Exploring the IMF of star clusters: a joint SLUG and LEGUS effort

G. Ashworth,[1,2]★ M. Fumagalli,[1,2] M. R. Krumholz,[3] A. Adamo,[4] D. Calzetti,[5]
R. Chandar,[6] M. Cignoni,[7] D. Dale,[8] B. G. Elmegreen,[9] J. S. Gallagher, III,[10]
D. A. Gouliermis,[11,12] K. Grasha,[5] E. K. Grebel,[13] K. E. Johnson,[14] J. Lee,[15]
M. Tosi[16] and A. Wofford[17]

[1]*Institute for Computational Cosmology, University of Durham, South Road, Durham DH1 3LE, UK*
[2]*Centre for Extragalactic Astronomy, University of Durham, South Road, Durham DH1 3LE, UK*
[3]*Research School of Astronomy & Astrophysics, Australian National University, Canberra, ACT 2611, Australia*
[4]*Department of Astronomy, Oskar Klein Centre, Stockholm University, Stockholm, SE-10691, Sweden*
[5]*Department of Astronomy, University of Massachusetts-Amherst, Amherst, MA 01003, USA*
[6]*Department of Physics and Astronomy, University of Toledo, Toledo, OH 43606, USA*
[7]*Department of Physics, University of Pisa, Largo Pontecorvo 3, I-56127 Pisa, Italy*
[8]*Department of Physics & Astronomy, University of Wyoming, Laramie, WY 82071, USA*
[9]*IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY 10598, USA*
[10]*Department of Astronomy, University of Wisconsin-Madison, Madison, WI 53706, USA*
[11]*Institut für Theoretische Astrophysik, Zentrum für Astronomie der Universität Heidelberg, D-69120 Heidelberg, Germany*
[12]*Max Planck Institute for Astronomy, D-69117 Heidelberg, Germany*
[13]*Astronomisches Rechen-Institut, Zentrum für Astronomie der Universität Heidelberg, D-69120 Heidelberg, Germany*
[14]*Department of Astronomy, University of Virginia, 530 McCormick Road, Charlottesville, VA 22904, USA*
[15]*Space Telescope Science Institute, Baltimore, MD 21218, USA*
[16]*INAF–Osservatorio Astronomico di Bologna, Bologna, I-40127, Italy*
[17]*Instituto de Astronomía, UNAM, Ensenada, CP 22860, Baja California, Mexico*

ABSTRACT

We present the implementation of a Bayesian formalism within the Stochastically Lighting Up Galaxies (SLUG) stellar population synthesis code, which is designed to investigate variations in the initial mass function (IMF) of star clusters. By comparing observed cluster photometry to large libraries of clusters simulated with a continuously varying IMF, our formalism yields the posterior probability distribution function (PDF) of the cluster mass, age and extinction, jointly with the parameters describing the IMF. We apply this formalism to a sample of star clusters from the nearby galaxy NGC 628, for which broad-band photometry in five filters is available as part of the Legacy ExtraGalactic UV Survey (LEGUS). After allowing the upper-end slope of the IMF (α_3) to vary, we recover PDFs for the mass, age and extinction that are broadly consistent with what is found when assuming an invariant Kroupa IMF. However, the posterior PDF for α_3 is very broad due to a strong degeneracy with the cluster mass, and it is found to be sensitive to the choice of priors, particularly on the cluster mass. We find only a modest improvement in the constraining power of α_3 when adding Hα photometry from the companion Hα-LEGUS survey. Conversely, Hα photometry significantly improves the age determination, reducing the frequency of multi-modal PDFs. With the aid of mock clusters, we quantify the degeneracy between physical parameters, showing how constraints on the cluster mass that are independent of photometry can be used to pin down the IMF properties of star clusters.

Key words: methods: statistical – stars: luminosity function, mass function – galaxies: individual: NGC 628 – galaxies: star clusters: general.

1 INTRODUCTION

A core parameter that is common to many astrophysical problems is the stellar initial mass function (IMF), which describes how stellar

★ E-mail: greg.ashworth@durham.ac.uk

masses are distributed across a group of stars at the time of their birth. Constraining the form and evolution of the IMF as a function of galaxy properties is of great importance in our theoretical understanding of how star formation proceeds in galaxies and star clusters. Moreover, as the form of the IMF will dictate the make-up of a given star cluster, and hence its mass-to-light ratio, any inference an observer makes about the properties of a given unresolved stellar population via photometry will be strongly dependent on the assumed shape of the IMF (see e.g. Bell et al. 2003).

The IMF is traditionally accepted to be universal, invariant both in time and with respect to the properties of the local environment. Two commonly adopted functional forms that describe the IMF are the broken power law of Kroupa (2001) and the lognormal form of Chabrier (2003), which features a power-law tail at the high-mass end ($M \geq 1\,\mathrm{M_\odot}$). Both these functional forms possess a high-mass slope that is consistent with that of the power-law IMF originally presented by Salpeter (1955), with index $\alpha = -2.35$. This form of the IMF appears to hold in most nearby environments observed to date (Massey, Johnson & Degioia-Eastwood 1995; Bastian, Covey & Meyer 2010).

Although theoretical predictions exist for the variation of the IMF with, for example, metallicity and gas temperature (Adams & Fatuzzo 1996; Bonnell, Clarke & Bate 2006; Krumholz 2011), at present there is no widely accepted evidence of IMF variation. There is, however, an increasing body of literature supporting the notion of a non-universal IMF, with variations across different environments. For instance, van Dokkum & Conroy (2010) found evidence suggesting a bottom-heavy IMF in early-type galaxies based on the strength of distinctive absorption lines (Na I doublet and Wing–Ford molecular FeH band), which can be modelled with an overabundance of low-mass stars in massive galaxies. Similar detections have been reported in the more recent literature (Conroy & van Dokkum 2012), with further evidence in support of a steeper IMF in massive elliptical galaxies coming from dynamical measurements (Cappellari et al. 2012). While both methods consistently imply more bottom-heavy IMFs in massive galaxies, they appear to disagree in detail when applied to individual galaxies rather than to a galaxy population (see Smith 2014). However, a more recent study by Lyubenova et al. (2016) using consistent data sets has not found this disagreement to be apparent.

At the other end of the galaxy mass function, recent observations have also revealed apparent correlations between the Hα-to-FUV luminosity ratio and the star formation rate (SFR; Lee et al. 2009, 2016) in dwarf galaxies, which can be attributed to variation in the high-mass end of the IMF. Similar suggestions of variation in the IMF have been made based on comparisons between the Hα-to-FUV luminosity ratio and Hα and *R*-band surface brightness in starburst galaxies (Meurer et al. 2009). Other indications of a variable IMF have been gleaned using techniques involving the Hα equivalent width and galaxy colours (Gunawardhana et al. 2011), and chemical evolution arguments (Matteucci 1994; Thomas 1999).

Attempts to explain these observations include the framework of the integrated galactic initial mass function (IGIMF; Kroupa 2001; Kroupa & Weidner 2003; Weidner, Pflamm-Altenburg & Kroupa 2011; Kroupa et al. 2013). In this theory, the most massive stars allowed to form in a cluster (corresponding to the high-mass cutoff of the IMF) are correlated with the cluster mass, which in turn depends on the galaxy SFR. These correlations suppress the formation of massive stars in dwarf galaxies, resulting in a lower Hα-to-FUV ratio, in line with the aforementioned observations.

However, the very same observations can be explained by alternative models that do not resort to deterministic variations of the IMF. For example, stochasticity has been shown to suitably describe the Hα-to-FUV ratio correlation in dwarf galaxies without the need for an explicit correlation between star and cluster masses. Indeed, a combination of the random sampling of clusters from a cluster mass function (CMF) and a random sampling of stars from the IMF (Fumagalli, da Silva & Krumholz 2011) appears to satisfactorily reproduce the observed Hα-to-FUV ratio in dwarf galaxies. Similar results have been reported by Calzetti et al. (2010) and Andrews et al. (2013), who have studied the ratio of Hα to cluster mass, finding that the IMF upper-end slope has no clear mass dependence (see also Corbelli et al. 2009). Andrews et al. (2014) find that stochasticity is similarly effective in explaining observations of the spiral galaxy M83. Additional challenges to the IGIMF theory are discussed, for instance, in Oey et al. (2013) and Lamb et al. (2016). These authors present observations of what are likely to be solitary O and B stars with no visible evidence of escape from nearby clusters. Such an occurrence would violate the requirement of the IGIMF that the maximum stellar mass correlates with cluster mass. However, previous indications of the possibility of solitary OB stars (de Wit et al. 2004, 2005; Lamb et al. 2010) have been disputed by Gvaramadze et al. (2012), and more recently Stephens et al. (2017) have made arguments in opposition to stochastic IMF sampling as a rebuttal of the IGIMF through observations of compact clusters of pre-main-sequence stars around massive stars previously thought to be isolated.

This ambiguity surrounding the universality of the IMF calls for the development of new techniques and methods that explore the properties of the IMF in different environments. To this end, Bayesian techniques have recently been applied to stellar population synthesis (SPS) models, in an attempt to infer the parameters that describe the IMF. A particularly appealing feature of these Bayesian approaches is the possibility to marginalize over other 'nuisance' parameters, and to explore to what extent the IMF is degenerate with respect to other parameters used in SPS models.

For example, Weisz et al. (2015) use Bayesian techniques combined with colour–magnitude diagrams of resolved star clusters to derive a constant IMF high-end slope that is both steeper than the canonical Kroupa value of $\alpha_3 = -2.3$ and constant in Andromeda. Dries, Trager & Koopmans (2016) have also recently presented an example of how Bayesian methods can be applied to this problem, successfully recovering the input parameters describing the IMFs in a selection of mock simple stellar populations produced with varying IMFs, provided however that a representative set of stellar templates is available during the analysis.

Following a similar idea, in this work we aim to combine Bayesian techniques with the stochastic SPS code SLUG (Stochastically Lighting Up Galaxies; da Silva, Fumagalli & Krumholz 2012, 2014; Krumholz et al. 2015a) to explore the IMF properties of star clusters. This analysis builds on the method described in Krumholz et al. (2015b), who combined SLUG with Bayesian inference to constrain the mass, age and extinction of clusters from broad-band photometry. In this work, expanding on their previous analysis, we develop a formalism to also investigate whether the high-mass end of the IMF can be constrained through photometry of individual clusters, while tracking at the same time degeneracies between the other parameters that regulate the mass-to-light ratio of the stellar populations. To achieve this goal, we first need to extend the current capabilities of SLUG, as we describe briefly in Section 2 and in Appendix A. Specifically, we enable a new mode in SLUG through which libraries of simulated clusters can

be constructed using a continuous distribution of arbitrary IMFs. These libraries can then be compared to actual observations using Bayesian techniques to infer the posterior probability distribution function (PDF) for the parameters that describe the cluster physical properties, including the IMF, within SPS models.

Following the implementation of this extension in the SLUG code, in Section 3, we apply this extended Bayesian formalism to photometric data on star clusters from the Legacy ExtraGalactic UV Survey (LEGUS). LEGUS is a Cycle 21 *Hubble Space Telescope* Treasury programme (Calzetti et al. 2015), designed to investigate star formation and its relation with the galactic environment across a representative sample of nearby galaxies within 12 Mpc, on scales ranging from individual stars up to kiloparsec-scale structures. In addition to this application, in Section 4 we construct a library of mock star clusters, which we use to explore from a theoretical point of view how accurately the upper-end slope of the IMF (hereafter the α_3 parameter) can be recovered using photometry from star clusters. In doing this, we also quantify the degeneracy between IMF parameters and other physical quantities, such as mass and age. Summary and conclusions follow in Section 5, where we discuss future improvements to our method and describe future experiments that could yield tight constraints on the IMF of star clusters. Throughout this work, unless otherwise specified, magnitudes are in the Vega system.

2 THE SLUG SOFTWARE SUITE

The core of the SLUG suite is an SPS code that, unlike the majority of traditional SPS codes, such as STARBURST99 (SB99; Leitherer et al. 1999), includes the effects of stochastic sampling. Stochastic sampling of the IMF was pioneered in the MASSCLEAN software of Popescu & Hanson (2009). In SLUG, this stochasticity is achieved through generating galaxies by drawing stars and clusters from underlying PDFs that define the IMF and the CMF in the simulated galaxy. These populations are then evolved through time to the desired age, accounting for cluster disruption. We refer the reader to previous work for a detailed description of the algorithms adopted in SLUG (e.g. da Silva et al. 2012; Krumholz et al. 2015a). Briefly, in its basic 'cluster mode', SLUG constructs star clusters by drawing stars from an IMF until a desired target mass is reached. This sampling method allows SLUG to accurately capture the effects of stochasticity, which are particularly relevant for low-mass clusters (e.g. Cerviño et al. 2002; Elmegreen 2002; Haas & Anders 2010). Moreover, the ability to populate star clusters (or galaxies) by randomly drawing from the IMF and CMF provides a convenient way to generate large libraries of realistic clusters that are representative of the observed population in nearby galaxies (e.g. Krumholz et al. 2015a).

Along with the core SPS code, the SLUG distribution includes the CLUSTER_SLUG software, which is a package designed for Bayesian analysis of observed cluster photometry (Krumholz et al. 2015a). Specifically, CLUSTER_SLUG uses Bayesian inference to recover the posterior PDFs for physical parameters (e.g. mass, age, extinction) that influence the photometric properties of a given star cluster. To this end, the CLUSTER_SLUG package makes use of a large 'training set' of SLUG models that, through Bayesian analysis combined with kernel density estimation (KDE) techniques (da Silva et al. 2014; Krumholz et al. 2015a), yields the PDFs for the parameters of interest given the observed broad-band photometry of a stellar cluster.

In this work, we extend this technique to the case of varying IMFs. Within the Bayesian framework of CLUSTER_SLUG, the posterior PDFs of parameters describing the functional form of the IMF (for instance, the high-end slope α_3 in the Kroupa IMF) can

be derived from cluster photometry, provided that SLUG libraries containing a continuous distribution of clusters with respect to the parameter of interest (e.g. α_3) are available. To this end, we have implemented a variable IMF capability in SLUG, with which libraries of simulated clusters can be generated with continuously variable IMF parameters. This modification is discussed in detail in Appendix A.

3 APPLYING SLUG TO OBSERVED STAR CLUSTERS

Having enabled a variable IMF capability in SLUG, we now combine a large library of simulated star clusters with LEGUS broad-band photometry. This photometry is taken from a star cluster catalogue of the galaxy NGC 628. We use this combination to investigate whether it is possible to constrain the upper-end slope of the IMF with our Bayesian technique.

3.1 The NGC 628 photometric catalogues

For this analysis, we make use of the multiwavelength imaging available from the LEGUS project. As part of the LEGUS survey (see Calzetti et al. 2015), local volume galaxies have been imaged with five broad-band filters using the Wide-Field Camera 3 (WFC3), with parallel observations from the Advanced Camera for Surveys (ACS). Of the 50 galaxies in the LEGUS sample, we select NGC 628 for our analysis. The filters used for NGC 628 are as follows: WFC3 *F*275*W*, *F*336*W* and *F*555*W*, along with ACS *F*435*W* and *F*814*W*. These may be generally thought of as the NUV, *U*, *B*, *V* and *I* bands. In our study, besides broad-band photometry, we also make use of narrow-band Hα data, which have been obtained in the WFC3 UVIS *F*657*N* filter as part of the Hα-LEGUS programme (Chandar et al., in preparation).

NGC 628 is a nearby face-on grand-design spiral galaxy of type SA(s)c[1] at a distance of 9.9 Mpc (Calzetti et al. 2015). Currently the most well-studied LEGUS galaxy, it is one of the two galaxies discussed in Krumholz et al. (2015b). Thus, the properties of star clusters in NGC 628 have already been analysed using the SLUG suite for a constant IMF, which provides us with a reference point to which we can compare our results when variations of the IMF are included. Two pointings are available for this galaxy (NGC 628 East and NGC 628 Centre), and here we focus on the East pointing (hereafter NGC 628E) as done in Krumholz et al. (2015b). Moreover, Hα data are available for the NGC 628E pointing, which allows us to explore the effects of including photometry that is sensitive to massive stars in our analysis.

We refer the readers to relevant LEGUS papers for a detailed description of the data processing techniques, and particularly on how cluster catalogues are generated (e.g. Calzetti et al. 2015; Grasha et al. 2015; Adamo et al. 2017). Here we provide only a brief description of the most relevant procedures. We make use of clusters that have been detected in at least four filters. For all clusters, both an aperture correction and correction for foreground Galactic extinction have been applied following standard procedures. All the clusters have been visually classified, according to the LEGUS definition of classes 1–3. Class 1 objects are compact and centrally concentrated clusters, class 2 objects are clusters with elongated

---

[1] The NASA/IPAC Extragalactic Database (NED) is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

surface-brightness profiles and class 3 objects loosely comprise associations or systems with multiple resolved peaks in surface brightness. We omit class 4 objects from our analysis, as they represent for most part isolated stars, background galaxies or artefacts.

With this selection, we are left with a cluster catalogue comprised of 259 clusters from NGC 628E, as in Krumholz et al. (2015b) and Adamo et al. (2017). Of these, we further select 241 clusters with a detection in the WFC3 UVIS *F657N* filter. Of the clusters with positive detections in this filter, 16 are discarded from the following analysis as visual inspection shows that their morphology is indicative of emission not arising from H II regions (Lee et al., in preparation). Altogether, our selection results in a final catalogue containing 225 clusters in the LEGUS classes 1–3.

3.2 Building a library of 2×10^8 simulated clusters

The first step towards applying our formalism to the LEGUS cluster catalogue is the generation of a library that provides the training set required for Bayesian analysis in CLUSTER_SLUG. Having enabled the variable IMF mode in SLUG, we now trivially generate an extended library of simulated star clusters, including a continuous range of IMFs with a varying upper-end slope (α_3). Our focus on the massive end of the IMF is justified by the availability of NUV and Hα photometry in LEGUS and Hα-LEGUS.

To construct this reference library with a variable IMF, which we use in all subsequent calculations unless stated otherwise, we extend the fiducial library described in Krumholz et al. (2015b)[2] to include a new dimension for the variable upper-end slope α_3. This library, which we dub pad_020_vkroupa_MW, contains 2×10^8 clusters distributed in mass as

$$p_M(\log M) \propto \begin{cases} 1, & 2 < \log M < 4 \\ 10^{-(\log M - 4)}, & 4 \leq \log M < 8 \, , \end{cases} \quad (1)$$

where the mass M is in units of M$_\odot$. As in Krumholz et al. (2015b), this choice is purely dictated by computational efficiency, as low-mass clusters are cheaper to simulate. Furthermore, with this weighting, we generate the largest number of clusters in the mass range where stochasticity is dominant so as to appropriately sample the dispersion in the population arising from random sampling. As described below, however, the choice of weighting does not affect our analysis as we can apply appropriate weighting schemes to produce arbitrary mass distributions. Similarly, we assume an age distribution defined by

$$p_T(\log T) \propto \begin{cases} 1, & 5 < \log T < 8 \\ 10^{-(\log T - 8)}, & 8 \leq \log T < \log T_{max} \, , \end{cases} \quad (2)$$

which results in the generation of more young clusters, in which stochastic effects associated with massive stars are more prominent. In these equations, the time T is in units of years, and T_{max} is the maximum age for which stellar tracks are available.

The clusters are populated with stars drawn from a variable IMF. Here, we assume a Kroupa-like IMF, defined as a broken power law comprised of three segments:

$$\alpha_i = \begin{cases} -0.3, & 0.01 \leq M < 0.08, i = 1 \\ -1.3, & 0.08 \leq M < 0.5, i = 2 \\ -3.0 \leq \alpha_3 \leq -1.5, & 0.5 \leq M < 120, i = 3 \, , \end{cases} \quad (3)$$

---

[2] The library is publicly available at www.slugsps.com and contains simulated cluster photometry in the selection of filters used in the LEGUS observations of NGC 628.

where α_i is the power-law index of the *i*th segment and the stellar masses are in units of M$_\odot$. The value of the third segment slope (α_3) is drawn from a flat distribution between -3.0 and -1.5 at the start of each realization.

Photometry is computed using Padova tracks including thermally pulsating AGB stars (Vassiliadis & Wood 1993; Girardi et al. 2000; Vázquez & Leitherer 2005), which sets the maximum age $T_{max} = 15$ Gyr. We assume a metallicity $Z = 0.02$, consistent with the solar value, and we apply a uniform extinction to each cluster assuming a Milky Way extinction curve (Cardelli, Clayton & Mathis 1989), which we normalize by drawing the visual extinction A_V from a flat distribution in the range $0 < A_V < 3$. In this work, we use the default SB99 spectral synthesis mode of SLUG, which emulates the behaviour of SB99 in choosing which atmosphere model is used for a given star. Details of this implementation can be found in appendix A2 of Krumholz et al. (2015a). While this set of atmospheres is commonly adopted in SPS modelling, we note that our results will naturally depend on this choice.

We further compute the nebular contribution to the broad-band photometry, assuming that 50 per cent of ionizing photons are converted into nebular emission (see also Section 3.6). The choice of these parameters is dictated by the properties of NGC 628 (see also Section 3.6.1). Krumholz et al. (2015b) demonstrate that the results of CLUSTER_SLUG analysis are mostly insensitive to the choice of tracks, extinction curve and metallicity, and so for this study we make use of only one choice for each of these parameters. Metallicity, choice of tracks and ionization efficiency do become important when considering clusters with ages less than 3 Myr however.

We do not include a treatment of binary stars or stellar rotation in this study. The presence of binaries (especially among massive stars) has been shown to produce appreciable effects on the integrated light of star clusters (Eldridge & Stanway 2009), resulting in a less luminous and bluer stellar population (Li & Han 2008). However, as their inclusion requires several additional parameters, binaries are currently not treated in SLUG. Rotation, which likewise affects results of SPS models (Vázquez et al. 2007; Levesque et al. 2012; Leitherer et al. 2014), can be included through the use of the latest Geneva tracks (Ekström et al. 2012) available in SLUG. However, we choose to consider the non-rotating Padova tracks to include the AGB treatment currently not present for the Geneva tracks with rotation. The effects of rotation can be as much as a 50 per cent increase in luminosity of O stars, along with a factor of 2 increase in luminosity between the UV and NIR (Vázquez et al. 2007; Levesque et al. 2012). However, the effects are only significant if all the stars have initial rotation velocities of ≈ 300 km s^{-1} (Eldridge & Stanway 2009).

Throughout this analysis, we also make use of the fiducial library of Krumholz et al. (2015b) with a constant Kroupa IMF to compare our results to the case of a non-variable IMF. A summary of the properties of these reference libraries is given in Table 1. Our library contains 2×10^8 clusters, which is a factor of 20 larger than the fiducial library of Krumholz et al. (2015b). This choice is dictated by the need to increase the number of clusters to account for the extra free parameter of the IMF slope α_3, while keeping the library size commensurate to the available computational resources. We have performed tests using smaller libraries of 10^8 clusters, finding that we have reached a satisfactory convergence.

3.3 Bayesian techniques and choice of priors

With a library of 2×10^8 clusters featuring variable IMFs in hand, we next make use of the Bayesian analysis tool CLUSTER_SLUG to

Table 1. Table of the parameters used in generating the two reference libraries of SLUG star clusters used in this work.

Name	Tracks	Base IMF	Z	Extinction[a]	ϕ[b]	$\log M$ ($M_\odot$)	$\log T$ (yr)	A_V (mag)	No. of realizations	Variable parameters
pad_020_vkroupa_MW[c]	Padova AGB	Kroupa	0.020	MW	0.5	2–8	5–10.18	0–3	2×10^8	$-3.0 \leq \alpha_3 \leq -1.5$
pad_020_kroupa_MW[d]	Padova AGB	Kroupa	0.020	MW	0.5	2–8	5–10.18	0–3	1×10^7	–

Notes. [a] MW = Milky Way extinction curve.

[b] ϕ is the fraction of ionizing photons that produce nebular emission within the aperture, combining the effects of a covering fraction less than unity and some portion of the ionizing photons being absorbed directly by dust.

[c] Variable IMF model.

[d] Constant IMF model (Krumholz et al. 2015b).

handle the computation of the posterior PDFs of the α_3 parameter. A description of the algorithms adopted in CLUSTER_SLUG has been presented in da Silva et al. (2014) and Krumholz et al. (2015a), and we refer the readers to this work for details. Briefly, the Bayesian analysis tool CLUSTER_SLUG takes a set of absolute magnitudes M_F over a given set of filters, along with the associated errors ΔM_F. By means of a large library of SLUG models, CLUSTER_SLUG exploits KDE to map the observed photometry into a PDF for the cluster physical parameters. In previous work, the only parameters under consideration were the cluster mass M, age T, extinction A_V. Here, we further extend CLUSTER_SLUG to also output PDFs for IMF parameters that are chosen to be variable, which in our case is the α_3 parameter. During this calculation, we set the KDE bandwidth to $h = 0.1$ dex for the physical variables and to $h = 0.1$ mag for the photometric variables, as in Krumholz et al. (2015b). During testing, a shift to a smaller bandwidth of $h = 0.05$ did not appreciably affect the results of the analysis.

As discussed in the previous section, the simulated clusters in our library are not weighted evenly in parameter space for reasons of computational efficiency and to better sample the range of parameters where stochasticity is most relevant. It therefore becomes necessary to 'flatten' the input library, prior to feeding it into CLUSTER_SLUG. This step is accomplished simply by inverting the weights chosen during the calculation of the library of SLUG clusters. After the library has been flattened, we further need to make a choice of the priors to be used during the Bayesian analysis. Throughout this work, we explore the effects of two different choices of prior. Our first choice is a flat prior, with $p_{prior} = 1$. Our second choice is a prior in the form of

$$p_{prior}(\boldsymbol{x}) \propto M^{-1} T^{-0.5}, \tag{4}$$

for $T > 10^{6.5}$ yr. This functional form is physically motivated by the cluster mass distribution, which is observed to follow a power law with index -2. The time dependence accounts instead for the cluster disruption (see Krumholz et al. 2015b). For young ages of $T \leq 10^{6.5}$ yr, where cluster disruption is largely ineffective, we switch to a prior $p_{prior}(\boldsymbol{x}) \propto M^{-1}$. As evident from the above equations, we elect to adopt a flat prior for both the IMF high-end slope α_3 and the extinction A_V in both cases.

3.4 Analysis of NGC 628E with broad-band photometry

Our next step after constructing a library of simulated clusters with variable IMF is to apply our CLUSTER_SLUG analysis to the LEGUS broad-band photometry of NGC 628E. With our extension of the SLUG code that now handles the case of a variable IMF, we use the CLUSTER_SLUG package to calculate the one-dimensional (1D) posterior PDFs of all the physical parameters of interest (mass, age, extinction and α_3) for the 225 clusters we have selected from NGC


Figure 1. Example of a triangle plot for cluster ID 56 in NGC 628E, computed assuming the physically motivated priors from Krumholz et al. (2015b). The 1D posterior PDFs for $\log M$, $\log T$, A_V and α_3 that are constructed marginalizing over all the other parameters are shown in the top panels of each column. The contour plots show instead the joint posterior PDFs, where the intensity reflects the probability density as indicated by the colour bar. In each panel, contours are spaced in intervals of 0.2 unit element. All the PDFs are normalized to have a unit integral.

628E. Two-dimensional (2D) PDFs are also computed, which are useful to explore correlations among parameters. At this stage, we apply both flat priors and the physically motivated priors previously presented in Krumholz et al. (2015b, see equation 4). In this section, we begin our analysis by considering the results obtained with the physically motivated priors, but in Section 3.5 we explicitly investigate the effects of the choice of priors.

Before considering the cluster population as a whole, it is instructive to inspect the corner plots for two example clusters, which we choose to bracket the range of behaviour seen in our analysis. As a first example, we show the corner plot for cluster ID 56 in Fig. 1, where the posterior PDFs are obtained using the physically motivated priors. This is the same cluster that is presented as an example in fig. 8 of Krumholz et al. (2015b), and having it as an example here allows us to make a direct comparison with previous work. However, we note that it is not part of the reduced 225 cluster catalogue as it has no observed Hα emission. The shapes of the posterior PDFs for mass, age and extinction computed for a variable IMF slope are quite similar to those recovered by Krumholz

Figure 2. As Fig. 1, but for cluster ID 591.

Figure 3. The cumulative posterior PDFs for the star clusters in our sample from NGC 628E, computed assuming both flat priors (in red, labelled as Flat) and the physically motivated priors of Krumholz et al. (2015b, in blue, labelled as K+15). Also included are the cumulative PDFs calculated using the fiducial library of Krumholz et al. (2015b) with a constant Kroupa IMF and physically motivated priors (in purple, labelled as Kroupa IMF K+15). The distributions are comparable overall. There is significant deviation between the two priors in the case of α_3, with the Krumholz et al. (2015b) priors having a significant preference towards more shallow IMF slopes.

et al. (2015b) with a fixed IMF, with the exception of a small excess of probability towards lower masses and ages. This cluster also happens to be representative as to the behaviour of a significant fraction of the clusters in our sample, which exhibit multiple peaks in the posterior PDF for mass and age. Furthermore, and most significantly, the 1D posterior PDF for α_3 is very broad, spanning the full range of possible values with a modest preference for shallow slopes. An inspection of the joint posterior PDFs further reveals that the IMF slope parameter is largely insensitive to age and extinction, and is degenerate with respect to mass.

At the opposite end of the spectrum, we show in Fig. 2 the case of a cluster (ID 591) for which mass, age and extinction are well constrained, as is clear from the sharp posterior PDFs of these quantities. Despite the fact that these physical parameters are well constrained, the posterior PDF of α_3 is still quite broad, with most of the probability being contained in the interval $\alpha \approx (-2.3, -1.5)$. As was seen in cluster 56, the corner plot reveals a modest preference of shallow IMF exponents, with $\alpha \gtrsim -2.2$.

The features highlighted in these examples appear to be quite general of the entire sample. This is shown in Fig. 3, where we display the cumulative posterior PDFs for the entire sample, which we obtain by combining the 1D posterior PDFs for each cluster. It must be noted that we do not analyse the clusters jointly for this, and that the PDFs for the clusters are combined following our analysis.

In addition to the two choices of prior, which we will discuss below, we also include in this figure the cumulative PDFs calculated using the fiducial library of Krumholz et al. (2015b), assuming a constant IMF. A comparison between the posterior PDFs for age and extinction obtained for a constant versus variable IMF (the purple and blue lines, respectively) reveals excellent agreement between the two. Besides validating our procedure, the fact that the age and extinction PDFs computed in these two cases are virtually indistinguishable implies that α_3, A_V and T have modest covariance, as suggested by the shape of the joint PDFs in Figs 1 and 2.

A similar conclusion holds for the posterior PDF of the mass, although differences between the constant and variable IMF cases become more noticeable. Indeed, there is some divergence both at very low and very high masses, particularly with the variable IMF
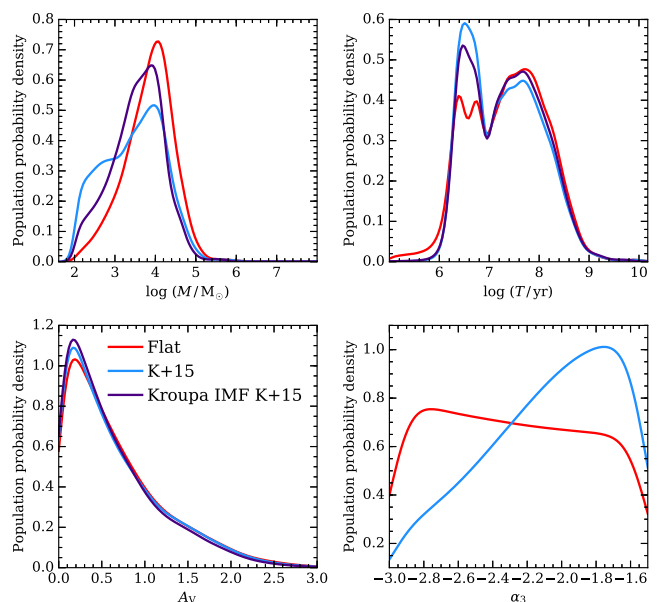
case being skewed on average towards low masses. The origin of this shift, albeit modest, can be attributed to the degeneracy between mass and IMF highlighted above. Indeed, compared to the case of a fixed IMF with $\alpha_3 = -2.3$, the analysis that includes a variable IMF seems to prefer, on average, lower cluster masses, which are populated by stars that are on average more massive. This effect is indeed visible in the cumulative posterior PDF for α_3, which is skewed towards α_3 values that are shallower than the canonical Salpeter value.

Moreover, Fig. 3 shows that, in line with what is found when examining the PDFs for individual clusters, the cumulative posterior PDF for α_3 is quite broad, implying that only weak constraints on the slope of the IMF can be obtained when comparing broad-band photometry with simulated clusters. As noted above, however, the posterior PDF appears skewed towards a shallow IMF, with $\alpha_3 \gtrsim -2.5$. This effect is modest, and not particularly statistically significant. We can nevertheless inspect the median α_3 recovered for each cluster, and study whether systematic trends can be found as a function of mass and LEGUS class. This analysis is shown in Fig. 4. When considering the physically motivated priors (squares and diamonds), a noticeable correlation with mass appears. In this figure, we have also separated clusters by their LEGUS class. We see that the majority of the class 3 objects (squares), which include clusters with multiple peaks and what are likely to be associations, are recovered at low masses (see also Adamo et al. 2017). These objects are characterized by shallow slopes of the upper end of the IMF. Conversely, class 1 and class 2 clusters (diamonds) are on average more massive, although a fit through the data still indicates a correlation between the median α_3 and the cluster mass.

The correlation between the value of α_3 and the cluster mass can be explained as follows. Due to the degeneracy between mass

(a) Physically motivated priors

(b) Age prior only

Figure 4. Panel (a) shows the correlation between the medians of the posterior PDFs for $\log M$ and α_3, with class 3 objects (light colours) and class 1–2 objects (dark colours) separated for both flat (triangles) and the Krumholz et al. (2015b) priors (quadrilaterals). Note the more apparent correlation between mass and α_3 for the physically motivated priors, with a lower limit of α_3 around the canonical Kroupa value (-2.3). Panel (b) shows the same but with the prior on the mass removed, leaving only a prior on age. This results in a distribution very similar to a purely flat prior.

and IMF, the observed light can be modelled either by a relatively massive cluster with canonical $\alpha_3 = -2.3$ or by a lower mass cluster with a more shallow IMF. The choice of a physical prior of the form $\propto M^{-1}$ skews the mass PDF towards low values (see also Section 3.5), making the latter case (lower cluster mass and hence a more shallow IMF) preferred. Therefore, the combination of imposing a physically motivated prior with a varying upper-end slope of the IMF moves the clusters in the α_3/M plane along a diagonal in the direction of lower mass and shallower slopes compared to the canonical Kroupa IMF. From a physical point of view, correlations between the cluster mass and the IMF slope are expected. For instance, a truncation of the IMF in excess of the effect induced by random sampling is predicted by the IGIMF theory, in which low-mass clusters have a bottom-heavy IMF (Kroupa 2001; Kroupa & Weidner 2003; Weidner et al. 2011; Kroupa et al. 2013). However, the effect we see in panel (a) of Fig. 4 is in the opposite direction to the effect predicted by the IGIMF. The fact that the posterior PDFs for individual clusters are very broad and that, as discussed below, the median values are quite sensitive to the choice of prior cautions against far-reaching conclusions of the nature of this correlation. Indeed, at this time, we cannot exclude the possibility that the observed trend arises from a mass-dependent correlation at second order that is not explicitly captured in our analysis.

Finally, we produce a median-centred stack of the 1D and 2D posterior PDFs for the LEGUS clusters, which is shown in Fig. 5. Degeneracies between the cluster mass and age, and between the cluster mass and IMF slope (as seen in panel (a) of Fig. 4), are apparent from the shape of the contours in their respective joint posterior PDFs. The broad shapes of the posterior PDFs in α_3 are also visible. The mass–age degeneracy is, at least in part, caused by the need for a higher mass to produce the same luminosity in older clusters.

3.5 Effects of the choice of prior

The analysis of clusters from NGC 628E using the CLUSTER_SLUG tools presented in the previous section shows that age and extinc-
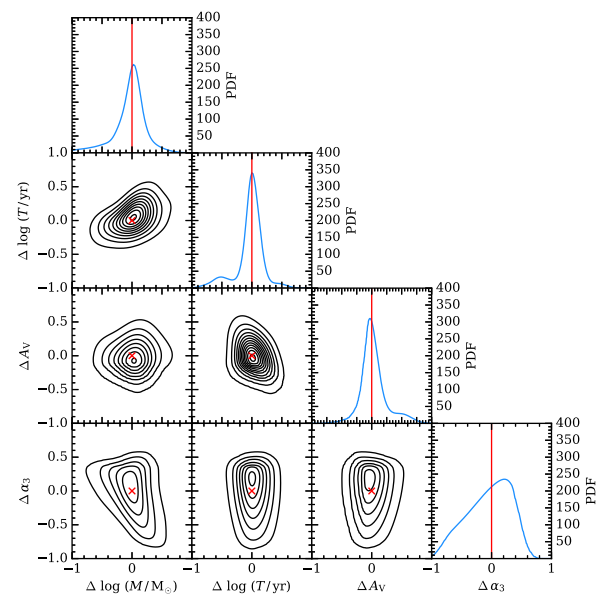

Figure 5. Stacks of the 1D and 2D PDFs (median-centred) for the clusters in NGC 628E. The 2D PDF contours are spaced in steps of 50, beginning at 50. The red crosses and lines mark the zero-points.

tion are largely insensitive to the parameters describing the IMF. Moreover, we have shown that the posterior PDF of α_3 is weakly constrained by broad-band photometry. When choosing a physically motivated prior on mass, the recovered PDFs are skewed towards lower masses and shallower IMF slopes compared to the case of a constant Kroupa IMF. As shown above, this effect arises from a degeneracy between mass and α_3, which results in diagonal shifts in the M/α_3 plane. It is therefore expected that the choice of prior on the mass will have a consequential effect on the parameters describing the IMF (and vice versa), as we show in this section.

We start by examining again the corner plot for the example cluster ID 56, this time computed assuming a flat prior on all the parameters, which is shown in Fig. 6. A comparison with Fig. 1, in which a physically motivated prior was chosen, highlights how

Figure 6. As Fig. 1, but with a flat prior on all physical parameters.

Figure 7. Comparison between the median values for each parameter returned by CLUSTER_SLUG for flat priors and the physically motivated priors by Krumholz et al. (2015b). The grey line marks the one-to-one correspondence. While extinction and age are insensitive to the choice of priors, mass and α_3 depend on the assumed prior, owing to a degeneracy between these two parameters.

a flat prior shifts the probability from a locus of low mass and young age ($M \sim 10^{3.5}$ M$_\odot$, $T \sim 10^{6.7}$ yr) to a locus of higher mass and age ($M \sim 10^5$ M$_\odot$, $T \sim 10^{7.6}$ yr). As a consequence, in our model based on stochastic sampling, the observed photometry can be realized at higher cluster masses with steeper IMFs ($\alpha_3 \lesssim -2.3$) compared to the case of low cluster mass. In other words, to produce a comparable UV luminosity to that of a higher mass cluster, a low-mass cluster needs a larger number of massive stars, something that can be achieved with a more shallow IMF slope.

By repeating the analysis for the full sample but using a flat prior, we see that the behaviour described above is indeed general. Compared to a physically motivated prior, a flat prior induces a shift in the cumulative PDFs (red lines in Fig. 3) towards higher masses and steeper IMF slopes. Conversely, there is no noticeable difference in the extinction and only a modest change occurs in the age cumulative distribution, with some probability flowing from $T \gtrsim 10^{6.5}$ yr to $T \lesssim 10^{6.5}$ yr. As a consequence of this shift in M/α_3, the correlation between the cluster mass and the IMF slope disappears when considering a flat prior (triangles in panel (a) of Fig. 4).

Panel (b) of Fig. 4 shows the same plot but with the analysis performed using only the age component of the Krumholz et al. (2015b) prior. Here we see a distribution of points that is almost identical to that of a flat prior, suggesting that the mass prior is dominant in its effect on the IMF that we recover. On these plots, we note that the fact that the medians are centred around $\alpha_3 \approx -2.3$ is not indicative that the data prefer a Kroupa IMF. Rather, this is a mere reflection of the fact that the posterior PDFs on the IMF slope are very broad and, by construction, are centred around the canonical value of $\alpha_3 \approx -2.3$. Based on this result, it may be tempting to conclude that the choice of a flat prior is preferable over the physically motivated prior of Krumholz et al. (2015b). However, while we do not wish to overinterpret the trend visible Fig. 4, the choice of our physically motivated prior is preferable given the power-law nature of the CMF. This implies that either our SLUG simulations are not fully describing the properties of clusters in NGC 628E or that indeed there is (weak) preference for more shallow IMF slopes in low-mass clusters.

Finally, rather than considering the cumulative PDFs, we compare the medians of the posterior PDFs for the physical parameters of individual clusters, as shown in Fig. 7. Consistent with the previous discussion, the medians of the age and extinction PDFs are insensitive to the choice of prior, while the medians for the mass are systematically offset in the direction of lower masses for a physically motivated prior. This shift is modest ($\sim$0.2 dex) for clusters above $\sim$10$^{3.5}$ M$_\odot$, but it becomes more significant (up to $\sim$0.5 dex) at the low-mass end. This figure also shows that shifts of a factor of 2 in mass that are modest on a logarithmic scale result in more appreciable variations of a power-law exponent on a linear scale, with $\Delta\alpha_3 \sim$ 0.2–0.3.

One possible way to avoid the IMF slope having such a dependence on the mass prior would be to analyse all the clusters as a single unit with a common mass function between them. At present, a single cluster is considered more likely to be less massive, as suggested by our choice of prior. However, for an ensemble of clusters that follow the CMF that forms the prior on the mass, we would retrieve a distribution of masses that follows this CMF. In this situation, there would not be a preference for every cluster to be less massive, rather for the ensemble of clusters to have a preference for less massive clusters.

3.6 Adding Hα photometry to the Bayesian analysis

3.6.1 Preparation of the Hα photometry

As shown in Section 3.4, our analysis of star clusters from NGC 628E based on the five broad-band LEGUS filters provides only modest constraints on the IMF slope. We also find that several clusters have broad or multi-peaked posterior PDFs in some of the physical parameters, especially age. As massive OB stars dominate the Hα emission in clusters at younger ages, the inclusion of Hα photometry is expected to improve some of the constraints on

Figure 8. As Fig. 1, but for cluster ID 292, which has observed Hα emission. Panel (a) is for the standard LEGUS broad-band filters only, whereas panel (b) includes the *F*657*N* filter to represent Hα. We see significant improvement in the recovery of the mass, age and extinction.

the physical parameters of interest. To this end, we include in our analysis data obtained as part of Hα-LEGUS (Chandar et al., in preparation) that are collected through the WFC UVIS *F*657*N* filter that, at the redshift of the LEGUS galaxies, encompasses the Hα line.

To derive aperture photometry for Hα data, we follow the method developed by Lee et al. (in preparation), which we briefly summarize here. During the processing of the *F*657*N* data, the standard approach for applying aperture corrections to broad-band photometry cannot be trivially applied to the nebular emission, given that the stellar continuum and the nebular emission differ in their spatial extent. To account for this difference and to produce photometry that can be compared to SLUG simulations, we follow a two-step procedure.

First, we synthesize an aperture-corrected continuum flux in the *F*657*N* filter, $f_{657,\,c}$. This is achieved by linearly interpolating the aperture-corrected continuum flux measured in the *F*547*M* and *F*814*W* filters, with a weight tuned to recover the observed aperture-corrected flux in the *F*657*N* filter for clusters without Hα emission. This procedure yields a synthesized *F*657*N* aperture-corrected continuum flux. Next, we repeat the above procedure, but for fluxes that are not aperture corrected, thus obtaining an interpolated continuum flux in the *F*657*N* filter that is not aperture corrected. This is needed to compute the line emission, as explained next. To derive the net Hα line flux $f_{657,\,H\alpha}$, we then subtract the synthesized continuum flux (without aperture correction) from the measured flux in the *F*657*N* filter.

Combining the two steps, we produce the aperture-corrected total flux (continuum plus line) in the *F*657*N* filter as $f_{657,\mathrm{corr}} = f_{657,H\alpha} + f_{657,c}$. Thus, at the end of this procedure, we have created a synthetic aperture correction from the weighted average of the aperture corrections on the bracketing filters and applied this to the measured *F*657*N* flux. Finally, as for the LEGUS broad-band photometry, we apply a correction for Galactic extinction. This results in a total flux value for the *F*657*N* filter that can be used in our

analysis and is comparable to the *F*657*N* photometry generated by SLUG. This total flux (hereafter referred to as Hα flux for simplicity) is the quantity that we compare directly with the output of SLUG.

The SPS module at the core of the SLUG package predicts the number of ionization photons produced by stars in each cluster, which we convert into nebular emission following the methods described in Krumholz et al. (2015a). At this stage, we assume that only a fraction ($\phi = 0.5$) of the ionizing photons are converted into nebular emission. After convolving the simulated cluster spectral energy distribution (SED) with the *F*657*N* filter transmission curve, SLUG generates a total flux like the $f_{657,\,\mathrm{corr}}$ we synthesize for the observations as described above. Throughout this work, we do not include Hα emission from stars themselves (e.g. from Be stars), as they are found to contribute only a few per cent compared to the nebular emission.

We note that, due to the extended nature of the Hα emission, there is a possibility that we are missing some fraction of the flux in the *F*657*N* filter from extended (but typically low surface brightness) emission. As we will show below, however, we are able to reproduce all the findings of our analysis using a set of mock clusters that, by definition, are not affected by loss of flux on large scales. This implies that any potential loss of flux has a minor effect on our results. Moreover, the choice of a nebular fraction of $\phi = 0.5$ already mitigates this effect. We have also conducted tests to verify that our results are not sensitive to the exact choice of this coefficient. Indeed, we find that repeating our analysis using a library with $\phi = 0.7$ does not have a noticeable effect on the medians of the PDFs we return for the physical parameters.

3.6.2 Results including Hα photometry

To illustrate the effect of including Hα in our analysis, we turn to an example cluster with observed Hα emission (ID 292), which is shown in Fig. 8. Looking at the differences between panels (a) and (b), we see a significant improvement in the age determination,

Figure 9. The cumulative posterior PDFs for a subset of young clusters (with median age $T \leq 10^{6.5}$ yr) computed for both choices of priors and for a constant Kroupa IMF. The left-hand panels show the PDFs based on broad-band photometry only, while the right-hand panels show the PDFs computed including the Hα flux. The inclusion of Hα in the analysis results in much sharper PDFs in age, which also lifts some of the degeneracy in extinction. There is also a small shift towards steeper IMF slopes for both priors.

with the locus at $T \sim 10^{6.5}$ yr being clearly preferred over the older ages. In this particular case, the inclusion of Hα further improves the constraints that can be derived on the mass and extinction, removing the bimodality seen in both. In turn, as the lower mass at $M \sim 10^3$ M$_\odot$ is now preferred, the posterior PDF for α_3 sharpens compared to the case in which no Hα was included.

As we will show in the following section, the improvement in age determination obtained when including Hα in the analysis is common to much of the cluster sample, and is most noticeable in the case of younger clusters. However, the variation in the median of the age PDFs is modest, and we see no appreciable difference in the shape of the cumulative PDFs when considering our entire cluster catalogue. Thus, given that the shapes of the cumulative PDFs for both mass and age are primarily dependent on the scatter among clusters rather than the width of the PDFs for the individual clusters, the shapes of the PDFs shown in Fig. 3 do not show significant change following the inclusion of Hα. Similarly, we see no appreciable difference in the cumulative PDF for α_3. This is primarily caused by the broad shape of the α_3 PDFs in the individual clusters, even following the inclusion of Hα.

Finally, we examine the effect of including Hα in the analysis of young star clusters by limiting our choice of clusters to those whose posterior PDF in age has a median of $T \leq 10^{6.5}$ yr. It is for this subset of clusters that the NUV and Hα photometry of the LEGUS and Hα-LEGUS surveys is expected to yield the tightest constraints. The cumulative PDFs for this subset of clusters, computed with both choices of priors, a fixed Kroupa IMF and with or without Hα, are shown in Fig. 9. This subset contains 21 clusters for the case of flat priors with no Hα, 24 with Hα, 21 for physical priors with no Hα, 18 with Hα and 20 for the case of the fixed Kroupa IMF with no Hα with 21 when Hα is included.

Several interesting features can be seen in this figure. Considering first the PDFs derived assuming the physically motivated priors (blue curves), the inclusion of Hα significantly improves the age constraint (compare the left- and right-hand panels), thus sharpening

the age posterior PDFs. However, we see no appreciable shift in the positions of the medians of the PDFs. A more subtle difference is visible when comparing the mass and α_3 PDFs. Conversely, more substantial differences appear when examining the posterior PDFs for the extinction, with the inclusion of Hα favouring extinctions up to ~0.5 dex higher. Comparisons with the PDFs computed with a constant Kroupa IMF (purple line) reveal that the differences observed when including Hα are independent of the choice of the IMF.

Moreover, in line with the discussion presented in the previous section, Fig. 9 shows how critical the mass–α_3 degeneracy is in this type of analysis. At fixed age, a physically motivated prior skews the mass determination towards smaller values, which is compensated by a more shallow slope on the upper end of the IMF. Conversely, a flat prior (particularly when including Hα) prefers steeper IMF slopes that result in higher masses. In between, the constant Kroupa IMF case with $\alpha_3 = -2.3$ places the preferred masses at intermediate values between those inferred with flat and physically motivated priors. In summary, a key result of this work is that, due to the mass–IMF degeneracy, the choice of prior on the mass strongly affects the resultant shape of the posterior PDF of α_3, preventing robust constraints on the upper-end slope of the IMF even when including Hα. Turning this argument around, the assumption of a fixed Kroupa IMF, which is equivalent to imposing a strong prior on α_3, affects the mass determination by up to ~0.5 dex for young and low-mass clusters.

4 ANALYSING THE PERFORMANCE OF THE TECHNIQUE USING MOCK CLUSTERS

Having applied our Bayesian formalism to the study of star clusters in NGC 628E, we have found that a degeneracy between mass and the slope of the upper end of the IMF prevents us from reliably constraining the value of α_3 both when using broad-band photometry alone and when Hα is included. We have also shown

Table 2. Table of the parameters used in SLUG for generating the two grids of mock star clusters used in this work. Within the parentheses, we list the adopted step size.

Name	Tracks	Base IMF	Z	Extinction[a]	ϕ[b]	$\log M$ ($M_\odot$)	$\log T$ (yr)	A_V (mag)	#[c]	α_3
MocksRun1[d]	Padova AGB	Kroupa	0.020	MW	0.5	2–6 (1)	6–9 (0.5)	0.1,0.5,1,1.5,2	3	$-2.8, -1.6(0.2)$
MocksRun2[e]	Padova AGB	Kroupa	0.020	MW	0.5	2–6 (1)	6–9 (0.5)	0.1,0.5,1,1.5,2	3	-2.3

Notes. [a]MW = Milky Way extinction curve.
[b]ϕ is the fraction of ionizing photons that produce nebular emission within the aperture, combining the effects of a covering fraction less than 1 and some portion of the ionizing photons being absorbed directly by dust.
[c]Number of realisations.
[d]Variable IMF mocks.
[e]Constant IMF mocks.

how the choice of prior on the mass has a knock-on effect on the shape of the posterior PDF for α_3. Equivalently, the choice of a constant IMF may lead to systematic errors in the mass determination of more than 0.5 dex (depending on age).

In this section, we exploit mock clusters to take a more detailed look at our ability to recover the posterior PDFs for physical parameters of clusters using our Bayesian formalism. We also characterize the significance of the mass–α_3 degeneracy, providing forecasts on the ability to constrain the IMF in future experiments when the cluster mass can be determined via techniques that are independent of photometry.

To achieve these goals, we generate a grid of 700 mock star clusters using SLUG, with masses in the range $\log(M/M_\odot) = 2\text{–}6$ in steps of 1 dex, ages in the range $\log(T/\text{yr}) = 6\text{–}9$ in steps of 0.5 dex and five choices of extinction, $A_V = 0.1, 0.5, 1.0, 1.5, 2.0$ mag. For the IMF upper-end slope (α_3), we select values ranging from -2.8 to -1.6 in steps of 0.2. These parameters, together with the other SLUG parameters that define the SPS calculation, are listed in Table 2.

We also generate a second grid of 175 clusters with a constant Kroupa IMF ($\alpha_3 = -2.3$) spanning the same ranges of mass, age and extinction. The parameters for this grid are also listed in Table 2. Due to the stochastic nature of the SLUG code, clusters simulated with a given choice of parameters do not necessarily have identical properties. For this reason, we generate three realizations of each cluster across both of our grids to help prevent the inclusion of rare outliers in our analysis. Using the variable library generated for the analysis of NGC 628E in Section 3.2, we apply CLUSTER_SLUG to our grid of mock clusters with a variable IMF. Likewise, we apply CLUSTER_SLUG to the fixed IMF grid using the fiducial library of Krumholz et al. (2015b) as a training set. We use flat priors during this analysis as the clusters are drawn from a flat distribution of physical parameters, without following, for example, a CMF.

At first, we examine the precision with which CLUSTER_SLUG can recover the physical parameters of the clusters. This is achieved by studying the intrinsic width of the posterior PDFs recovered by our Bayesian analysis. A sharp PDF not only means that physical parameters can be recovered to high precision, but also ensures that the global properties of a sample (e.g. via the cumulative PDFs) reflect the intrinsic scatter of the population. Conversely, very broad PDFs reflect not only a low precision in recovering parameters, but also the inability to characterize the underlying scatter of physical parameters within the cluster population. In practice, we stack both the 1D and 2D posterior PDFs returned by CLUSTER_SLUG for all 700 mock clusters with a variable IMF. The PDFs are median-centred before stacking, and to select a representative group of objects, we choose for each cluster the realization corresponding to the second

Figure 10. Stacks of the 1D and 2D PDFs (median-centred) for 700 mock clusters simulated with a variable IMF and the choice of parameters listed in Table 2. 2D PDF contours are spaced in steps of 100, beginning at 100. The red crosses and lines mark the zero-points.

best 'fit' in mass (i.e. second smallest residual). These stacks are shown in Figs 10 and 11.

Inspecting the 1D posterior PDFs recovered from broad-band photometry alone (Fig. 10), we see that the PDFs for the mass, age and extinction have a well-defined peak. Conversely, the posterior PDF for α_3 is broad and flat, confirming that broad-band photometry alone has weak constraining power on the upper end of the IMF. Additional insight is offered by the analysis of the joint PDFs. The well-known degeneracy between age and extinction is recovered by our analysis, as is the degeneracy between mass and α_3 discussed in the previous sections. Further, the contours on the joint PDFs of age and extinction with α_3 indicate that T and A_V are largely insensitive to variations in the upper-end slope of the IMF.

Restricting our analysis to a subset of mock clusters with young ages ($T \leq 10^{6.5}$ yr; Fig. 11), we see that the performance of our Bayesian technique improves, with the posterior PDFs of all the parameters becoming narrower. Qualitatively, however, the majority of the trends we see in Fig. 10 remain. By comparing the posterior PDFs computed with and without the inclusion of Hα (right- and left-hand panels, respectively), we see that Hα is critical in improving the age determination, breaking the age–extinction degeneracy. The broadness in the age PDF seen in the left-hand panel of Fig. 11
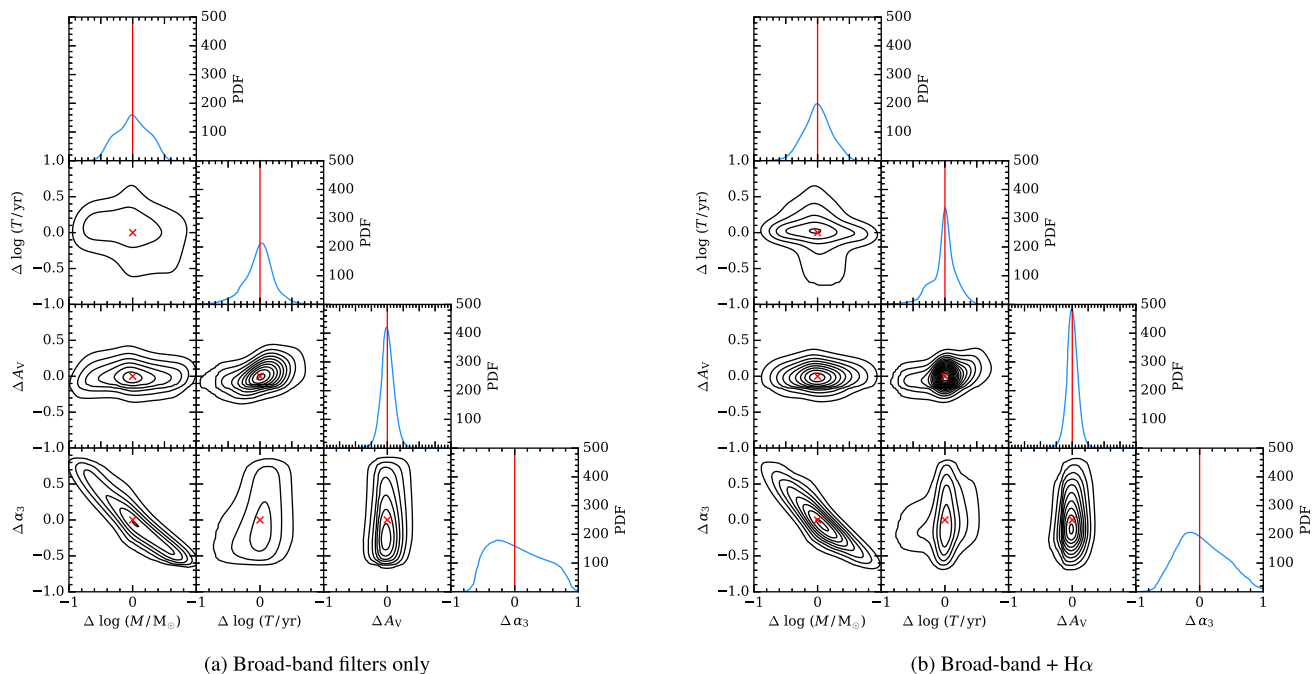
Figure 11. As Fig. 10, but for a subset of young mock clusters, with age $T \leq 10^{6.5}$ yr. The left- and right-hand panels show, respectively, the PDFs recovered from broad-band photometry alone, and with the inclusion of Hα. 2D PDF contours are spaced in steps of 50, beginning at 50. The red crosses and lines mark the zero-points.

is reduced greatly following the inclusion of Hα, which is seen in the narrowing of the age PDF in the panel on the right. In turn, a better constraint in age (and extinction) is reflected in a narrower mass PDF, although it remains broad due to the mass–IMF degeneracy. Indeed, with refined PDFs on age and extinction, an evident 'banana' shape emerges for the joint PDF for M–α_3, which clearly highlights the degeneracy between mass and IMF that has shaped much of the discussion in this work. By means of our Bayesian analysis, we can also accurately quantify this degeneracy (see below), concluding that a variation in $\Delta\alpha_3 \sim \pm 0.5$ maps into a mass variation of $\Delta \log(M/M_\odot) \sim \pm 1$ dex.

Next, we consider how accurate our Bayesian procedure is in recovering the underlying physical parameters for each cluster. This is achieved by comparing the medians of the inferred posterior PDFs to the true value used when computing the mocks. As was done for the stacks, in an effort to apply our analysis to a representative sample of clusters, we consider for each set of parameters the realization that yields the second best 'fit' to the mass. These results are presented in Fig. 12. Looking at all four physical parameters, we see that the medians of the inferred PDFs are a trustworthy indicator for the true value of the mass, age and extinction when considering the average population. For the case of constant IMF (purple diamonds), the scatter between individual clusters is modest, of the order of $\sim$0.2–0.3 dex (mag). Conversely, the addition of a variable IMF parameter greatly increases the scatter of residuals on a cluster by cluster basis when compared to the grid of Kroupa IMF clusters. This increase in scatter is especially noticeable in mass, and originates from the M–α_3 degeneracy described above. As the mass of the cluster and its IMF slope are highly dependent on each other, the wide scatter in mass is further reflected in the inaccurate values inferred for α_3. Indeed, given the flat and broad posterior PDFs for the IMF slope, the medians of α_3 have no physical meaning and simply reflect the median of the prior on α_3 [in this case, a flat distribution between $\alpha_3 = (-3.0, -1.5)$].

Given this result, it is clear that obtaining refined constraints on the upper-end slope of the IMF requires a different approach. It is readily apparent that the ability to apply a prior to the cluster mass that is independent of the broad-band photometry may prove very useful in analysis such as this. Independent constraints on the masses of clusters that may then be used as priors may come, for instance, from dynamical measurements (e.g. Ho & Filippenko 1996; Gerhard 2000). The era of 30 m telescopes will make such measurements more feasible.

To illustrate this point in more quantitative terms, we present in Fig. 13 a study of the accuracy with which α_3 can be recovered as a function of the width of a prior on the mass. Results of three choices of IMF are presented, when leaving the mass unconstrained, and when applying a top-hat prior with width 0.5, 0.2 and 0.05 dex about the target cluster mass, for both broadband filters only and broad-band filters combined with Hα. This analysis shows that, as expected, the medians of the PDFs on α_3 converge to the true value as the precision with which the mass is constrained increases. Due to the difference between logarithmic intervals in mass and the linear scale in the IMF upper-end exponent, there is a weak scaling of α_3 with the prior width, with the underlying values being recovered to within $\Delta\alpha_3 \sim \pm 0.2$ when the mass is constrained to better than 10 per cent. However, more realistic constraints on the mass (to within a factor of 2), for instance via dynamical measurements, appear to be useful to rule out extreme variations in the IMF upper-end slope through multiwavelength broad-band photometry, given the assumptions of our model. The inclusion of Hα results in a wider scatter towards zero for the unconstrained case. This advantage becomes much less evident as the mass is constrained however. Finally, in Fig. 14, the mode rather than the median of the α_3 posterior PDF is used to represent the recovered α_3. Here we see a much faster convergence than we see in Fig. 13. The scatter also reduces rapidly as we tighten the constraints on the mass, although this improvement is

Figure 12. Comparisons of the medians of the inferred posterior PDFs for all physical parameters to the true values adopted when simulating the mock clusters. Results for individual clusters are shown with crosses, while the median residuals and the corresponding 10th and 90th percentiles are shown with solid points. We plot analysis including broad-band photometry alone with (red, circles) and without (purple, diamonds) a variable IMF. We also show results of the variable IMF case when including Hα (blue, squares). Points are shifted along the *x*-axis for visualization purposes, with the leftmost point situated at the true *x*-axis value.

negligible in the $\alpha_3 = -1.6$ case, where the scatter remains very large.

5 SUMMARY AND CONCLUSIONS

In this paper, we have developed a Bayesian formalism to study the properties of the IMF in star clusters with multiwavelength broad-band photometry combined with SLUG stochastic SPS models. To investigate variation in the IMF, we first extended the capabilities of SLUG, implementing a flexible algorithm by which the PDFs defining physical functions such as the IMF can be set to variable mode. With this variable mode, it is now possible to perform simulations of clusters with an IMF for which specific parameters can be varied continuously across a desired range.

Exploiting this variable mode, which is now part of the main distribution of SLUG, we constructed a library of 20 million star clusters, simulated with a Kroupa-like IMF for which the slope of the high-mass end (α_3) was randomly chosen for each cluster, evenly selecting between -3.0 and -1.5. With this library, we then exploited the CLUSTER_SLUG Bayesian analysis framework described in Krumholz et al. (2015b) to compute the posterior PDF of α_3 for clusters with observed multiwavelength broad-band photometry.

As a proof of concept, we applied this formalism to a catalogue of 225 star clusters from the local galaxy NGC 628, for

which broad-band and Hα photometry is available from the LEGUS (Calzetti et al. 2015) and Hα-LEGUS (Chandar et al., in preparation) programmes. After assuming a physically motivated prior for the mass and age of star clusters, we found that our formalism recovers the PDFs of the core physical parameters of mass, age and extinction for the entire sample, consistent with those resulting from the analysis of the same clusters with a constant IMF (Krumholz et al. 2015b). Conversely, we found very broad and flat PDFs for α_3, indicating that broad-band photometry alone is unable to tightly constrain the slope of the upper end of the IMF. This result is primarily driven by a noticeable degeneracy between mass and α_3.

Due to the interlink between mass and IMF slope, we found that the posterior PDF for α_3 is quite sensitive to the choice of priors. Compared to the assumption of a flat prior in mass, a physically motivated choice derived from the power-law nature of the CMF leads to posterior PDFs for the mass that are systematically skewed towards lower cluster masses and with a preference for more shallow IMF slopes, in excess of the canonical Kroupa value of $\alpha_3 = -2.3$. However, the fact that the posterior PDF is very broad and that it is sensitive to the choice of prior cautions against far-reaching conclusions on the nature of the IMF with present data. Conversely, applying a prior to the cluster ages alone has a negligible effect on our analysis.

Figure 13. Comparison between the IMF slopes α_3 measured from the inferred posterior PDF (taking the median) and the true value as a function of width of a prior on the mass. Results from three IMF slopes are shown. The filled regions represent the 10th and 90th percentile range. We omit clusters with a target mass of $100\,M_\odot$ due to the relative variation in the resultant generated masses. Panel (a) was produced using only the five broad-band filters, whereas panel (b) was produced using the additional Hα filter. We see some improvement for the case with no constraints, with marginal improvement for the tighter constraints.

Adding Hα fluxes to our analysis does not significantly narrow the posterior PDF for α_3, although it reduces the degeneracy between age and extinction. As a consequence, the posterior PDFs for age and extinction, and to some extent mass, sharpen. This effect is particularly noticeable for the subsets of young star clusters, with ages $T \leq 10^{6.5}$ yr.

We also quantified the accuracy and precision with which our procedure is able to recover the input physical parameters for a grid of mock clusters, simulated across a wide range of physical parameters. We find that the posterior PDFs for mass, age and extinction have, on average, well-defined medians. The widths of the peaks in the PDFs are particularly narrow when considering a subset of young clusters and when including Hα in the analysis. We find that we accurately recover the input parameters of our mock clusters, on average, with the medians of the mass, age and extinction agreeing well with the input parameters. However, we still experience difficulty in recovering our input α_3, finding a very broad posterior PDF across many mock clusters.

Finally, we have presented future avenues for improvement in our ability to constrain the slope of the upper end of the IMF via broad-band photometry. One possible path, rather than performing the analysis of the photometry on a cluster-by-cluster basis as we do in this study, is to perform a joint analysis of the star clusters with an underlying mass function that is common to all. This ensemble method could lessen the impact of the mass prior on our results. Furthermore, the ability to break the degeneracy between mass and α_3 via priors on the mass that are independent of the photometry (such as dynamical cluster mass measurements) is expected to provide us with the best way forward in our attempts to constrain the IMF slope by photometry. Indeed, we have explicitly shown that,

while mass constraints of 10 per cent are required to reliably recover α_3, priors that restrict the allowed mass range to within a factor of 2 are sufficient to rule out extreme variations in the allowed range of α_3.

In conclusion, while our ability to recover the slope of the upper end of the IMF via broad-band photometry is limited at present time, there are good prospects to further develop the formalism we have presented here to obtain improved constraints on α_3 with current data, or via future spectroscopic observations in the era of 30 m telescopes.

ACKNOWLEDGEMENTS

Figure 14. Comparison between the IMF slopes α_3 measured from the inferred posterior PDF (taking the mode rather than the median) and the true value as a function of width of a prior on the mass. Results from three IMF slopes are shown. The filled regions represent the 10th and 90th percentile range. We omit clusters with a target mass of $100\,\mathrm{M}_\odot$ due to the relative variation in the resultant generated masses.

REFERENCES

Adamo A. et al., 2017, ApJ, preprint (arXiv:1705.01588)
Adams F. C., Fatuzzo M., 1996, ApJ, 464, 256
Andrews J. E. et al., 2013, ApJ, 767, 51
Andrews J. E. et al., 2014, ApJ, 793, 4
Bastian N., Covey K. R., Meyer M. R., 2010, ARA&A, 48, 339
Bell E. F., McIntosh D. H., Katz N., Weinberg M. D., 2003, ApJS, 149, 289
Bonnell I. A., Clarke C. J., Bate M. R., 2006, MNRAS, 368, 1296
Calzetti D., Chandar R., Lee J. C., Elmegreen B. G., Kennicutt R. C., Whitmore B., 2010, ApJ, 719, L158
Calzetti D. et al., 2015, AJ, 149, 51
Cappellari M. et al., 2012, Nature, 484, 485
Cardelli J. A., Clayton G. C., Mathis J. S., 1989, ApJ, 345, 245
Cerviño M., Valls-Gabaud D., Luridiana V., Mas-Hesse J. M., 2002, A&A, 381, 51
Chabrier G., 2003, PASP, 115, 763
Conroy C., van Dokkum P. G., 2012, ApJ, 760, 71
Corbelli E., Verley S., Elmegreen B. G., Giovanardi C., 2009, A&A, 495, 479
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 Wit W. J., Testi L., Palla F., Vanzi L., Zinnecker H., 2004, A&A, 425, 937
de Wit W. J., Testi L., Palla F., Zinnecker H., 2005, A&A, 437, 247
Dries M., Trager S. C., Koopmans L. V. E., 2016, MNRAS, 463, 886
Ekström S. et al., 2012, A&A, 537, A146
Eldridge J. J., Stanway E. R., 2009, MNRAS, 400, 1019
Elmegreen B. G., 2002, ApJ, 577, 206
Fumagalli M., da Silva R. L., Krumholz M. R., 2011, ApJ, 741, L26
Gerhard O., 2000, in Lançon A., Boily C. M., eds, ASP Conf. Ser. Vol. 211, Massive Stellar Clusters. Astronomical Society of the Pacific., San Francisco, p. 12
Girardi L., Bressan A., Bertelli G., Chiosi C., 2000, A&AS, 141, 371
Grasha K. et al., 2015, ApJ, 815, 93
Gunawardhana M. L. P. et al., 2011, MNRAS, 415, 1647
Gvaramadze V. V., Weidner C., Kroupa P., Pflamm-Altenburg J., 2012, MNRAS, 424, 3037
Haas M. R., Anders P., 2010, A&A, 512, A79
Ho L. C., Filippenko A. V., 1996, ApJ, 466, L83
Kroupa P., 2001, MNRAS, 322, 231
Kroupa P., Weidner C., 2003, ApJ, 598, 1076
Kroupa P., Weidner C., Pflamm-Altenburg J., Thies I., Dabringhausen J., Marks M., Maschberger T., 2013, Planets, Stars and Stellar Systems. Springer-Verlag, Berlin, p. 115
Krumholz M. R., 2011, in Treyer M., Wyder T., Neill J., Seibert M., Lee J., eds, ASP Conf. Ser. Vol. 440, UP2010: Have Observations Revealed a Variable Upper End of the Initial Mass Function?. Astron. Soc. Pac., San Francisco, p. 91
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
Lamb J. B., Oey M. S., Werk J. K., Ingleby L. D., 2010, ApJ, 725, 1886
Lamb J. B., Oey M. S., Segura-Cox D. M., Graus A. S., Kiminki D. C., Golden-Marx J. B., Parker J. W., 2016, ApJ, 817, 113
Lee J. C. et al., 2009, ApJ, 706, 599
Lee J. C., Veilleux S., McDonald M., Hilbert B., 2016, ApJ, 817, 177
Leitherer C. et al., 1999, ApJS, 123, 3
Leitherer C., Ekström S., Meynet G., Schaerer D., Agienko K. B., Levesque E. M., 2014, ApJS, 212, 14
Levesque E. M., Leitherer C., Ekstrom S., Meynet G., Schaerer D., 2012, ApJ, 751, 67
Li Z., Han Z., 2008, ApJ, 685, 225
Lyubenova M. et al., 2016, MNRAS, 463, 3220
Massey P., Johnson K. E., Degioia-Eastwood K., 1995, ApJ, 454, 151
Matteucci F., 1994, A&A, 288, 57
Meurer G. R. et al., 2009, ApJ, 695, 765
Oey M. S., Lamb J. B., Kushner C. T., Pellegrini E. W., Graus A. S., 2013, ApJ, 768, 66
Popescu B., Hanson M. M., 2009, AJ, 138, 1724
Salpeter E. E., 1955, ApJ, 121, 161
Smith R. J., 2014, MNRAS, 443, L69
Stephens I. W. et al., 2017, ApJ, 834, 94
Thomas D., 1999, MNRAS, 306, 655
van Dokkum P. G., Conroy C., 2010, Nature, 468, 940
Vassiliadis E., Wood P. R., 1993, ApJ, 413, 641
Vázquez G. A., Leitherer C., 2005, ApJ, 621, 695
Vázquez G. A., Leitherer C., Schaerer D., Meynet G., Maeder A., 2007, ApJ, 663, 995
Weidner C., Pflamm-Altenburg J., Kroupa P., 2011, in Treyer M., Wyder T., Neill J., Seibert M., Lee J., eds, ASP Conf. Ser. Vol. 440, UP2010: Have Observations Revealed a Variable Upper End of the Initial Mass Function?. Astron. Soc. Pac., San Francisco, p. 19
Weisz D. R. et al., 2015, ApJ, 806, 198

APPENDIX A: IMPLEMENTING A CONTINUOUSLY VARIABLE IMF IN SLUG

As described in Section 2, SLUG is a stochastic SPS code that we make use of in our study of IMF variations. This requires a large library of simulated clusters with a wide variety of IMF shapes.

While in its previous version SLUG already handled arbitrary IMF shapes, the parameters specifying the functional form of the IMF were fixed by an input parameter file, which prevented the user from running large simulations with a continuous distribution of clusters with respect to the parameters defining the shape of the IMF. As an example, consider the Kroupa IMF constructed with three power-law segments, having slopes $\alpha_1 = -0.3$, $\alpha_2 = -1.3$ and $\alpha_3 = -2.3$. If we wished to investigate variation in α_3, it would have been necessary to run many libraries of SLUG simulations, each configured with a different input value of α_3. This is clearly not desirable, especially when dealing with simulations of millions of clusters. To circumvent this limitation, we develop a new feature of SLUG that allows users to vary the parameters that control the IMF continuously, in a similar way to what is done with the other physical parameters (e.g. mass, age and extinction).

To this end, we implement an extension of SLUG's PDF capabilities, introducing nested PDFs. In SLUG, PDFs are generated and handled in a very general way (da Silva et al. 2012, 2014; Krumholz et al. 2015a). Each PDF is described by an arbitrary number of segments, with each segment having a functional form of choice, such as a lognormal distribution or a power-law distribution (a full list of functional forms is given by Krumholz et al. 2015a). In the code, these PDFs are used to represent the IMF, the CMF, star formation history, distribution of extinctions (A_V) and the cluster lifetime function.

Our extension is implemented by enabling a hierarchy of PDFs, by which a 'master' PDF can control, at run time, the behaviour of 'slave' PDFs. With this new feature, users can define the functional form of the IMF (the 'slave' PDF) by means of an arbitrary number of segments as before. However, one or more parameters defining these segments can now be set to a variable mode. The user further specifies the form of one 'master' PDF for each variable parameter, from which the numerical values for the variable parameter are drawn during the simulation. Specifically, at run time, SLUG draws new realizations for each of the variable parameters from the 'master' PDFs and reinitializes the 'slave' PDFs (the IMF in our case) for each cluster accordingly. The end result of this implementation when applied to the IMF case is that we can easily construct libraries of simulations with continuously varying IMFs, using only an input parameter file.

Operationally, the use of this variable feature merely requires users to set the correct flag when defining the PDF segments in the SLUG IMF definition file, and then to create a companion file to define the PDF from which the parameter is drawn. Detailed instructions for the use of this new feature are available in the latest release of the SLUG user manual,[3] along with examples. The source code featuring this extension is publicly available from the SLUG repository, and it is now part of the main SLUG release.[4]

Here we provide examples of the capabilities of this variable mode, together with tests to validate the new version of the code. As a first example, we consider the simple case of a Salpeter-like IMF (Salpeter 1955), which is defined by a single power-law


Figure A1. Values of the power-law index α of a variable Salpeter-like IMF (blue histogram), obtained from 10^4 cluster realizations generated using the newly developed variable mode in SLUG. The requested normal distribution for α has a mean of -2.35 and a dispersion of 0.2, and is plotted in red. The best-fitting normal distribution is displayed in black, and closely agrees with the input PDF.


Figure A2. Comparison of the maximum stellar mass for two sets of realizations of a $500\,M_\odot$ cluster at 10^6 yr. In the first set (blue histogram), the stars are drawn from a canonical Kroupa IMF in 5000 realizations. In the second set, we generate 10^5 realizations using the variable IMF mode, with the upper-end slope α_3 drawn from a flat distribution between -3.0 and -1.5. The subset of realizations with a value of α_3 within ± 0.3 of the Kroupa value is shown in red, while the full range is shown in grey. The histograms are rescaled such that there are an equal number of clusters in the bins for the Kroupa and constrained variable cases, such that their size is relative to the complete variable set.

---

[3] Available at http://slug2.readthedocs.io
[4] Available at http://www.slugsps.com

segment with slope α. In this run, we set α to variable mode, further providing a Gaussian PDF with mean $\alpha = -2.35$ and dispersion of 0.2 as the 'master' PDF. The values of α drawn in 10^4 cluster realizations of a single SLUG run are shown to follow closely the requested normal distribution in Fig. A1.

Next, we demonstrate the use of the variable PDFs for the case of an IMF with multiple segments. Due to our flexible implementation, we can choose to vary only one of the parameters of this IMF. In this demonstration, we vary only the high-mass slope (α_3) of a Kroupa-like IMF, with its slope generated by evenly selecting between -3 and -1.5 for each cluster realization. We generate 10^5 realizations of a $500\,\mathrm{M_\odot}$ star cluster, and we plot histograms of the most massive star in each simulated cluster in Fig. A2. Given enough realizations, we can extract from the resulting library of SLUG models a subset of simulations with $\alpha_3 \approx -2.3$, which we then compare to simulations that have been run in the non-variable

mode using a canonical Kroupa IMF. We find excellent agreement when comparing the shapes of the two distributions, which validates our new implementation. Fig. A2 also shows that, as expected, the inclusion of shallow IMF slopes in the library, with α_3 up to values of -1.5, increases the probability of drawing massive stars compared to a normal Kroupa IMF, skewing the maximum stellar mass distribution to high values.

In summary, our extension to the way SLUG handles PDFs enables users to create large libraries of clusters simulated using a continuous distribution of parameters defining the functional form of the IMF. These libraries can then be used jointly with Bayesian analysis to derive posterior PDFs for the IMF parameters via comparisons with broad-band photometry.

This paper has been typeset from a TeX/LaTeX file prepared by the author.