

Cluster population demographics in NGC 628 derived from stochastic population synthesis models

Jianling Tang ^{1,2}★, Kathryn Grasha ^{1,2,3}†‡ and Mark R. Krumholz ^{1,2}

¹Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia

²ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia

³Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

Accepted 2024 July 22. Received 2024 July 22; in original form 2023 January 15

ABSTRACT

The physical properties of star cluster populations offer valuable insights into their birth, evolution, and disruption. However, individual stars in clusters beyond the nearest neighbours of the Milky Way are unresolved, forcing analyses of star cluster demographics to rely on integrated light, a process fraught with uncertainty. Here, we infer the demographics of the cluster population in the benchmark galaxy NGC 628 using data from the Legacy Extra-galactic UV Survey (LEGUS) coupled to a novel Bayesian forward-modelling technique. Our method analyses all 1178 clusters in the LEGUS catalogue, ~ 4 times more than prior studies severely affected by completeness cuts. Our results indicate that the cluster mass function is either significantly steeper than the commonly observed slope of -2 or is truncated at $\approx 10^{4.5} M_{\odot}$; the latter possibility is consistent with proposed relations between truncation mass and star formation surface density. We find that cluster disruption is relatively mild for the first ≈ 200 Myr of cluster evolution; no evidence for mass-dependent disruption is found. We find suggestive but not incontrovertible evidence that inner galaxy clusters may be more prone to disruption and outer galaxy clusters have a more truncated mass function, but confirming or refuting these findings will require larger samples from future observations of outer galaxy clusters. Finally, we find that current stellar track and atmosphere models, along with common forms for cluster mass and age distributions, cannot fully capture all features in the multidimensional photometric distribution of star clusters. While our forward-modelling approach outperforms earlier backward-modelling approaches, some systematic differences persist between observed and modelled photometric distributions.

Key words: methods: statistical – techniques: photometric – software: data analysis – galaxies: individual: NGC 628 – galaxies: star clusters: general – galaxies: star formation.

1 INTRODUCTION

The distribution of star clusters’ properties (mass, age, and physical size) and the variation of this distribution with galactic environment provides crucial clues to the physics of star formation and star cluster evolution. Moreover, because of the processes responsible for setting this distribution may depend on galactic environment, cluster demographics can also trace the history of galaxy assembly and evolution, acting as a fossil record of the environments that existed when clusters formed (e.g. Adamo et al. 2015; Menon et al. 2021). Given the importance of cluster demographics, it is not surprising that there have been many attempts to measure them, and that cluster demographics figure prominently in the science cases for a number of large surveys of nearby galaxies, such as PHAT (Dalcanton et al. 2012), PHANGS-HST (Lee et al. 2022), and LEGUS (Calzetti et al. 2015); see

Krumholz, McKee & Bland-Hawthorn (2019a) for a detailed review.

Different aspects of cluster demographics probe different physics. By studying the shape of the cluster mass function (CMF), and in particular the initial CMF (ICMF) that applies to the youngest clusters, we can constrain star formation theories and provide an observational check on simulations. Studies to date show that over much of its range the ICMF is well described by a power-law $dN/dM \propto M^{\alpha_M}$ with a slope $\alpha_M \approx -2$ across a wide range of galaxy properties, corresponding to equal mass per logarithmic bin and suggesting a scale-free formation process (e.g. Fall & Chandar 2012; Chandar et al. 2014). However, the shape of the ICMF at its high-mass end ($\gtrsim 10^4 M_{\odot}$) remains uncertain. Some authors report that a truncated distribution such as a Schechter function describes the data better than the pure power law (e.g. Bastian et al. 2012; Adamo et al. 2017; Johnson et al. 2017), while others question the statistical significance of claimed detections and instead suggest that the dearth of massive clusters is simply a size-of-sample effect (e.g. Larsen 2009; Mok, Chandar & Fall 2020). If there is a real truncation in the mass function, its location must depend somehow on the galactic environment, since rapidly star-forming galaxies with

* E-mail: janett.jianling@gmail.com

† ARC DECRA Fellow.

‡ Visiting Fellow.

large cluster populations often harbour clusters with masses larger than the reported truncation masses in more modestly star-forming galaxies (e.g. Chandar, Fall & Whitmore 2015; Linden et al. 2017). The possible absence of a high-mass break in the ICMF, as well as its value and variations with the galactic environment, can provide important clues about the process of star formation.

As clusters evolve and disperse from their birthplaces, they experience mass-loss. Various theories govern how they disrupt, ranging from the ‘infant mortality’ stage of rapid gas removal after star formation (Hills 1980; Baumgardt & Kroupa 2007; Dinbier & Walch 2020) to long-time-scale processes such as two-body relaxation (Fall & Zhang 2001; Lamers, Baumgardt & Gieles 2010; Kruijssen et al. 2012) to processes such as external tidal shocking that operate on intermediate time-scales (Bastian et al. 2005; Lamers et al. 2005b; Gieles et al. 2006; Gieles & Renaud 2016; Howard, Pudritz & Harris 2017; Fensch et al. 2019; Webb, Reina-Campos & Kruijssen 2019). To place observational constraints on these theories, we study the cluster age function (CAF). As with the ICMF, observations suggest that the CAF is reasonably well described by a power law form $dN/dT \propto T^{\alpha_T}$, where T is cluster age. However, the value of α_T remains highly controversial, partly due to disagreement among observational groups about what constitutes a cluster (Krumholz et al. 2019a). Despite this uncertainty, the CAF encodes the time-scales of diverse cluster disruptive processes. One question of particular interest is the form of the joint mass–age distribution, which gives a complete description of cluster formation and disruption. If the mass and age distributions are separable, i.e. if $d^2N/dM dT \propto (dN/dM)(dN/dT)$, this implies mass-independent disruption (MID) of clusters (Fall, Chandar & Whitmore 2005, 2009; Chandar et al. 2015), while if they are not this implies mass-dependent disruption (MDD; Lamers, Gieles & Portegies Zwart 2005a; Gieles, Lamers & Portegies Zwart 2007; Grudić et al. 2021). Knowing whether MID or MDD holds in a particular galaxy or sub-galactic region would in turn place strong constraints on the mechanisms by which clusters disrupt.

Part of the reason that observations have not yet settled debates about the existence and value of a cutoff in the ICMF, the slope of the CAF, and whether cluster disruption is mass-dependent or mass-independent, is that determining the properties of clusters from observations is not trivial. To sample a wide range of environments, studies of cluster populations must work with integrated light rather than resolved stellar populations, since in the crowded environments of star clusters, it is generally only possible to resolve individual stellar sources for a handful of the most nearby galaxies. The traditional approach to extracting cluster demographics from this type of data is to convert the integrated light measurements for each cluster to physical properties such as mass and age by comparing the observed photometry to a grid of simple stellar population (SSP) models and generating a set of best-fitting parameters. One then bins the clusters by mass and age to obtain mass and age distributions.

However, this approach encounters several difficulties. First, the binning process usually entails the loss of useful information, and as a result parameters determined by fitting data to histograms are heavily biased (Maschberger & Kroupa 2009; El-Badry, Weisz & Quataert 2017). Secondly, the conventional approach of using χ^2 fitting to convert photometry to masses and ages implicitly assumes that the uncertainties on cluster mass and age can be approximated as Gaussian. This is often a poor assumption, because the mapping between physical properties and photometry is non-linear and non-monotonic, particularly once one adds the additional complication of dust extinction. Consequently, a particular set of photometric measurements may plausibly fit two (or more) widely separated

loci in physical parameter space, yielding posterior mass and age distributions with complex non-symmetrical and multimodal shapes (de Meulenaer et al. 2014; Fouesneau et al. 2014; Krumholz et al. 2015b; Bialopetravičius, Narbutis & Vansevicius 2019). The problem is exacerbated for low-mass clusters, where stochastic sampling of the IMF leads to a wide range of possible photometry even for clusters of fixed mass and age (Piskunov et al. 2009).

However, the most severe issue for measuring cluster demographics is completeness. The problem is that star clusters become fainter as one moves to both lower mass and older age, so a magnitude limit corresponds to a complex shape in the parameter space of age and mass, meaning that both the mass and age distributions are subject to large potential biases. The conventional way of handling this is to discard clusters outside of a limited range of both mass and age over which completeness is expected to be $\gtrsim 90$ per cent, for example, to retain only clusters with estimated masses $5000 M_\odot$ and estimated ages < 200 Myr in LEGUS. Such a drastic truncation of the sample avoids bias, but at the cost of a large loss of statistical power at both lower masses and old ages. The former is particularly concerning, since the steep mass function means that low-mass clusters form the majority of the available sample; the sample truncation required to avoid bias from survey magnitude limits therefore often involves discarding the majority of the data.

The main objective of this paper is to present a complete analysis pipeline and preliminary results for cluster demographics in the benchmark galaxy NGC 628 using the catalogue of clusters measured by the *Hubble Space Telescope* Treasury Program Legacy Extragalactic Ultraviolet Survey (LEGUS; Calzetti et al. 2015) coupled to the novel Bayesian forward modelling method proposed by Krumholz et al. (2019b) that overcomes the problems identified above. Specifically, this method naturally copes with complex and non-deterministic mappings between photometric measurements and physical properties, and it enables us to use a total of 1178 clusters catalogued by LEGUS with non-zero completeness values, as compared to earlier modelling where severe completeness cuts reduced the sample to ≈ 300 (Adamo et al. 2017). This paper is structured as follows: In Section 2, we introduce our target galaxy, cluster catalogue, and analysis of completeness. In Section 3, we summarize our analysis method. We present our results for cluster demographics in Section 4, and discuss their implications in Section 5. We summarize the findings and discuss future prospects in Section 6.

2 DATA DESCRIPTION

Here, we summarize the properties of our target galaxy (Section 2.1), the star cluster catalogue we use as the basis for our study (Section 2.2), and our treatment of survey completeness (Section 2.3).

2.1 NGC 628

Our target NGC 628, also known as Messier 74, is a well-studied grand-design spiral galaxy, located at 9.9 Mpc (Olivares et al. 2010). We choose NGC 628 for this study as its large size (apparent radius of 5.2 arcmin; Gusev, Egorov & Sakhilov 2014) and face-on orientation allow for a detailed study of the stellar cluster population. We use observations of NGC 628 taken from the LEGUS survey, a Cycle 21 HST Treasury programme that targeted 50 local galaxies ($\lesssim 15$ Mpc) with the *Hubble Space Telescope* with broad-band filter coverage from the UV to the near IR. All targets were imaged with either the Wide Field Camera 3 (WFC3) or the Advanced Camera for Surveys (ACS) in the NUV (WFC3 *F275W*), *U* band (*F336W*), *B* band

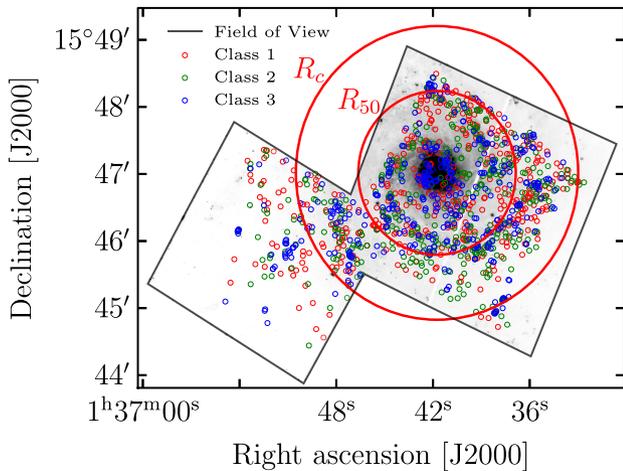


Figure 1. V -band ($F555W$) mosaic image of the centre and east fields of NGC 628 from LEGUS-HST. Symbols show the location of the star clusters, coloured by the cluster classifications: class 1 (red circle), class 2 (green circle), and class 3 (blue circle). The black boxes outline the observational field of view. For a detailed description of cluster classification method, see Section 2.2. The two large red circles represent the co-rotation radius R_c (Section 4.3.2), and the median galactocentric radius of the LEGUS star cluster catalogue R_{50} (Section 4.3.1).

(ACS/ $F435W$ or WFC3/ $F438W$), V band (ACS or WFC3/ $F555W$ or $F606W$), and I band (ACS or WFC3 $F814W$). For the remainder of this paper, we follow the conventional Johnson passband naming convention UV , U , B , V and I without converting to the Johnson system. As part of the LEGUS survey, NGC 628 is covered by two pointings; one at the galactic centre (NGC 628c) and the other at the eastern edge (NGC 628e). We show the LEGUS V -band image of NGC 628 in Fig. 1. Detailed descriptions of the survey and data reduction are provided in Calzetti et al. (2015).

2.2 Cluster catalogue

The photometric star cluster catalogue for NGC 628 that we use in this study comes from the LEGUS survey, and we refer readers to Grasha et al. (2015) and Adamo et al. (2017) for a complete description of how it is constructed. Here, we provide a brief description of the parts of the process most pertinent to this study. LEGUS uses a two-step pipeline for cluster classification. First, an automated cluster extraction tool is used to extract potential cluster candidates with at least a 3σ detection in a minimum of five contiguous pixels from the white light images. The cluster candidates are then further refined and required to have a V -band concentration index (CI; the difference in magnitude between a radius of 1 and 3 pixels) of > 1.4 mag for the centre pointing and > 1.3 mag for the east pointing. The CI cut serves to separate point-like sources (i.e. stars) from more extended sources (i.e. star clusters). The star cluster candidates are required to be detected in the V band and a minimum of at least three of the other filters, with photometric error $\sigma_\lambda \leq 0.3$ mag in each band.

To create the final cluster catalogue, the LEGUS team then visually inspects all candidates brighter than -6 in the V band and classifies them into four morphological categories: (1) centrally concentrated clusters with spherically symmetric profiles; (2) clusters with asymmetric radial profiles; (3) multi-peaked clusters with underlying diffuse emission; and (4) non-cluster contaminants, such as background galaxies, stars, bad pixels, or edge artefacts. In this work,

we limit our analysis to Class 1 (426), 2 (437), and 3 (413) clusters, for a total sample of 1276; we reduce this to 1275 by removing one cluster that we found to be a duplicate entry in the LEGUS NGC 628c and 628e catalogues, located in the small region where the fields overlap. Fig. 1 shows the mosaic V -band image of NGC 628 overlaid with the cluster positions, coloured by their morphological classification.

2.3 Observational completeness limits

Because we intend to forward model the cluster population, we require knowledge of the completeness of the observational catalogue. Specifically, we require knowledge of the function $P_{\text{obs}}(\mathbf{m})$, which describes the probability that a hypothetical star cluster with a vector of magnitudes \mathbf{m} in the various LEGUS filters would be included in the cluster catalogue. To compute this function, we employ the completeness limits reported by Adamo et al. (2017), who carry out artificial cluster tests to derive the completeness of the LEGUS automated catalogue generation procedure in each filter independently. For each filter, F and a range of magnitudes m_F in that filter, they determine the probability $P_{\text{obs},F}(m_F)$ that a cluster would be recovered by the automated extraction procedure. In what follows, we linearly interpolate these tabulated data to obtain a continuous function $P_{\text{obs},F}(m_F)$ that gives the probability that a cluster of arbitrary magnitude m_F in filter F will be recovered.

To determine the completeness for a cluster with a vector of magnitudes \mathbf{m} , we use 10 000 Monte Carlo trials. In each trial, we randomly assign each filter F to be a detection or a non-detection with probability $P_{\text{obs},F}(m_F)$ as determined from the interpolated artificial cluster test results. We then determine from this set of detections and non-detections if the cluster would be catalogued following the same criteria used in construction of the actual LEGUS catalogue, i.e. the cluster is catalogued only if it is (1) detected in V band and at least three other bands, and (2) has a visual magnitude $M_V \leq -6$. We then take $P_{\text{obs}}(\mathbf{m})$ to be equal to the fraction of the Monte Carlo trials in which the cluster is catalogued.

We use this method to calculate the completeness both of our library of synthetic clusters (see Section 3) and of the actual LEGUS catalogue. For the latter, we find a total of 97 clusters for which we estimate $P_{\text{obs},F}(\mathbf{m}) = 0$; these can be present in the catalogue because a small number of clusters with $m_V > -6$ were added by hand. To avoid introducing errors in our completeness estimate, we remove these from the sample. We remove a total of 97 clusters from the sample of NGC 628 due to their zero completeness values, resulting in a final sample size of 1178 clusters.

3 METHODS

In this section, we describe the pipeline used to derive the cluster demographics, motivated by the forward Bayesian modelling approach demonstrated in Krumholz et al. (2019b). For reasons of brevity, we only summarize the method, and refer readers to Krumholz et al. for details.

3.1 Overview of the method

Given a set of unresolved photometric measurements of a cluster population, how do we infer their underlying demographics? To answer this question, we first propose a joint distribution of mass M , age T , and extinction A_V for the cluster population, which we denote as $f(M, T, A_V | \theta)$, where θ represents a vector of parameters describing the joint distribution f . For example, if we assume the

mass distribution of clusters is described by a Schechter function $f(M, T, A_V | \theta) \propto M^{\alpha_M} \exp(-M/M_{\text{break}})$, then θ contains the slope α_M and break mass M_{break} of the Schechter function. We describe the functional forms we consider for $f(M, T, A_V | \theta)$ and the parameters they involve in Section 3.3.

We seek to derive the posterior distribution of θ . We do so in the usual way for a Bayesian method, but writing the posterior as the product of a prior and a likelihood function

$$p_{\text{post}}(\theta | \{\mathbf{m}\}) \propto \mathcal{L}(\{\mathbf{m}\} | \theta) p_{\text{prior}}(\theta). \quad (1)$$

where $\mathcal{L}(\{\mathbf{m}\} | \theta)$ is the likelihood for our set of N_{obs} photometric measurements. We defer a discussion of priors to Section 3.3, and present details of our method of calculating the likelihood function in Section 3.2, but to summarize the latter here, we compute the likelihood function using a Gaussian mixture model derived from a large library of synthetic star clusters; as we change θ , we adjust the weights applied to this library, which changes the value of $\mathcal{L}(\{\mathbf{m}\} | \theta)$. Our method can therefore be summarized into the following steps:

- (i) We generate a library composed of synthetic clusters with weights based on a proposed distribution of cluster physical properties and the observational completeness of the survey.
- (ii) Using the newly created library, we produce a synthetic distribution in photometric space, which will be used to compare with the observations.
- (iii) We adjust the model parameters θ to maximize the resemblance between the observations and the photometric distribution, as parameterized by the likelihood function. As we do so, we map out the posterior distribution.

We carry out the final step of this procedure using a Markov chain Monte Carlo (MCMC) method as implemented in the software package EMCEE (Foreman-Mackey et al. 2013). For all the calculations presented in this paper, we use 4000 iterations of 100 walkers, discarding the first 300 iterations for burn-in; both visual inspection of the posteriors and quantitative evaluation of the autocorrelation time indicates that the chains are well converged. We describe the first two steps of the method – generating and re-weighting the library, in Section 3.2.

3.2 Calculation of the likelihood function

As discussed above, we compute the likelihood function using the Gaussian mixture model described in Krumholz et al. (2019b), which operates on a library of N_{lib} synthetic clusters, each of which is characterized by a mass M , age T , extinction A_V , and a vector of photometric magnitudes \mathbf{m} in each of the filters used in the observations. The full library is further described by a vector of photometric bandwidths \mathbf{h} , which we set to 0.1 mag in all filters; see Krumholz et al. (2015b, 2019b) for detailed discussion of the meaning of the bandwidth and the motivation for choosing this value. The synthetic clusters in the library are generated using the Padova tracks including models for asymptotic giant branch (AGB) stars (Vázquez & Leitherer 2005), and are integrated with the ‘STARBURST99’ treatment of stellar atmospheres (Leitherer et al. 1999). Details of the library construction are provided in Appendix A. We note that the choice of stellar tracks does affect some of our results for cluster demographics. We consider both Padova-AGB and MIST (Choi et al. 2016) models, and choose to use the former because they yield model luminosity functions that more closely match the observations. See Appendix B for details. To model extinctions, we adopt a Milky Way extinction curve in which the optical and UV (UV , U , B and V bands) extinctions are taken from Fitzpatrick (1999), and

the IR (I band) extinctions are obtained from Landini et al. (1984). Given the library, Krumholz et al. (2019b) show that the likelihood function for the parameters θ can be written as (their equation 12)

$$\mathcal{L}(\{\mathbf{m}\} | \theta) \propto \prod_{i=1}^{N_{\text{obs}}} \left[\mathcal{A}(\theta) \sum_{j=1}^{N_{\text{lib}}} w_j(\theta) \mathcal{N}(\mathbf{m}_i | \mathbf{m}_j, \mathbf{h}'_i) \right], \quad (2)$$

where $w_j(\theta)$ is the statistical weight of the j th library cluster (which depends on the parameters θ as described below), \mathbf{m}_i and \mathbf{m}_j are the vectors of magnitudes for the i th observed and j th library clusters, respectively, $\mathcal{A}(\theta) = [\sum_{j=1}^{N_{\text{lib}}} w_j(\theta)]^{-1}$ is a normalization factor, and $\mathcal{N}(\mathbf{x} | \mathbf{x}_0, \sigma)$ is the standard multidimensional Gaussian distribution centred at \mathbf{x}_0 and with width σ , evaluated at position \mathbf{x} . The quantity \mathbf{h}'_i is given by $\mathbf{h}'_i = (\mathbf{h}^2 + \sigma_i^2)^{1/2}$, where σ_i the observational error on the magnitude of cluster i ; the quantities \mathbf{h}'_i , \mathbf{h} , and σ_i are vectors with one element per filter, and the expression for \mathbf{h}'_i should be understood as applying separately to each filter. If the magnitude of a specific filter of the cluster is undetected, the product of the observed and library cluster magnitudes becomes zero, thus not contributing to the likelihood function. In practice, we evaluate equation (2) numerically using the CLUSTER_SLUG module of the SLUG software suite (Krumholz et al. 2015a), which implements a tree-based order $N_{\text{obs}} \ln N_{\text{lib}}$ algorithm for performing the calculation that is much faster than a naive brute force evaluation, which would have a computational cost of the order of $N_{\text{obs}} N_{\text{lib}}$. The CLUSTER_SLUG module and the entire SLUG (da Silva, Fumagalli & Krumholz 2012, 2014; Krumholz et al. 2015a) software suite are freely available from <http://www.slugsp.com>.

The quantity in square brackets in equation (2) is the distribution of photometric magnitudes for the cluster library evaluated at the vector of magnitudes \mathbf{m}_i for the i th observed cluster. This distribution, and thus the likelihood function as whole, depends on the parameters describing the cluster population θ only through the weight functions $w_j(\theta)$ that describe the statistical weight of each library cluster. The relationship between weights and θ is given by (equation 9 of Krumholz et al. 2019b)

$$w_j(\theta) = P_{\text{obs}}(\mathbf{m}_j) \frac{f(M_j, T_j, A_{V,j} | \theta)}{p_{\text{lib}}(M_j, T_j, A_{V,j})}. \quad (3)$$

The denominator $p_{\text{lib}}(M_j, T_j, A_{V,j})$ is the distribution function describing the sampling density of the library (see Appendix A), while the numerator $f(M_j, T_j, A_{V,j} | \theta)$ is distribution of the physical properties given θ , and the pre-factor $P_{\text{obs}}(\mathbf{m}_j)$ is the probability that a cluster with vector of magnitude \mathbf{m}_j would be included in the LEGUS catalogue. The terms in equation (3) can be intuitively understood as follows. The factor $P_{\text{obs}}(\mathbf{m}_j)$ accounts for the fact that only a fraction of clusters with magnitudes \mathbf{m}_j will be observed due to the completeness limits; we compute this probability as described in Section 2.3. The denominator is the probability density for drawing clusters with a particular combination of physical parameters $(M_j, T_j, A_{V,j})$ while constructing the library. Finally, the numerator represents the true probability density for a cluster population described by the parameter set θ , so that the ratio $f(M_j, T_j, A_{V,j} | \theta) / p_{\text{lib}}(M_j, T_j, A_{V,j})$ represents the factor by which we must up- or down-weight the library so that clusters in the library have the same mass, age, and extinction distribution as clusters in reality; if this weight factor is unity, then our library is sampled from the same distribution of cluster properties as the real population.

Intuitively, then, our method consists of iteratively adjusting the parameters θ and thus the weights $w_j(\theta)$ to bring the predicted photometric distribution into as close agreement as possible with the observed one. This will, in turn, adjust the mass and age distributions,

since these are determined by the same vector of parameters θ as the luminosity distribution. Again, we remind readers that this is just an intuitive description of the underlying process; formal proof that equation (2) is the correct likelihood function to accomplish this adjustment, along with some details of how we handle complications like clusters where some filters are missing due to the fields of view in the different filters not being perfectly overlapping, is provided in Krumholz et al. (2019b).

3.3 Demographic models and priors

Cluster demographics depend on cluster formation and destruction mechanisms, and thus the set of candidate parametric model distributions $f(M, T, A_V | \theta)$ we consider is necessarily informed by theoretical expectation. In this work, we consider the two most prominent models: MID (Section 3.3.1; Fall et al. 2005, 2009; Chandar et al. 2015) and mass-dependent destruction/disruption (MDD; Section 3.3.2; Lamers et al. 2005a; Gieles et al. 2007; Grudić et al. 2021). We couple both to a parametric model for the distribution of extinctions (Section 3.3.3).

3.3.1 The mass-independent disruption model

In the MID model, the rate at which star clusters are destroyed is independent of cluster mass. The CMF therefore has the same shape at all ages, and the distribution function $f(M, T, A_V | \theta)$ can be separated into two distinct functions, one describing the mass distribution $p_M(M)$ and one the age distribution $p_T(T)$. Given the observational evidence that the mass function is a (possibly) truncated power law, we will adopt a functional form for $p_M(M)$ given by a Schechter function,

$$p_M(M) \propto M^{\alpha_M} \exp\left(-\frac{M}{M_{\text{break}}}\right). \quad (4)$$

For the age distribution, photometry cannot differentiate between a bound unbound cluster, so the age distribution is required to be flat for times that are shorter than the physical time required for the stars in a cluster to disperse; of course the distribution can also be flat out to older ages if the mechanisms responsible for disruption do not begin until some time after a cluster forms. Regardless of its physical, origin, we call the time at which cluster disruption begins T_{MID} . After this time, clusters will disrupt, and we approximate the age distribution as a power law. We therefore have

$$p_T(T) \propto \begin{cases} 1, & \text{if } T < T_{\text{MID}} \\ \left(\frac{T}{T_{\text{MID}}}\right)^{\alpha_T}, & \text{if } T > T_{\text{MID}} \end{cases}. \quad (5)$$

Thus, the joint mass-age distribution in the MID model is

$$\frac{d^2 N}{dM dT} \propto M^{\alpha_M} \exp\left(-\frac{M}{M_{\text{break}}}\right) \max(T, T_{\text{MID}})^{\alpha_T}. \quad (6)$$

The MID model therefore has four free physical parameters that we place in our vector θ : α_M , $\log M_{\text{break}}$, α_T , and $\log T_{\text{MID}}$. We adopt flat priors on α_M from -4 to 0 and α_T from -3 to 0 , reflecting a broad range around previous literature values; we will see that these choices have little effect, as our MCMC never approaches these boundaries. The priors on $\log M_{\text{break}}$ and $\log T_{\text{MID}}$ require somewhat more thought. As for $\log M_{\text{break}}$, we impose a flat prior from 2 to 7 because with we barely see clusters more massive than $10^{6.5} M_{\odot}$. The lower mass limit of 2 in log scale is to ensure a reasonable MCMC walker range. For $\log(T_{\text{MID}}/\text{yr})$, we impose flat priors from 5 to 10 based on physical plausibility. At the lower end, clusters cannot

disperse on less than a crossing time-scale, and even the densest clusters detected in LEGUS have crossing time-scales well above 10^5 yr. The upper limit is roughly the age of the Universe.

3.3.2 The mass-dependent disruption model

For the MDD model, clusters lose mass at a rate that varies as a pure power-law function of their current mass (Gieles et al. 2007; Grudić et al. 2021), $dM/dT \propto -M^{\gamma_{\text{MDD}}}$. For such a mass-loss rate, the present-day mass of a cluster born with initial mass M_i at age T is

$$M = M_i \left[1 - \gamma_{\text{MDD}} \left(\frac{M_0}{M_i} \right)^{\gamma_{\text{MDD}}} \frac{T}{T_{\text{MDD},0}} \right]^{1/\gamma_{\text{MDD}}}, \quad (7)$$

where $T_{\text{MDD},0}$ represents the time required for a cluster with mass M_0 to have fully disrupted, defined as having reached a present-day mass $M = 0$. The joint mass–age distribution therefore obeys

$$\frac{d^2 N}{dM dT} \propto \frac{d^2 N}{dM_i dT} \frac{dM_i}{dM}. \quad (8)$$

If the distribution of initial masses M_i follows the Schechter function form given by equation (4), then we have

$$\frac{d^2 N}{dM dT} \propto M^{\alpha_M} \eta^{\alpha_M+1-\gamma_{\text{MDD}}} \exp\left(-\eta \frac{M}{M_{\text{break}}}\right), \quad (9)$$

where

$$\eta(M, T) = \left[1 + \gamma_{\text{MDD}} \left(\frac{M_0}{M} \right)^{\gamma_{\text{MDD}}} \frac{T}{T_{\text{MDD},0}} \right]^{1/\gamma_{\text{MDD}}} \quad (10)$$

is the ratio of the initial and present-day cluster masses for a cluster of present-day mass M and age T .

The MDD model therefore has four free parameters: α_M , $\log M_{\text{break}}$, γ_{MDD} , and $T_{\text{MDD},0}$; note that M_0 is not a separate parameter, because the mass–age distribution depends only on the combination of parameters $M_0^{\gamma_{\text{MDD}}}/T_{\text{MDD},0}$. We therefore without loss of generality choose $M_0 = 100 M_{\odot}$ in what follows. However, for reader convenience we will also report the commonly used t_4 parameter, which is simply the disruption time for a cluster of mass $10^4 M_{\odot}$; this is given by $t_4 = 10^{2\gamma_{\text{MDD}}} T_{\text{MDD},0}$. We adopt the same priors on α_M and $\log M_{\text{break}}$ as in the MID model (Section 3.3.1). For γ_{MDD} , previous observational estimates and N -body simulations give values in the range 0.6 – 0.7 (Lamers, Baumgardt & Gieles 2010), and we adopt broad priors that include this range: we take γ_{MDD} to be flat from 0 to 1 . These limits stem from physical considerations: if $\gamma_{\text{MDD}} < 0$ then low-mass clusters lose mass more slowly than massive ones, contrary to the physical expectations of the model, while if $\gamma_{\text{MDD}} \geq 1$ then no clusters ever disrupt because there is no T for which $M = 0$. Finally, we adopt flat priors on $\log(T_{\text{MDD},0}/\text{yr})$ from 4 to 10 ; these limits are broad enough not to matter, because none of our walkers ever approach them.

3.3.3 Dust extinction

The exact functional shape of the distribution of dust extinctions $p_{A_V}(A_V)$ is unknown, so we model it as non-parametrically as possible. Following Krumholz et al. (2019b), we adopt a simple piece-wise linear form over the range $A_V = 0 - 3$ mag characterized by six free parameters ($p_{A_V,i}$, $0 \leq i \leq 6$) to be fit, representing the value of the probability density function (PDF) at $A_V = i \Delta A_V$ mag with $\Delta A_V = 0.5$ mag. Thus, the functional form we adopt for the

extinction distribution is

$$p_{A_V}(A_V) \propto \begin{cases} p_{A_V,i} + (p_{A_V,i+1} - p_{A_V,i}) \left(\frac{A_V}{\Delta A_V} - i \right), & i < \frac{A_V}{\Delta A_V} \leq i+1 \\ 0, & A_V > 3 \text{ mag} \end{cases} \quad (11)$$

We treat $p_{A_V,i}$ for $i = 0 - 5$ as parameters of our model to be fit, with $p_{A_V,6}$ fixed by the requirement that the total area under the PDF be unity. We set the priors on $p_{A_V,i}$ to be flat for all values > 0 , subject to the requirement that $p_{A_V,i} > 0$ remain positive for all A_V .

Combining this with the mass and age distributions, our final functional form to be fit is

$$f(M, T, A_V | \theta) \propto \frac{d^2 N}{dM dT} p_{A_V}(A_V), \quad (12)$$

with $d^2 N/dM dT$ given by equation (6) or equation (9) for the MID or MDD models, respectively.

3.4 Model selection using Akaike weights

The method described thus far allows us to compute the posterior PDFs of the model parameters θ for both the MID and MDD models. To determine whether the MID or MDD model provides a better fit and more accurate description of the data, we use the Akaike information criterion (AIC) to assess the fit quality provided by non-nested models. For models tested using AIC, we first identify the walker with the highest likelihood, then compute,

$$\text{AIC}_{(\text{MID}, \text{MDD})} = 2k - 2 \ln \hat{\mathcal{L}}_{(\text{MID}, \text{MDD})}, \quad (13)$$

where k represents the number of free parameters in the models, and $\hat{\mathcal{L}}$ is the maximum of the likelihood function. In our case, k is 11 for both the MID and MDD models, with four parameters describing the joint mass–age distribution (α_M , M_{break} , and either α_T and T_{MID} or γ_{MDD} and $T_{\text{MDD},0}$, six parameters representing the dust extinction shape, and one extra parameter to describe the total number of clusters present in the galaxy.

For the model comparison, we calculate the Akaike weights for either the MID or MDD model as

$$w_{\text{MID}/\text{MDD}} = \frac{\frac{e^{-\Delta_{\text{MID}/\text{MDD}}}}{2}}{\frac{e^{-\Delta_{\text{MID}}}}{2} + \frac{e^{-\Delta_{\text{MDD}}}}{2}}. \quad (14)$$

The AIC measures the amount of information in the data preserved by a given model, with the relative Akaike weight w of one model indicating the confidence level at which we can claim that it preserves more information than the other models considered. The model with the highest Akaike weight is our best fit and thus the preferred model.

4 RESULTS

In this section, we report our fit parameters and model comparison results (Section 4.1), along with comparisons between models and observed photometry to verify that our best-fitting models do a reasonable job at reproducing the observations (Section 4.2). To search for variations in cluster population demographics with galactocentric radius, we also separately analyse clusters in the inner and outer galaxy (Section 4.3).

4.1 Full catalogue fits

We first analyse the full LEGUS NGC 628 catalogue using the method described in Section 3. We summarize the marginal posteriors

we derive on all model parameters in the first block in Table 1. We also report results for the nuisance parameters describing the dust distribution in Appendix C. Our model comparison yields $w(\text{MID}) \approx 1$ and $w(\text{MDD}) < 10^{-11}$, suggesting that the MID model does a substantially better job capturing the variations in the data, and we will therefore focus on this as our preferred model from this point forward.

Fig. 2 shows the posterior PDFs of α_M , $\log M_{\text{break}}$, α_T , and $\log T_{\text{MID}}$ we obtain from our fit. In this plot, the extent of the axes reflects the full range of model parameters allowed by our priors, so we can immediately read off where parameters are well constrained by the data, versus where they are unconstrained and occupy the full range of values allowed by our priors. These figures have a few noteworthy features. First, they suggest that cluster disruption in NGC 628 is weak to non-existent out to ages of ≈ 200 Myr – the posterior PDF suggests that either $T_{\text{MID}} \approx 2 \times 10^8$ yr and that disruption becomes fairly rapid only after this point ($\alpha_T \approx -1.4$) or that disruption begins early, $T_{\text{MID}} \approx 2$ Myr but is very mild ($\alpha_T \approx -0.3$); in either scenario, disruption is limited at the young to moderate ages where our sample has most of its statistical power. Secondly, the fits provide evidence that the mass function is better described by either a truncated Schechter-form CMF with a slope close to the commonly found $\alpha_M = -2$ but then an exponential cutoff at $M_{\text{break}} \approx 10^{4.5} M_{\odot}$ or by a steeper power law with $\alpha_M \approx -2.3$.

4.2 Photometric comparisons

Before accepting the results of our fits, we must validate that our models adequately recover the observed luminosity distribution, since matching this distribution is the goal of our forward model.

4.2.1 Luminosity functions

To visualize the comparisons between the distributions of observations and the library, we first plot the 1D cluster luminosity distributions in the five LEGUS bands, showing both the measured distribution and the luminosity distribution we predict using the model parameters. The model prediction follows immediately from our expression for the likelihood function in terms of our Gaussian mixture model (equation 2), and is simply

$$p(m) \propto \sum_{j=1}^{N_{\text{lib}}} w_j(\theta) \mathcal{N}(m | m_j, h), \quad (15)$$

where m is the magnitude in the filter of interest, m_j is the magnitude of the j^{th} library cluster in that filter, and h is the library bandwidth. For the observations, we marginalize over the observational scatter using a bootstrap resampling method. Specifically, the LEGUS catalogue reports a central value and an uncertainty (assumed to be Gaussian) on the magnitude of each cluster in each photometric band; to generate the observed distributions, we draw 20 000 samples from this Gaussian for each cluster, and generate an observed luminosity distribution for each realization. We then plot the median and 5th to 95th percentile range of these 20 000 trails.¹

¹We caution that near the completeness limit of the LEGUS sample this resampling procedure is subject to Malmquist bias. The nature of the bias is as follows: for clusters near the $V = -6$ mag lower limit on the LEGUS catalogue, there is a significant chance in any given realization that they will wind up being assigned a V magnitude fainter than -6 , leading to them being excluded from the catalogue we produce for that realization. If the original LEGUS catalogue from which we started included clusters fainter

Table 1. Fitting results.

Catalogue	Model	w	α_M	$\log(M_{\text{break}}/M_\odot)$	α_{M_4}	α_T	$\log(T_{\text{MID}}/\text{yr})$	$\log(T_{\text{MDD},0}/\text{yr})$	γ_{MDD}	$t_4(\text{Myr})$
<i>Full sample</i>										
All	MID	1.00	$-2.16^{+0.32}_{-0.04}$	$5.61^{+0.49}_{-1.24}$	$-2.85^{+0.10}_{-0.05}$	$-1.41^{+1.19}_{-0.40}$	$8.22^{+0.07}_{-1.86}$	–	–	–
	MDD	$<10^{11}$	$-1.40^{+0.07}_{-0.08}$	$4.00^{+0.05}_{-0.05}$	–	–	–	$6.65^{+0.10}_{-0.07}$	$0.98^{+0.02}_{-0.04}$	268^{+37}_{-29}
<i>Radial division</i>										
$\{r \leq R_{50}\}$	MID	0.6027	$-1.97^{+0.51}_{-0.08}$	$5.07^{+0.34}_{-0.79}$	$-2.74^{+0.34}_{-0.07}$	$-1.84^{+1.47}_{-0.75}$	$8.28^{+0.07}_{-1.90}$	–	–	–
	MDD	0.3973	$-2.24^{+0.04}_{-0.04}$	$7.24^{+0.51}_{-0.57}$	–	–	–	$7.61^{+0.19}_{-0.22}$	$0.40^{+0.10}_{-0.07}$	218^{+47}_{-38}
$\{r > R_{50}\}$	MID	0.2390	$-1.97^{+0.51}_{-0.08}$	$5.07^{+0.34}_{-0.79}$	$-2.74^{+0.34}_{-0.07}$	$-1.84^{+1.47}_{-0.75}$	$8.28^{+0.07}_{-1.90}$	–	–	–
	MDD	0.7610	$-1.79^{+0.10}_{-0.10}$	$4.47^{+0.25}_{-0.12}$	–	–	–	$6.83^{+0.44}_{-0.23}$	$0.85^{+0.11}_{-0.23}$	238^{+62}_{-47}
$\{r \leq R_c\}$	MID	0.9175	$-2.18^{+0.04}_{-0.03}$	$7.20^{+0.56}_{-0.71}$	$-2.73^{+0.04}_{-0.05}$	$-1.13^{+0.20}_{-0.22}$	$8.07^{+0.10}_{-0.16}$	–	–	–
	MDD	0.0825	$-2.17^{+0.03}_{-0.03}$	$7.40^{+0.41}_{-0.54}$	–	–	–	$8.06^{+0.13}_{-0.16}$	$0.22^{+0.07}_{-0.05}$	288^{+50}_{-40}
$\{r > R_c\}$	MID	0.8131	$-2.04^{+0.63}_{-0.38}$	$5.00^{+2.28}_{-0.81}$	$-2.78^{+0.53}_{-0.35}$	$-0.75^{+0.47}_{-1.47}$	$7.74^{+1.66}_{-1.68}$	–	–	–
	MDD	0.1869	$-2.23^{+0.38}_{-0.27}$	$5.47^{+1.68}_{-1.08}$	–	–	–	$7.96^{+1.47}_{-1.19}$	$0.61^{+0.29}_{-0.42}$	685^{+15845}_{-582}
<i>Subpopulations</i>										
$\{p_{\text{obs}} \geq 90 \text{ per cent}\}$	MID	1.00	$-2.16^{+0.17}_{-0.04}$	$5.32^{+0.35}_{-0.57}$	$-2.90^{+0.06}_{-0.04}$	$-1.34^{+1.17}_{-0.65}$	$8.22^{+0.08}_{-1.85}$	–	–	–
	MDD	$<10^{-4}$	$-2.14^{+0.08}_{-0.05}$	$4.91^{+0.15}_{-0.21}$	–	–	–	$7.11^{+0.56}_{-0.20}$	$0.85^{+0.11}_{-0.29}$	450^{+141}_{-96}
Class 3 excluded	MID	1.00	$-1.99^{+0.16}_{-0.08}$	$5.16^{+0.34}_{-0.76}$	$-2.76^{+0.08}_{-0.06}$	$-1.05^{+0.98}_{-0.48}$	$8.21^{+0.11}_{-2.37}$	–	–	–
	MDD	$<10^{-4}$	$-1.74^{+0.06}_{-0.06}$	$4.90^{+0.08}_{-0.09}$	–	–	–	$6.19^{+0.06}_{-0.05}$	$0.98^{+0.02}_{-0.03}$	91^{+10}_{-9}
OGCs excluded	MID	0.0364	$-1.50^{+0.07}_{-0.34}$	$4.26^{+0.19}_{-0.34}$	$-2.44^{+0.07}_{-0.30}$	$-0.40^{+0.04}_{-1.01}$	$6.39^{+1.76}_{-0.05}$	–	–	–
	MDD	0.9636	$-1.79^{+0.06}_{-0.05}$	$4.33^{+0.07}_{-0.06}$	–	–	–	$6.64^{+0.19}_{-0.11}$	$0.93^{+0.05}_{-0.10}$	217^{+32}_{-26}
$\{D_{\text{norm}}^5 > 2\}$	MID	1.00	$-2.17^{+0.33}_{-0.04}$	$5.58^{+0.39}_{-1.21}$	$-2.86^{+0.11}_{-0.04}$	$-1.52^{+1.31}_{-0.70}$	$8.24^{+0.08}_{-1.91}$	–	–	–
	MDD	$<10^{-11}$	$-2.16^{+0.05}_{-0.04}$	$5.42^{+0.39}_{-0.24}$	–	–	–	$8.11^{+0.43}_{-0.46}$	$0.27^{+0.25}_{-0.23}$	400^{+78}_{-58}

Notes. The first column indicates the subset of the LEGUS catalogue used in the fit, the second indicates the type of model fit (MID or MDD), the third gives the Akaike weight for that model, and the remaining columns give marginal posterior PDFs derived for each model parameter, reported as $q_{50}^{+(q_{84}-q_{50})} / (q_{50}-q_{16})$, where q_N denotes the estimated N^{th} percentile. Thus, our central values are the median of the PDF, and the ranges shown correspond to the 68 per cent confidence interval. The co-rotation radius is represented by R_c , while the median galactocentric radius is denoted as R_{50} . In the subpopulation section, we present trimmed catalogue fits based on four different trimming criteria. First, we consider the cluster population denoted by $p_{\text{obs}} \geq 90$ per cent. Additionally, we provide three additional trimmed catalogue fits where we exclude class 3 objects, old globular clusters, and clusters with a normalized distance to the nearest 5th neighbour in the library greater than 2σ .

We compare the observed and model-predicted luminosity functions in Fig. 3. As the figure shows, our best-fitting model does an excellent job of reproducing the full luminosity function in the bluer bands, and reproduces the shape of the peak of the luminosity function very well in every band. The only place where our model lies significantly outside the observed confidence interval is in the I band, where the observed luminosity function drops off somewhat more steeply than our model fit at both the bright and dim ends. At the dim end, this is likely due to our completeness model not being perfectly accurate – one can see hints of this in the V band as well, where our model cuts off sharply at -6 mag, the nominal catalogue

than $V = -6$ mag this would be mostly compensated by clusters just below the magnitude limit of the original catalogue scattering upward in some realizations – that is, for every cluster with $V = -6.001$ mag in the original LEGUS catalogue that scatters below the magnitude cut in almost half the realizations, there should be another cluster with $V = -5.999$ mag that is excluded from the original catalogue but scatters above the magnitude cut in almost half the realizations. However, because clusters for which the central estimate for $V = -5.999$ mag do not appear in the LEGUS catalogue at all, this compensation does not happen, and as a result we end up undercounting clusters near the magnitude limit by as much as a factor of ≈ 2 . This effect is significant for magnitudes within a few σ of the catalogue cutoff, and since the typical photometric error in LEGUS is $\sigma \approx 0.1$ mag, we expect significant bias for $V \gtrsim -6.3$ mag.

limit, but the data fall off slightly more smoothly. The Malmquist bias discussed above may contribute slightly to this discrepancy, but given that the observed luminosity function in V band turns over near $V = -6.5$ mag, roughly 5σ above the catalogue limit for the typical $\sigma \approx 0.1$ mag uncertainty in LEGUS, it is unlikely to be the dominant effect. At the bright end, the discrepancy may be a result of our simple Schechter functional form being too simple to fully capture the full shape of the mass distribution. A second possibility is that this mismatch is a result of errors in our assumed values of metallicity or nebular covering fraction, a topic we address further momentarily.

It is also interesting to compare the fidelity with which we are able to reproduce the luminosity function with the results of prior studies that derived cluster mass and age distributions from a more traditional backward-modelling approach rather than our Bayesian forward-modelling method. Adamo et al. (2017) used the cluster fitting code YGGDRASIL to assign masses and ages to each cluster in NGC 628, and find that the population is best fit by an MDD model with $t_4 = 190$ Myr, $\gamma_{\text{MDD}} = 0.65$, $\alpha_M = -2.03$, and $M_{\text{break}} = 2.03 \times 10^5 M_\odot$. These results are derived using the same stellar tracks as we use in our library, but come from a fit to the ‘exclusive’ catalogue that excludes morphologically complex class 3 clusters. To compare to these results, we therefore fit the LEGUS catalogue excluding class 3 sources using our method (see Section 5.1 for details), and then generate predicted luminosity functions for both

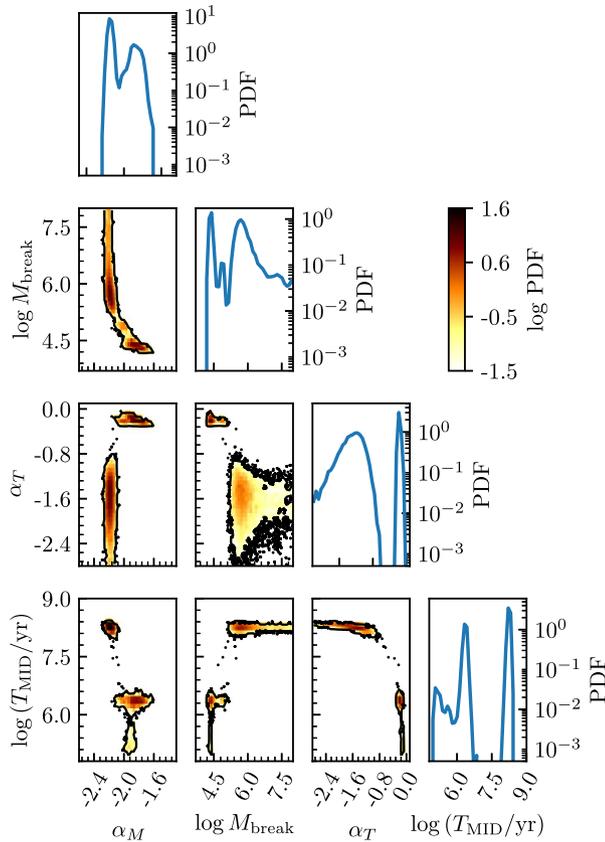


Figure 2. Corner plot showing the one-dimensional and two-dimensional histograms of the posterior PDFs of the parameters α_M , $\log M_{\text{break}}$, α_T , and $\log T_{\text{MID}}$ for the full LEGUS NGC 628 catalogue. We omit the nuisance parameters describing the distribution of extinctions A_V . The 1D blue histograms show one-dimensional marginal PDFs for each parameter, while the colour maps show log probability densities in various 2D cuts through parameter space. The outermost contour level in the 2D plots is set to enclose 99 per cent of the samples, and black dots show individual MCMC samples outside this threshold. The extent of each axis is rounded according to the prior range to achieve optimal visualization.

Adamo et al. (2017)’s best-fitting model and our best fit, using the same method as in Fig. 3, and compare to the observed LEGUS catalogue excluding class 3 clusters. We show the results in Fig. 4. It is clear that our method performs substantially better at reproducing the observed luminosity function, particularly in the redder bands.

4.2.2 Colour distributions

Since our method attempts to fit the full five-dimensional photometric distribution at once, we can also check for agreement between model and observations in multiple dimensions. To this end, we present as examples the $UV - U$ versus U colour–magnitude diagram in Fig. 5, and the $UV - U$ versus $U - B$ colour–colour diagram in Fig. 6; we show the colour–magnitude and colour–colour diagrams for other bands in Appendix D. As with the luminosity distribution, we compute the colour distribution predicted by our fit directly from the Gaussian mixture model, using an expression analogous to equation 15, and we compute the observed distribution via a Monte Carlo resampling of the LEGUS measurements including their uncertainties. Specifically, to generate the distribution of the observations in colour–magnitude or colour–colour space, we draw

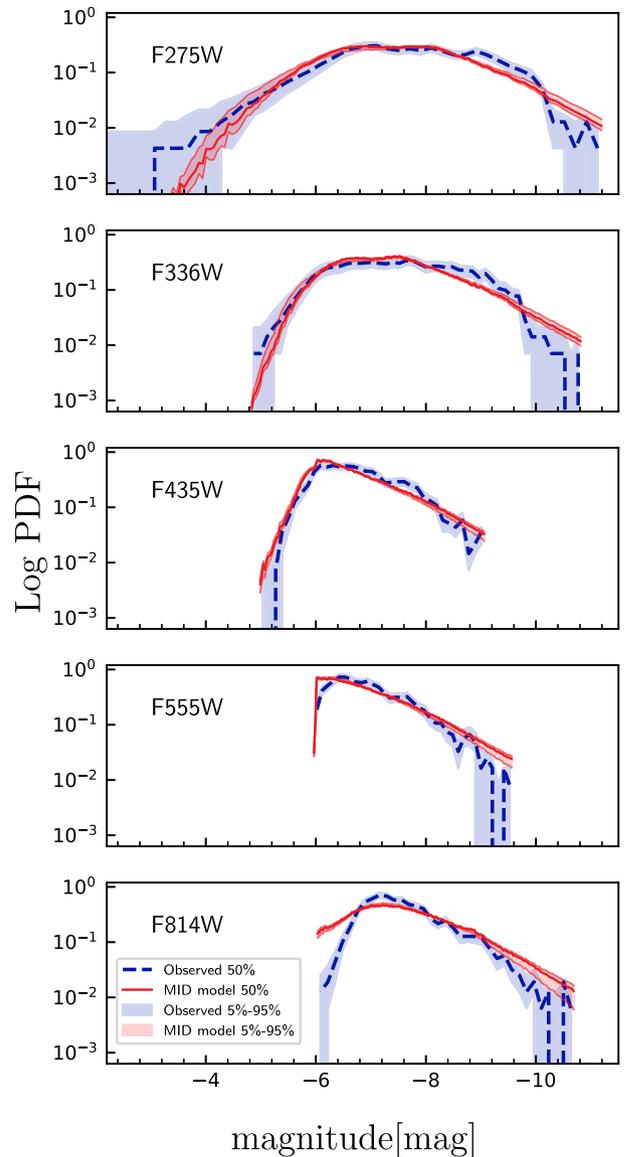


Figure 3. Comparison between observed and model-predicted luminosity functions in UV , U , B , V , and I bands. Blue-dashed lines show the luminosity functions of the LEGUS catalogue, divided into uniformly spaced bins; distributions are normalized to have unit integral, and blue-shaded regions show the 5th–95th percentiles in normalized bin counts derived from 20 000 Monte Carlo re-samplings of the observations – see the main text for details of the procedure. The red solid lines illustrate the predicted photometric distributions of the best MID models, computed using the 50th percentile values post burn-in. The shaded bands surrounding these lines indicate the 5th–95th percentiles for each bin. The blue-dashed lines show the 50th percentiles of the photometric distributions in each individual bin, using bootstrapped samples of the observed luminosities. The blue-shaded bands surrounding these dashed lines indicate the 5th–95th percentiles for each bin.

20 000 samples from the Gaussian error distribution for each cluster and band, thereby generating a sample of $\sim 10^7$ points in each possible colour–magnitude and colour–colour diagram. We then plot contour lines that enclose 50, 84, and 95 per cent of these samples to indicate the range covered by the observations. These figures also show reasonably good agreement between the observed and model-predicted colour–magnitude and colour–colour distributions, and

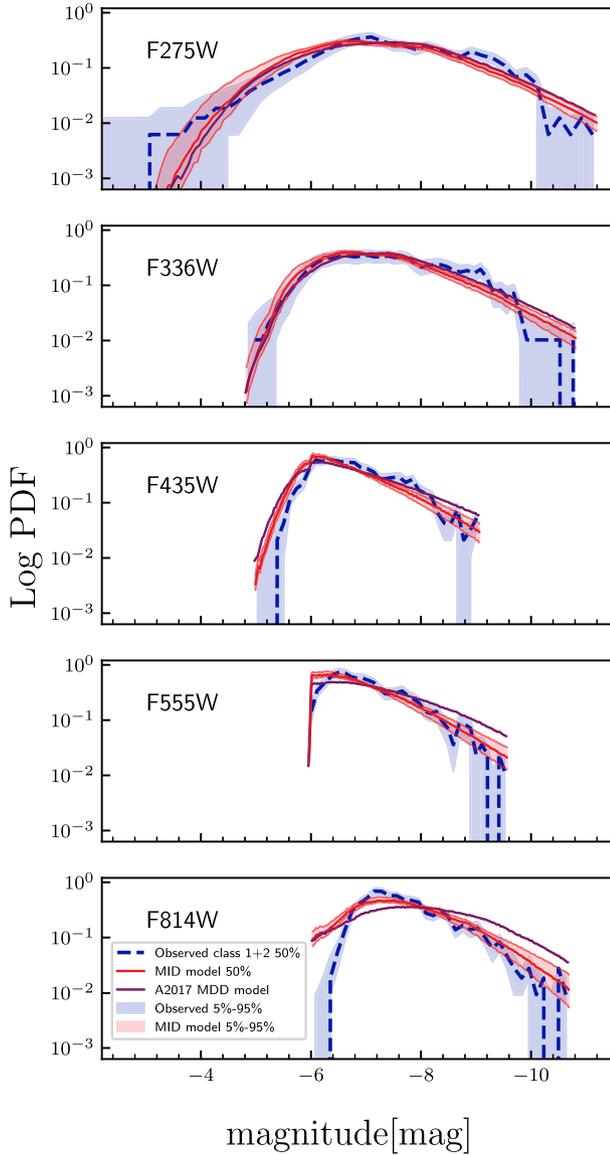


Figure 4. Same as Fig. 3, but now showing the luminosity function only for class 1 and 2 clusters, and including a comparison to the luminosity distribution predicted using the best-fitting parameters obtained by Adamo et al. (2017) (black line).

suggest that our best-fitting model provides a reasonable representation of the distribution of observations in 5D photometric space.

The largest deviations between our model and the observations occur for clusters with the bluest colours. Fig. 5 illustrates that the observations include a tail at a UV magnitude of approximately -6 , featuring the most negative $UV - U$ colour index (bluest colour), which is not well reproduced by the models. This feature is also present in the $U - B$ colour distributions shown in Fig. 6, where our model exhibits a cutoff in $U - B$ colour space at the bluest end, indicating limitations in reproducing the bluest clusters in the observed samples.

One possible explanation for this discrepancy is that our model library assumes Solar metallicity, while NGC 628 is slightly sub-Solar (e.g. Moustakas et al. 2010). Consistent with this hypothesis, Deger et al. (2022) compare PHANGS-*HST* clusters to tracks through

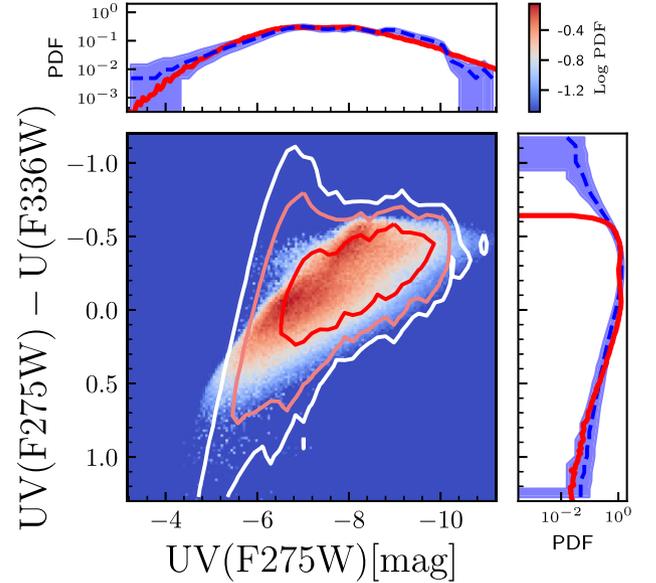


Figure 5. Comparison between observed and model-predicted cluster distributions in UV magnitude versus $UV - U$ colour. The heatmap in the central panel shows the predicted colour-magnitude distribution for our model, and is computed using the 50th percentile values of all parameters. The red, coral, and white contour lines enclose the 50, 84, and 95 percent of a bootstrap resampling of the observed distribution in the LEGUS catalogue – see the main text for details of the bootstrapping procedure. The flanking panels compare 1D PDFs of the observed (blue-dashed lines) and predicted (red solid lines) colour and magnitude. To visualize the uncertainties around photometric measures, we represent the 5th–95th percentiles on the observational distribution as blue-shaded regions.

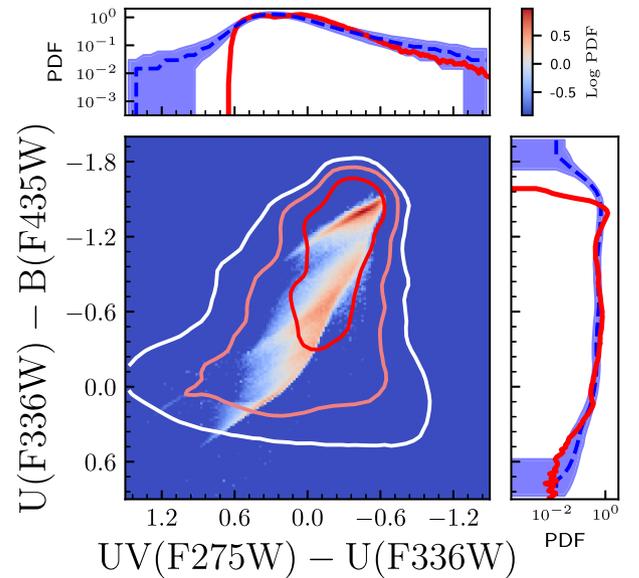


Figure 6. Same as Fig. 5, but showing the $U - B$ versus $UV - U$ colour-colour distributions of the observations and the best-fitting model.

Table 2. Goodness-of-fit log p -values for one-sample KS test comparisons of between the observed and best-fitting model luminosity functions (cf. Fig. 3) in the UV , U , B , V , and I bands, for both the best-fitting MID and MDD models.

	UV	U	B	V	I
$\ln p_{\text{MID}}$	−1.80	−4.03	−10.04	−3.62	−8.43
$\ln p_{\text{MDD}}$	−3.09	−4.50	−8.38	−5.13	−11.92

colour-space predicted by non-stochastic stellar population synthesis models using a range of metallicity and nebular covering fraction. They find the bluest colours achievable by the models are sensitive to these two parameters, with plausible variations in them having relatively little effect over most of the track, but yielding variations at the extreme blue end by ~ 0.3 mag, comparable to the offset we find between our bluest observed clusters and our bluest SLUG models. However, Krumholz et al. (2015b) show that metallicity variations at this level have relatively little effect when inferring cluster masses or ages. In any event, the effects of this discrepancy on population-level statistics should be small, since they involve only a handful of clusters out of our full catalogue of > 1000 .

While Figs 3, 5, and 6 allow qualitative comparison between the best-fitting model and the observed photometric distributions, it is also possible to make a quantitative comparison. The conventional approach for assessing whether a model is a good fit typically involves measures such as p -values derived from χ^2 or Kolmogorov–Smirnov (KS; Smirnov 1939) tests. However, it is important to keep in mind what these tests assess: the p -value indicates the confidence with which we can rule out the null hypothesis that measurement error (for a χ^2 test) or finite sample size (for a KS test) are the *only* sources of disagreement between model and data, i.e. the null hypothesis for a χ^2 or KS test is that the model is perfect and that discrepancies between it and the data arise only because of limitations in the measurement. In our study, we have identified distinct features in the data that are not replicated by any of the models, as demonstrated by the comparisons of colour–colour and colour–magnitude between the models, so we do not expect even our best model to achieve a ‘good fit’ in the sense of a high p -value, and thus most of the value in the test is in comparing relative p -values. With this limitation understood, we proceed to carry out KS test comparisons between the observed and model-predicted 1D luminosity distributions shown in Fig. 3, and the analogous distributions for the best-fitting MDD model. We present the resulting p -values in Table 2.

As expected, the MID model in the UV band has a p -value high enough to be consistent with the hypothesis that the model is perfect at the few σ level. Among the five available bands, the MID model is preferred in all bands except B based on higher p -values, consistent with the finding of CLUSTER_SLUG favouring the MID model. However, we caution that one should *not* perform model comparisons just by examining relative p -values, since, among other defects, this approach ignores the higher dimensional correlations present in the real data, which our comparison based on Akaike weights properly captures. In general, we also see that model–observation agreement is better in the bluer bands and worse in the redder ones. None the less, the comparison reinforces the point that, while our models capture the major qualitative features of the observed photometric distributions well, they are clearly not perfect. Given the much larger sample of clusters are we can use compared to previous work, we have the statistical power to detect discrepancies that would have been invisible to earlier analyses on much smaller data sets.

4.3 Inner versus outer galaxy clusters

Having presented the full catalogue fits, we now search for radial variations in cluster demographics. For this purpose, we assign every cluster in the LEGUS catalogue a galactocentric radius, taking the centre of NGC 628 to be at RA = 24.17°, Dec. = 15.78°, and adopting an inclination $i = 25.2^\circ$ and position angle PA = 25° from North (Grasha et al. 2015, 2017). We use the centre location, inclination, and position angle to deproject the cluster coordinates from RA/Dec. to a galactocentric coordinate system. Once clusters have been assigned radii, we divide the catalogue in two ways: first simply by making two sub-catalogues of equal size containing the inner and outer halves of the sample, and second by making two sub-catalogues containing clusters inside and outside the galactic co-rotation radius. The advantage of the former approach is that it ensures that we have equal statistical power in both regions, and thus do not miss radial variations because our inner or outer galaxy sample is too small to see them. The advantage of the latter approach is that co-rotation marks a physically motivated radius where we might expect to see a change in cluster behaviour, as opposed to the galactocentric radius that marks an equal number division, which is solely a function of the size and location of the LEGUS pointings.

4.3.1 Equal number division

For our first analysis, we divide the LEGUS cluster sample into two radial bins each containing an equal number of clusters. The median radius that makes this even division is $R_{50} = 3.53$ kpc. We fit MID and MDD models to both the subset of clusters with $r \leq R_{50}$ and those with $r > R_{50}$ and report the results in the second and third blocks of rows in Table 1. Our model selection statistics strongly favours the MID model in the inner region, whereas the outer region prefers the MDD model. Comparing these two sub-divided regions, we present the fitted parameters in Table 1 and plotted the posteriors in 2D parameter spaces of the CMF and CAF in Figs 7 and 8, respectively. For the outer galaxy age distribution, even though the MDD model is (slightly) preferred, we show results for the slightly less favoured MID model so that we can make a direct comparison.

Examining Fig. 7, its upper panel shows the marginal posterior PDFs of α_M and M_{break} of inner and outer regions divided by R_{50} . We notice that the posterior mass distributions of the inner bin (indicated by the red contours) are unimodal, corresponding to a single probability peak in the parameter space centred at $\alpha_M \approx -2.2$ and $\log(M_{\text{break}}/M_\odot) \gtrsim 6$, while the posterior for the outer region (indicated by blue contours) is bimodally distributed with two distinct probability peaks, one at $\alpha_M \approx -2$ and $\log(M_{\text{break}}/M_\odot) \approx 5$, the second at $\alpha_M \approx -1.5$ and $\log(M_{\text{break}}/M_\odot) \sim 4.3$. Despite this difference, however, the 2σ confidence contours for the two regions slightly overlap, and the 95 per cent confidence contours clearly lie along a continuous ridge of probability running from low α_M and high M_{break} to higher α_M and lower M_{break} . Thus, we formally detect a difference between the inner and outer regions, but at low confidence. To the extent that there is a real difference, it is that there is stronger evidence for a truncation in the mass function in the outer galaxy, an effect consistent with the correlation between truncation mass and star formation rate surface density identified by Wainer et al. (2022).

The upper panel of Fig. 8 provides an analogous plot for the two parameters $-\alpha_T$ and T_{MID} – that describe the CAF. The qualitative conclusion to be drawn from Fig. 8 is similar to that for the mass PDFs as shown in the upper panel of Fig. 7, i.e. the posterior PDFs of the age parameters for the inner and outer radial bins are mainly separated at 2σ levels, but are clearly tracing out the same underlying ridge of

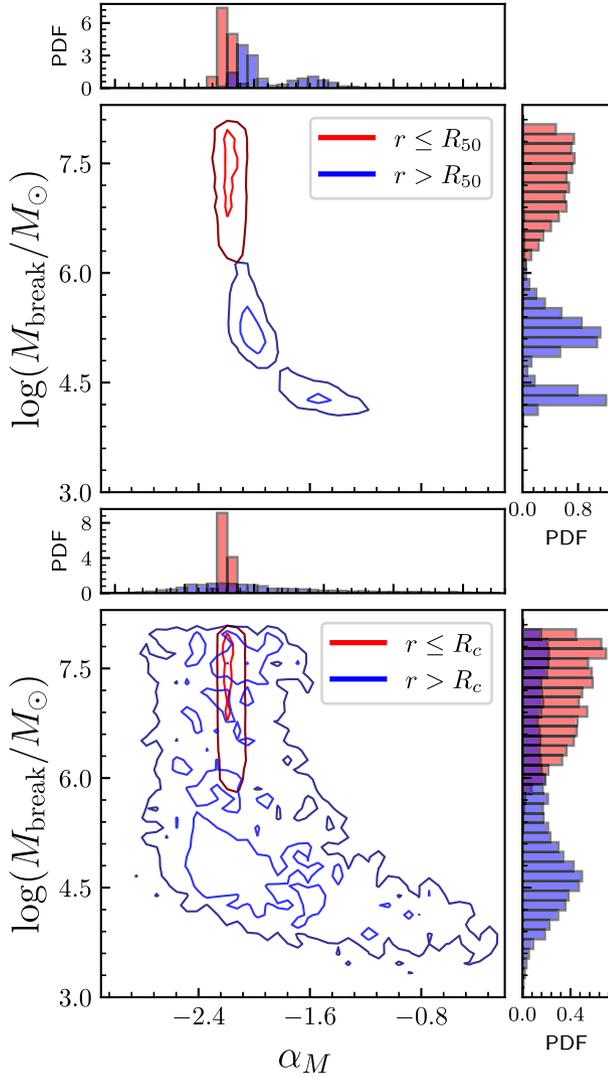


Figure 7. 2D marginal posterior joint PDF of the parameters α_M and $\log M_{\text{break}}$ that describe the CMF. Upper panel: PDFs for clusters inside (red) and outside (blue) the co-rotation radius R_c . Blue and red contour lines in the central panel correspond to loci that enclose 50 and 95 per cent of the MCMC samples, and thus correspond approximately to the median and 2σ significance confidence level. The histograms flanking the central panel show the corresponding 1D marginal PDFs of α_M and $\log M_{\text{break}}$ for the inner (red) and outer (blue) galaxy samples. Lower panel: same as upper panel, except we show the parameters α_M and $\log M_{\text{break}}$ for clusters inside (red) and outside (blue) the median cluster galactocentric radius R_{50} .

probability. The contours would overlap if we extended our plots to higher than 95per cent confidence; we again therefore consider this a low-confidence detection of a difference. The clusters in the inner region of the galaxy (represented by red contours) begin to disrupt around $T_{\text{MID}} \approx 30$ Myr, with a relatively rapid disruption rate of $\alpha_T \approx -0.8$ thereafter. Conversely, the PDFs of the age parameters in the outer region (represented by blue contours) appear to be more spread out. The PDF peak for the clusters in the outer region suggests either moderate disruption starting at 1–5 Myr ($\alpha_T \approx -0.4$), or no disruption until ~ 200 Myr. In either case, to the extent that our detection is real, it suggests that disruption always faster in the inner galaxy at ages $\gtrsim 100$ Myr, but that whether disruption in the inner galaxy is faster or slower at younger ages depends on which of the

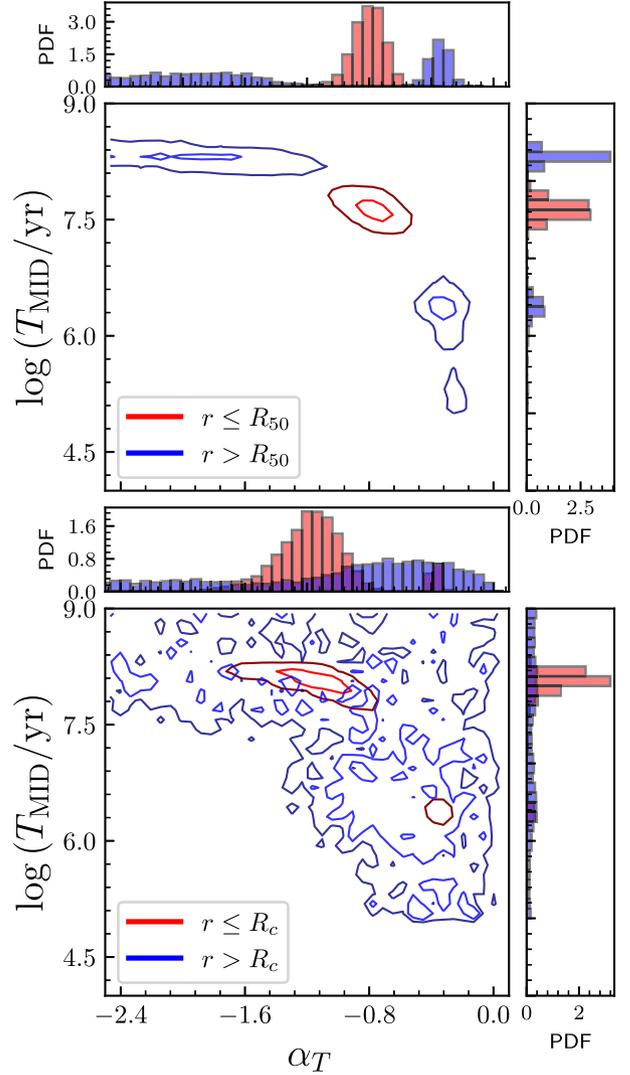


Figure 8. Same as Fig. 7, except we show the marginal PDFs of the parameters describing the CAF, α_T and $\log T_{\text{MID}}$, rather than the parameters describing the CMF.

two possibilities identified by the MCMC we favour for the outer galaxy – for the case of late disruption, $T_{\text{mid}} \sim 200$ Myr, the inner galaxy has faster disruption at all cluster ages, while for the early but moderate disruption scenario in the outer galaxy, $T_{\text{mid}} \sim 3$ Myr and $\alpha_T \sim -0.4$, disruption is initially faster in the outer galaxy, but the inner galaxy catches up and overtakes after ~ 100 Myr, when ~ 30 – 40 per cent of clusters still remain intact.

4.3.2 Co-rotation radius division

An alternative way of separating the cluster population is into those inside and outside the galactic co-rotation radius at $R_c = 6.3$ kpc where the spiral pattern and orbiting clusters move at equal speed (Sakhibov, Gusev & Hemmerich 2021). To the extent that encounters with molecular clouds or other structures associated with the spiral arms influence cluster formation or disruption, we might expect to see changes in cluster population demographics across co-rotation. Making this division, we have 1159 clusters inside R_c and 117 clusters outside R_c .

The bottom panels of Figs 7 and 8 show the inferred posteriors for the clusters located inside and outside the co-rotation radius; we again concentrate on the MID model, because model comparison prefers it to MDD. Here, we see differences between the inner and outer galaxy posterior PDFs that are qualitatively consistent with those we observed for a division of the sample into two equal parts. However, the small number of clusters in the outer galaxy sample when we divide at co-rotation ensures that the posterior PDFs are very broad, and, particularly for the age distribution, largely just reflect our priors. Given the very broad outer galaxy posteriors, we cannot rule out the possibility that the cluster demographics inside and outside co-rotation are the same. Doing so would likely require significantly more outer-galaxy clusters than the LEGUS catalogue provides.

4.4 Individual cluster ages and masses

In addition to the population-level inference of the marginal posteriors, we can also compute individual cluster-by-cluster fits for cluster masses and ages in NGC 628 using the method described in Krumholz et al. (2015a, b), with the important difference that since we have now fit for the mass and age distribution of NGC 628 we can use this result as an informative prior. We briefly summarize the method here and refer readers to Krumholz et al. (2015b) for a full description. The method is analogous to that described in Section 3: we consider a single cluster with a vector of magnitudes \mathbf{m} and corresponding uncertainties σ , and invoke Bayes' theorem to write the posterior probability distribution for the cluster mass M , age T , and extinction A_V as

$$p_{\text{post}}(\boldsymbol{\psi} \mid \mathbf{m}, \sigma) \propto \sum_{j=1}^{N_{\text{lib}}} w_j(\boldsymbol{\psi}_j) \mathcal{N}(\mathbf{m} \mid \mathbf{m}_j, \mathbf{h}'_j) \mathcal{N}(\boldsymbol{\psi} \mid \boldsymbol{\psi}_j, \mathbf{h}_\boldsymbol{\psi}), \quad (16)$$

where $\boldsymbol{\psi} = (\log M, \log T, A_V)$ is the vector of parameters we wish to infer, and quantities subscripted by j are properties of the j th library cluster: \mathbf{m}_j and \mathbf{h}'_j are the vectors of photometric magnitudes and bandwidths (identical to those used in equation 2), $\boldsymbol{\psi}_j$ is the (log mass, log age, extinction) vector, $\mathbf{h}_\boldsymbol{\psi}$ is the bandwidth in (log mass, log age, and extinction), and $w_j(\boldsymbol{\psi}_j)$ is the weight. The prior enters in the calculation of w_j : Library clusters whose properties are a priori more likely are weighted more highly (e.g. lower mass clusters will be assigned more weight than higher mass ones, since most clusters are low mass; cf. equation 18 of Krumholz et al. 2015a). The difference between our calculation here and that presented in Krumholz et al. (2015b) is that, whereas Krumholz et al. adopted a range of candidate uninformative priors, here we use the best-fitting model from our population-level analysis as our prior.

In Figs 9 and 10, we show how the marginal posterior mass and age distributions we derive by this procedure compare to those derived using the deterministic stellar population fitting code YGGDRASIL (Zackrisson et al. 2011) and reported for the clusters in NGC 628 (Adamo et al. 2017) (upper panels), and to those derived using the uninformative priors from Krumholz et al. (2015b) (lower panels); for the latter, we use their $\alpha_M = -1.5$, $\alpha_T = -0.5$ case. We construct these figures using bootstrap resampling: for each cluster in the NGC 628 catalogue, we randomly draw a CLUSTER_SLUG mass from the marginal posterior PDF of mass derived from equation (16) using informative priors; to construct the top panel we then draw corresponding masses from the central value and uncertainty (assumed to be Gaussian in $\log M$) for the clusters reported by Adamo et al. (2017), while for the bottom panel we also draw the corresponding mass from equation (16), but with $w(\boldsymbol{\psi}_j)$

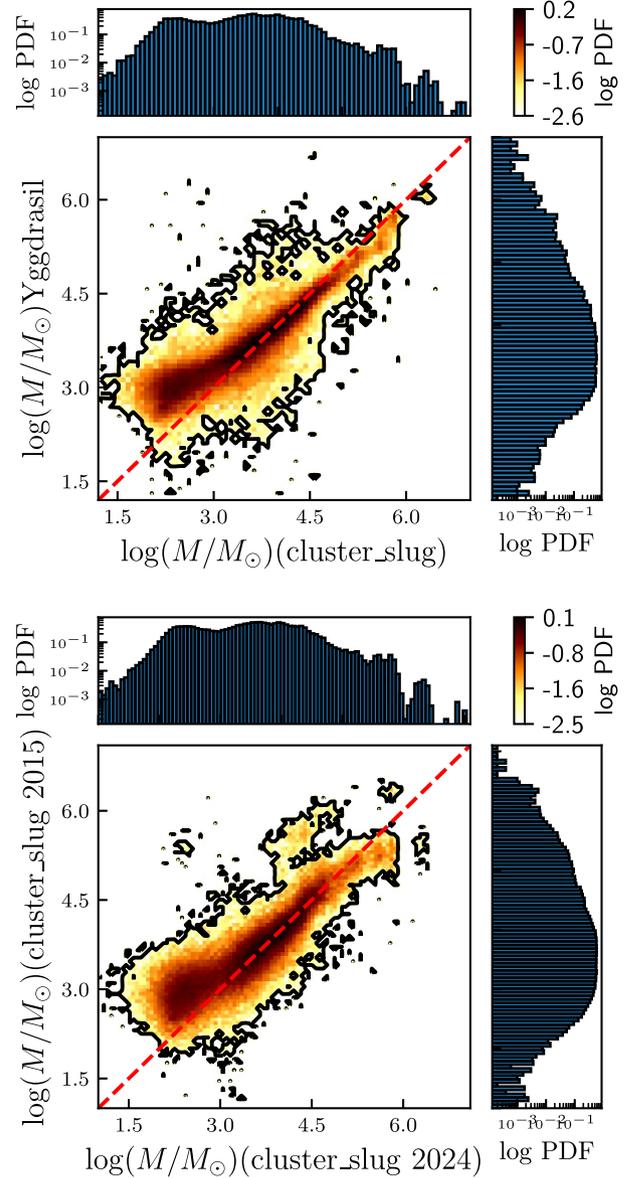


Figure 9. Comparison of NGC 628 cluster mass estimates computed in this paper (equation 16; horizontal axis) to those computed in alternative ways (vertical axes); the top panel compares to deterministic YGGDRASIL models (as reported in Adamo et al. 2017), while the bottom panel compares to results derived using the informative priors from Krumholz et al. (2015b); see the main text for details. The 2D density maps display the joint PDF of masses derived by the two methods, estimated using the bootstrap resampling procedure described in the main text (Section 4.4). The top and right flanking panels present the 1D distribution of masses computed by each model.

evaluated using the uninformative priors. We repeat this procedure until we have a sample of 35 000 pairs for each panel, and we plot the distributions of these points in Fig. 9; the procedure is identical for ages in Fig. 10.

Examining these figures, we see that overall YGGDRASIL tends to produce somewhat higher cluster masses (i.e. the CLUSTER_SLUG predicted masses $< 1000 M_\odot$ are biased above the one-to-one lines) and younger ages compared to CLUSTER_SLUG (i.e. the points are biased slightly below the one-to-one lines), particularly at the lowest masses and youngest ages; agreement is closer at higher

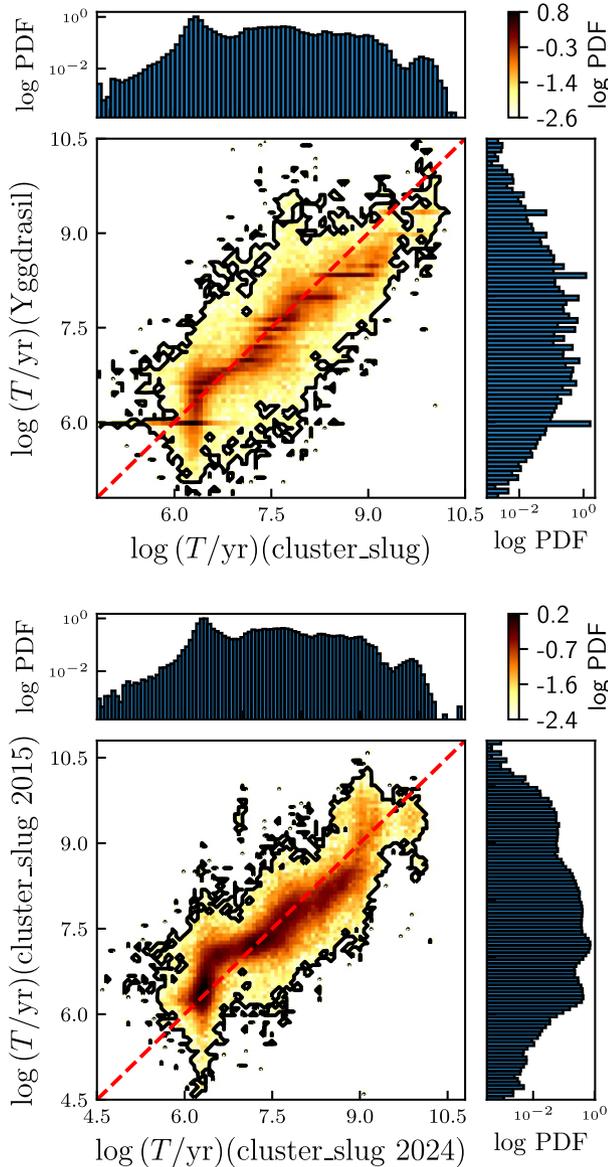


Figure 10. Same as Fig. 9, except we show age estimates instead of mass estimates. Note that the horizontal streaks in the top panel are an artefact of the YGGDRASIL model, which uses a relatively coarse grid of ages.

mass and age (i.e. the points cluster more closely around the one-to-one lines).² One explanation for this asymmetry at $M \sim 10^3 M_\odot$ is, for a given brightness of a cluster, YGGDRASIL assigns this cluster a unique mass based on a fixed mass-to-light ratio. However, CLUSTER_SLUG consider another possibility that it can be a low-mass star cluster that happens to contain a massive star that dominates the light profile due to the impartial sampling of the IMF. Because we use an informative prior on mass with a slope ~ -2 , which assigns a high probability to less massive clusters, the scenario mentioned above becomes prominent. Consequently, we expect CLUSTER_SLUG to assign systematically lower masses to dim

²The horizontal streaks visible in Fig. 10 are not physically meaningful; they are artefact of YGGDRASIL using a discrete, coarsely sampled grid of ages for its fits.

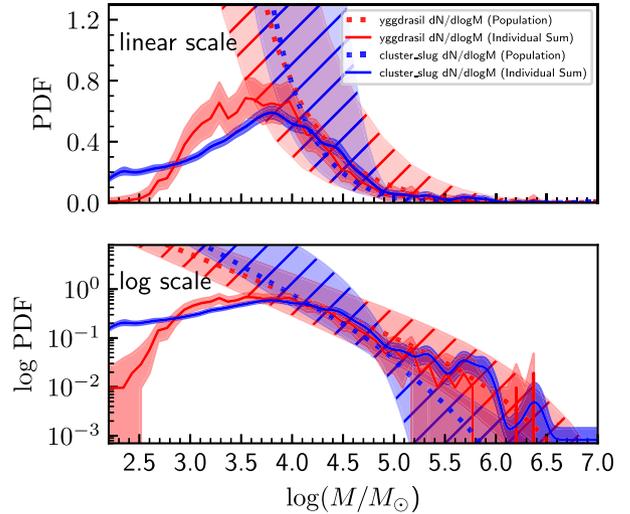


Figure 11. Comparison of the inferred posterior PDFs in linear (upper panel) and log scales (lower panel) for the mass distribution of the entire population of star clusters using the CLUSTER_SLUG and YGGDRASIL SSP models. The dotted blue line is the mass PDF obtained by summing the marginal mass PDFs derived with CLUSTER_SLUG over all clusters; the shaded region around this line is the 5–95 per cent confidence interval derived from bootstrap resampling via the procedure described in the main text. The red solid line and shaded region show the same quantities derived from YGGDRASIL deterministic models following Adamo et al. (2017). The dashed lines and hatched regions show the corresponding fits to the population mass distribution, computed as described in the main text.

clusters than YGGDRASIL, as shown in Fig. 9. Similarly, deterministic models interpret a blue colour as requiring a young age, while a stochastic model allows for the possibility that it might also be the result of an older cluster containing a single bright star undergoing a blue loop. Consequently, CLUSTER_SLUG’s uncertainty intervals tend to extend to larger masses and ages than YGGDRASIL’s, and this effect is largest in the least massive and youngest clusters, which are subject to the largest stochastic fluctuations. As noted above, this phenomenon was previously reported by Krumholz et al. (2015b), so the main new thing to take from our analysis here is that switch from uninformative to informative priors does not appear to alter it significantly.

Consistent with this finding, we see that there are somewhat smaller differences between the CLUSTER_SLUG models with the two different priors. Compared to the uninformative prior adopted in Krumholz et al. (2015b), our population fit has a steeper slope at low mass, since here we have $\alpha_M \approx -2$ as compared to $\alpha_M = -1.5$, and so we again find that the current fit favours somewhat lower masses at the low-mass end. Conversely, there are clearly some clusters for which the uninformative prior results in a mass $\gtrsim 10^5 M_\odot$, while the informative one reduces the estimated mass to $\approx 10^{4.5} M_\odot$. This is clearly a result of our more sharply truncated mass function.

To explore this further, in Fig. 11, we plot the inferred posterior distribution of cluster masses, marginalized over extinction and age, for all the clusters in NGC 628. The blue solid line in this figure is computed by summing the individual cluster posterior PDFs of mass, and the shaded blue band around it is the 5–95 per cent confidence interval, which we derive from our bootstrap samples; the red solid line and shaded band are the result of applying the same procedure to the YGGDRASIL results. (We omit the corresponding line for uninformative priors to avoid clutter, but it would show the same

qualitative effects.) Comparing these lines, two models agree very well from 10^4 – $10^7 M_\odot$. At lower masses, the CLUSTER_SLUG mass distribution is flatter than the YGGDRASIL one due to the individual mass PDFs having large dispersions due to CLUSTER_SLUG’s inclusion of uncertainties from stochasticity, as described above.

To investigate the consistency between the PDF derived from individual cluster masses and the best-fitting Schechter mass functions we derive, in Fig. 11 we have also overplotted dashed lines representing those best-fitting models; we do so both for the Schechter function derived in this paper, and using the YGGDRASIL-based provided in Adamo et al. (2017); the dotted line is the most likely model, and the hatched region around it is the 5–95 per cent confidence interval derived from bootstrap resampling of the posterior PDFs. Interestingly, we see that while our best-fitting CLUSTER_SLUG model matches the distribution relatively well near the peak of the PDF, it is more sharply truncated at $10^5 M_\odot$ than the mass function we get by adding up the individual cluster posterior PDFs. The individual cluster PDFs have a small population of clusters above $10^6 M_\odot$, which is absent in our Schechter function fit to the whole population.

This comparison highlights an important feature of our method, which is that it is most sensitive to the peak of the distribution, in contrast to methods based on fitting binned data, which assign equal weight (except for Poisson errors) to nearly unpopulated mass bins as to highly populated ones. Indeed, there is no a priori reason why we should expect to obtain the same result from a method that weights all *clusters* equally versus one that weights all *mass bins* equally, nor why we should expect a forward-modelling method that emphasizes reproducing the observed luminosity distribution to match a backward modelling method that has no such constraint. In an ideal world, where our stellar population synthesis models and the functional forms we use to parameterize the cluster population were both accurate, these approaches would agree, but clearly that is not the case given our current models. None the less, this comparison does suggest that in future work, our method might be improved by adopting a more general functional form for the mass distribution that allows independent adjustment of the shape at lower and higher masses, rather than relying on the more prescriptive functional forms (such as the Schechter function) that have traditionally been used. It also means that we should apply the informative priors with caution for very luminous clusters, since our priors may be unfairly discounting the possibility of these being very massive.

5 DISCUSSION

Here, we discuss some implications of our findings. We begin with a series of additional verification checks on our method (Section 5.1), and then discuss our findings regarding truncation of the cluster mass distribution (Section 5.2) and the cluster age distribution (Section 5.3).

5.1 Method robustness against outliers

All our results in Section 4 are derived from fitting the full LEGUS cluster catalogue. However, there are plausible reasons why one might consider excluding various parts of the catalogue from the analysis, for example because one is worried that particular parts of the catalogue are unreliable, or exercise unreasonable amounts of leverage on the result. For this reason, it is important to assess the robustness of our method against various potential outlier sub-populations by repeating our analysis on trimmed catalogues where we remove some subset of the clusters. In all these cases, once we

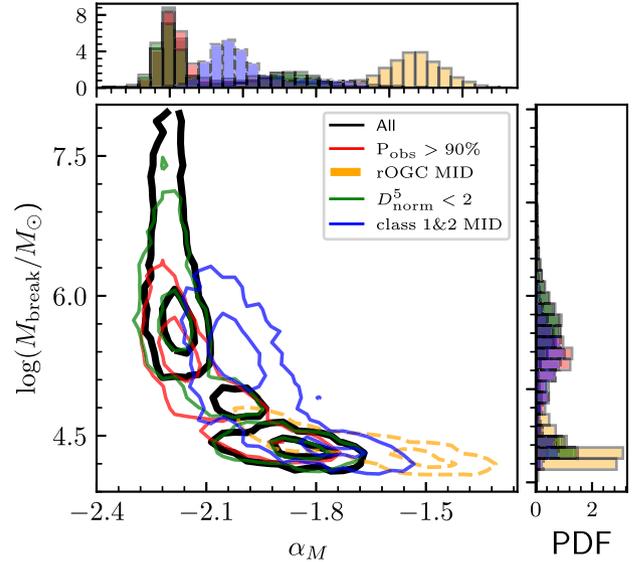


Figure 12. 2D marginal posterior joint PDF of the parameters α_M and $\log M_{\text{break}}$ that describe the CMF, comparing the results derived from the full LEGUS catalogue (black contours) to those derived for five trimmed catalogues where we have removed clusters with low P_{obs} (red, Section 5.1.1), candidate metal-poor old globular clusters (orange, Section 5.1.2), class 3 clusters (green, Section 5.1.3), and clusters with $D_{\text{norm}}^5 > 2$ (blue, Section 5.1.4). In all cases, the inner contour encloses 50 per cent of the MCMC samples, and the outer contour encloses 95 per cent of the samples. The histograms flanking the central panel show the corresponding 1D marginal PDFs of the full and trimmed catalogues.

remove the target subset of clusters, the remainder of our analysis pipeline is identical to the one used for the full catalogue in Section 4.

We compare the mass and age parameters we derive using the reduced catalogues to those we obtain from the full catalogue in Figs 12 and 13, respectively, and report the marginal posterior PDFs for all parameters from the reduced catalogues in Table 1. We also include a table presenting the cluster number statistics of each trimmed catalogue in Table 3.

5.1.1 Trimming low-completeness clusters

In Section 2.3, we explain our method to derive the completeness function $P_{\text{obs}}(\mathbf{m})$ for the data and the library. Knowledge of this function, coupled to our forward-modelling techniques, allows us to avoid the traditional method of excluding clusters with masses $\lesssim 5000 M_\odot$ and ages $\gtrsim 200$ Myr, and in turn allowing us to use essentially all of the LEGUS catalogue, rather than only ≈ 25 per cent of it as in previous work (e.g. Adamo et al. 2017). However, our approach is sensitive to our completeness estimates, which therefore represent a potential source of systematic error. To test this possibility, we repeat the numerical experiments with a more conservative approach of only including clusters where we estimate the completeness is ≥ 90 per cent. This reduces the sample to 1120 clusters, but ensures that $P_{\text{obs}}(\mathbf{m})$ never varies by more than 10 per cent, and thus errors in it have little effect.

Examining Figs 12 and 13, we find that the regions of allowed parameter space for both the parameters describing mass (α_M , $\log M_{\text{break}}$) and those describing age (α_T , $\log T_{\text{MID}}$) are slightly more concentrated for the high completeness sample, but are almost entirely overlapping with the probability maxima we obtain for the

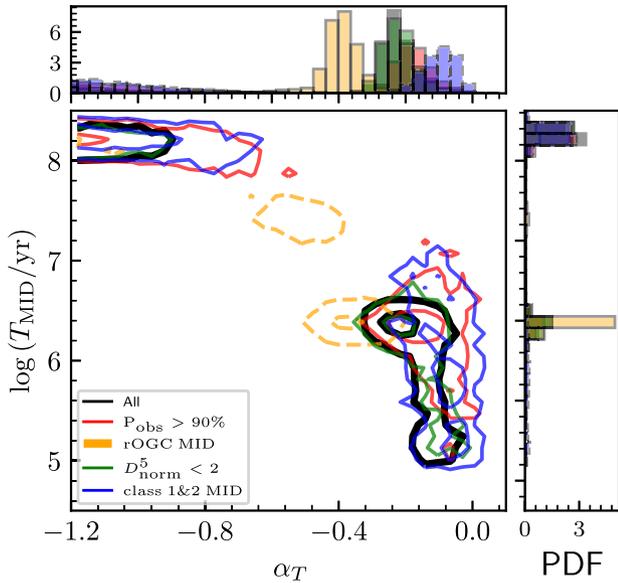


Figure 13. Same as Fig. 12, except showing the PDF of the parameters α_T and $\log(T_{\text{MID}}/\text{yr})$ that describe the CAF for the MID model. For completion, we include the trimmed catalogue comprising only class 1 and 2 clusters as dashed contours. However, since this case the catalogue containing only class 1 and 2 clusters that excludes the old globular clusters (rOGC) is better fit by the MDD model, which cannot be directly compared with the MID fits.

Table 3. Number statistics and best-fitting model for each trimmed catalogue.

Catalogue	N_c	Best model	Section
All	1178	MID	Section 4.1
$\{p_{\text{obs}} \geq 90 \text{ per cent}\}$	1120	MID	Section 5.1.1
OGCs excluded	991	MID	Section 5.1.2
Class 3 excluded	813	MDD	Section 5.1.3
$\{D_{\text{norm}}^5 > 2\}$	1160	MID	Section 5.1.4

Notes. The last column points to the corresponding section of each trimmed catalogue. N_c represents the number of clusters included in the count that have a non-zero completeness value.

full catalogue; the results are consistent at the 1σ level. We therefore conclude that possible systematic errors in the completeness function have at most a trivial impact on the outcome.

5.1.2 Exclusion of candidate metal-poor old globular clusters

Whitmore et al. (2023) argue that using a single metallicity for all the clusters in NGC 628 may yield misleading results, because some of the clusters are old globular clusters (OGCs) with substantially lower metallicities than the remainder of the population. Since our method assumes a single (Solar) metallicity, this represents another source of potential bias for us. To assess the robustness of our results against this bias, we re-fit the NGC 628 catalogue after removing candidate OGCs; following Whitmore et al. (2023), we identify these candidates based on their extremely red colours, $V - I > 0.95$ mag. Examining the comparison between the results derived from this trimmed catalogue and those from the full catalogue in Figs 12 and 13, we find that the results derived after removing the OGCs are consistent with those of the full catalogue in the sense that the contours largely overlap, but

that exclusion of the OGCs appears to eliminate the possibility of a purely power law but steep mass function that is present in the full catalogue. Instead, the OGC-cleaned catalogue strongly favours a truncation in the mass function at $M_{\text{break}} \approx 10^{4.5} M_{\odot}$. Similarly, of the two possibilities for cluster disruption allowed by the full catalogue – no disruption to ≈ 200 Myr and strong disruption thereafter, or weak disruption starting at ≈ 2 Myr, the OGC-cleaned catalogue favours the latter possibility. We therefore conclude that the possible inclusion of metal-poor clusters in our catalogue, and our approach of fitting the data using a single metallicity, represents a small potential source of systematic error, in that these clusters appear to be responsible for driving our full-catalogue fit toward one of the two possible ‘islands’ of probability that it finds. At the same time, however, this effect is not strong enough to skew the results dramatically, in the sense that the posteriors for the OGC-removed catalogue are almost entirely contained within the range of posteriors allowed for the full catalogue.

5.1.3 Exclusion of class 3 clusters

Our next test is to remove from our catalogue clusters identified as LEGUS class 3. These objects are characterized by irregular, diffuse, or multi-peaked morphologies, likely indicative of a population that is either unbound or is bound but is so young that it has not dynamically relaxed. There is considerable debate in the literature about whether these objects should be categorized as clusters at all (see Krumholz et al. 2019a, and references therein), and previous studies indicate that their mass and age distributions differ from those of the more dynamically relaxed class 1 and 2 objects (Grasha et al. 2015; Adamo et al. 2017). To investigate the outcome without class 3 objects, we re-fit the NGC 628 catalogue after removing them.

Consulting Fig. 12, the resulting constraints on the mass distribution are only slightly different from those derived from the full catalogue; Table 1. The best-fitting break mass is similar to full catalogue fits, and the best-fitting index α_M is very slightly shallower, ≈ -2 rather than ≈ -2.1 . The best-fitting parameters describing cluster ages are nearly identical to those of the full catalogue. Neither of these changes is very significant. Thus, we again see that the differences from the full catalogue, while statistically significant, are not particularly physically significant.

5.1.4 Exclusion of clusters with poor photometric matches

In addition to the sub-populations discussed above that have physical meanings, we are also interested in examining potential outliers in photometric space, where our library may not contain any examples capable of reproducing the data. We wish to ensure that these outliers are not biasing our results. To identify photometric outliers, for each catalogue cluster, we compute the normalized photometric distance to each library cluster, as defined by Krumholz et al. (2015a):

$$D_{\text{norm},i} \equiv \sqrt{\frac{1}{N_F} \sum_{j=1}^{N_F} \left(\frac{m_{j,\text{obs}} - m_{j,i}}{\sigma m_{j,\text{obs}}} \right)^2}, \quad (17)$$

where $m_{j,\text{obs}}$ and $\sigma_{j,\text{obs}}$ are the magnitude and uncertainty of the observed cluster in the j th of N_F filters, and $m_{j,i}$ is the magnitude of the i th library cluster in the same filter. We then compute the fifth nearest-neighbour distance D_{norm}^5 for each observed cluster, and remove from the catalogue clusters with $D_{\text{norm}}^5 > 2$, i.e. those for which our library includes fewer than five cluster that fall within

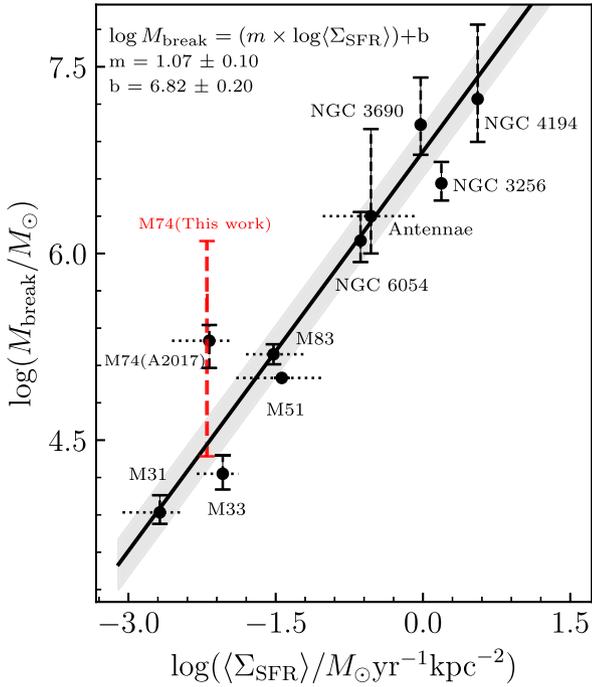


Figure 14. Comparison of break masses M_{break} and star formation rate surface densities Σ_{SFR} for a sample of nearby galaxies. Black points show the compilation of Wainer et al. (2022), while the black line shows the best-fitting linear relation obtained by Johnson et al. (2017); the grey band shows the 84 per cent confidence interval around the relation. ‘M74(A2017)’ denotes the truncated mass derived in Adamo et al. (2017). The red-dashed line shows our fitted truncated mass for the full NGC 628 catalogue at the 84 per cent confidence interval. We find that the CMF is well described by either a Schechter function with a slope $\alpha_M \approx -2$ and a truncation at $M_{\text{break}} \approx 10^{4.5} M_{\odot}$ or a steeper slope $\alpha_M \approx -2.2$ but a substantially higher truncation mass; this uncertainty leads to the extended range of possible M_{break} values in our study. We take our value of Σ_{SFR} for NGC 628 from Sun et al. (2023).

the measured photometric uncertainty interval. This criterion ensures that we exclude clusters that few or no good matches in our catalogue.

Examining Figs 12 and 13, which show the results of fitting the resulting trimmed catalogue, we see that the results align well with the fits from the full sample. The mass parameters are nearly identical, and the age ones differ by less than the uncertainties. This indicates that the method remains robust against outliers in photometric space that exhibit poor correspondence with the library clusters.

5.1.5 Summary of robustness tests

It is worth pausing at this point to summarize the results of our robustness tests. Specifically, we conclude that, for most possible ways we trim the catalogue, the results favour a bimodal shape of the posterior PDFs, where it can either be a truncated mass function with a break mass $M_{\text{break}} \sim 10^{4.5} M_{\odot}$ and a slope close to -2 , experiencing early and mild disruption at $\sim 10^{6.5}$ yr, or a pure power law with a slope of ~ -2.2 , with either no or very late disruption. Examining Figs 12 and 13, it is clear that the 95 per cent confidence regions for the full catalogue (black) in overlap with the 95 per cent confidence intervals of all the trimmed catalogues, and that the 95 per cent confidence intervals of all trimmed catalogues overlap with each other. Thus, there is no strong evidence for differences. The exclusion we consider that has the largest effect is if we remove OGCs; doing so shifts the posteriors in favour of the first of the two probability

peaks and away from the second, but even in this case the 95 per cent confidence intervals overlap and thus there is no convincing evidence for a change.

5.2 CMF with a high-mass truncation

Our results clearly favour a CMF that at high masses declines more sharply than the $\alpha_M = -2$ power-law commonly found to describe CMFs; our fits are ambivalent about what functional form best describes this, with one possibility being a power law with a steeper slope $\alpha_M \approx -2.2$ and no exponential truncation, and the other being a power law with slope $\alpha_M \approx -2$ but an exponential truncation at $M_{\text{break}} \approx 10^{4.5}$. Given this finding, it is interesting to compare our M_{break} to values reported in the literature. Previous studies of CMF in spiral galaxies such as M83 (Adamo et al. 2015), NGC 628 (Adamo et al. 2017), and M51 (Messa et al. 2018) all found that Schechter functions provided better fits to the CMF than pure power laws, but with exponential truncations in mass at $\approx 10^{5.6} M_{\odot}$, which agrees with our best-fitting truncation mass within an order of magnitude. However, given the bimodal nature of our posterior PDF, we cannot rule out either the possibility of a pure power-law shape of the CMF, or that there is a truncation at a break mass ≈ 0.5 –1 dex smaller than this value.

It is also worth noting that we do not necessarily expect all galaxies to have the same truncation masses for their CMFs. Johnson et al. (2017) propose a relation between the break mass and the star formation rate density ($\langle \Sigma_{\text{SFR}} \rangle$), perhaps due to both varying systematically with mean interstellar pressure; Wainer et al. (2022) extend this proposal by adding new data. To place our fitted break mass into this context, we add our value for M_{break} in NGC 628 to Wainer et al.’s compilation in Fig. 14; for this purpose, we adopt a star formation rate density for NGC 628 from Sun et al. (2023). As the figure shows, our best-fitting break mass is consistent with the relation fit by Johnson et al.,

$$\log \frac{M_{\text{break}}}{M_{\odot}} = (1.07 \pm 0.10) \log \frac{\langle \Sigma_{\text{SFR}} \rangle}{M_{\odot} \text{ yr}^{-1} \text{ kpc}^{-2}} + (6.82 \pm 0.20), \quad (18)$$

but given the extremely broad uncertainty implied by our bimodal posterior PDF, this is not a particularly strong statement. If the lower M_{break} turns out to be correct, the resulting value would agree well with the Johnson et al. relationship, which in turn would provide a natural explanation for the small truncation mass: NGC 628 has a lower star formation rate per unit area, presumably indicative of a lower interstellar pressure. Notably, Wainer et al. (2022) find a significant underestimation of the upward bias in the truncated mass of NGC 628 fitted by Adamo et al. (2017). Wainer et al.’s revised truncation mass associated with a larger error bar of approximately 1.1 dex, aligns more closely with our fitted truncated mass and wide range of uncertainties.

5.3 Implications for cluster disruption

We find that MID is a better description of the cluster age distribution than MDD in all regions except for clusters with galactocentric radii $< R_{50}$, and that disruption is weak at young ages – either no disruption at all occurs for ≈ 200 Myr and only then becomes rapid ($\alpha_T \approx -1.4$), or disruption starts from young ages but is very mild ($\alpha_T \approx -0.3$). As with our results for the CMF, it is interesting to put this result into the context of previous studies. As summarized by Krumholz et al. (2019a), the results of these studies have tended to depend strongly on the method of cluster

catalogue construction, with authors who use ‘exclusive’ catalogues that exclude objects that do not have compact, round morphologies favouring MDD and weak disruption, while those who use ‘inclusive’ catalogues without such morphological filters favouring MID and more rapid cluster disruption, $\alpha_T \approx -1$. This difference arises because the morphological filters tend to preferentially affect the youngest clusters, so applying them flattens the age distribution. There is a great deal of debate in the literature, which we shall not revisit here, about the extent to which morphological filters artificially remove real star clusters from the sample, versus the extent to which a failure to filter leads to catalogues including large numbers of unbound associations that were never bound to begin with.

Because we include LEGUS class 3 sources in our catalogue, and thus are not applying morphological filters, one would naively have expected our results to match more closely those coming from the inclusive catalogues. However, our actual result – MID but weak disruption, at least at ages below a few hundred Myr – sits partway between the inclusive and exclusive results, and this is likely at least in part because our method allows us to investigate significantly older ages than in previous studies. This makes us much less sensitive to the young ages where the uncertainty about morphological filtering is most acute. Our statistical signal is not dominated by the small number of clusters with ages $\lesssim 100$ Myr, but by the much larger number of clusters at somewhat older ages, where identification of clusters is much less ambiguous since at these ages all bound stellar systems are relaxed.

It is also interesting to compare our results to those for NGC 628 specifically, for which Adamo et al. (2017) found MID of clusters in the inner pointing NGC 628c, but MDD at the outskirts in pointing NGC 628e. Since we choose to analyse the cluster population as a whole or divide it by galactocentric radius, rather than by pointing, the samples we analyse are not completely identical. However, since the clusters belonging to NGC 628c (the inner region) constitute the great majority of the population, our model selection result for NGC 628 as a whole is consistent with the findings of Adamo et al. There is possibly some tension in the fact that we favour MID even when we divide our catalogues by galactocentric radius, while Adamo et al. find MDD in NGC 628e, but since the divisions are not exactly the same we cannot make this statement more quantitatively. However, it does reinforce the point that our method has the advantage that, since we are forward modelling the completeness in each field, we can easily break up the clusters by any criterion of our choice to derive the demographics in different parts of the galaxy.

6 SUMMARY AND CONCLUSION

We use a novel method described by Krumholz et al. (2019b) to derive the demographics of the cluster population of the nearby spiral galaxy NGC 628 from unresolved photometry measured by the LEGUS survey (Calzetti et al. 2015). This method has the advantage of integrating diverse observational data imaged using different filters or depths, without the need for binning or imposing extreme completeness cuts. As a result, our method allows us to analyse the roughly 1200 clusters in NGC 628 catalogued by the LEGUS survey, as opposed to the past work that was restricted to include only the most massive and youngest ~ 30 per cent of the population. Our method also obviates the need to assign a unique mass and age to every individual cluster of the population, and therefore avoids the inevitable information loss incurred by reducing complex, sometimes

multiply peaked posterior probability distributions for individual cluster masses and ages to single values or even Gaussians. We validate the method by comparing the photometric distributions that emerge from our fits to the observed ones and finding good agreement, indicating that our approach can identify models for the cluster population that reproduce observations. Our analysis results in three primary findings.

First, we use the AIC to compare two classes of parametric models for the time-evolution of the cluster population, MID and MDD. We find that the MID model fits the data better, both for the entire NGC 628 cluster catalogue and for all sub-divisions of it except one; the exception is if we consider only the half of the catalogue closer to the galactic centre. We thus conclude that cluster mass and age distributions are likely at least approximately separable and that the disruption time-scales of the clusters have no dependence on the cluster masses.

Secondly, we find that CMFs are well described by Schechter functions $dN/dM \propto M^{\alpha_M} \exp(-M/M_{\text{break}})$, and that there is good evidence that the mass function at $\sim 10^5 M_\odot$ is steeper than the slope $\alpha_M = -2$, most commonly found in previous studies. However, we are at present unable to determine if the data are better fit by an untruncated power law with a steeper slope of $\alpha_M \approx -2.2$, or by a shallower slope but a truncation mass $M_{\text{break}} \approx 10^{4.5} M_\odot$. The latter possibility is consistent with the relationship between CMF truncation mass and star formation surface density proposed by Johnson et al. (2017) and Wainer et al. (2022).

Thirdly, we find that the cluster age distribution in NGC 628 indicates that cluster disruption is weak at ages up to ≈ 200 Myr; it either only begins at such ages and then becomes rapid, or begins early but then is very mild, with only a factor of ≈ 2 decrease in cluster number per decade in age (index $\alpha_T \approx -0.3$). In past studies, the value of α_T has proven highly sensitive to the cluster catalogue type used, with studies using inclusive cluster catalogues (those that do not include filters based on morphology) generally finding $\alpha_T \sim -1$ while those using exclusive catalogues (which are filtered for morphology) finding $\alpha_T \sim -0.3$. The significance of our study is that even though we are using an inclusive catalogue, we still find relatively weak disruption at young ages. One possible reason is that we can use a much larger sample and probe to considerably larger ages than previous studies such as Adamo et al. (2017).

The three conclusions above are derived from our analysis of all the clusters in NGC 628 catalogued by LEGUS. However, we also group clusters by galactocentric radius and search for differences in cluster demographics with radius. We find marginal evidence for radial variations, with suggestive hints that inner galaxy clusters are subject to more severe disruption than those further out, and that inner galaxy clusters have a mass function that is better described by a pure power law while outer galaxy ones have a mass function that is more strongly truncated. However, our ability to search for radial variations is limited by the fairly small number of clusters available at larger galactocentric radii, since the great majority of the available sample comes from the LEGUS NGC 628c pointing, which targets the galactic centre. Given the suggestive hints in our analysis, increasing the sample of outer-galaxy clusters would be a worthwhile effort.

Looking forward, now that we have demonstrated the performance of our pipeline on NGC 628, we will apply it to other cluster catalogues, both those originating from LEGUS and from other studies. Analysis of a diverse sample of galaxies will allow us to compare the demographics of cluster populations in various galactic environments. This, in turn, opens up the possibility of testing a range

of theoretical models for cluster formation and disruption and their variation with the galactic environment.

ACKNOWLEDGEMENTS

JT and KG are supported by the Australian Research Council through the Discovery Early Career Researcher Award (DECRA) Fellowship DE220100766 funded by the Australian Government. MRK acknowledges support from the ARC Future Fellowship and Laureate Fellowship funding schemes, awards FT180100375 and FL220100020. This work is supported by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013. Based on observations made with the NASA/ESA Hubble Space Telescope, obtained at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555. These observations are associated with program #13364.

DATA AVAILABILITY

The cluster catalogue data underlying this paper are publicly available online at [MAST](#). The software suite, tools, and LEGUS catalogues required to perform the MCMC optimization, as well as the output chains, are publicly accessible on [Bitbucket](#).

REFERENCES

- Adamo A., Kruijssen J. M. D., Bastian N., Silva-Villa E., Ryon J., 2015, *MNRAS*, 452, 246
- Adamo A. et al., 2017, *ApJ*, 841, 131
- Ashworth G. et al., 2017, *MNRAS*, 469, 2464
- Bastian N., Hempel M., Kissler-Patig M., Homeier N. L., Trancho G., 2005, *A&A*, 435, 65
- Bastian N. et al., 2012, *MNRAS*, 419, 2606
- Baumgardt H., Kroupa P., 2007, *MNRAS*, 380, 1589
- Bialopetravičius J., Narbutis D., Vansevičius V., 2019, *A&A*, 621, A103
- Calzetti D. et al., 2015, *AJ*, 149, 51
- Chandar R., Whitmore B. C., Calzetti D., O’Connell R., 2014, *ApJ*, 787, 17
- Chandar R., Fall S. M., Whitmore B. C., 2015, *ApJ*, 810, 1
- Choi J., Dotter A., Conroy C., Cantiello M., Paxton B., Johnson B. D., 2016, *ApJ*, 823, 102
- Cioni M. R. L., Girardi L., Marigo P., Habing H. J., 2006, *A&A*, 452, 195
- Dalcanton J. J. et al., 2012, *ApJS*, 200, 18
- da Silva R. L., Fumagalli M., Krumholz M., 2012, *ApJ*, 745, 145
- da Silva R. L., Fumagalli M., Krumholz M. R., 2014, *MNRAS*, 444, 3275
- de Meulenaer P., Narbutis D., Mineikis T., Vansevičius V., 2014, *A&A*, 569, A4
- Deger S. et al., 2022, *MNRAS*, 510, 32
- Dinnbier F., Walch S., 2020, *MNRAS*, 499, 748
- El-Badry K., Weisz D. R., Quataert E., 2017, *MNRAS*, 468, 319
- Fall S. M., Chandar R., 2012, *ApJ*, 752, 96
- Fall S. M., Zhang Q., 2001, *ApJ*, 561, 751
- Fall S. M., Chandar R., Whitmore B. C., 2005, *ApJ*, 631, L133
- Fall S. M., Chandar R., Whitmore B. C., 2009, *ApJ*, 704, 453
- Fensch J. et al., 2019, *A&A*, 628, A60
- Fitzpatrick E. L., 1999, *Publ. Astron. Soc. Pac.*, 111, 63
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *Publ. Astron. Soc. Pac.*, 125, 306
- Fouesneau M. et al., 2014, *ApJ*, 786, 117
- Gieles M., Renaud F., 2016, *MNRAS*, 463, L103
- Gieles M., Portegies Zwart S. F., Baumgardt H., Athanassoula E., Lamers H. J. G. L. M., Sipiør M., Leenaarts J., 2006, *MNRAS*, 371, 793

- Gieles M., Lamers H. J. G. L. M., Portegies Zwart S. F., 2007, *ApJ*, 668, 268
- Grasha K. et al., 2015, *ApJ*, 815, 93
- Grasha K. et al., 2017, *ApJ*, 842, 25
- Grudić M. Y., Kruijssen J. M. D., Faucher-Giguère C.-A., Hopkins P. F., Ma X., Quataert E., Boylan-Kolchin M., 2021, *MNRAS*, 506, 3239
- Gusev A. S., Egorov O. V., Sakhibov F., 2014, *MNRAS*, 437, 1337
- Hills J. G., 1980, *ApJ*, 235, 986
- Howard C. S., Pudritz R. E., Harris W. E., 2017, *MNRAS*, 470, 3346
- Johnson L. C. et al., 2017, *ApJ*, 839, 78
- Kruijssen J. M. D., Pelupessy F. I., Lamers H. J. G. L. M., Portegies Zwart S. F., Bastian N., Icke V., 2012, *MNRAS*, 421, 1927
- Krumholz M. R., Fumagalli M., da Silva R. L., Rendahl T., Parra J., 2015a, *MNRAS*, 452, 1447
- Krumholz M. R. et al., 2015b, *ApJ*, 812, 147
- Krumholz M. R., McKee C. F., Bland-Hawthorn J., 2019a, *ARA&A*, 57, 227
- Krumholz M. R., Adamo A., Fumagalli M., Calzetti D., 2019b, *MNRAS*, 482, 3550
- Lamers H. J. G. L. M., Gieles M., Portegies Zwart S. F., 2005a, *A&A*, 429, 173
- Lamers H. J. G. L. M., Gieles M., Bastian N., Baumgardt H., Kharchenko N. V., Portegies Zwart S., 2005b, *A&A*, 441, 117
- Lamers H. J. G. L. M., Baumgardt H., Gieles M., 2010, *MNRAS*, 409, 305
- Landini M., Natta A., Oliva E., Salinari P., Moorwood A. F. M., 1984, *A&A*, 134, 284
- Larsen S. S., 2009, *A&A*, 494, 539
- Lee J. C. et al., 2022, *ApJS*, 258, 10
- Leitherer C. et al., 1999, *ApJS*, 123, 3
- Linden S. T. et al., 2017, *ApJ*, 843, 91
- Maschberger T., Kroupa P., 2009, *MNRAS*, 395, 931
- Menon S. H. et al., 2021, *MNRAS*, 507, 5542
- Messa M. et al., 2018, *MNRAS*, 473, 996
- Mok A., Chandar R., Fall S. M., 2020, *ApJ*, 893, 135
- Moustakas J., Kennicutt Robert C. J., Tremonti C. A., Dale D. A., Smith J.-D. T., Calzetti D., 2010, *ApJS*, 190, 233
- Olivares E. F. et al., 2010, *ApJ*, 715, 833
- Piskunov A. E., Kharchenko N. V., Schilbach E., Röser S., Scholz R. D., Zinnecker H., 2009, *A&A*, 507, L5
- Sakhibov F., Gusev A. S., Hemmerich C., 2021, *MNRAS*, 508, 912
- Smirnov N. V., 1939, *Bull. Math. Univ. Moscou*, 2, 3
- Sun J. et al., 2023, *ApJ*, 945, L19
- Vázquez G. A., Leitherer C., 2005, *ApJ*, 621, 695
- Wainer T. M. et al., 2022, *ApJ*, 928, 15
- Webb J. J., Reina-Campos M., Kruijssen J. M. D., 2019, *MNRAS*, 486, 5879
- Whitmore B. C. et al., 2023, *MNRAS*, 520, 63
- Zackrisson E., Rydberg C.-E., Schaerer D., Östlin G., Tuli M., 2011, *ApJ*, 740, 13

APPENDIX A: THE CLUSTER LIBRARY

The Krumholz et al. (2019b) approach described in the main text requires a library of synthetic clusters. We construct a library of 4×10^7 clusters for this purpose using the SLUG stochastic stellar population synthesis code (da Silva et al. 2012; Krumholz et al. 2015a). Construction of the library involves two steps. In the first, we choose a mass M , age T , and extinction A_V for each cluster from a specified library distribution $p_{\text{lib}}(M, T, A_V)$. We choose this library distribution to be

$$p_{\text{lib}}(M, T, A_V) \propto p(M)p(T)p(A_V), \quad (\text{A1})$$

with

$$p(M) \propto \begin{cases} (M/10^2 M_\odot)^{-1}, & 10^2 < M/M_\odot \leq 10^5 \\ 10^{-3}(M/10^5 M_\odot)^{-2}, & 10^5 < M/M_\odot \leq 10^7 \\ 0, & \text{otherwise} \end{cases}$$

$$p(T) \propto \begin{cases} T^{-1}, & 10^5 < T/\text{yr} < 1.5 \times 10^{10} \\ 0, & \text{otherwise} \end{cases}$$

$$p(A_V) \propto \begin{cases} \text{const}, & 0 < A_V < 3 \text{ mag} \\ 0, & \text{otherwise} \end{cases}$$

This distribution maximizes sampling of clusters with low masses and young ages, where stochastic effects are strongest and thus the largest number of samples is needed to fully characterize the luminosity distribution.

The second step is to generate photometric magnitudes for each cluster. To do so, we run SLUG using the Padova tracks with TP-AGB stars (Cioni et al. 2006), coupled to the ‘starburst99’ treatment of stellar atmospheres (Leitherer et al. 1999). We draw stars from a Chabrier IMF using SLUG’s ‘stop nearest’ sampling policy. Full explanations of these choices are provided in Krumholz et al. (2015a), and in the SLUG documentation available at <https://slug2.readthedocs.io/>. The SLUG parameter file used to generate the library is available at https://bitbucket.org/janet_jianling_tang/legus_slug23/src/master/. This file includes all the necessary inputs and parameters described above. The library itself is too large to post, but is available upon request to the authors.

APPENDIX B: THE EFFECTS OF STELLAR TRACKS

To investigate how the choice of stellar tracks affects our analysis, we repeat our fit to the full catalogue, which in the main text we perform using a library generated from Padova–AGB stellar tracks (Vázquez & Leitherer 2005), with a library generated from MIST stellar tracks (Choi et al. 2016). We leave all other aspects of the procedure unchanged. We then use the best-fitting results to generate 1D luminosity functions for the MIST models via the procedure described in Section 4.2.1, which we compare to the fiducial Padova results and to the observations in Fig. B1. In this figure, Padova–AGB and observed curves shown here are identical to those shown in Fig. 3, so the only difference is the addition of the MIST model results. Examining the figure, it is evident that the best-fitting Padova models reproduce the observed luminosity functions substantially better than the best-fitting MIST models, especially at high luminosity and in the redder *V* and *I* bands. It is this difference that leads to us adopting the Padova–AGB models as our preferred library for the analysis in the main text.

If we compare the best-fitting parameter values themselves, we find that, similar to the Padova fits, the fits using MIST tracks also prefer the MID over the MDD model. However, for MIST, we find a significantly lower truncation mass, $M_{\text{break}} \sim 10^4 M_\odot$, and a shallower mass slope of $\alpha_M \sim -1$. These changes combine so that the Padova and MIST models actually yield relatively little difference in the slope of the mass function near $M \sim 10^4 M_\odot$, where the LEGUS catalogue is most complete and sensitive – that is, if we evaluate the slope $d \log p_M / d \log M$ at $M = 10^4 M_\odot$, then we obtain similar values close to -2 for both $\alpha_M \approx -2.2$, $M_{\text{break}} \approx 10^{5.6} M_\odot$, the 50th percentile values we find for the Padova models, and for the $\alpha_M \sim -1$, $M_{\text{break}} \sim 10^4 M_\odot$ we obtain from MIST. The fits differ mainly away from this mass, where sensitivity is limited by small number statistics at the massive end and catalogue completeness

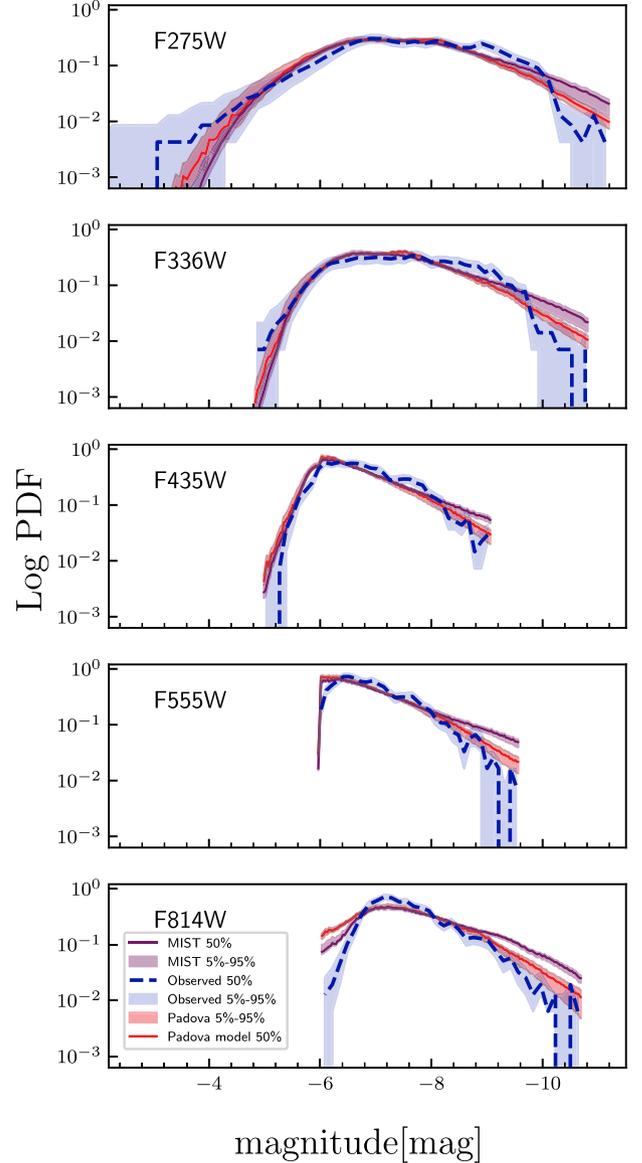


Figure B1. Same as Fig. 3, but now also showing luminosity functions generated using MIST stellar tracks (purple). The observed (blue) and Padova track fit (red) curves are identical to those shown in Fig. 3.

at the low-mass end. The MIST result suggests a mass function that becomes relatively flat for masses $\ll 10^4 M_\odot$ compared to the Padova CMF, but that drops off sharply at $\gg 10^4 M_\odot$ due to its lower break mass.

APPENDIX C: DISTRIBUTION OF EXTINCTIONS

As discussed in the main text, in addition to posteriors for our parameters of interest, our fit also necessarily produces posteriors for the distribution of dust extinctions. While this is a nuisance parameter, it is important for us to verify that the results is reasonable. We show the 1D marginal posterior PDF of dust extinctions that we derive for our MID model fit to the full LEGUS catalogue in Fig. C1. We see that the results are in reasonable agreement with expectations: the extinction is small to moderate for most clusters, $A_V \lesssim 1$, with a

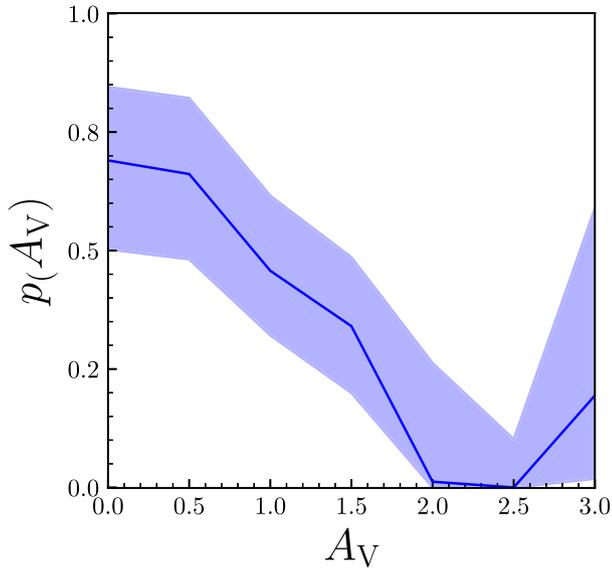


Figure C1. Posterior PDF of the distribution of dust extinctions A_V . The solid line shows the 50th percentile result, and the shaded band shows the 5th to 95th percentile range.

minority having higher extinction. We have confirmed that the upturn in A_V is not due to outliers in photometric space.

We note that uncertainty about the shape of the extinction distribution does contribute to the overall uncertainty of our fits, since the posterior PDFs for all quantities that we care about must be marginalized over the nuisance A_V distribution. In particular, as discussed in Krumholz et al. (2015a), star clusters can exhibit degeneracies between mass, age, and extinction, leading to multiple fits of comparable likelihood with varying sets of physical parameters. In such cases, utilizing a full PDF as done in this work, provides more realistic estimate of the errors as compared to traditional SED fitting methods that return a single best fit with a Gaussian error estimate. However, there is certainly still room for improvement. In particular, Ashworth et al. (2017) have demonstrated that the inclusion of further photometric data, particularly narrow-band $H\alpha$, can help break some of these degeneracies in forward-modelling techniques. Presumably, the addition of such data to our population-level fits would also help reduce the uncertainties.

APPENDIX D: ADDITIONAL COMPARISONS IN PHOTOMETRIC SPACE

In addition to the colour–colour and colour–magnitude plots of $UV-U$ versus $U-B$ and $UV-U$ versus U presented in Section 4.2, where the most significant disagreements occur, we also present colour–magnitude comparisons for all the other bands in Figs D2–D4 and colour–colour comparison in Fig. D1.

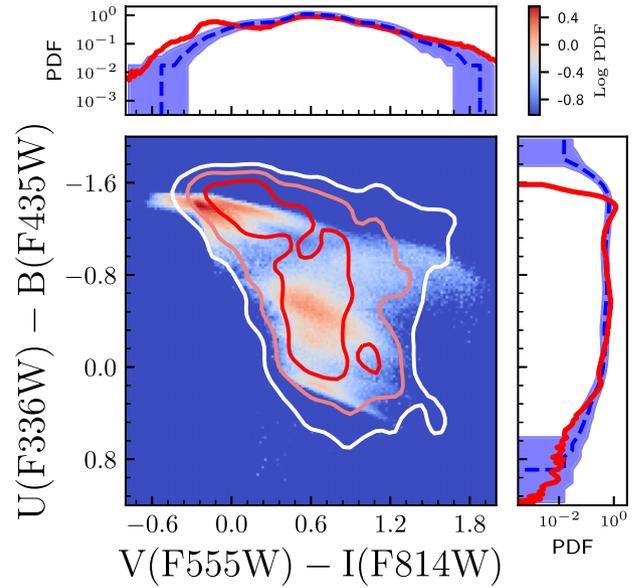


Figure D1. Same as Fig. 5, but showing the $U-B$ versus $V-I$ colour–colour distributions of the observations and the best-fitting model. The y-axis is inverted to align with conventions used in star cluster evolutionary tracks, such as fig. 20 of Deger et al. (2022).

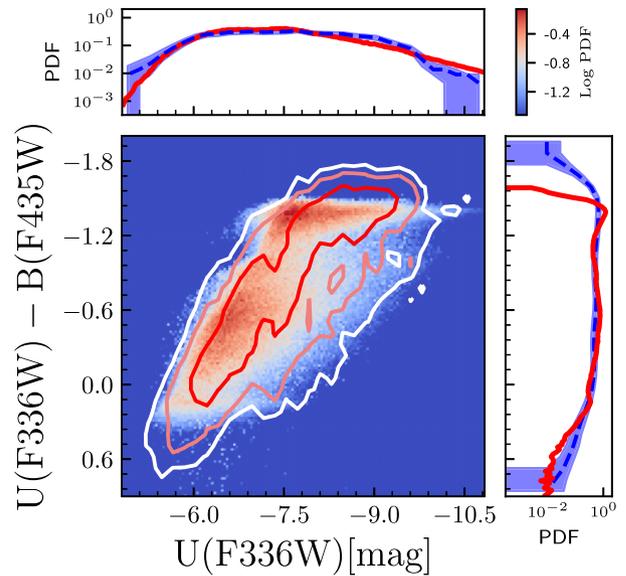


Figure D2. Same as Fig. 5, except we show U versus $U-B$ in colour–magnitude space here.

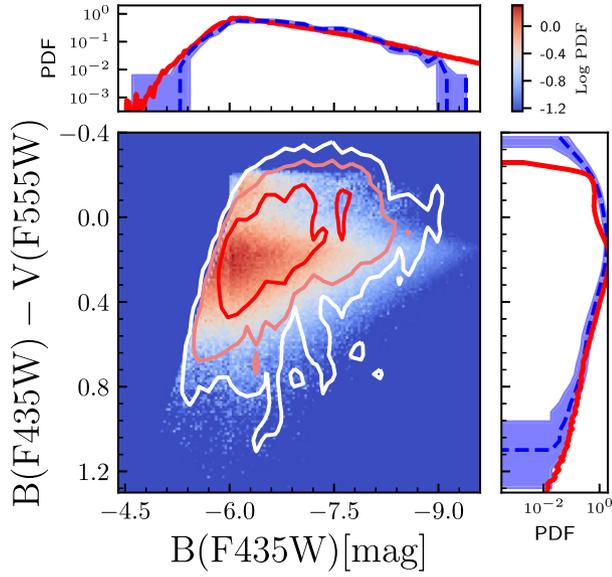


Figure D3. Same as Fig. 5, except we show B versus $B - V$ in colour-magnitude space here.

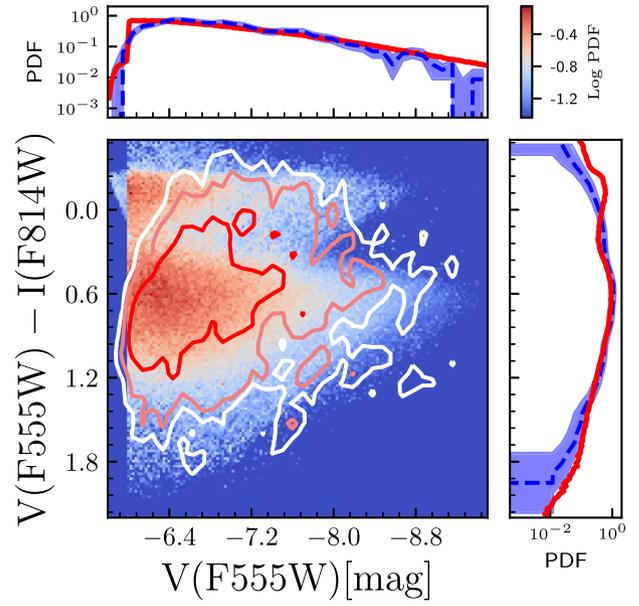


Figure D4. Same as Fig. 5, except we show V versus $V - I$ in colour-magnitude space here.

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