidispersion which implies the equality of mean and variance, and can be quantified by the Pearson χ^2 dispersion statistic (see Section 3.2.). The Poisson distribution function is typically displayed as

$$f(y; \mu) = \frac{\mu^y e^{-\mu}}{y!}, \quad (3)$$

where the mean and variance are given by

$$\text{Mean} = \mu, \quad \text{Variance} = \mu, \quad (4)$$

representing a particular case of equation (1) with $V(\mu) = \mu$ and $a(\phi) = 1$. Thus, a regression equation derived from equation (1) may be used as a GLM for a count response, y . The usual link function, $g(\mu)$, is the natural log function such that $\mu = e^\eta$ (see e.g. Hardin & Hilbe 2012). It is worth noting that GLMs are not simple log transforms of the response variable, but rather, the expected counts from a Poisson regression is an exponentiated linear function of η , thereby keeping the response variable on its original scale. Often, count data do not enjoy the Poisson assumption of equidispersion resulting in a Poisson dispersion statistic (see Section 3.2) with a value greater than one.

3.2 Overdispersion

Overdispersion in Poisson models occurs when the response variance is greater than the mean. It may arise when there are violations in the distributional assumptions of the data such as when the data are clustered, thereby violating the likelihood requirement of the independence of observations. Overdispersion may cause standard errors of the estimates to be deflated or underestimated, i.e. a variable may appear to be a significant predictor when it is in fact not. A key approach for checking overdispersion is by means of the dispersion statistic, \mathcal{D} ,

$$\mathcal{D} = \frac{\chi^2}{N - N_p}, \quad (5)$$

where N is the number of observations and N_p is the number of parameters in the model. Then $N - N_p$ represents the residual degrees of freedom. For a Poisson GLM, the Pearson χ^2 value is

$$\chi^2 = \sum_{i=1}^N \frac{(Y_i - \mu_i)^2}{\mu_i}, \quad (6)$$

where Y_i represents the observed values, and μ_i is the mean and variance of Y_i . Poisson overdispersion occurs when the variation in the data exceeds the expected variability based on the Poisson distribution, resulting in \mathcal{D} being greater than 1. Small amounts of overdispersion are of little concern; a rule of thumb is: if $\mathcal{D} > 1.25$, then a correction may be warranted (Hilbe 2014).

If overdispersion is observed, then there are several corrective measures in common practice. Options are adjusting the standard errors by scaling, applying sandwich or robust standard errors, or bootstrapping standard errors for the model. However, only the standard errors will be adjusted and not the regression coefficients, β , which often can be affected by overdispersion as well (e.g. Hilbe 2011). This paper examines the efficacy of using Bayesian estimation methods on a more general discrete distribution known as the NB. The NB distribution contains a second parameter called the dispersion or heterogeneity parameter which is used to accommodate Poisson overdispersion as described below.

3.3 Negative binomial regression

The NB distribution has long been recognized as a full member of the exponential family, originally representing the probability of observing y failures before the r th success in a series of Bernoulli trials. It can also be formulated as a Poisson model with gamma heterogeneity (Hilbe 2011). The NB model, as a Poisson–gamma mixture model, is appropriate to use when the overdispersion in an otherwise Poisson model is thought to take the form of a gamma shape or distribution, i.e. $a(\phi) = 1/k$, with $k > 0$. The NB probability distribution function is then given by

$$f(y; k, \mu) = \frac{\Gamma(y+k)}{\Gamma(k)\Gamma(y+1)} \left(\frac{k}{\mu+k}\right)^k \left(1 - \frac{k}{\mu+k}\right)^y. \quad (7)$$

The distribution function has two parameters, μ and k , allowing more flexible models than the Poisson distribution. The symbol Γ represents the gamma function.³ The mean and variance are given by

$$\text{Mean} = \mu; \quad \text{Variance} = \mu + \frac{\mu^2}{k} = \mu + \alpha\mu^2. \quad (8)$$

The NB distribution has distributional assumptions similar to the Poisson distribution with the exception that it has a dispersion parameter $\alpha = 1/k$ to accommodate wider count distribution shapes than allowed by the Poisson model. As the dispersion parameter, α , approaches 0, $\lim_{\alpha \rightarrow 0} \alpha\mu^2 = 0$ or $\lim_{k \rightarrow \infty} \mu^2/k = 0$, then the variance equals the mean which recovers the Poisson distribution.

It should be noted that if different clusters of counts have different gamma shapes, indicating differing degrees of correlation within data, and if the NB Pearson χ^2 dispersion statistic is greater than one, then the NB model may itself be overdispersed, i.e. the data may be both Poisson and NB overdispersed. Random effects and mixed effects Poisson and NB models are then reasonable alternatives (Hilbe 2014).

³ If n is a positive integer, $\Gamma(n) = (n-1)!$

Table 1. Main assumptions of each regression model family.

	Normal	Log-normal	Poisson	Negative binomial
Response variable	Real	Positive	Non-negative integer	Non-negative integer
Null values	✓	✗	✓	✓
Sample variance	Homoscedastic	Homoscedastic	Heteroscedastic	Heteroscedastic
Overdispersion	✗	✗	✗	✓

Table 2. Summary of the parameters used in this work from the catalogue of GCs compiled by Harris et al. (2013).

Parameter	Definition
N_{GC}	Number of globular clusters
M_V	Absolute visual magnitude
σ	Bulge velocity dispersion
M_{BH}	Central black hole mass
M_{dyn}	Dynamical mass
$\epsilon_{N_{GC}}$	Uncertainty in N_{GC}
ϵ_{M_V}	Uncertainty in M_V
ϵ_{σ}	Uncertainty in σ
$\epsilon_{M_{BH}}$	Uncertainty in M_{BH}

An additional situation should also be mentioned. If the Poisson dispersion statistic is less than one, this is evidence of Poisson underdispersed data. The NB model is not appropriate for handling Poisson underdispersion; however, the generalized Poisson model is. We do not discuss underdispersed data in this paper, but the subject warrants future study as to how it applies to astrophysical data. To guide the reader, Table 1 displays the main assumptions of the OLS, OLS with a log-transformed response variable, Poisson and NB regression models discussed in the previous sections.

4 DATA SET

As a study case, we use the catalogue of GCs presented in Harris et al. (2013) (see also Harris et al. 2014).⁴ The data are composed of 422 galaxies with published measurements of their GC populations. There is a range of galaxy morphologies from which we indexed 247 as elliptical (E), 94 as lenticular (S0), 55 as spirals (S) and 26 as irregulars (Irr) galaxies for illustrative purposes. Note that the original catalogue presents 69 different subcategories of morphological classifications which will be discussed in Section 6. This is a compilation of literature data from a variety of sources obtained with the *Hubble Space Telescope* as well as a wide range of other ground-based facilities. Beyond N_{GC} , we select the following properties for our analysis: central black hole mass, dynamical bulge mass, bulge velocity dispersion and absolute visual magnitude as described in Table 2.

5 MODELLING THE POPULATION SIZE OF GLOBULAR CLUSTERS

Within this section we demonstrate the application of Bayesian GLM regression for modelling the relationship between N_{GC} and the following galaxy properties: M_{BH} , σ , M_V and M_{dyn} . Hereafter, unless otherwise stated, the analysis is made using a subsample of 45 objects from which we have observations for all the property

predictors. In Section 5.4 an additional analysis uses the entirety of the available data.

A few common terms in statistical modelling need to be reviewed to facilitate our model applications. The analysis focus is the prediction of N_{GC} as a function of the global galaxy properties. Therefore, N_{GC} represents the count (i.e. a non-negative integer) response variable, while M_V , M_{BH} and M_{dyn} are interchangeably called covariates, explanatory variables or predictors. If included in the model, the galaxy morphological type is also considered a nominal categorical predictor (see Section 3). The whole analysis is performed using Just Another Gibbs Sampler (JAGS),⁵ a program for analysis of Bayesian hierarchical models using a Markov chain Monte Carlo (MCMC) framework.⁶ For each regression case, we initiate three Markov chains by starting the Gibbs samples at different initial values sampled from a normal distribution with zero mean and standard deviation of 10. The initial adapting and burning phases were set to 22 000 steps followed subsequently by 50 000 steps, which was sufficient to guarantee convergence of each chain for all studied cases.

We now use the relationship between M_{BH} and N_{GC} as an example to illustrate how the statistical model is built. To motivate the use of the more general NB distribution, we start the analysis assuming a GLM Poisson regression model neglecting the uncertainties in measurements at this stage for simplicity.⁷ This leads to the following model:

$$\begin{aligned}
 N_{GC,i} &\sim \text{Poisson}(\mu_i), \\
 \mu_i &= e^{\eta_i}, \\
 \eta_i &= \beta_0 + \beta_1 M_{BH,i}, \\
 \beta_0 &\sim N(0, 10^6), \\
 \beta_1 &\sim N(0, 10^6), \\
 i &= 1, \dots, N.
 \end{aligned}
 \tag{9}$$

This set of equations reads as follows: each galaxy in the data set, composed of N objects, has its GC population sampled from a Poisson distribution whose expected value, μ , relates to the central black hole mass through a linear relation expressed by η . Since we do not have previous information about the values of the coefficients β_0 and β_1 , we assigned non-informative Gaussian priors with zero mean and standard deviation equal to 10^6 . We refer the reader to Appendix A for an example of how to implement a Poisson GLM in JAGS. The fitted curve for this model is displayed in Fig. 1. The grey shaded areas represent 50, 95 and 99 per cent prediction intervals, which are the regions where a future observation will fall with these

⁵ <http://CRAN.R-project.org/package=rjags>

⁶ Note that count models can be approached by other methods, such as a full maximum likelihood algorithm (see Hilbe 2011, for a review).

⁷ Neglecting the errors at this point does not affect the conclusions regarding the level of Poisson overdispersion.

⁴ The complete catalogue can be obtained at http://www.physics.mcmaster.ca/harris/GCS_table.txt.

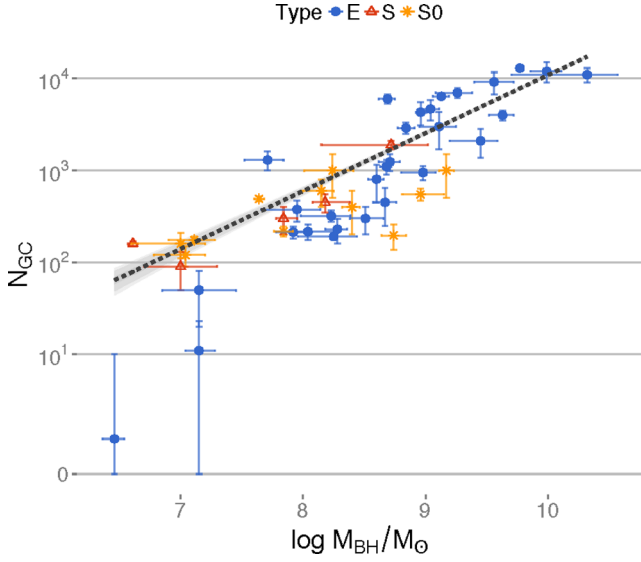


Figure 1. Total number of GCs, N_{GC} , plotted versus the central black hole mass, M_{BH} . The dashed line represents the expected value of N_{GC} for each value of M_{BH} using Poisson GLM regression, while the shaded areas depicts 50, 95 and 99 per cent prediction intervals. Galaxy types are coded by shape and colour as follows: ellipticals (E; blue solid circles), spirals (S; red open triangles) and lenticulars (S0; orange asterisks). An ArcSinh transformation is applied in the y-axis for better visualization of the whole range of N_{GC} values, including the null ones.

given probabilities.⁸ Note that the areas in the plot are too narrow to be visually discriminated. A visual inspection clearly indicates that the Poisson model is not adequate to explain the data variability since most of the data fall outside the three prediction intervals. Also, the dispersion statistic for this model is $\mathcal{D} = 1039$, which is a strong indication of an inadequate model. All other covariates, σ , M_V and M_{dyn} , lead to models with similarly high levels of Poisson overdispersion. Hence, hereafter we discuss construction of the full model based on the NB family to mitigate overdispersion and to include the uncertainties in the observational quantities. Unlike the Poisson model, by employing a NB distribution we allow the incidence rate of GCs to be itself a random variable.

Continuing with our working example, we keep the discussion using the relationship between N_{GC} and M_{BH} , but see Appendix B for descriptions of the other models. The first step is to understand how to include information about the uncertainties in the measurements (see e.g. Andreon & Hurn 2013, for a review of measurement errors in astronomy). Measurement errors in the response count variable are the trickiest part to be modelled. The classical model with an additive error term $y = y^* \pm \varepsilon$ is inappropriate since it does not ensure that the observed value y is non-negative. The appropriate model is described below and its graphical representations are displayed in Fig. 2:

$$N_{GC;i} \sim \text{NB}(p_i, k),$$

$$p_i = \frac{k}{k + \mu_i},$$

⁸ Not to be confused with the commonly used confidence interval in frequentist statistics. A 95 per cent confidence interval will contain the sample mean with 95 per cent probability. In other words, a larger number of repeated samples from the data would contain the sample mean 95 per cent of the time.

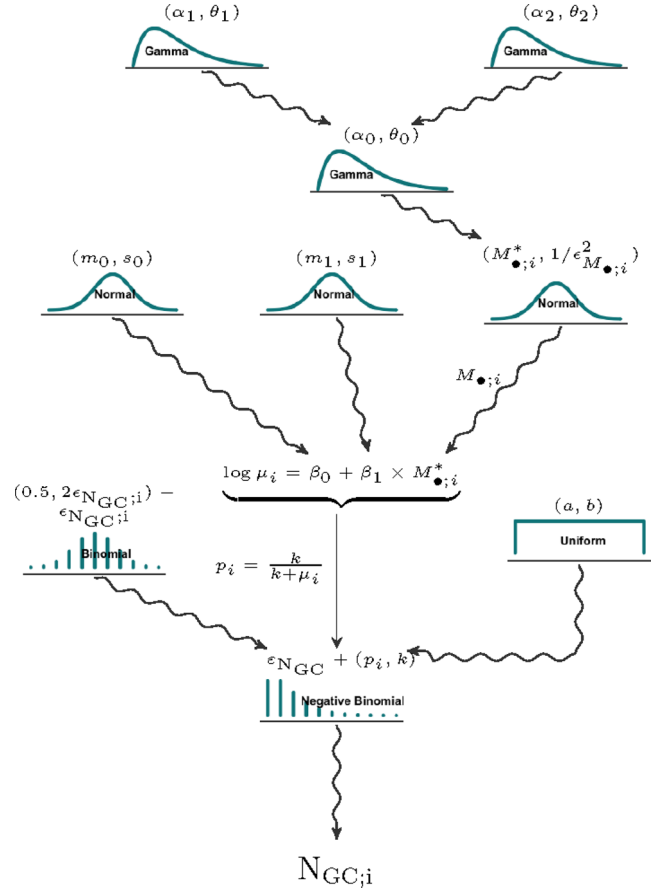


Figure 2. A graphical model of equation (10) representing the hierarchy of dependencies for a data set of galaxies indexed by the subscript i . The sinusoidal curves represent stochastic dependencies, while straight arrows a deterministic ones. To save space, we replace M_{BH} by M_{\bullet} in the diagram.

$$\mu_i = e^{\eta_i} + \epsilon_{N_{GC};i},$$

$$\eta_i = \beta_0 + \beta_1 M_{BH;i}^*,$$

$$k \sim \mathcal{U}(0, 5),$$

$$M_{BH;i} \sim \mathcal{N}(M_{BH;i}^*, e^2_{M_{BH;i}^*}),$$

$$\epsilon_{N_{GC};i} \sim \mathcal{B}(0.5, 2e_{N_{GC};i}) - e_{N_{GC};i},$$

$$\beta_0 \sim \mathcal{N}(0, 10^6),$$

$$\beta_1 \sim \mathcal{N}(0, 10^6),$$

$$M_{BH;i}^* \sim \Gamma(\alpha_0, \theta_0),$$

$$\alpha_0 \sim \Gamma(0.01, 0.01),$$

$$\theta_0 \sim \Gamma(0.01, 0.01),$$

$$i = 1, \dots, N. \quad (10)$$

The above is slightly more complex than the model displayed in equation (9) and reads as follows. Each galaxy in the data set with N objects, has its GCs population sampled from a NB distribution whose expected value, μ , relates to the central black hole mass through the linear predictor η . The additional transformation $p_i = k/(k + \mu_i)$ is required due to how the NB distribution is parametrized in JAGS. The uncertainties related to the counts, $\epsilon_{N_{GC};i}$, are taken to be associated with the mean, μ , of the NB distribution and are modelled using a shifted binomial distribution,

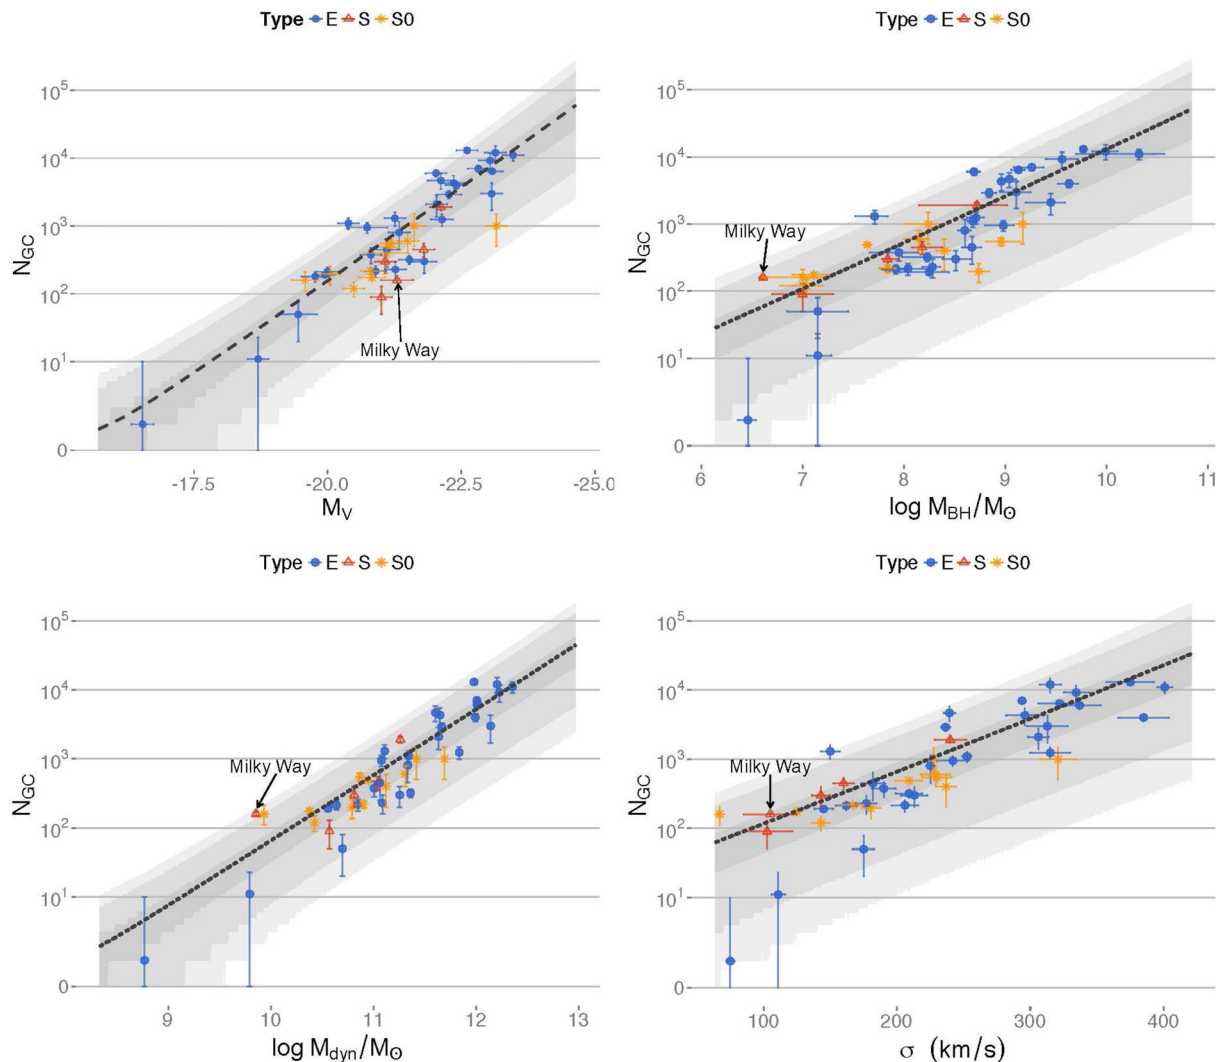


Figure 3. GC population, N_{GC} , plotted against visual absolute magnitude (M_V ; top left-hand panel), black hole mass (M_{BH} ; top right-hand panel), dynamical mass (M_{dyn} ; bottom left-hand panel) and bulge velocity dispersion (σ ; bottom right-hand panel). In each panel the dashed line represents the expected value of N_{GC} for each value of the covariate using NB GLM regression, while the shaded areas depicts 50, 95 and 99 per cent prediction intervals. Galaxy types are coded by shape and colour as follows: ellipticals (E; blue solid circles), spirals (S; red open triangles) and lenticulars (S0; orange asterisks). An ArcSinh transformation is applied in the y-axis for better visualization of the whole range of N_{GC} values, including the null ones.

\mathcal{B} , with zero mean and taking on integer values in the range $[-e_{N_{GC,i}}, +e_{N_{GC,i}}]$ (see e.g. chapter 13 from Cameron & Trivedi 2013, from which this approach is loosely based). Uncertainties associated with the observed predictor $M_{BH,i}$ are modelled using a Gaussian distribution with unobserved mean given by the ‘true black hole mass’, $M_{BH,i}^*$, and standard deviations given by the reported uncertainties in the observed black hole mass, $e_{M_{BH,i}}$. Since $M_{BH,i}^*$ is itself an unobserved variable, we add a non-informative Γ prior on top of which we added non-informative hyperpriors for the shape, α_0 , and rate, θ_0 , parameters of the Γ distribution. The choice of a Γ prior is motivated by the fact that the black hole mass is a continuous, but non-negative quantity which makes Γ a more suitable distribution. For the shape parameter k , we assigned a non-informative uniform prior, \mathcal{U} , as suggested in Zuur, Hilbe & Ieno (2013). For the coefficients β_0 and β_1 we assigned non-informative Gaussian priors with zero mean and standard deviation equal to 10^6 .

Adapting the model above for each combination of N_{GC} and a given galaxy property generates the fitted curves displayed in Fig. 3.

The grey shaded area represents 50, 95 and 99 per cent prediction intervals, while the dashed line represents the expected value of N_{GC} for each value of the covariate. Note the remarkable agreement between the model and the observed values with prediction intervals enclosing the entirety of the data, including objects that have been previously declared outliers and even removed from analysis, such as our own Milky Way (e.g. Burkert & Tremaine 2010; Harris & Harris 2011; Harris et al. 2014).

5.1 Fit diagnostics

If the Markov chains are all representative of the posterior distribution of the fitted parameters, they should overlap each other. Traceplots, Fig. 4, and density plots, Fig. 5, are two useful visual diagnostics that are commonly used to test for chain convergence. We can see that the chains do mix well after the burn-in period, suggesting that the chains are producing representative values from the posterior distribution for β_0 , β_1 and k . Additionally, we used a more quantitative check, viz., the so-called Gelman–Rubin




Figure 4. Illustration of MCMC diagnostics. Three chains were generated by starting the Gibbs algorithm at different initial values sampled from a normal distribution with zero mean and standard deviation 10. Steps 42 000–72 000 are shown here. The figure displays the results for the model N_{GC} versus M_{BH} , with the traceplots for β_0 , β_1 and k displayed from top to bottom.

statistic (Gelman & Rubin 1992). The underlying idea is that if the chains have reached convergence, the average difference between the chains should be similar to the average difference across steps within the chains. The statistic equals unity if the chains are fully converged. As a rule, values above 1.1 indicate that the chains have failed to properly converge. The Gelman–Rubin statistic fell below 1.05 for all estimated parameters in our analysis. Hence, once we convince ourselves that the model is working properly, the next step in the analysis is to add interpretations to the fitted coefficients as we discuss now.

5.2 Interpretation of the coefficients

The exponentiated coefficients e^{β_i} of Poisson and NB regressions are also known as rate ratios, or incidence rate ratios, which quantify how an increase of unity in the predictor variable affects the number of occurrences of the response variable. From Table 3, displaying the means and respective 95 per cent credible intervals of the posterior distribution for each parameter, the exponentiated coefficient $\beta_1 = 1.59$ of the M_{BH} predictor gives a rate ratio of $e^{1.59} = 4.9$. Therefore, according to the model, a galaxy whose central black hole has a mass of e.g. $\approx 10^8 M_\odot$ has on average approximately five times more GCs than a galaxy whose $M_{BH} \approx 10^7 M_\odot$.⁹ In other words, 1 dex¹⁰ variation increase in the M_{BH} leads to an

⁹ Note that the analysis was made using $\log M_{BH}$.

¹⁰ A dex difference of a given quantity x is a change by a factor of 10^x .




Figure 5. Overlapped density plots with different colours by chain. The plot is a comparison of the target distribution by each chain, representing a visual test for convergence. The figure displays the results for the model N_{GC} versus M_{BH} , with the posteriors for β_0 , β_1 and k displayed from top to bottom.

Table 3. β_i coefficients and scale parameter, k , from Bayesian NB regression analysis with N_{GC} as the response variable and M_{BH} , M_{dyn} , σ and M_V as predictors. The upper and lower limits enclose 95 per cent of the credible intervals around the posterior means.

Predictor	β_0	β_1	k
M_{BH}	-6.49 ± 2.6	1.59 ± 0.30	1.53 ± 0.60
M_{dyn}	-17.72 ± 2.75	2.19 ± 0.24	2.46 ± 0.97
σ	2.99 ± 0.78	0.02 ± 0.003	1.52 ± 0.59
M_V	-20.50 ± 3.9	-1.28 ± 0.17	2.23 ± 1.1

approximately five times increase in the incidence of GCs in a given galaxy. Likewise, an increase of 1 dex in M_{dyn} leads to an increase of $e^{2.19} = 8.9$ times in the population size of GCs. Another way to state this, given two galaxies with a difference in dynamical mass of 1 dex, the more massive one has a production rate of GCs 8.9 times more efficient on average. Similar interpretation can be made on the other parameters. Another question of interest is how to determine the best predictor of N_{GC} . In the following, we discuss how to address this problem from a Bayesian perspective.

5.3 Model comparison

To find the best predictors for the GC population, we compare the models using the dispersion statistics \mathcal{D} defined in Section 3.2, and the deviance information criterion (DIC; Spiegelhalter et al. 2002).

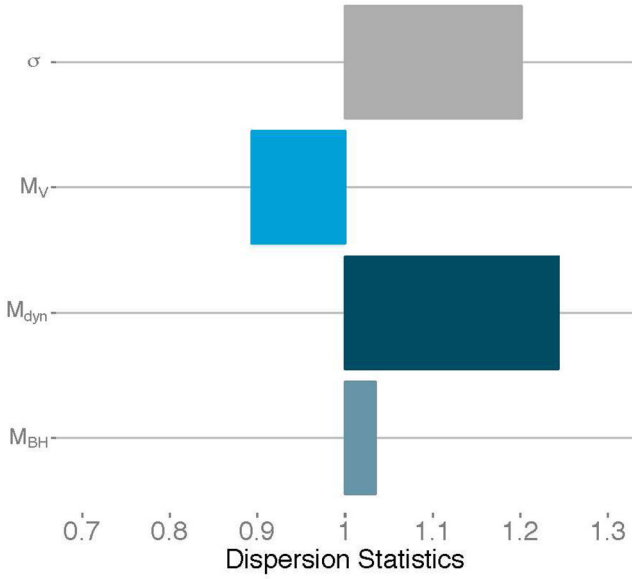


Figure 6. Dispersion statistics, \mathcal{D} , for each model. Values above 1 represent overdispersion, while values below 1 indicate underdispersion.

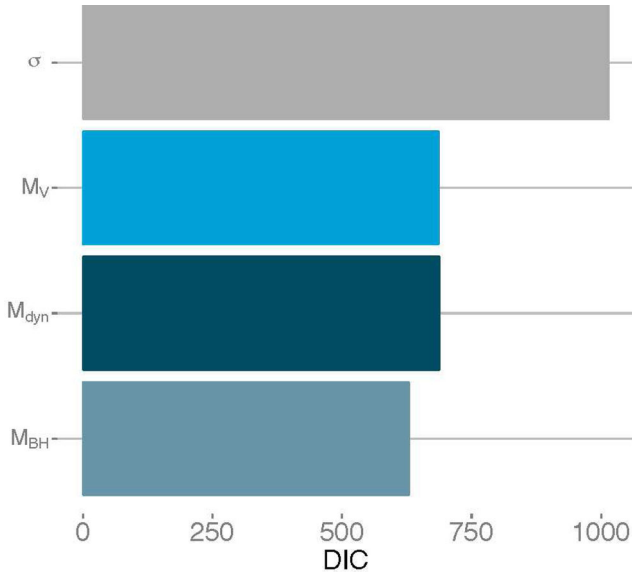


Figure 7. Deviance information criterion, DIC, for each model. Smaller DIC values correspond to preferred models.

The latter represents a compromise between the goodness of fit and model complexity. It is defined as

$$\text{DIC} = \overline{\text{Dev}} + p_D, \quad (11)$$

where the $\overline{\text{Dev}}$ is the average of the deviance $\text{Dev}(\theta)$ defined as $\text{Dev}(\theta) = -2 \log \mathcal{L}(\text{data}|\theta)$, with \mathcal{L} representing the likelihood function. The effective number of parameters, p_D , is calculated as

$$p_D = \overline{\text{Dev}} - \text{Dev}(\theta), \quad (12)$$

where θ is the vector of model parameters (β_0 , β_1 , k for the case in study here). The preferred model has the smallest value for the DIC statistic. Figs 6 and 7 depict the results for the model comparison using the same data set. The black hole mass displays the lowest values for \mathcal{D} and DIC, with dispersion statistics as low as $\mathcal{D} = 1.05$. Although derived from an independent analysis, these

findings corroborate previous claims about the tight connection between the central black hole mass and GC population (Burkert & Tremaine 2010). Nevertheless, it is worth noting that this is not in agreement with a previous analysis performed by Harris et al. (2013) using the same catalogue, where they found M_{dyn} as a better predictor for N_{GC} than M_{BH} .¹¹

5.4 Further analysis with the entire data set

Hereafter, we provide a more extensive analysis using the entire catalogue of 422 galaxies. The only quantities available for all objects are the N_{GC} , galaxy morphological type and M_V (Harris et al. 2013). The advantage of using count models for this type of analysis is apparent from the six galaxies for which no GCs were detected. Such a scenario is naturally accommodated by discrete likelihoods while avoiding the failings of logarithmic transformations to the response. The statistical model we use is the same as that discussed in the beginning of this section and can be described as

$$N_{\text{GC};i} \sim \text{NB}(p_i, k),$$

$$p_i = \frac{k}{k + \mu_i},$$

$$\mu_i = e^{\eta_i} + \epsilon_{N_{\text{GC};i}},$$

$$\eta_i = \beta_0 + \beta_1 M_{V;i}^*,$$

$$k \sim \mathcal{U}(0, 5),$$

$$M_{V;i} \sim \mathcal{N}(M_{V;i}^*, e_{M_{V;i}}^2),$$

$$\epsilon_{N_{\text{GC};i}} \sim \mathcal{B}(0.5, 2e_{N_{\text{GC};i}}) - e_{N_{\text{GC};i}},$$

$$\beta_0 \sim \mathcal{N}(0, 10^6),$$

$$\beta_1 \sim \mathcal{N}(0, 10^6),$$

$$M_{V;i}^* \sim \mathcal{U}(-26, -10),$$

$$i = 1, \dots, N.$$

(13)

Overall, the model is similar to the one described in equation (10). The difference is in the prior for the unobserved true absolute visual magnitude, $M_{V;i}^*$, to which we assigned a uniform prior over the range of magnitudes covered by the catalogue. The fitted model shows remarkable agreement with the data as displayed in Fig. 8. Very few objects fall outside the prediction intervals over a wide range of galaxy brightnesses. The dispersion statistics for this model is $\mathcal{D} = 1.15$, and the credible intervals for the fitted β coefficients and scaling parameter, k , are shown in Fig. 9. Likewise, as in the previous section, we can interpret the β coefficient as follows. The mean value of β_1 exponentiated is ≈ 0.4 , which implies that a galaxy whose absolute visual magnitude is one unit greater than another reference galaxy has on average 0.4 times less GCs, i.e. a galaxy brighter by 1 mag over another has on average 2.5 times more GCs. Likewise, a galaxy with $M_V = -20$ has on average $2.5^5 \approx 100$ times more GCs than a galaxy with $M_V = -15$, which is consistent with a visual inspection of Fig. 8. Another advantage of our approach is the possibility to extrapolate the regression solution without making non-physical predictions. The fitted model predicts a nearly zero occurrence of GCs for galaxies with $M_V \geq -11$. Considering the total galaxy luminosity, $L = 10^{0.4(M_V \odot - M_V)} L_{\odot}$, with $M_{V \odot} = 4.83$, the

¹¹ It is important to note that we are not modelling the same relationship as Harris et al. (2013) who modelled $\log N_{\text{GC}}$, the logarithm transformation of N_{GC} , while we model N_{GC} in the original scale.

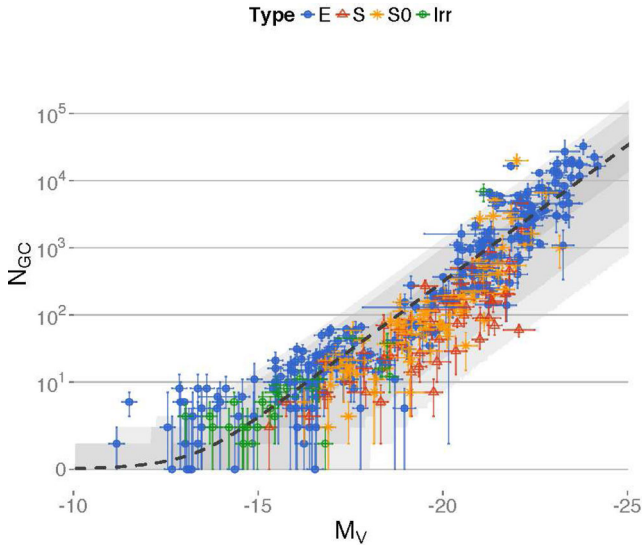


Figure 8. GC population, N_{GC} plotted against visual absolute magnitude M_V . The dashed line represents the expected value of N_{GC} for each value of M_V , while the shaded areas depicts 50, 95 and 99 per cent prediction intervals. Galaxy types are coded by shape and colour as follows: ellipticals (E; blue solid circles), spirals (S; red open triangles), lenticulars (S0; orange asterisks) and irregulars (Irr; green open circles). An ArcSinh transformation is applied in the y-axis for better visualization of the whole range of N_{GC} values, including the null ones.

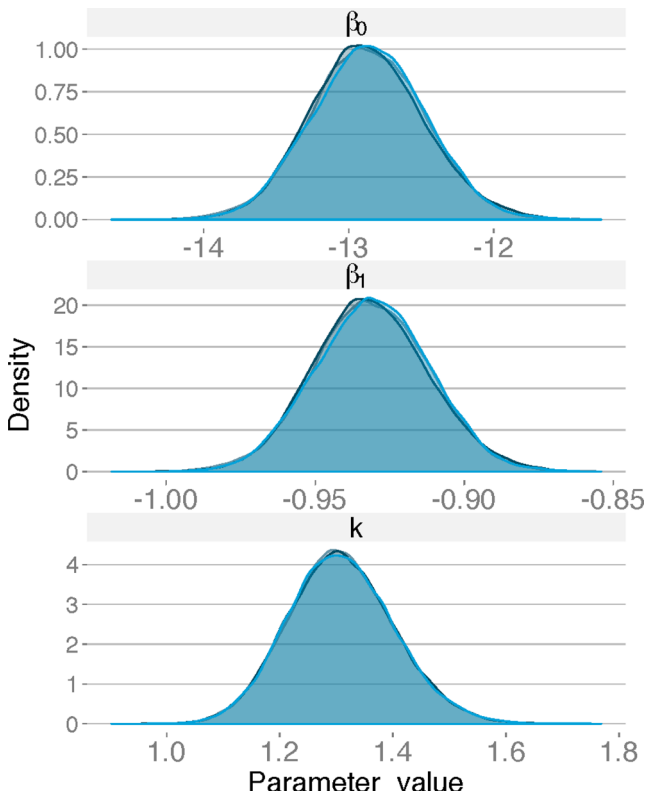


Figure 9. Overlapped density plots with different colours by chain. The plot is a comparison of the target distribution by each chain, representing a visual test for convergence. The figure displays the results for the model N_{GC} versus M_V , with the posteriors for β_0 , β_1 and k displayed from top to bottom.

model suggests that galaxies with $L \leq 2 \times 10^6 L_{\odot}$ are unlikely to host populations of GCs, thus agreeing with the literature (e.g. Harris et al. 2013). The use of Bayesian prediction intervals allow us to make some interesting predictions: for instance from Fig. 8, we can state that galaxies with luminosities $L \leq 8.5 \times 10^7 L_{\odot}$ (or $M_{V_{\odot}} \geq -15$) should not contain more than 10 GCs with 99 per cent probability.

The analysis performed so far did not account for information regarding different galaxy morphological types. Therefore, we are implicitly assuming a pooled estimate (e.g. Gelman & Hill 2007): all different galaxy types are sampled from the same common distribution ignoring any possible variation among them. On the other extreme, performing an independent analysis for each class would mean making the assumption that each morphological type is sampled from independent distributions and that variations between them cannot be combined. In the next section we discuss a more flexible approach together with a brief overview of GLMMs.

6 GENERALIZED LINEAR MIXED MODELS

As our final analysis, we introduce one of the most important extensions of the GLM methodology known as GLMMs. In particular, we focus on one of the simplest GLMM incarnations known as the random intercepts model. The random intercepts model, in our context, includes an additional term ζ_j to account for a class (galaxy type) specific deviation from the common intercept β_0 :

$$\eta_{ij} = \beta_0 + \beta_1 M_{V,i} + \zeta_j, \quad (14)$$

where the index j runs from 1 to 69 representing each of the different galaxy subtypes reported in Harris et al. (2013). A standard approach to modelling ζ_j in a standard linear mixed regression model is to assume the conditional normality of the random intercepts with $\zeta_j \sim \mathcal{N}(0, 1/\tau)$, and $\tau \sim \Gamma(0.01, 0.01)$. Our intention in incorporating this extra term into the model is not to simply adjust the data, but rather the aim is to identify any particular galaxy subtype which deviates from the overall population mean. For this purpose, we employed a popular method for variable selection from a Bayesian perspective known as least absolute shrinkage and selection operator (LASSO) which is discussed in the following section.

6.1 Bayesian LASSO

The original LASSO regression was proposed by Tibshirani (1996) to automatically select a relevant subset of predictors in a regression problem by shrinking some coefficients towards zero (see also Uemura et al. 2015, for a recent application of LASSO for modelling Type Ia supernovae light curves). For a typical linear regression problem:

$$y_i = \beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip} + \epsilon, \quad (15)$$

with ϵ denoting Gaussian noise, LASSO estimates linear regression coefficients $\beta = \beta_0 + \beta_1 x_1 + \dots + \beta_p x_p$ by imposing a L_1 -norm penalty in the form

$$\arg \min_{\beta} \left\{ \sum_{i=1}^N \left(y_i - \sum_{j=1}^p \beta_j x_{ij} \right)^2 + \kappa \sum_{j=1}^p |\beta_j| \right\}, \quad (16)$$

where $\kappa \geq 0$ is a tunable constant that controls the level of sparseness of the solution. The number of zero coefficients thereby increases as κ increase. Tibshirani (1996) also noted that the LASSO estimate




Figure 10. Illustrative comparison between Laplace and Gaussian priors. The Gaussian distribution is represented by dashed lines, while the Laplace distribution by solid lines. For all curves we assign a zero mean, and the scale (or standard deviation, σ , for the Gaussian case) parameters 0.25 (dark blue lines) and 0.5 (cyan lines).

has a Bayesian counterpart when the β coefficients have a double-exponential prior (i.e. a Laplace prior) distribution:

$$f(\zeta; \tau) = \frac{1}{2\tau} \exp\left(-\frac{|\zeta_j|}{\tau}\right), \quad (17)$$

where $\tau = 1/\kappa$. The idea was further developed and is known as Bayesian LASSO (see e.g. Park et al. 2008). Hereafter, we use the LASSO formulation for a slightly different purpose, viz., variable selection for random intercept models (see e.g. Bernardo et al. 2011, p. 165). The underlying idea is to discriminate between galaxy types that follow the overall population mean, i.e. $\zeta_1 = 0$, and galaxies that require an additional adjustment in the intercept, i.e. $\zeta_i \neq 0$. In order to include this information, we replace the linear predictor η by equation (14) and add the following equations in the model described by equation (13):

$$\begin{aligned} \zeta_j &\sim \text{Laplace}(0, \tau), \\ \tau &= 1/\kappa, \\ \kappa &\sim \Gamma(0.01, 0.01), \\ j &= 1, \dots, 69. \end{aligned} \quad (18)$$

The role of the Laplace prior is to assign more weight to regions either near to zero or in the distribution tails as compared to a normal prior. A visual inspection of Fig. 10 confirms this notion. For the parameter κ , we assigned a diffuse (non-informative) gamma hyperprior in the form $\kappa \sim \Gamma(0.01, 0.01)$, which avoids the need of an ad hoc choice of κ . Note that other possibilities exist such as e.g. iteratively finding κ via cross-validation to maximize predictive power.

Analysis results are displayed in Fig. 11. Overall, it suggests that we do not need to add an additional intercept for predicting N_{GC} from M_V . This is consistent with the fact that prediction intervals in Fig. 8 enclose ~ 98.8 per cent of the data set without any need of a random intercept. Nevertheless, the following galaxy types require systematic adjustments: spirals galaxies with moderate size




Figure 11. Caterpillar plot for the random intercepts ζ_i versus the subcategories of galaxy morphological classifications. The thick and thin horizontal lines represent 90 and 95 per cent credible intervals, respectively.

of nuclear bulge (Sb), barred lenticulars (SB0), lenticulars (S0) and dwarf elliptical galaxies (dE0N and dE1N). MG represents one single object. Also, UGC 3274 is the brightest galaxy of the galaxy cluster ACO 539 (Lin & Mohr 2004). Fig. 12 shows that the dE0N and dE1N objects have a large number of GCs on average when compared to other galaxy types with similar luminosities, while the lenticulars have systematically fewer GCs than expected for the overall galaxy population. This can be quantified by looking at the mean value of ζ in Fig. 11. For S0 galaxies the mean value of ζ is -0.42 indicating that, on average, S0 galaxies have 34 per cent ($1 - e^{-0.42}$) fewer GCs than other galaxy types in the same range of luminosities. Generally speaking, galaxy types with 95 per cent credible intervals falling on the right-hand side of the dashed grey vertical line in Fig. 11 have more GCs than the overall population mean, while galaxy types on the left-hand side have fewer GCs than the population mean. While a detailed investigation of the causes of this behaviour is beyond the scope of this work, it is important to stress the ability of hierarchical Bayesian models to explore the multilevel statistical properties of the objects under study in a unified way.

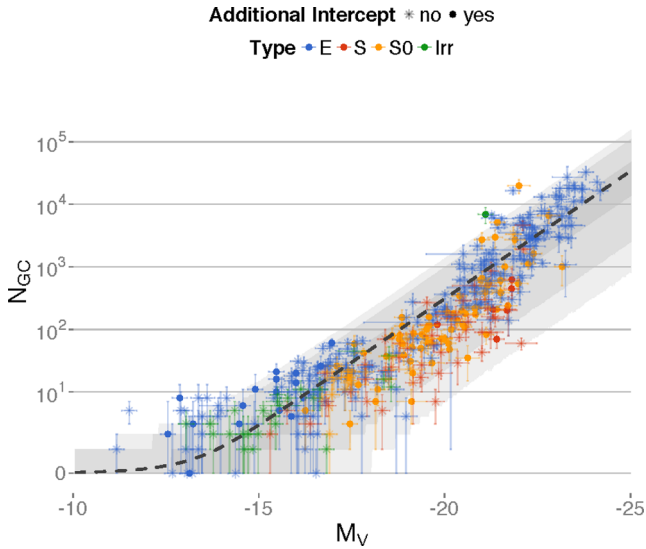


Figure 12. GC population, N_{GC} plotted against visual absolute magnitude M_V . The dashed line represents the expected value of N_{GC} for each value of M_V , while the shaded areas depicts the 50, 95 and 99 per cent prediction intervals. Galaxy types are coded by colours as follows: ellipticals (E; blue), spirals (S; red), lenticulars (S0; orange) and irregulars (Irr; green). Asterisks represent galaxies belonging to subtypes whose random intercept ζ is consistent with zero, while circles represent the ones with $\zeta \neq 0$. An ArcSinh transformation is applied in the y-axis for better visualization of the whole range of N_{GC} values, including the null ones.

7 CONCLUSIONS

We employed a Bayesian NB regression model to analyse the population size of GCs in the presence of galactic attributes such as central black hole mass, brightness and morphological type. Hence, demonstrating how GLMs designed to represent count data provide reliable outcomes and interpretations. The main scientific results and features of our analysis can be summarized as follows.

(i) The population size of GC is on average 35 per cent lower on S0 galaxies if compared to other galaxies with similar luminosities.

(ii) The relationship between the number of GCs and other galaxy properties has more variation than expected by a Poisson process, but can be well modelled by a NB GLM.

(iii) The Bayesian modelling herein employed naturally accounts for heteroscedasticity, intrinsic scatter and errors in measurements in both axes (either discrete or continuous).

(iv) Predicted intervals around the trend for expected N_{GC} envelope the data, including the Milky Way, which was previously considered an outlier.

(v) The random intercepts model (with a Bayesian LASSO) applied to the correlation between GC population and brightness allows us to account for the presence of 69 different galaxy subcategories of morphological classifications, and automatically identifies particular types not following the overall population mean. Galaxy types dE1N, dE0N, E/cD, S0, Sb0 and Sb show significant deviations from the general trend. Based on the sample studied here, we advise these types to be further scrutinized in order to clarify if there is any physical mechanism behind such deviations or merely an observational bias.¹²

¹² Type MG also shows significant deviation, but this is probably a consequence of small sample size (MG corresponds to only one object).

(vi) By employing a hierarchical Bayesian model for the random intercepts and unobserved covariates (e.g. true black hole mass), we allow the model to borrow strength across units. This happens via their joint influence on the posterior estimates of the unknown hyperparameters.

(vii) If extrapolated, the fitted model predicts a suppression in the presence of GCs for galaxies with luminosities $L \lesssim 2 \times 10^6 L_{\odot}$.

(viii) The central black holes mass is in fact a good predictor of the number of GCs. 1 dex increase in M_{BH} leads to an approximate five times increase in the incidence of GCs. The origin of such correlation is still a matter of debate. One possible explanation is that both properties are associated with a common event such as major mergers, thus galaxies experimenting a recent major merger should have a large M_{BH} mass and GC populations (e.g. Jahnke & Macciò 2011). The total mass of GCs and the central black hole mass can also correlate with the bulge binding energy in elliptical galaxies (e.g. Snyder et al. 2011; Saxton, Soria & Wu 2014). Rapid growth of the nuclear black hole of a galaxy might be fuelled by a massive inflow of cold gas towards the centre of the galaxy. The gas inflow would trigger star formation and the formation of GCs. Hence, leading to an indirect correlation between the total number of GCs and the M_{BH} . Scrutinizing which one among these and other possibilities, if any, are responsible for this correlation (causal or not) is beyond the purposes of this work. However, it does provide a clear example on how the adoption of modern statistical methods can point to intriguing astrophysical questions.

A statistical model is based on an appropriate probability distribution function assumed to generate or describe a data set. Hence, the parameter estimating likelihood function must specify a probability distribution on the appropriate scale under study. Discrete data, and count data in particular, are not continuous as are data described by the Gaussian distribution. The most appropriate way to model count data is by using a discrete probability distribution, e.g. a Poisson or NB likelihood, otherwise the model will likely be biased and misspecified – the price to be paid for employing the wrong likelihood estimator for the data of interest.

GLMs are a cornerstone of modern statistics, and an invaluable instrument for astronomical investigations given their potential application to a variety of astronomical problems beyond Gaussian assumptions. A prompt integration of these methods into astronomical analyses will allow contemporary statistical techniques to become common practice in the research of 21st century astronomy.

ACKNOWLEDGEMENTS

We thank Johannes Buchner for careful revision and constructive comments. The IAA Cosmostatistics Initiative (COIN)¹³ is a non-profit organization whose aim is to nourish the synergy between astrophysics, cosmology, statistics and machine learning communities. EEOI is partially supported by the Brazilian agency CAPES (grant number 9229-13-2). ALC-S acknowledges funding from a CNPq, BJT-A fellowship (400857/2014-6). MK acknowledges support by the DFG project DO 1310/4-1. This work was written on the collaborative OVERLEAF platform,¹⁴ and made use of the GITHUB¹⁵ repository web-based hosting service and GIT version control software. Work on this paper has substantially benefited from using the

¹³ <https://asaip.psu.edu/organizations/iaa/iaa-working-group-of-cosmostatistics>

¹⁴ www.overleaf.com

¹⁵ www.github.com

collaborative website AWOB (<http://awob.mpg.de>) developed and maintained by the Max-Planck Institute for Astrophysics and the Max-Planck Digital Library. The bibliographic research was possible thanks to the tools offered by the NASA Astrophysical Data Systems.

REFERENCES

- Andreon S., Hurn M. A., 2010, MNRAS, 404, 1922
 Andreon S., Hurn M., 2013, Stat. Analysis Data Mining, 6, 15
 Ata M., Kitaura F.-S., Müller V., 2015, MNRAS, 446, 4250
 Bernardo J., Bayarri J., Berger J., Dawid A., Heckerman D., 2011, Bayesian Statistics 9. Oxford Univ. Press, Oxford
 Brodie J. P., Strader J., 2006, ARA&A, 44, 193
 Burkert A., Tremaine S., 2010, ApJ, 720, 516
 Cameron A. C., Trivedi P. K., 2013, Regression Analysis of Count Data, 2nd edn. Cambridge Univ. Press, Cambridge
 De Souza R. S. et al., 2015, Astron. Comput., 12, 21
 Durrell P. R. et al., 2014, ApJ, 794, 103
 Elliott J., de Souza R. S., Krone-Martins A., Cameron E., Ishida E. E. O., Hilbe J., 2015, Astron. Comput., 10, 61
 Gebhardt K. et al., 2000, ApJ, 539, L13
 Gelman A., Hill J., 2007, Data Analysis Using Regression and Multi-level/Hierarchical Models (Analytical Methods for Social Research). Cambridge Univ. Press, New York
 Gelman A., Rubin D. B., 1992, Stat. Sci., 7, 457
 Griswold M., Parmigiani G., Potosky A., Lipscomb J., 2004, Biostatistics, 1, 1
 Hardin J. W., Hilbe J. M., 2012, Generalized Linear Models and Extensions, 3rd edn. Stata Press, College Station, TX
 Harris G. L. H., Harris W. E., 2011, MNRAS, 410, 2347
 Harris W. E., Harris G. L. H., Alessi M., 2013, ApJ, 772, 82
 Harris G. L. H., Poole G. B., Harris W. E., 2014, MNRAS, 438, 2117
 Hilbe J. M., 2011, Negative Binomial Regression, 2nd edn. Cambridge Univ. Press, Cambridge
 Hilbe J., 2014, Modeling Count Data. Cambridge Univ. Press, Cambridge
 Jahnke K., Macciò A. V., 2011, ApJ, 734, 92
 Jong P. D., Heller G. Z., 2008, Generalized Linear Models for Insurance Data. Cambridge Univ. Press, Cambridge
 Kruijssen J. M. D., 2014, Classical Quantum Gravity, 31, 244006
 Lansbury G. B., Lucey J. R., Smith R. J., 2014, MNRAS, 439, 1749
 Lin Y.-T., Mohr J. J., 2004, ApJ, 617, 879
 Lindsey J. K., 1999, Stat. Medicine, 18, 2223
 Marley J. K., Wand M. P., 2010, J. Stat. Software, 37, 1
 Marschner I. C., Gillett A. C., 2012, Biostatistics, 13, 179
 Nelder J. A., Wedderburn R. W. M., 1972, J. R. Stat. Soc., Ser. A, 135, 370
 O'Hara R. B., Kotze D. J., 2010, Methods Ecol. Evolution, 1, 118
 Park Trevor Casella George, 2008, J. Am. Stat. Assoc., 103, 681
 Raichoor A., Andreon S., 2012, A&A, 543, A19
 Raichoor A., Andreon S., 2014, A&A, 570, A123
 Rhode K. L., 2012, AJ, 144, 154
 Saxton C. J., Soria R., Wu K., 2014, MNRAS, 445, 3415
 Snyder G. F., Hopkins P. F., Hernquist L., 2011, ApJ, 728, L24
 Spiegelhalter D. J., Best N. G., Carlin B. P., Van Der Linde A., 2002, J. R. Stat. Soc. Ser. B, 64, 583
 Tibshirani R., 1996, J. R. Stat. Soc. Ser. B, 58, 267
 Tremaine S. et al., 2002, ApJ, 574, 740
 Uemura M., Kawabata K. S., Ikeda S., Maeda K., 2015, PASJ, 67, 55
 Zuur A., Hilbe J., Ieno E., 2013, A Beginner's Guide to GLM and GLMM with R. Highland Statistics, Newburgh, UK

APPENDIX A: JAGS MODEL

A1 Poisson GLM

The basic JAGS syntax for a Poisson GLM model:

```
GLM.pois <- model{
#Priors for regression coefficients

beta.0 dnorm(0, 0.000001)
beta.1 dnorm(0, 0.000001)

#Poisson GLM Likelihood

for (i in 1 : N){
eta[i] <- -beta.0 + beta.1 * x[i]
log(mu[i]) <- eta[i]
y[i] dpois(mu[i])
}
}
```

A2 Negative binomial GLM

The basic JAGS syntax for a NB GLM model:

```
GLM.NB <- model{
#Priors for regression coefficients
beta.0 dnorm(0, 0.000001)
beta.1 dnorm(0, 0.000001)
k dunif(0.001, 10)

#NB GLM Likelihood

for (i in 1 : N){
eta[i] <- -beta.0 + beta.1 * x[i]
log(mu[i]) <- eta[i]
p[i] <- -k/(k + mu[i])
y[i] dnegbin(p[i], k)
}
}
```

Another approach to fit a NB model in JAGS is via a combination of a Gamma distribution with a Poisson distribution in the form (see e.g. Marley & Wand 2010; Hilbe 2011)

```
GLM.NB <- model{
#Priors for regression coefficients
beta.0 dnorm(0, 0.000001)
beta.1 dnorm(0, 0.000001)
k dunif(0.001, 10)
```

```
#NB GLM Likelihood

for (i in 1 : N){
eta[i] < -beta.0 + beta.1 * x[i]
log(mu[i]) < -eta[i]
rateParm[i] < -k/mu[i]
g[i] dgamma(k, rateParm[i])
y[i] dpois(g[i], k)
}
}
```

APPENDIX B: BAYESIAN MODEL FOR EACH COVARIATE

B1 Dynamical mass versus globular cluster population

Bayesian NB GLM model for the relationship between N_{GC} and galaxy dynamical mass M_{dyn} . Since, there is no information about the uncertainties in the measurements of M_{dyn} , we neglect this in this particular model:

$$\begin{aligned}
N_{GC;i} &\sim \text{NB}(p_i, k), \\
p_i &= \frac{k}{k + \mu_i}, \\
\mu_i &= e^{\eta_i} + \epsilon_{N_{GC};i}, \\
\eta_i &= \beta_0 + \beta_1 M_{dyn;i}, \\
k &\sim \mathcal{U}(0, 5), \\
\epsilon_{N_{GC};i} &\sim \mathcal{B}(0.5, 2e_{N_{GC};i}) - e_{N_{GC};i}, \\
\beta_0 &\sim \mathcal{N}(0, 10^6), \\
\beta_1 &\sim \mathcal{N}(0, 10^6), \\
\alpha_0 &\sim \Gamma(0.01, 0.01), \\
\theta_0 &\sim \Gamma(0.01, 0.01), \\
i &= 1, \dots, N.
\end{aligned}
\tag{B1}$$

B2 Bulge velocity versus globular cluster population

Bayesian NB GLM model for the relationship between N_{GC} and bulge dispersion velocity σ :

$$\begin{aligned}
N_{GC;i} &\sim \text{NB}(p_i, k), \\
p_i &= \frac{k}{k + \mu_i}, \\
\mu_i &= e^{\eta_i} + \epsilon_{N_{GC};i}, \\
\eta_i &= \beta_0 + \beta_1 \sigma_i^*, \\
k &\sim \mathcal{U}(0, 5), \\
\sigma_i &\sim \mathcal{N}(\sigma_i^*, e_{\sigma_i}^2), \\
\epsilon_{N_{GC};i} &\sim \mathcal{B}(0.5, 2e_{N_{GC};i}) - e_{N_{GC};i}, \\
\beta_0 &\sim \mathcal{N}(0, 10^6), \\
\beta_1 &\sim \mathcal{N}(0, 10^6), \\
\sigma_i^* &\sim \Gamma(\alpha_0, \theta_0), \\
\alpha_0 &\sim \Gamma(0.01, 0.01), \\
\theta_0 &\sim \Gamma(0.01, 0.01), \\
i &= 1, \dots, N.
\end{aligned}
\tag{B2}$$

(B1) This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.