

Theoretical predictions for IMF diagnostics in UV spectroscopy of star clusters

G. Ashworth,^{1,2} Michele Fumagalli,^{1,2★} Angela Adamo³ and Mark R. Krumholz⁴

¹*Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK*

²*Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, UK*

³*Department of Astronomy, The Oskar Klein Centre, AlbaNova University Centre, Stockholm University, SE-106 91 Stockholm, Sweden*

⁴*Research School of Astronomy & Astrophysics, Australian National University, Canberra ACT 2611, Australia*

Accepted 2018 July 29. Received 2018 July 29; in original form 2018 April 5

ABSTRACT

We explore the possibility of using UV spectroscopy in combination with broad-band photometry as diagnostic tools for understanding the shape of the initial mass function (IMF) in unresolved stellar populations. Building on our previous work, we extend the Stochastically Lighting Up Galaxies (SLUG) code to include a high-resolution UV spectral synthesizer and equivalent width calculation capabilities. We first gain a qualitative understanding of how UV spectral features behave as the parameters that define a star cluster in SLUG (mass, age, extinction, and IMF slope α_3) are changed. We then exploit Bayesian inference techniques to recover the α_3 values for clusters simulated with SLUG, using mock observations of these clusters comprised of broad-band photometry and equivalent width measurements of a selection of UV spectral features. We find some improvement when compared to attempts using broad-band photometry alone (with the interquartile range of the α_3 posterior PDF shrinking by $\simeq 32\%$), although we still do not yet fully break the known degeneracy between the cluster mass and α_3 . Finally, we make predictions about how effective real observations will be by quantifying our ability to constrain α_3 as a function of limiting equivalent width. We find that observations sensitive to a modest equivalent width of $\simeq 9 \text{ \AA}$ are sufficient to improve the recovery of the IMF slope parameter by $\simeq 32\%$ (interquartile range of posterior PDF median residuals), moving to $\simeq 39\%$ when we include all the significant spectral features in the wavelength range $900 - 3000 \text{ \AA}$.

Key words: methods: statistical – stars: luminosity function, mass function – galaxies: clusters: general.

1 INTRODUCTION

Young, bright, massive stars are responsible for much of the ultra-violet (UV) flux of a stellar population, and it is from these stars that many distinctive features in the UV region of the spectrum originate. The number of these stars that exist in a cluster is dependent on the upper end slope of the initial mass function (IMF) (e.g. Salpeter 1955). The IMF describes the distribution of stars by mass in a stellar population, and it is a core ingredient in our understanding of many astrophysical phenomena. As these massive stars contribute strongly to the UV region of the spectrum, this region holds many clues which can aid investigations into the shape of the IMF. Accurately modelling the UV region of a star cluster's spectrum is therefore important if we are to improve our understanding of the IMF from observations of unresolved stellar

populations. It is such a task that we undertake over the course of this paper.

Describing the shape of the IMF is an important goal as it is the IMF that defines the make-up of stellar populations. Its shape is generally considered to be universal (Bastian, Covey & Meyer 2010), with most observations suggesting that, at its simplest, it takes the form of a power law of slope ~ -2.3 (Salpeter 1955), assuming a more shallow slope at low stellar masses (Massey, Johnson & Degioia-Eastwood 1995; Kroupa 2001; Chabrier 2003; Bastian, Covey & Meyer 2010).

However, the concept of a universal IMF is not without challenge. There exist many studies which imply variation in the shape of the IMF, and many of these concentrate on the upper end slope. These studies range from theoretical arguments based on temperature and metallicity (Adams & Fatuzzo 1996; Bonnell, Clarke & Bate 2006; Krumholz 2011) through to the integrated galactic initial mass function (IGIMF) (Kroupa 2001; Kroupa & Weidner

★ E-mail: michele.fumagalli@durham.ac.uk

2003; Weidner, Pflamm-Altenburg & Kroupa 2011; Kroupa et al. 2013), which is based on the assumption that the most massive star that can form in a star cluster is related to the mass of that cluster (i.e. the upper limit of the IMF). In a recent study, Dib, Schmeja & Hony (2017) find suggestions of a variable IMF in the Milky Way through ‘lonely’ O star fractions, and ? infer a shallow high-mass end IMF slope for 30 Doradus. However, no evidence of IMF variation has been widely accepted in the literature, and a variety of alternative explanations have been proposed to many of these variable IMF indications, from the application of stochastic sampling of the IMF (Fumagalli, da Silva & Krumholz 2011) to the reported observations of isolated massive OB stars (Oey et al. 2013; Lamb et al. 2016).

In this paper, we continue the work started in Ashworth et al. (2017). In that study, we investigated the possibility of constraining the IMF from broad-band photometry using the Stochastically Light Up Galaxies (SLUG; da Silva, Fumagalli & Krumholz 2012, 2014; Krumholz et al. 2015a; Ashworth et al. 2017) stellar population synthesis (SPS) code. Using *Hubble Space Telescope* photometry from the Legacy ExtraGalactic UV Survey (LEGUS; Calzetti et al. 2015), combined with Bayesian analysis exploiting a large library of SLUG models, we demonstrated that accurate recovery of the IMF was non-trivial. Using mock clusters generated with SLUG, we went on to quantify the degeneracy between the mass and the IMF, and showed that using broad-band photometry alone, without accurate knowledge of the cluster mass, we could not put tight constraints on the slope of the IMF.

In this work, we include the UV region of the spectrum in addition to the broad-band photometry used previously. As UV features arise in the spectra of young, hot, and massive stars, including this region of the spectrum in the analysis should allow us to constrain the IMF’s high-end slope more closely than was shown to be possible using broad-band photometry alone. Other studies have made use of this fact, for example, in the recent work of Wofford et al. (2014), the authors use UV spectroscopy to constrain the high-mass end of the IMF (finding it not to be shallower than the Salpeter value). Previous work that has combined SPS modelling with UV spectral features to determine the IMF includes, e.g. Pettini et al. (2000), Steidel et al. (2004), and Quider et al. (2009), although these studies tend to perform more traditional statistical analysis. For instance, comparing deterministic SPS model spectra by eye to galaxy spectra calculated from discrete sets of physical parameters.

As in our previous work, we make use of SLUG to perform this study. The paper is structured in the following way. In Section 2, we describe new extensions made to the SLUG code in order to produce high-resolution UV spectra. These modifications are then used to construct several libraries of simulated clusters with a wide coverage of the mass-age-extinction-IMF parameter space. In Section 3, we continue with a qualitative exploration how the behaviour of the physical parameters that define the stellar populations maps onto UV spectral features and broad-band photometry. Next, in Section 4, we produce photometry and equivalent width measurements for a selection of simulated clusters that form a set of mock observations. These observations are then analysed using the Bayesian inference methods used in previous studies (Krumholz et al. 2015b; Ashworth et al. 2017), which we now extend to include spectral features as additional input data. Finally, we make predictions as to what spectral features are of importance when planning observations, and whether observations would be able to recover the high-mass IMF slope of an unresolved stellar population using the methods presented in this paper.

2 IMPLEMENTING HIGH-RESOLUTION UV SPECTRA IN SLUG

To generate our mock star clusters, we use the SLUG code (da Silva et al. 2012, 2014; Krumholz et al. 2015a; Ashworth et al. 2017). SLUG is a stochastic stellar population synthesis (SPS) code that allows the user to generate star clusters and galaxies by randomly drawing stars from an input IMF, as well as drawing star clusters themselves from a cluster mass function when populating galaxies. SLUG also has the capability to randomly sample the star formation history when populating galaxies. As part of the SLUG suite of software¹, we distribute a Bayesian analysis toolset by the name of BAYESPHOT. The BAYESPHOT tool uses Bayesian inference techniques to calculate posterior probability distribution functions (PDFs) of physical parameters that define the cluster from a set of input observables (generally broad-band photometry; Krumholz et al. 2015a). This is accomplished through the use of kernel density estimation (KDE) techniques, combined with a large library (which we refer to as the ‘training set’) of SLUG models.

2.1 Stellar atmosphere grids

Up until now, SLUG has been capable of producing spectra with a resolution of $\sim 10 \text{ \AA}$. However, this is not a high enough resolution to allow us to make use of the equivalent widths of absorption line features in our analysis. As it is these features that we wish to use to infer the slope of the high-mass end of the IMF, we have implemented a high-resolution spectral capability in the SLUG code. Along with this new capability, the associated machinery required to calculate equivalent widths for arbitrary line features and use them in Bayesian inference has also been implemented.

The highest spectral resolution that SLUG is capable of producing is limited to the resolution of the stellar atmospheres to which it has access. Therefore, in order to produce high-resolution spectra, SLUG must have access to high-resolution stellar atmosphere models in the wavelength region of interest. To this end, we modified SLUG to handle the high-resolution ($\sim 0.4 \text{ \AA}$) atmospheres of Leitherer et al. (2010), which are distributed as part of STARBURST99 (Leitherer et al. 1999). These atmospheres cover a wavelength range of 900 – 3000 $\text{\AA}$.

This new spectral synthesiser option, dubbed SB99HRUV in SLUG, covers the UV region of the spectrum (900 – 3000 $\text{\AA}$) using the following atmospheres:

- (i) Stars classified as Wolf–Rayet stars use the Potsdam PoWR atmospheres (Gräfener, Koesterke & Hamann 2002; Hamann & Gräfener 2004; Sander, Hamann & Todt 2012), specifically the revision shipped with STARBURST99.
- (ii) We use Kurucz (Kurucz 1992; Lejeune, Cuisinier & Buser 1997) atmospheres for all stars with $M < 5 M_{\odot}$, for stars with $T < 170\,000\text{K}$ in the mass range $5 - 10 M_{\odot}$, and for stars with $T < 22\,000\text{K}$ at masses $> 10 M_{\odot}$.
- (iii) For all other stars, we use the high-resolution IFA atmospheres (Leitherer et al. 2010).
- (iv) Stars that fall outside of these atmosphere grids are treated as black bodies, although this is a rare occurrence given our choice of tracks.

For the rest of the wavelength range of the spectrum, we use the standard lower resolution atmospheres (equivalent to selecting

¹ Available at www.slugsps.com

the default SB99 spectral synthesizer option, which mimics the atmosphere selection of STARBURST99 that are employed by SLUG. Further information is available in Krumholz et al. (2015a) and the SLUG user manual. The core differences between this and the high-resolution region described above are the use of the Hillier atmospheres (Hillier & Miller 1998) for Wolf-Rayet stars, and the use of Pauldrach OB atmospheres (Pauldrach et al. 1998) instead of the new IFA atmospheres.

All the atmospheres used during the work performed in this paper are those that are distributed with either SLUG, or in the case of the PoWR atmospheres and the high-resolution IFA UV atmospheres, those that are distributed with STARBURST99. Finally, we stitch the high- and low-resolution spectra together, replacing the 900 – 3000 Å region of the low-resolution spectrum with the high-resolution spectrum calculated using the atmospheres listed above.

In our calculation, we assume the Padova stellar tracks including thermally pulsating AGB stars (Vassiliadis & Wood 1993; Girardi et al. 2000; Vázquez & Leitherer 2005). We must however consider the effects of the limits of these tracks on the IMF we have selected. Indeed, we show in Section 3.1 how there can be significant differences between clusters generated with a $30 M_{\odot}$ mass cut-off compared to an IMF which cuts off at $120 M_{\odot}$, for example. There are, however, some observational indications for stars with masses greater than the assumed cut-off (for example, the recent work of Schneider et al. 2018, who find evidence for stars with masses up to $200 M_{\odot}$ in 30 Doradus) which warrant further exploration before applying our formalism to real data.

2.2 Selecting atmospheres

When selecting the high-resolution IFA atmospheres, we have implemented both the nearest-neighbour algorithm used by STARBURST99, and a new interpolation scheme whereby the atmosphere for a particular star is calculated through a barycentric interpolation. This can be toggled by simply changing a Boolean value in the high-resolution UV spectral synthesizer before compiling SLUG. The nearest-neighbour scheme allows direct comparison with STARBURST99, and is also less computationally intensive.

There are several steps to our new interpolation scheme. We begin with the parameters of the star for which we wish to obtain an atmosphere. First, we find the two nearest neighbour atmospheres in the temperature–surface gravity plane that exist within the atmosphere library. We then iterate through the next nearest neighbours until we satisfy the conditions required to enclose the star within a triangle of three atmospheres within this plane. Having formed this triangle of nearest neighbour atmosphere models, we find the barycentric coordinates of the star of interest within this triangle, and then use these to weight the atmospheres on each vertex of the triangle. The weighted combination of these atmospheres gives us an interpolated model for the star. If the triangle cannot be closed, we default to the nearest-neighbour method. We use this interpolation scheme during this paper rather than the nearest neighbour approach due to the sometimes large gaps in the parameter space. The definition of ‘nearest neighbour’ is the same one used by STARBURST99 for these atmospheres, with each direction (in the surface gravity–temperature plane) weighted differently (with the surface gravity axis weighted by a factor of 5; Leitherer et al. 2010).

To test the scheme, a test star was selected, and then removed from the atmosphere grid. The spectrum of the nearest-neighbour star and the interpolated atmosphere were then compared to the true atmosphere of the missing star. The interpolated star lay notice-

ably closer to the true value across most absorption features.² Our work is based on mock observations constructed with a consistent choice of interpolation scheme. These mock observations are then analysed using libraries of SLUG models constructed using the same interpolation scheme. Thus, we defer the study of which scheme performs best when applied to observations to future work.

We further note that, due to the low density of points in some regions of the temperature–surface gravity plane, certain combinations of parameters can lead to significant differences in flux between the high-resolution spectrum and the low-resolution spectrum that we stitch into, as in STARBURST99. This leaves discontinuities at the borders of the high-resolution region in some cluster parameter configurations, and affects photometry for filters that cross into this region. To ensure smooth transitions, we integrate a region in both the low- and high-resolution spectra (covering 2400 – 2900 Å) and scale the high-resolution spectrum by the ratio between these integrals, so that we can eliminate any large discontinuities at the points where the two spectra join. It is recommended, however, that users planning to compare SLUG photometry alone to real data use the low-resolution spectrum.

2.3 Equivalent width calculations

As the high-resolution IFA spectra are provided in both line and continuum form, we are able to normalize the spectrum without the difficulties of attempting to create a pseudo-continuum around each line feature. Therefore, we use only the rectified spectrum (as is output by STARBURST99) to calculate equivalent widths, which gives us a normalized spectrum consisting of features belonging only to the stars covered by the new high-resolution IFA atmospheres (the massive O and B type stars). This ensures that a true continuum is adopted in the equivalent width calculations rather than an empirically determined pseudo-continuum. As O and B stars dominate the signal for the features analysed in this work, this treatment provides an excellent approximation overall. In simple testing using an approximate normalization of the spectra for the other stellar types as performed by STARBURST99, we find that the majority of our ‘independent’ line list (see Table 2) is not strongly affected. The strongest effects are apparent around the ages where Wolf-Rayet stars contribute to the spectrum. Even then, only our Si II and O V features are strongly affected (with the latter present in the representative line list and the former in the independent line list). Between 10^5 yr and 10^7 yr, we see only minor differences for other lines in our independent line list (less than 10 percent when moving to the oldest clusters). The exception is for Mg I and Mg II, where we see differences of approximately 20–30%. There will also be an effect on any photometric filters which overlap the high-resolution region. However, we must note that these differences will change if newer versions of the Wolf-Rayet atmospheres are used. Overall, the use of O and B type stars alone is an acceptable approximation for this theoretical study, and the inclusion of more information in the spectrum in the form of additional stellar types may in fact assist in the recovery of the cluster physical parameters. It should also be

²Even when using the basic nearest-neighbour atmosphere selection scheme, results produced by SLUG will not be directly comparable to STARBURST99 due to the differing track methods that SLUG employs. However, using an older revision of the SLUG code, and using the nearest-neighbour atmosphere selection method, we were able to validate our implementation by reproducing the rectified and non-rectified spectra produced by STARBURST99 for a selection of test clusters.

Table 1. Table of the parameters used in SLUG for generating the mock star clusters used within this work. The base IMF used for all these model sets is the three component broken power law of the Kroupa IMF. The break points between the segments of this power law are situated at 0.01, 0.08, 0.5, $M_{*,\text{Max}}$. Z is the metallicity. The (D) represents discrete sampling, with # realizations at each value, and (C) represents continuous sampling, with # realizations drawn from a flat distribution across the given range.

Name	$M_{*,\text{Max}}$	Z	Log M ($M_{\odot}$)	Log t (yr)	# ^a	α_3
Library	120	0.020	2.7–5.7 (C)	4.7–7.7 (C)	5×10^7	−3.2 – −1.3 (C)
Mock Observations	120	0.020	3–5 (C)	5–7 (C)	5×10^2	−3.0 – −1.5 (C)
Set 1: High cut-off, solar Z	120	0.020	3,6 (D)	5,6,7 (D)	10^3	−3.0, −2.3, −1.5 (D)
Set 2: High cut-off, high Z	120	0.050	3,6 (D)	5,6,7 (D)	10^3	−2.3
Set 3: High cut-off, low Z	120	0.004	3,6 (D)	5,6,7 (D)	10^3	−2.3
Set 4: Low cut-off, solar	30	0.020	3,6 (D)	5,6,7 (D)	10^3	−3.0, −2.3, −1.5 (D)
Set 5: Low cut-off, high Z	30	0.050	3,6 (D)	5,6,7 (D)	10^3	−2.3
Set 6: Low cut-off, low Z	30	0.004	3,6 (D)	5,6,7 (D)	10^3	−2.3
Squares	120	0.020	3,6 (D)	5–7 (C)	10^5	−3.0 – −1.5 (C)

Note: ^aNumber of realizations.

noted that these rectified spectra are not extinguished, and contain no nebular emission. In this work, we run SLUG with nebular emission calculations enabled, so that both nebular emission and extinction are applied to the photometric filters and the raw non-rectified spectrum. Overall, we must stress that, even with the limitations detailed above, as we are comparing SLUG mock observations to libraries of other SLUG models, this approach is internally consistent for this study.

Finally, the equivalent width is calculated by interpolating linearly between each point in the spectrum (as they are the bin-centres of variable width bins, this is a best approximation), and then integrating under the resulting function across the integration region for a given spectral feature. Further information on the use of these new features is available in the SLUG documentation³. The code is currently publicly available, including the high-resolution IFA atmospheres. The PoWR Wolf–Rayet atmospheres will be added to the distribution when new versions from the PoWR group are made available (Hamann, private communication) as these supersede the versions provided with STARBURST99, which are used here.

3 GRIDS OF HIGH-RESOLUTION SLUG MODELS

In order to carry out the work described in this paper, we must generate a large selection of SLUG models. These models will be used for empirical study, to form mock observations, and to form the large library of simulations required by BAYESPHOT. The parameters that define these different models are detailed in Table 1.

The ‘Library’ is the large set of models that makes up the ‘training set’ we use for our Bayesian analysis. The Library contains 5×10^7 model star clusters, with their masses drawn from a flat distribution between $5 \times 10^2 M_{\odot}$ and $5 \times 10^5 M_{\odot}$, and their ages drawn from a flat distribution between 5×10^4 yr and 5×10^7 yr. The extinction, A_V is drawn from a flat distribution between 0 and 3, and the IMF high-mass slope α_3 is drawn from a flat distribution between −3.2 and −1.3. The actual shape of the distribution is not important during the Bayesian analysis, but a flat distribution makes sense in this context as the clusters we use to generate our input photometry are also generated using flat distributions. Nebular emission is calculated using an ionizing photon fraction $\phi = 0.73$, the default in SLUG. It must be noted once again, however, that neb-

ular emission (much like extinction) is not included in the rectified spectrum used for calculating equivalent widths from the combinations of high-resolution stellar atmospheres. This does mean that any nebular emission features are absent from the spectra we use to calculate the equivalent widths. However, much of the increased flux from the nebular emission, and the effects of extinction, would be negated by the normalization process. In addition, as the observations are generated with the same restrictions as the models used in the Bayesian analysis the work is self-consistent.

The ‘Mock Obs.’ models are those whose photometry and equivalent widths form our set of mock observations that we use in our Bayesian analysis. Their masses are drawn from a slightly more narrow set of distributions in mass, age, and IMF than the library, as shown in Table 1. This means we have no mock clusters situated right at the edge of the library in any part of parameter space (save for extinction, where we cannot go negative), or even outside of it. This is a desirable property when studying clusters with BAYESPHOT. ‘Sets’ 1–6 are models used for studies of both individual lines and the behaviour of the equivalent width for varying age, mass, α_3 , and metallicity. Each set is comprised of 1000 realizations, enough for the median equivalent widths to be acceptably converged. Finally, the ‘squares’ set is a group of models intended to be used to investigate the changes in equivalent width in the age– α_3 plane. These models are drawn from flat distributions in log t and α_3 .

All of these sets of models have both photometry, along with equivalent widths for the features listed in Table 2. A selection of the features are taken from Fanelli et al. (1992). Others were defined by inspection of test model spectra combined with reference to NED⁴ and Kramida et al. (2015). The photometric filters are one of the standard five-band filter sets used in the LEGUS project (Calzetti et al. 2015): F275W, F336W, and F555W filters from Wide Field Camera 3 (WFC3 UVIS), and the F435W and F814W filters from the Advanced Camera for Surveys (ACS). In addition to photometry and equivalent widths, we also generate full spectra for the models referred to as ‘Sets’ 1–6.

We note that when they form, star clusters are embedded in gaseous clouds, and are therefore unable to be probed by observations as required by this study. It can take up to a few Myr for these clouds to be cleared, even for low-mass clusters (as detailed by, for example Hollyhead et al. 2015). However, in this theoretical study,

⁴The 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.

³<http://slug2.readthedocs.io>

Table 2. Table of the line features used in this work, where we list the integration regions used for calculating the equivalent width of each feature, along with an ionization energy if appropriate. The ionization potentials given were obtained from Kramida et al. (2015). The lines are sorted by increasing leading edge wavelength of their integration region. Spectral features taken from Fanelli et al. (1992) are marked with an asterisk (*). In C IV, the letters ‘a’, ‘c’, ‘e’ represent the absorption, centre, and emission, respectively (as given in Fanelli et al. 1992). For other P-Cygni features such as N v, the ‘a’ and ‘b’ represent the absorption and emission components, respectively. For other features such as Si III, the ‘a’ and ‘b’ represent two occurrences of the same feature, with ‘a’ being the stronger in equivalent width. Features marked with a † mark are those used in the full-set analysis as they are sufficiently independent of one another. Sub-regions of larger spectral features are indented.

Feature name	Integration range (Å)	ϕ^a (eV)
O VI	1015.0–1050.0	138.1189
O viat†	1015.0–1037.0	–
O vib	1037.0–1050.0	–
C III†	1174.0–1177.0	47.88778
Si IIIb	1205.0–1207.0	33.49300
N v	1220.0–1250.0	97.89013
N vat†	1220.0–1240.0	–
N vb	1240.0–1250.0	–
BI1302*	1292.0–1312.0	–
C II†	1334.0–1338.0	24.3845
O v	1355.0–1375.0	113.8989
O vat†	1355.0–1370.0	–
O vb	1370.0–1375.0	–
Si IV* †	1387.0–1407.0	45.14179
BI1425*	1415.0–1435.0	–
Si IIIat†	1416.0–1418.0	33.49300
Fe1453 * †	1440.0–1466.0	–
C IV	1527.0–1569.0	64.49351
C iva* †	1530.0–1550.0	–
C ivc*	1540.0–1560.0	–
C ive*	1550.0–1570.0	–
BI1617*	1604.0–1630.0	–
He II†	1639.5–1641.0	54.4177650
BI1664*	1651.0–1677.0	–
N IV†	1703.0–1727.0	77.4735
BI1719*	1709.0–1729.0	–
BI1853*	1838.0–1868.0	–
Si I†	1843.0–1846.0	8.15168
Fe IIa*	2382.0–2422.0	16.19920
BI2538*	2520.0–2556.0	–
Fe IIb*	2596.0–2622.0	16.19920
MgWide*	2670.0–2870.0	–
Mg II* †	2784.0–2814.0	15.035271
Mg I* †	2839.0–2865.0	7.646236

Note: ^aIonization potential.

the simulations at the age point of 0.1 Myr are an illustrative extreme low-age case, and in the online Appendix A we will consider the effects of narrowing our age range to > 1 Myr.

3.1 Effects of physical parameters on spectral features

To investigate the effects of changing the variable physical parameters on star cluster spectra in a qualitative manner, we run several sets of SLUG models for a variety of IMF slopes, metallicities, cluster masses, and cluster ages. Their parameters are listed in Table 1 as Sets 1–6.

As a first example, we look into the effects of changing the value of α_3 on the equivalent widths of the spectral features, and how these equivalent widths then evolve with time. This is done for both low-mass ($10^3 M_\odot$) clusters and high-mass ($10^6 M_\odot$) clusters to get a handle on the effects of stochasticity on the equivalent widths we observe. Results are shown in Fig. 1 for the N IV feature, where the points correspond to the median value from the set of 1000 realizations, and the error bars cover the extent of the 5th to 95th percentile range. We stress that the size of the error bars is not Poisson error, but arises from the stochastic sampling of the IMF (an effect most pronounced in low-mass star clusters and clusters with bottom heavy IMFs).

The effects of stochasticity are readily apparent when comparing the low- and high- mass clusters [panels (a) and (b) to panels (c) and (d), respectively], with the low-mass clusters having a wide scatter in equivalent width at all times. This stochasticity is mostly absent for the large $10^6 M_\odot$ clusters, with some scatter being apparent for the extremely steep IMF, as expected, due to the small possibility of generating a large star instead of the much more probable large number of small stars.

In the case of the IMF with the $120 M_\odot$ cut-off [panels (a) and (c)], we see an upward trend in equivalent width for all three of the α_3 values, although it is somewhat obscured by the scatter in the case of the low-mass cluster with a steep IMF. We see much the same behaviour for the truncated IMF in panels (b) and (d), with an overall increase in equivalent width compared to the IMF with the higher mass cut-off. This truncation results in a reduction of extremely hot and massive stars, which in turn leads to less of the Nitrogen being highly ionized, and therefore results in more absorption at young ages for shallower IMF slopes. Already we are able to infer that including lines in our analysis helps to constrain α_3 , although this may be hampered by stochasticity at low masses.

We now move to examine the shape of the spectra themselves, a study which relates to the analysis of stacked spectra. We also look at how metallicity affects the spectrum, and how this compares with the effects caused by variations in the IMF’s parameters. A sample of the spectrum around the N v features is given in Fig. 2, with panel (a) showing how the spectrum varies at a fixed metallicity (solar), fixed age (10^5 yr), and fixed mass ($10^6 M_\odot$) for changing values of the IMF slope ($\alpha_3 = -3, -2.3, -1.5$). Panel (b) shows how the spectrum varies at the same fixed age and mass, but with a fixed IMF slope of $\alpha_3 = -2.3$, and metallicities of $Z = 0.004, 0.020, 0.050$. In the changing IMF case, we see that the strength of the absorption line feature increases as the IMF becomes more shallow, due to the increased probability of producing hot, massive stars. We see that changing metallicity has a similar effect on the N v line that changing the IMF does, and a high-metallicity environment appears to be degenerate with a top heavy IMF in the ‘N va region’ of the feature. However, the ‘N vb integration region’ (the emission feature) is not appreciably affected by the changing metallicity, whereas it does differentiate between the three IMF slopes.

Unfortunately, we are at present unable to vary the metallicity continuously in SLUG because we lack a set of atmosphere models that is sufficiently densely sampled to allow for effective interpolation. Should a continuous variation of metallicity be implemented in SLUG, it would be possible to introduce it as a fifth physical variable in the Bayesian analysis. This is a degeneracy which should be explored in future work when the necessary features have been implemented in SLUG. For the rest of this study, however, we maintain a constant solar metallicity.

Fig. 3 shows how the equivalent width of two spectral features changes through the age– α_3 plane. Each panel contains the same set




Figure 1. Evolution of the equivalent width of N IV with cluster age for a range of IMF upper-end slopes. Panels (a) and (b) are clusters with a mass of $10^3 M_{\odot}$, whereas panels (c) and (d) are clusters with masses of $10^6 M_{\odot}$. Panels (a) and (c) have the IMF cut-off at $120 M_{\odot}$, whereas panels (b) and (d) have the IMF cut-off at $30 M_{\odot}$. The points correspond to the median value from the set of realizations, and the error bars cover the 5th to 95th percentile range. Blue squares correspond to $\alpha_3 = -3.0$, purple triangles correspond to $\alpha_3 = -2.3$, and red circles correspond to $\alpha_3 = -1.5$. The leftmost point of each trio is positioned at the correct age value, and the others are shifted slightly to aid readability.

of 10^6 cluster realizations drawn randomly from the distributions listed in Table 1 (the ‘Squares’ set of models). The equivalent width assigned to each grid element is the median of all the realizations that fall within that element. Panel (a) shows the plane for the C III feature, and panel (b) shows the plane for the C IV feature, chosen as they both illustrate well how the equivalent width changes with the two variables, the latter being a strong line and the former being weak. We see that both C III and C IV allow us to differentiate the IMF slope at young ages, whereas for older clusters we lose this distinguishing power as the massive stars that are responsible for the line features evolve outside the main sequence. Note that,

unlike the C IV line which increases in strength as the IMF becomes shallow, the C III line becomes weaker as this happens. This effect occurs because a shallower IMF results in a higher proportion of the cluster’s light coming from stars that lack the C III feature due to them being hot enough to ionise carbon past the C^{+2} state responsible for producing the line.

We once again note that observations are limited in how young an age they are able to probe, and would be likely limited to ages of a few Myr or greater. However, even at these ages, we still see useful differences in the equivalent widths in Fig. 3.




Figure 2. Renormalized stacked spectra of 1000 model star clusters of mass $10^6 M_{\odot}$ at an age of 10^5 yr. In panel (a), each set of 1000 was generated at a different IMF ($\alpha_3 = -3.0, -2.3, -1.5$), at solar metallicity. In panel (b), each set of 1000 was generated using a different metallicity, with $Z = 0.004, 0.020, 0.050$. All models were generated using the canonical Kroupa IMF as a base IMF shape. The spectrum shows the region around the N v line, with the integration region used for the calculation of the equivalent width outlined in a light grey colour. We see the line increase in strength as the IMF becomes shallower. Note also the degeneracy in the absorption region between the IMF and the metallicity. The emission region is unaffected by changing metallicity.




Figure 3. 2D plots of equivalent width versus changes in cluster age or IMF slope for two spectral regions (C III and C IV λ), where the equivalent width reported for a grid cell is the median of the equivalent widths for the clusters that reside in that area of parameter space. In both panels, we see that the equivalent width of these lines may be used to differentiate the various possible values of α_3 at young ages. In addition, the combination of the two lines provides us with a way to also differentiate the age of the cluster in conjunction with the value of α_3 .

4 ANALYSIS OF MOCK OBSERVATIONS

Having discussed the qualitative behaviour of the equivalent width as a function of physical parameters, we now move to study more quantitatively how combining the equivalent widths of spectral features with broad-band photometry may be used to constrain the slope of the IMF, α_3 .

The mock observations are produced from the photometry and equivalent widths of a selection of SLUG model clusters. These clus-

ters were generated by drawing their physical parameters randomly, as listed in Table 1 (‘Mock Obs.’). It is using this set of mock observations that we attempt to recover the underlying physical properties of the corresponding simulated clusters. For each set of mock observations we extract the photometry and equivalent widths, assigning each a nominal error of 0.05. We must stress that the mock clusters are not taken from the library, and are generated independently. However, as we use the same tracks and atmospheres to generate

the mock clusters, the result is somewhat optimistic compared to observations.

The first step is to generate posterior PDFs (1D and 2D) for the physical parameters of each mock cluster that we have generated. An example (cluster 19) is plotted in Fig. 4, with red markings corresponding to the true values of each parameter. The analysis was performed four times for each cluster. First, using one of the LEGUS filter sets (as used for NGC 628e, see Adamo et al. 2017), which is comprised of the F275W, F336W, and F555W filters from WFC3 UVIS, and the F435W, and F814W filters from ACS. Next, the procedure is repeated using the following spectral features in place of the photometry: C III, C IVa, N IVa, O IVa, O Vb. These features were selected as a representative set of five, with reference to their appearance in the IMF–age plane (c.f. Fig. 3), which we use to identify those lines most able to distinguish the cluster age or the IMF slope whilst also giving us a range of line strengths and including an emission feature. Next, the analysis is repeated using both the photometry and the equivalent widths of these spectral features together, first using this representative set of five lines, and then finally using all the spectral features listed in Table 2 marked with a † symbol. We cannot include all the line features in the table in our analysis at the same time, as not all the features are independent. Some of our features are sub-features of a single feature we also list, are blends containing many features, or are even multiple occurrences of the same line. For example, we list both the full line feature and the separate components of the lines with P-Cygni profile. For the purposes of this analysis, we choose to use only the absorption component when using the ‘full’ independent line list, and indeed the shoulder features may depend greatly on the assumed wind model and therefore may include more uncertainties. We also discard the weaker feature when a line occurs more than once, and we remove the blend features which contain multiple lines.

Panel (a) of Fig. 4 shows the results of using the broad-band photometry alone. We see quite good recovery of the cluster age and the extinction. However, recovery of the cluster’s mass and the value of α_3 is poor, due to the underlying degeneracy between the two. Conversely, in panel (b) we see that lines alone recover the age and the IMF slope much more successfully, at the expense of the mass and the extinction. This is expected as A_V is unconstrained by construction, and the equivalent width (having been calculated from a normalized spectrum) is weakly dependent on the total cluster mass. In panel (c), we see the combination of the two results in good recovery for all the physical parameters. Indeed, the inclusion of the UV lines further refines our determination of the cluster’s age, and helps to alleviate the mass–IMF degeneracy that effects analysis performed with broad-band photometry alone (Ashworth et al. 2017), although it does not fully break it. Panel (d) is the same as panel (c), but the analysis was performed with all the line features. The results are mostly comparable, with slight shifts towards more accurate mass, age, and IMF, and a spreading out in the extinction direction. We conclude therefore, that the representative set of five lines is sufficient to greatly improve the recovery of the cluster’s physical parameters.

In Fig. 5, we plot the residuals (the difference between either the median or mode of the posterior PDFs and the true value of the corresponding parameter) for 10 of the mock observation clusters. We see much the same effect as we see in Fig. 4 in most cases, and there is some improvement overall in the recovery of the IMF. Note also how the medians and modes move much closer together, as the posterior PDFs become narrower. The overall trends seen for Cluster 19 in Fig. 4 apply more generally across our mock observations.

Finally, we quantify the overall behaviour of the PDFs across our entire input data set by looking at stacks of the all the posterior PDFs for 500 mock observation clusters, which we show in Fig. 6. In panels (a) and (b) the posterior PDFs are centred on their medians, whereas those in panels (c) and (d) are centred on the true values of each parameter for each of the clusters in the stack. Panels (a) and (c) are produced using only the five LEGUS filters in the analysis, whereas panels (b) and (d) are produced using a combination of the both the photometry and the equivalent widths of the lines in our representative set of five. The 5th and 95th percentiles are marked with orange dot–dashed lines, and the first and third quartiles are marked with indigo dashed lines. The zero marker, on which the PDFs are centred as they are stacked, is the solid red line (or cross in the 2D PDFs).

As seen in the previous examples, there is significant improvement overall in our recovery of the physical parameters once the spectral features are included. Indeed, we see for the true-value centred stacks that we move from quite wide distributions around the value of interest to quite narrow peaks. This can be seen more clearly by looking at the changes in the interpercentile (5th to 95th) and interquartile (first to third) ranges, as listed in Table 3 for both the true-centred and median-centred cases. While the median-centred case (marked MC) is a useful companion to panels (a) and (b) of Fig. 6, the true-centred rows (marked TC) illustrate the improvement (or worsening) of our recovery of the parameters. In effect, we see a reduction in the overall scatter (as represented by the 5–95 range) across all parameters, save for the extinction, with improvements having a mean of $\sim 25\%$. For the interquartile range, we see a mean reduction in width of $\sim 23\%$, with a narrowing of $\sim 32\%$ for α_3 . The broadening of the posterior PDFs for extinction is caused by the lack of extinction information in the rectified spectra we use for calculating the equivalent widths of our spectral features. This results in the almost flat posterior PDFs that we see in Fig. 4. It is the addition of the lines, which are insensitive to the extinction, to the extinction sensitive photometry which causes the spreading out of our distributions here.

While we are unable to entirely break the mass–IMF degeneracy described by Ashworth et al. (2017) (curved contours in the joint $\log M$ – α_3 PDF are still visible even with the equivalent widths included in the analysis), we are able to make significant improvements on constraining the value of not only α_3 , but the mass and age of the clusters as well. We also note that some residual scatter is intrinsic due to the stochasticity in our models, and cannot be reduced by adding more clusters. Indeed, during testing we found no significant difference in scatter between stacks of 10, 100, or 500 clusters for a fixed Kroupa IMF case. Our analysis has therefore shown, when introducing additional information in the form of UV spectral lines, that it is possible to tighten the constraints on α_3 , and therefore possible to rule out extreme variations in α_3 as well.

4.1 Observational constraints

To aid the planning of new observations, in this section we aim to predict what spectral features are required to gain sufficient improvement in recovery of the IMF when compared to the case of using pure broad-band photometry. What is of interest is specifically what information is needed to get the improved recovery of α_3 that we desire. To this end, we take the empirical approach of adding in features by their overall strength (equivalent width) until there is a point of diminishing returns, where the addition of further spectral features does not improve the posterior PDFs by an appreciable amount.




Figure 4. Example corner plot for one of the mock observation clusters, cluster 19. The 1D posterior PDFs for $\log M$, $\log T$, A_V , and α_3 are shown in the top panels of each column. The joint posterior PDFs are represented by contours (beginning at 0, with increments of 0.5 unit element). They are also represented by a 2D map, where the intensity of the colour reflects the probability density as indicated on the colour bar. All PDFs are normalized to have unit integral. In panel (a), are the results when the analysis is performed using only the LEGUS filters, in panel (b) we show results from using the representative selection of lines only, and in panel (c) we show results from using a combination of the two. In panel (d), we present the result from analysis using all the line features that are independent (see Table 2). Note the improvement in the recovery of both the cluster's age and IMF slope when including the spectral lines in the analysis. The red crosses and lines mark the positions of the true values for the parameters that describe this cluster.

The most basic method by which we can do this is to limit which spectral features we include in our Bayesian analysis by imposing limits on the equivalent width to which observations may be sensitive. This is a simple approach which allows us,

to first approximation, to gauge the sensitivity required by real observations, and thus the wavelength coverage, resolution, and exposure time needed to constrain α_3 in a given sample of clusters.




Figure 5. Residuals of the 1D posterior PDFs for the parameters $\log M$, $\log t$, A_V , and α_3 for 10 of the clusters generated to provide mock observations. The residuals are calculated from both the medians of the distributions (blue crosses) and the modes of the distributions (purple triangles), when compared to the true values for each of these parameters. We see that the inclusion of the lines into the analysis improves the recovery of all of the parameters, and that in most cases the median and mode move closer together, indicating a narrowing of the distributions as well.

To this end, we select 100 example clusters from the ‘Mock Obs.’ model set, and repeat the Bayesian analysis presented in Section 4. However, instead of using a fixed selection of filters for the entire data set, we limit the inclusion of spectral features to those which have an equivalent width that passes a particular threshold for each cluster individually. We choose limits of 12, 9, 6, $3 \text{ \AA}$, along with the case where we have all the line features included, and the case where we use broad-band photometry alone. In the case of the $9 \text{ \AA}$ limit, this gives us between zero and three line features across the input clusters.

In panel (a) of Fig. 7, we see the residuals of the 1D posterior PDFs for the four physical parameters for this set of input clusters.

The markers represent the median residual value for each equivalent width limit, with the error bars covering the interquartile range. The markers in red represent the residuals calculated from the difference between the median of the posterior PDFs and the true values, whereas the purple points represent the residuals calculated using the difference between the mode of the posterior PDFs and the true values.


We note that even at a sensitivity limit of $9 \text{ \AA}$, we see definite improvement in age, mass, and IMF determination. As sensitivity increases, we see a slow improvement in recovery for these parameters, although the median residuals are close to zero across the range. However, the residuals for the extinction increase. This is caused by the line features having no extinction applied to them, resulting in an almost flat PDF for the A_V parameter as seen in Fig. 4. As more and more spectral features are included, the posterior PDF for extinction is spread out from the narrow peak we obtain when using filters alone, as more and more noisy PDFs are combined with it. Over this sample of 100 clusters, we find that even with a limit of $9 \text{ \AA}$ we are able to significantly improve the accuracy at which we are able to infer the value of α_3 , and indeed that further improvements in sensitivity provide limited improvement over this. In both age and cluster mass, however, further increases in sensitivity continue to visibly improve results.

The improvement in our ability to constrain the IMF is more visible in panel (b) of Fig. 7. In this figure, we see the evolution of the absolute value of the median of the residuals, as well as the evolution of the interquartile range, as we increase the sensitivity. Looking at the interquartile ranges, we see that our overall ability to constrain the IMF improves significantly with the introduction of equivalent width data, with improvement in α_3 being $\sim 32\%$ (median residual) or $\sim 37\%$ (mode residual) at the $9 \text{ \AA}$ limit. However, further improvements gained by increasing our sensitivity are small, finally coming to $\sim 39\%$ (median residual) or $\sim 61\%$ (mode residual) when we include all the spectral features in the analysis. We see similar behaviour in mass, whereas the improvement in age is more consistent across the range of equivalent width limits. Our constraints on extinction loosen, for the reason explained earlier.


5 SUMMARY AND CONCLUSIONS

In this paper, we have developed an extension to the SPS code SLUG, whereby the resolution of the output spectra has been greatly increased in the UV region (from $\simeq 10 \text{ \AA}$ to $\simeq 0.4 \text{ \AA}$). Also, a new equivalent width calculation mechanism has been implemented to take advantage of these new higher resolution spectra. These new capabilities were then leveraged to perform a theoretical study of the high-mass end of the IMF and its relationship to observable quantities for unresolved stellar populations.


Having performed a qualitative exploration of how the equivalent widths of spectral features are affected by changes in the various physical parameters ($\log M$, $\log t$, A_V , α_3) that describe our model star clusters, we applied the Bayesian analysis capabilities of BAYESPHOT to mock observations generated with SLUG. These observations (both broad-band photometry and the equivalent widths of UV line features) were analysed using a large library of SLUG models which covered a wide parameter space. We found similar difficulty in tightly recovering the value of α_3 as in Ashworth et al. (2017) where degeneracy between the cluster mass and the IMF slope hampered this recovery. Repeating the analysis while including the equivalent widths of five representative spectral features resulted in a significant improvement in the recovery of the physical parameters for the clusters when averaged over the whole set




(a) LEGUS Filters, Median centred



(b) Lines and LEGUS filters, Median centred



(c) LEGUS Filters, True centred



(d) Lines and LEGUS filters, True centred

Figure 6. Stacks of the one and 2D PDFs for 500 clusters generated as mock observations, as described in Table 1. The stacked PDFs are normalized by the number of clusters being stacked, with contours spaced in steps of 0.5, beginning at 0.5. The red crosses and lines mark the zero points. Panels (a) and (c) are the results of analysis using the LEGUS filters only, and are median-centred and centred on the true values, respectively. Likewise, panels (b) and (d) show the results of analysis using both the lines (C_{III} , C_{IVa} , N_{IVa} , O_{IVa} , O_{VIB}) and the LEGUS filters, once again median-centred and true-value-centred, respectively. Note the significant improvement in the recovery of the cluster age when including the lines in the analysis. We also see improvement in the recovery of the cluster mass and IMF, although with a residual $\log M - \alpha_3$ degeneracy in the joint posterior PDF for those parameters. The 5th and 95th percentiles are marked with orange dot-dashed lines, and the first and third quartiles are marked with indigo dashed lines. The change in the interpercentile ranges [for the true-centred cases, panels (c) and (d)] are given in Table 3.

Table 3. Table of the difference in inter-quartile and inter-percentile ranges for the median-centred (MC) and true-centred stacks (TC) shown in Fig. 6, changing from the case with photometry alone to the case including both photometry and the selected UV features. The numbers in parentheses correspond to the percentage change in the ranges, compared to the original value.

Range (Percentiles)	$\Delta \log M$ ($M_{\odot}$)	$\Delta \log t$ (yr)	ΔA_V	$\Delta \alpha_3$
5–95 (MC)	−0.61 (−35%)	−0.63 (−49%)	+0.18 (+31%)	−0.34 (−23%)
25–75 (MC)	−0.30 (−42%)	−0.18 (−41%)	+0.08 (+36%)	−0.24 (−35%)
5–95 (TC)	−0.84 (−39%)	−0.87 (−55%)	+0.14 (+20%)	−0.50 (−26%)
25–75 (TC)	−0.33 (−41%)	−0.22 (−45%)	+0.07 (+27%)	−0.24 (−32%)




Figure 7. The plots in panel (a) show the residuals of the 1D posterior PDFs for the parameters $\log M$, $\log t$, A_V , and α_3 for the first 100 mock observation clusters. The residuals are calculated from both the medians of the distributions (red upward-pointing markers) and the modes of the distributions (purple downward-pointing markers), when compared to the true values for each of these parameters. For each cluster, we have set a value of equivalent width below which no spectral features are included in the analysis. The markers correspond to the median value of the residual at each equivalent width limit, and the error bars cover the interquartile range. The markers are offset by ± 0.2 around their limit value for visual clarity. Panel (b) illustrates the behaviour of both the interquartile range, and the absolute value of the median residual for both the cases where it is calculated from the modes of the PDFs (reds, downward triangles) and the medians (blues, upward triangles). The markers in panel (b) are not offset, as there are no error bars to overlap.

of 500 mock observations, with the interquartile range of the IMF posterior PDF shrinking by 32%. However, we still see evidence of the mass–IMF degeneracy noted in previous work.

Finally, we provided predictions for the required sensitivity in equivalent width needed to make improvements in recovery of α_3 over using broad-band photometry alone. We find that a modest limit of 9 $\text{\AA}$ provides sufficient improvement (with the interquartile range of the median-residuals shrinking by 32% for the IMF), with further improvement as more line features are included in the analysis.

The study presented in this work is purely theoretical, and as such provides an idealised situation for recovering the IMF from spectra. Although, due to SLUG’s stochastic nature, it is not simply a trivial problem of finding the exact match in the library, it would be instructive to attempt to recover the IMF for clusters with an independently calculated α_3 value. The real-world performance improvements of the method presented in that study will depend on how well the UV features of the spectrum are modelled. High-resolution UV spectral

observations of a full-resolved stellar population, the central cluster of 30 Doradus (Crowther et al. 2016), for example, would provide suitable input equivalent widths for Bayesian analysis with SLUG, whilst also allowing an independent calculation of the IMF through star counts.

A limitation of the models generated during this work is the absence of binary systems in the star clusters we build. Binary star systems are thought to make up a large proportion of the star systems present in galaxies, with perhaps all O and B stars residing in binary systems, as summarized by Eldridge et al. (2017). Indeed, the presence of binaries affects the evolution of a given stellar system due to events such as mergers and mass transfer between the two stars, and their presence results in a given population of stars appearing more blue than would be predicted by models that do not contain binaries (Eldridge 2012). The effect of mergers and mass transfer would be to push the IMF to a shallower slope, as massive stars would become even more massive with time as they accrete. In

addition, Eldridge & Stanway (2009) note that the inclusion of binaries can spread the Wolf–Rayet phase out across a longer age range, which would affect our analysis if all stellar types were included in the equivalent width calculation. Although capturing these effects in SLUG requires tracking multiple new parameters (which could be varied within our Bayesian framework at the expenses of computational efficiency), this extension may provide a more realistic library that can be used to constrain the IMF. Another limitation is the absence of stellar rotation in our models. This is due to the use of the Padova tracks with thermally pulsating AGB stars which do not include stellar rotation. However, as SLUG is distributed with the Geneva stellar tracks (Ekström et al. 2012), which do include the effects of rotation, it would be possible to study what effects (if any) including rotation has on the IMF slopes that we recover. The inclusion of rotation should result in an increase in the luminosity of O stars (Vázquez et al. 2007; Levesque et al. 2012), although the effect of this would be dependent on the speed of this rotation. Once more this would require the addition of further variables into our analysis framework.

In conclusion, we find that the inclusion of UV spectral features does improve recovery of the physical parameters of unresolved stellar populations in our theoretical study, over that possible with broad-band photometry alone. However, we do not find that it makes it possible to fully break the mass–IMF degeneracy that exists in these systems. Having applied cuts to the included line features, we find observations sensitive to a modest equivalent width are capable of improving the recovery of the IMF high-mass slope α_3 to a significant degree, with a scatter narrow enough to either rule out or indeed distinguish extreme variations in the IMF.

ACKNOWLEDGEMENTS

GA acknowledges support from the Science and Technology Facilities Council (ST/P000541/1 and ST/M503472/1). MF acknowledges support by the Science and Technology Facilities Council (ST/P000541/1). This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 757535). AA acknowledges the support of the Swedish Research Council (Vetenskapsrådet) and the Swedish National Space Board (SNSB). We also thank the LEGUS Collaboration, along with Claus Leitherer and Wolf-Rainer Hamann for helpful discussions. We wish to thank the referee, Prof. Nate Bastian, for insightful comments on this manuscript. This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grants ST/H008519/1 and ST/K00087X/1, STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the National E-Infrastructure. This research has made use of the NASA/IPAC Extragalactic Database (NED), which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. For access to the data and codes used in this work, please contact the authors or visit www.slugsps.com.

REFERENCES

Adamo A. et al., 2017, *ApJ*, 841, 131
Adams F. C., Fatuzzo M., 1996, *ApJ*, 464, 256

Ashworth G. et al., 2017, *MNRAS*, 469, 2464
Bastian N., Covey K. R., Meyer M. R., 2010, *ARA&A*, 48, 339
Bonnell I. A., Clarke C. J., Bate M. R., 2006, *MNRAS*, 368, 1296
Calzetti D. et al., 2015, *AJ*, 149, 51
Chabrier G., 2003, *PASP*, 115, 763
Crowther P. A. et al., 2016, *MNRAS*, 458, 624
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
Dib S., Schmeja S., Hony S., 2017, *MNRAS*, 464, 1738
Ekström S. et al., 2012, *A&A*, 537, A146
Eldridge J. J., 2012, *MNRAS*, 422, 794
Eldridge J. J., Stanway E. R., 2009, *MNRAS*, 400, 1019
Eldridge J. J., Stanway E. R., Xiao L., McClelland L. A. S., Taylor G., Ng M., Greis S. M. L., Bray J. C., 2017, *PASA*, 34, e058
Fanelli M. N., O’Connell R. W., Burstein D., Wu C.-C., 1992, *ApJS*, 82, 197
Fumagalli M., da Silva R. L., Krumholz M. R., 2011, *ApJ*, 741, L26
Girardi L., Bressan A., Bertelli G., Chiosi C., 2000, *A&AS*, 141, 371
Gräfener G., Koesterke L., Hamann W.-R., 2002, *A&A*, 387, 244
Hamann W.-R., Gräfener G., 2004, *A&A*, 427, 697
Hillier D. J., Miller D. L., 1998, *ApJ*, 496, 407
Hollyhead K., Bastian N., Adamo A., Silva-Villa E., Dale J., Ryon J. E., Gazak Z., 2015, *MNRAS*, 449, 1106
Kramida A., Yu. Ralchenko Reader J., NIST ASD Team, 2015, NIST Atomic Spectra Database (ver. 5.3), [Online]. Available at: <http://physics.nist.gov/asd> [2017, May 16]. National Institute of Standards and Technology, Gaithersburg, MD
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, *The Stellar and Sub-Stellar Initial Mass Function of Simple and Composite Populations*. Springer Science+Business Media, Dordrecht, 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
Kurucz R. L., 1992, in Barby B., Renzini A., eds, *Proc. IAU Symp. 149, The Stellar Populations of Galaxies*. Angra dos reis, Brazil, p. 225
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
Leitherer C. et al., 1999, *ApJS*, 123, 3
Leitherer C., Ortiz Otálvaro P. A., Bresolin F., Kudritzki R.-P., Lo Faro B., Pauldrach A. W. A., Pettini M., Rix S. A., 2010, *ApJS*, 189, 309
Lejeune T., Cuisinier F., Buser R., 1997, *A&AS*, 125, 229
Levesque E. M., Leitherer C., Ekstrom S., Meynet G., Schaerer D., 2012, *ApJ*, 751, 67
Massey P., Johnson K. E., Degioia-Eastwood K., 1995, *ApJ*, 454, 151
Oey M. S., Lamb J. B., Kushner C. T., Pellegrini E. W., Graus A. S., 2013, *ApJ*, 768, 66
Pauldrach A. W. A., Lennon M., Hoffmann T. L., Sellmaier F., Kudritzki R.-P., Puls J., 1998, in Howarth I., ed., *ASP Conf. Ser. Vol. 131, Properties of Hot Luminous Stars*. Astron. Soc. Pac., San Francisco, p. 258
Pettini M., Steidel C. C., Adelberger K. L., Dickinson M., Giavalisco M., 2000, *ApJ*, 528, 96
Quider A. M., Pettini M., Shapley A. E., Steidel C. C., 2009, *MNRAS*, 398, 1263
Salpeter E. E., 1955, *ApJ*, 121, 161
Sander A., Hamann W.-R., Todt H., 2012, *A&A*, 540, A144
Schneider F. R. N. et al., 2018, *Science*, 359, 69
Steidel C. C., Shapley A. E., Pettini M., Adelberger K. L., Erb D. K., Reddy N. A., Hunt M. P., 2004, *ApJ*, 604, 534
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
- Wofford A., Leitherer C., Chandar R., Bouret J.-C., 2014, *ApJ*, 781, 122

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

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