



HAL
open science

Modelling and interpreting spectral energy distributions of galaxies with BEAGLE

Jacopo Chevallard, Stéphane Charlot

► **To cite this version:**

Jacopo Chevallard, Stéphane Charlot. Modelling and interpreting spectral energy distributions of galaxies with BEAGLE. Monthly Notices of the Royal Astronomical Society, 2016, 462, pp.1415-1443. 10.1093/mnras/stw1756 . hal-03645047

HAL Id: hal-03645047

<https://hal.science/hal-03645047v1>

Submitted on 10 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Modelling and interpreting spectral energy distributions of galaxies with BEAGLE

Jacopo Chevallard¹*† and Stéphane Charlot²

¹Scientific Support Office, Directorate of Science and Robotic Exploration, ESA/ESTEC, Keplerlaan 1, NL-2201 AZ Noordwijk, the Netherlands

²Sorbonne Universités, UPMC-CNRS, UMR7095, Institut d'Astrophysique de Paris, F-75014, Paris, France

Accepted 2016 July 18. Received 2016 July 17; in original form 2016 March 9

ABSTRACT

We present a new-generation tool to model and interpret spectral energy distributions (SEDs) of galaxies, which incorporates in a consistent way the production of radiation and its transfer through the interstellar and intergalactic media. This flexible tool, named BEAGLE (for BayESian Analysis of GaLaxy sEds), allows one to build mock galaxy catalogues as well as to interpret any combination of photometric and spectroscopic galaxy observations in terms of physical parameters. The current version of the tool includes versatile modelling of the emission from stars and photoionized gas, attenuation by dust and accounting for different instrumental effects, such as spectroscopic flux calibration and line spread function. We show a first application of the BEAGLE tool to the interpretation of broad-band SEDs of a published sample of $\sim 10^4$ galaxies at redshifts $0.1 \lesssim z \lesssim 8$. We find that the constraints derived on photometric redshifts using this multipurpose tool are comparable to those obtained using public, dedicated photometric-redshift codes and quantify this result in a rigorous statistical way. We also show how the post-processing of BEAGLE output data with the PYTHON extension PYP-BEAGLE allows the characterization of systematic deviations between models and observations, in particular through posterior predictive checks. The modular design of the BEAGLE tool allows easy extensions to incorporate, for example, the absorption by neutral galactic and circumgalactic gas, and the emission from an active galactic nucleus, dust and shock-ionized gas. Information about public releases of the BEAGLE tool will be maintained on <http://www.jacopochevallard.org/beagle>.

Key words: methods: data analysis – methods: statistical – dust, extinction – H II regions – galaxies: evolution – galaxies: stellar content.

1 INTRODUCTION

Over the last 15 yr, our understanding of how galaxies form and evolve has improved substantially. The combination of technological and theoretical progress has brought this field into a new era: advances in observational techniques (e.g. multi-object spectroscopy, efficient near-infrared CCDs) have enabled multiwavelength observations of large samples of galaxies out to the highest redshifts, while the steady rise of computational power and refinement of numerical techniques have fostered new approaches (e.g. semi-analytic models, hydro-dynamic simulations) to model the formation and evolution of galaxies. This progress has led to a general consensus about the main physical ingredients required to describe the evolution of the galaxy population (e.g. Gonzalez-Perez et al. 2014; Lu et al. 2014; Vogelsberger et al. 2014; Henriques et al. 2015; Schaye et al. 2015): collapse and hierarchical growth of dark

matter haloes; accretion of baryons on to these haloes; conversion of baryons into stars; feedback of massive stars and active galactic nuclei (AGN) on star formation; supernova- and AGN-driven outflows of metal-enriched gas; infall of both pristine and metal-enriched gas on to galaxies. The large-scale environment can also affect galaxy properties, in particular, by providing quenching mechanisms (e.g. tidal or ram-pressure stripping, strangulation; e.g. Lagos et al. 2014; Rafieferantsoa et al. 2015), and through its influence on the merger rate (e.g. Lackner et al. 2012; Rafieferantsoa et al. 2015) and galactic spins (e.g. Hahn, Teyssier & Carollo 2010; Codis et al. 2012). Although these different ingredients are present in many galaxy formation models, we still lack a detailed quantification of their respective roles in shaping the properties of galaxies. This is because of the complexity inherent in galaxy physics, which combines gravity, radiation hydro-dynamics, magnetic fields and high-energy physics, acting on scales from less than a pc (e.g. for the formation of proto-stellar cores) to over a Mpc (e.g. for environmental effects). For this reason, ‘first-principles’ simulations of galaxy formation remain far beyond the reach of current computational capabilities. Instead, small-scale baryonic physics is generally subsumed into

* E-mail: jchevall@cosmos.esa.int

† ESA Research Fellow.

sub-grid prescriptions, which vary from model to model (e.g. Scannapieco et al. 2012; Haas et al. 2013a,b; Vogelsberger et al. 2013; Torrey et al. 2014; Crain et al. 2015). The appropriateness of such prescriptions, and hence, our ability to understand galaxy formation, must be assessed by comparing simulated and observed galaxy properties.

Comparing the predictions of galaxy formation models with observations requires one to relate properties pertaining to the evolution of baryons in dark-matter haloes, such as gas cooling and star formation, with observables, such as ultraviolet, optical and infrared spectral energy distributions (SEDs). This can be achieved using models of stellar population synthesis and of the transfer of starlight through the interstellar and intergalactic media (e.g. Tinsley 1978; Bruzual 1983; Arimoto & Yoshii 1987; Guiderdoni & Rocca-Volmerange 1987; Buzzoni 1989; Bressan, Chiosi & Fagotto 1994; Worthey 1994; Leitherer & Heckman 1995; Fioc & Rocca-Volmerange 1997; Bruzual & Charlot 2003; Maraston 2005; Conroy & Gunn 2010; Vazdekis et al. 2010; Maraston & Strömbäck 2011). This spectral modelling involves several additional components, such as the stellar initial mass function, prescriptions for the evolution, spectral properties and release of heavy elements by individual stars of different initial masses and chemical compositions, and prescriptions for the influence of the interstellar medium (ISM) and the intergalactic medium (IGM) on stellar radiation. Galaxy SEDs can be computed in this way from the histories of star formation and chemical enrichment predicted by galaxy formation models. The interpretation of photometric and spectroscopic galaxy observations with such model SEDs to constrain stellar masses, metallicities, star formation histories and ionized-gas properties is at the base of most galaxy evolution studies.

Two major limitations, often neglected, affect this type of analysis: the adoption of oversimplified models to describe the wide variety of observed galaxy SEDs and the presence of ‘systematic’ model uncertainties. This second limitation has been addressed in several studies already (e.g. Charlot, Worthey & Bressan 1996; Cerviño, Luridiana & Castander 2000; Conroy, Gunn & White 2009; Percival & Salaris 2009; Conroy & Gunn 2010; Conroy, White & Gunn 2010a). The difficulty of precisely quantifying systematic model uncertainties has led to mainly qualitative conclusions, leaving the problem unsolved. The first limitation is easier to tackle, for example, by using more physically realistic models of galaxy SEDs and combining these with advanced statistical techniques to extract physical constraints from data. This appears as the most promising route to fully exploit the information gathered by modern photometric and spectroscopic galaxy surveys. Yet, the several tools proposed so far to interpret galaxy SEDs in terms of physical parameters do not allow one to fully exploit the high quality of modern data. For example, most existing approaches rely on the adoption of a rigid physical model (e.g. analytic, two-parameter star formation histories combined with a standard dust attenuation curve and the assumption that all stars in a galaxy have the same metallicity) to describe galaxy SEDs (e.g. Bolzonella et al. 2010; Wuyts et al. 2011; Hernán-Caballero et al. 2013; Ilbert et al. 2013; Bauer et al. 2013; Muzzin et al. 2013; Lundgren et al. 2014; Kochiashvili et al. 2015; Mortlock et al. 2015; Kawinwanichakij et al. 2016). Even with the inclusion of superimposed bursts of star formation (e.g. Kauffmann et al. 2003; Gallazzi et al. 2005; Pozzetti et al. 2007; Gallazzi & Bell 2009; da Cunha et al. 2010), this does not allow a physically consistent description of the contributions by stars, gas and dust to the integrated emission from a galaxy, nor the inclusion of a potential AGN component (a notable exception is the approach of Pacifici et al. 2012, who incorporate star formation and chemical

enrichment histories from numerical simulations of galaxy formation and emission from photoionized gas). Also, current spectral analysis tools are generally optimized to interpret either photometric or spectroscopic observations of galaxies, but not arbitrary combinations thereof. Finally, most existing tools suffer from additional limitations: many focus on the selection of ‘best-fitting’ parameters rather than on the uncertainties associated with these parameters (e.g. chi-square minimization techniques; Arnouts et al. 1999; Bolzonella, Miralles & Pelló 2000; Kriek et al. 2009); when this is not the case, the number of free parameters that can be explored is generally limited (e.g. with grid-based Bayesian techniques; da Cunha, Charlot & Elbaz 2008; Noll et al. 2009; Pacifici et al. 2012); and when more sophisticated (e.g. Markov Chain Monte Carlo, hereafter MCMC) techniques allow the exploration of more parameters, instrumental effects are generally not incorporated in the analysis (e.g. Acquaviva, Gawiser & Guaita 2011; Serra et al. 2011; Han & Han 2014).

In this paper, we introduce a new-generation tool to interpret galaxy SEDs, *BEAGLE* (for BayESian Analysis of GaLaxy sEDs), which incorporates several main novelties. The modular design of this tool, written in FORTRAN 2003/08, allows one to easily combine, in a physically consistent way, different prescriptions for the production of starlight in galaxies and its transfer through the ISM (absorption and emission by gas, attenuation by dust) and the IGM (absorption by gas). Other modules to be implemented in the future include the infrared emission from dust and the emission from an AGN. The *BEAGLE* tool includes several possible prescriptions to describe the star formation and chemical enrichment histories of galaxies, ranging from simple analytic functions to the predictions of sophisticated galaxy formation models. The flexible parametrization of these and other model parameters allows one to build mock galaxy catalogues as well as to interpret any combination of photometric and spectroscopic observation of galaxies by adapting model complexity (i.e. number of free parameters) to the data, without sacrificing coherence. Moreover, the adopted Bayesian framework allows the characterization of complex, non-linear correlations among model parameters in high-dimensional parameters spaces, comparisons between competing models, the description within a hierarchical framework of single objects as well as populations of objects, and a reduction of the parameter space using informative priors based on high-quality observations. Finally, posterior predictive checks comparing model predictions with observations enable one to identify model failures, which can drive future model developments.

In Section 2 below, we outline the modular, highly versatile approach used to describe the physical properties of stars, gas and dust in the *BEAGLE* tool. In Section 3, we describe in detail the statistical approach and output products arising from the analysis of galaxy SEDs with this tool. In Section 4 we present an example application of *BEAGLE* to the interpretation of broad-band SEDs of galaxies in a wide range of redshifts and discuss the results of such modelling in Section 5. We also introduce also a PYTHON-based extension to post-process *BEAGLE* results, named *PYP-BEAGLE* (for PYthon Postprocessing of *BEAGLE*). We compare our approach with existing SED fitting codes in Section 6 and summarize our conclusions in Section 7. Throughout the paper, we adopt a present-day solar metallicity $Z_{\odot} = 0.01524$ (corresponding to a zero-age metallicity of 0.017, see table 3 of Bressan et al. 2012) and the latest constraints on cosmological parameters from *Planck*, i.e. $\Omega_{\Lambda} = 0.6911$, $\Omega_{\text{m}} = 0.3089$, and $H_0 = 67.74$ (see last column ‘TT,TE,EE+lowP+lensing+ext’ of table 4 of Planck Collaboration 2015).

2 ASTROPHYSICAL INGREDIENTS OF BEAGLE

Our main aim in this paper is to design a general, astrophysically sound framework to fit any combination of photometric and spectroscopic galaxy observations, as well as to produce and analyse mock catalogues of galaxy SEDs. The need for a single framework to study both true and mock galaxy data arises from the growing role played by simulations in the optimal preparation and exploitation of modern surveys. Simulations are important to optimize the observational strategy (e.g. signal-to-noise thresholds, spectral resolution) needed to answer a given scientific question, but also to characterize the performances and systematics of the increasingly complex instruments mounted on new ground-based and space-based telescopes.

The ultraviolet, optical and infrared SEDs of galaxies include contributions by stars, gas, dust and potentially an AGN. Since these components are physically linked to one another, they must be described in a consistent way in spectral models of galaxies. In this section, we present our formalism to model the production of starlight in galaxies and its transfer through the ISM and the IGM. We appeal to the isochrone synthesis technique introduced by Charlot & Bruzual (1991).¹ In this approach, the luminosity (in units of $\text{erg s}^{-1} \text{\AA}^{-1}$) emitted at wavelength λ by a galaxy at time t from the onset of star formation can be expressed as

$$L_\lambda(t) = \int_0^t dt' \psi(t-t') S_\lambda[t', \mathbf{Z}(t-t')] T_\lambda^{\text{ISM}}(t, t'), \quad (1)$$

where $\psi(t-t')$ is the star formation rate at time $t-t'$ (which traces the star formation history), $S_\lambda[t', \mathbf{Z}(t-t')]$ the luminosity emitted per unit wavelength per unit mass by a simple stellar population (SSP) of age t' and chemical composition $\mathbf{Z}(t-t')$ (each element of this vector corresponding to a different chemical element), and $T_\lambda^{\text{ISM}}(t, t')$ the transmission function of the ISM. The function $S_\lambda(t', \mathbf{Z})$ can be expressed as (e.g. Conroy et al. 2009)

$$S_\lambda(t', \mathbf{Z}) = \int_{m_{\text{low}}}^{m_{\text{up}}} dm \phi(m) \Lambda_\lambda[L_{\text{bol}}(m, \mathbf{Z}, t'), T_{\text{eff}}(m, \mathbf{Z}, t'), \mathbf{Z}], \quad (2)$$

where m is the stellar mass, $\phi(m)$ the stellar initial mass function [IMF; defined such that $\phi(m)dm$ is the number of stars born with masses between m and $m+dm$] with lower and upper mass cutoffs m_{low} and m_{up} , and $\Lambda_\lambda[L_{\text{bol}}(m, \mathbf{Z}, t'), T_{\text{eff}}(m, \mathbf{Z}, t'), \mathbf{Z}]$ the spectral energy distribution of a star with bolometric luminosity $L_{\text{bol}}(m, \mathbf{Z}, t')$, effective temperature $T_{\text{eff}}(m, \mathbf{Z}, t')$, and chemical composition \mathbf{Z} .

Following Charlot & Longhetti (2001), we express the transmission function in the ISM as

$$T_\lambda^{\text{ISM}}(t, t') = T_\lambda^+(t, t') T_\lambda^0(t, t'), \quad (3)$$

where $T_\lambda^+(t, t')$ and $T_\lambda^0(t, t')$ are the transmission functions of the ionized gas and the neutral ISM, respectively. Here $T_\lambda^+(t, t')$ accounts for both the absorption and the emission of photons in the ionized gas, i.e. $T_\lambda^+(t, t')$ will be close to zero at wavelengths shorter than the H-Lyman limit and greater than unity at those corresponding to emission lines. The transmission functions of the neutral interstellar gas and the IGM will be discussed separately in Sections 2.6 and 2.7 below.

¹ The term ‘isochrone’ refers to the location of coeval stars with homogeneous chemical composition and different masses in the Hertzsprung–Russell diagram.

For a galaxy at redshift z , the spectral flux density (in units of $\text{erg s}^{-1} \text{\AA}^{-1} \text{cm}^{-2}$) reaching the observer at wavelength $\lambda_{\text{obs}} = \lambda(1+z)$ is related to the luminosity L_λ emitted at rest wavelength λ (equation 1) by the relation

$$F_{\lambda_{\text{obs}}} = \frac{L_\lambda}{4\pi d_L(z)^2} \frac{1}{1+z} T_{\lambda,z}^{\text{IGM}}, \quad (4)$$

where d_L is the luminosity distance at redshift z , $T_{\lambda,z}^{\text{IGM}}$ the (redshift-dependent) transmission function of the IGM, and the factor $1/(1+z)$ ensures energy conservation, by accounting for the wavelength stretching $d\lambda_{\text{obs}} = (1+z)d\lambda$.

In the next sections, we describe our prescriptions for the emission from stars and the ionized gas, the star formation and chemical enrichment histories, changes in the α -element to iron abundance ratio (α/Fe), attenuation by dust and IGM absorption. We stress that a main feature allowed by the modular design of the BEAGLE tool is that the prescriptions adopted here for these different physical ingredients can be easily replaced by alternative ones. We do not discuss here our prescriptions for an AGN component and infrared dust emission, which are currently being incorporated and will be the subject of a future release (Section 7).

2.1 Stellar population synthesis code

We adopt the latest version of the Bruzual & Charlot (2003) code to describe the emission of stellar populations at wavelengths between 91 \AA and 190 μm (equation 2). This version of the code incorporates updated stellar evolutionary tracks (Bressan et al. 2012; Marigo et al. 2013) and the MILES library of observed optical stellar spectra (Sánchez-Blázquez et al. 2006) to describe the properties of stars in the Hertzsprung–Russell diagram [$\Lambda_\lambda(L_{\text{bol}}, T_{\text{eff}}, \mathbf{Z})$ in equation (2)]. We compute in this way SSPs in wide ranges of ages, $t' \in [10^5, 2 \times 10^{10}]$ yr, and metallicities, $Z \in [0.0001, 0.04]$ (i.e. $[\text{Fe}/\text{H}] \in [-2.18, 0.42]$).²

2.2 Stellar initial mass function

The stellar IMF [term $\phi(m)$ in equation 2] controls the relative contributions by stars of different masses to the integrated SED of a galaxy. Recently, several studies have claimed the presence of systematic variations of the IMF, such as a potential steepening in early-type galaxies (e.g. van Dokkum & Conroy 2010; Conroy & van Dokkum 2012b; Cappellari et al. 2012; La Barbera et al. 2013; Sonnenfeld et al. 2015), for which theoretical motivations exist (e.g. Chabrier, Hennebelle & Charlot 2014). The implication of these claims for stellar population synthesis modelling are being complicated by the fact that different studies produce inconsistent results (Smith 2014; Zieleniewski et al. 2015) and may depend on underlying assumptions (Tang & Worthey 2015; McConnell, Lu & Mann 2016). In this paper, we adopt a fixed IMF corresponding to the Galactic-disc IMF of Chabrier (2003), with lower and upper mass cutoffs $m_{\text{low}} = 0.1$ and $m_{\text{up}} = 100 M_\odot$. We note that the Chabrier (2003) IMF is very similar to the ‘universal’ IMF of Kroupa (2001) and differs from the classical Salpeter (1955) IMF [$\phi(m) \propto m^{-2.35}$] through the turn-over of the distribution $\phi(m)$ at stellar masses $m < 1 M_\odot$.

² Here Z stands for the sum of the elements of vector \mathbf{Z} , which corresponds to the mass fraction in all elements heavier than helium.

2.3 α element-to-iron abundance ratio

Many spectral evolution models available in the literature (e.g. Pégase, Fioc & Rocca-Volmerange 1997; Le Borgne et al. 2004; GALAXEV, Bruzual & Charlot 2003; FSPS, Conroy et al. 2009; ; Vazdekis et al. 2010; MaStro, Maraston & Strömbäck 2011; Conroy & van Dokkum 2012a) rely on libraries of observed spectra of solar-neighbourhood stars (e.g. ELODIE, Prugniel & Soubiran 2001; Prugniel et al. 2007; STELIB, Le Borgne et al. 2003; Indo-US, Valdes et al. 2004; MILES, Sánchez-Blázquez et al. 2006). This is because current theoretical spectral libraries suffer from persisting inaccuracies linked to, for example, the inclusion of line blanketing, incomplete line lists, the accuracy of quantum-mechanical calculations, the treatment of convection in stellar interiors (see e.g. Martins & Coelho 2007). The problem of appealing to observed spectral libraries is that the abundances of solar-neighbourhood stars are tied to the particular star formation and chemical enrichment history at our location in the Milky Way. This potentially biases the interpretation with such libraries of galaxies with different star formation and chemical enrichment histories, such as massive, early-type galaxies, which have typically larger α/Fe ratio than star-forming, late-type galaxies like the Milky Way (e.g. Milone, Barbuy & Schiavon 2000; Worthey, Tang & Serven 2014).

A way to overcome the limitations of observed spectral libraries is to adopt theoretical libraries to compute differential properties of stellar populations with non-scaled solar element abundance ratios relative to the well-calibrated scaled-solar case. Following Walcher et al. (2009, see also Vazdekis et al. 2015), we account for the dependence of the function S_λ on α/Fe ratio in equation (1) by writing

$$S_\lambda(t', Z, [\alpha/\text{Fe}]) = \frac{S_\lambda^{\text{theor}}(t', Z, [\alpha/\text{Fe}])}{S_\lambda^{\text{theor}}(t', Z)} S_\lambda(t', Z), \quad (5)$$

where $S_\lambda(t', Z)$ is the luminosity per unit wavelength of an SSP of age t' and metallicity Z computed using an empirical library of stellar spectra (Section 2.1), for which $\alpha/\text{Fe} \approx \alpha/\text{Fe}_\odot$ (i.e. $[\alpha/\text{Fe}] \approx 0$), and $S_\lambda^{\text{theor}}(t', Z)$ and $S_\lambda^{\text{theor}}(t', Z, [\alpha/\text{Fe}])$ are the theoretical predictions for SSPs of same age t' and metallicity Z for scaled-solar ($[\alpha/\text{Fe}] = 0$) non-scaled solar ($[\alpha/\text{Fe}] \neq 0$) α element-to-iron abundance ratio, respectively. We adopt for $S_\lambda^{\text{theor}}(t', Z)$ and $S_\lambda^{\text{theor}}(t', Z, [\alpha/\text{Fe}])$ the predictions of Coelho et al. (2007, as updated by Coelho 2014), which cover the wavelength range $\lambda \in [0.13, 100] \mu\text{m}$ at low resolution, and $\lambda \in [0.25, 0.9] \mu\text{m}$ at high resolution, ages between 30 Myr and 14 Gyr, iron abundances $[\text{Fe}/\text{H}] \in [-1, 0.2]$ and α elements-to-iron abundance ratios $[\alpha/\text{Fe}] \in [0, 0.4]$.

2.4 Nebular emission

We adopt the recent photoionization models of star-forming galaxies of Gutkin, Charlot & Bruzual (2016, who follow the prescription of Charlot & Longhetti 2001) to describe the transmission function of the ionized gas $T_\lambda^+(t, t')$ in equation (3). These models combine the stellar population synthesis code described in Section 2.1 above with the latest version of the photoionization code CLOUDY (version 13.3, last described by Ferland et al. 2013). In this approach, the ensemble of H II regions and the diffuse gas ionized by young stars throughout a galaxy are described by means of effective (i.e. galaxy-wide) parameters. The main adjustable parameters of the photoionized gas are the interstellar metallicity, Z_{ISM} , the typical ionization parameter of a newly ionized H II region, U_S (which characterizes the ratio of ionizing-photon to gas densities at the edge

of the Stroemgren sphere), and the dust-to-metal (mass) ratio, ξ_d (which characterizes the depletion of metals on to dust grains). We adopt the large grid of models computed by Gutkin et al. (2016) for $Z_{\text{ISM}} \in [0.0001, 0.04]$ (i.e. $[\text{Fe}/\text{H}] \in [-2.18, 0.42]$), $\log U_S \in [-4, -1]$ and $\xi_d \in [0.1, 0.5]$. These provide the nebular (lines+continuum) emission at all SSP ages less than $t' = 10$ Myr (equation 3), over which 99.9 per cent of the ionizing photons are released by a single stellar generation. Hence, in this prescription, $T_\lambda^+(t, t') = 1$ for $t' > 10$ Myr. We do not include in the present paper the nebular emission produced by faint, post-asymptotic-giant-branch stars, shocks and a potential AGN component (these will be implemented in future releases of the BEAGLE tool).

2.5 Star formation and chemical enrichment histories

The star formation history $\psi(t - t')$ and chemical enrichment history $Z(t - t')$ entering the expression of the spectral energy distribution $L_\lambda(t)$ of a galaxy (equation 1) depend on the complex physical processes affecting baryons trapped in dark-matter potential wells (see Section 1). In principle, one should be able to constrain ψ and Z by comparing the observed SED of the galaxy to the predictions of spectral evolution models corresponding to different star formation and chemical enrichment histories (e.g. Heavens et al. 2004; Cid Fernandes et al. 2007; Panter et al. 2007; Pacifici et al. 2013; Tojeiro et al. 2013). This requires an approach including the broadest possible ranges in ψ and Z , enabling one to identify the parametrization best describing the observations as well as potential biases linked to the adoption of less appropriate assumptions (e.g. Pacifici et al. 2015).

A common way to parametrize the star formation and chemical enrichment histories of galaxies is to adopt a smooth, analytic (e.g. constant, exponentially declining or rising) function to describe $\psi(t - t')$ and a fixed metallicity Z to describe the chemical composition of stars of all ages in a galaxy. Although such parametrization may be adequate to describe the average properties of different galaxy classes (e.g. massive galaxies of different morphological types in the nearby Universe), they cannot account for the complex star formation and chemical enrichment histories of individual galaxies in a hierarchical universe. As an example, the Milky Way is known to have experienced different bursts of star formation in the past (e.g. Rocha-Pinto et al. 2000b) and to currently harbour stars spanning a broad range of chemical compositions (e.g. Rocha-Pinto et al. 2000a).³

This complexity may be more appropriately modelled by appealing to non-parametric descriptions of $\psi(t - t')$ and $Z(t - t')$. This requires discretizing the evolution in a finite number of ‘star formation periods’ (or ‘bins’, each spanning a range of ages), adjusting the mass fraction and metallicity in each period to best reproduce a given observation (see, e.g.; STARLIGHT, Cid Fernandes et al. 2004, 2005; MOPED, Heavens, Jimenez & Lahav 2000; STECKMAP, Ocvirk et al. 2006 and VESPA, Tojeiro et al. 2007). A common limitation of this approach is the large number of free parameters involved and their correlations, which requires high-S/N spectroscopic observations across a broad wavelength range to produce useful

³ The superposition of ‘stochastic’ bursts on top of smooth star formation histories, has been shown to better reproduce the observed spectroscopic properties of individual galaxies (e.g. Kauffmann et al. 2003; Brinchmann et al. 2004), but still assuming that all stars in a given galaxy have a fixed metallicity.

constraints (e.g. Tojeiro et al. 2013; McDermid et al. 2015; Citro et al. 2016; López Fernández et al. 2016). An alternative approach was recently proposed by Pacifici et al. (2012), who adopt physically motivated prescriptions for $\psi(t - t')$ and $Z(t - t')$ extracted from the semi-analytic post-treatment of a cosmological dark-matter simulation. They show that this provides a powerful new means of interpreting observed SEDs of galaxies in terms of physical parameters (e.g. Pacifici et al. 2013, 2015). They also show how the dependence of $\psi(t - t')$ and $Z(t - t')$ on the adopted galaxy formation model can be alleviated by resampling the model star formation and chemical enrichment histories to widen the range of evolutionary stages probed at any redshift, at the expense of weakening the link between the resulting library of star formation and chemical enrichment histories and the input cosmological and galaxy formation model (Pacifici et al. 2012).

We incorporate in BEAGLE all of the above approaches to describe the functions $\psi(t - t')$ and $Z(t - t')$ in equation (1). We achieve this by means of a general parametrization of the star formation and chemical enrichment histories of a galaxy, described as a succession of a flexible number of star formation periods, and by interpolating on the fly the grid of input SSPs (Section 2.1) to compute $S_\lambda[t', Z(t - t')]$ in equation (2) for any combination of age and metallicity. The shapes of $\psi(t - t')$ and $Z(t - t')$ in these components can be drawn from analytic functions or from different flavours of galaxy formation models, such as phenomenological and semi-analytic models and hydro-dynamic simulations. We note that such a parametrization allows one to easily separate the past history of star formation, imprinted in the spectral signatures of low-mass, long-lived stars, from the ‘current’ (i.e. averaged over the last 10–100 Myr) star formation, traced by young massive stars, which controls the nebular and far-ultraviolet emission (ignoring the contribution from shocks and a potential AGN). This parametrization is also appropriate to describe the stochastic nature of the star formation process in galaxies. A main advantage of the flexible approach presented here is the possibility to adapt the complexity of the description of the star formation and chemical enrichment histories of a galaxy to the type and quality of the available data (e.g. photometry versus spectroscopy, low- versus high-S/N).

2.6 Dust attenuation

In the previous sections, we have presented our prescriptions for the emission from stars and ionized gas, changes in heavy-element abundance ratios and the star formation and chemical enrichment histories of a galaxy in the BEAGLE tool. We now turn to the effect of dust attenuation on the stellar and nebular emission, expressed by the function $T_\lambda^0(t, t')$ in equation (3). The signatures of this effect compete with those of age and metallicity in galaxy SEDs (e.g. Wise & Silva 1996; Papovich, Dickinson & Ferguson 2001; Guo et al. 2011). The importance of a careful description of dust properties is also that this may be used to constrain the mechanisms of dust production (e.g. winds from asymptotic-giant-branch stars, supernova explosions; Höfner 2009; Cherchneff 2010) and destruction (e.g. shocks; Jones 2004; Jones & Nuth 2011) from the analysis of different types of galaxies at various redshifts.

We follow the standard terminology and refer by attenuation (or ‘effective absorption’) to the combined effects of absorption and scattering in and out of the line of sight to a galaxy caused by both local and global geometric effects, the term extinction being reserved for photon absorption along and scattering out of a single line of sight (e.g. Charlot & Fall 2000). The dependence of

extinction on wavelength (i.e. the extinction curve) has been measured in the Milky Way, the Small and the Large Magellanic Clouds (MW, SMC and LMC; Prevot et al. 1984; Bouchet et al. 1985; Clayton & Martin 1985; Cardelli, Clayton & Mathis 1988, 1989). These studies show that, in the Milky Way, the ultraviolet-to-near-infrared extinction curve varies along different lines of sight, while the strength of the characteristic absorption bump near 2175 Å drops from the average MW, to the LMC, to the SMC extinction curves. These extinction curves are commonly used to model the effect of attenuation by dust on galaxy SEDs, although this amounts to assuming that dust is unrealistically distributed in a uniform screen between the source and the observer.⁴ Such an assumption may introduce unwanted biases in the interpretation of galaxy SEDs. It also neglects the fact that young stars in their birth clouds are typically more attenuated than older stars in galaxies (e.g. Silva et al. 1998; Charlot & Fall 2000).

We adopt a more physically motivated prescription for attenuation by dust and express the transmission function of the neutral ISM (equations 1 and 3) as a function of stellar age t' and galaxy inclination θ ,

$$T_\lambda^0(t, t') \equiv T_\lambda^0(t', \theta). \quad (6)$$

Here we have dropped the dependence of T^0 on galaxy age t , i.e. we assume that attenuation by dust depends only on the current ISM properties and on the age distribution of stellar populations in a galaxy. We can rewrite the transmission function as

$$T_\lambda^0(t', \theta) = \exp[-\hat{\tau}_\lambda(t', \theta)], \quad (7)$$

where $\hat{\tau}_\lambda(t', \theta)$ is the attenuation (or effective absorption) optical depth of the dust affecting photons emitted at wavelength λ in all directions by all stars and gas in the galaxy, which emerge in the direction θ from the normal to the equatorial plane of the galaxy (assuming azimuthal symmetry).

We use the general parametrization in equation (7) to implement in a flexible way in the BEAGLE tool different prescriptions for $\hat{\tau}_\lambda(t', \theta)$, summarized in Table 1. These include simple mean extinction curves (MW, LMC, SMC), the starburst attenuation curve of Calzetti, Kinney & Storchi-Bergmann (1994), the two-component dust model of Charlot & Fall (2000, hereafter CF00) and new prescriptions by Chevallard et al. (2013, hereafter C13). We indicate in Table 1 how these different prescriptions account for three major features affecting dust attenuation in galaxies: the distribution of dust relative to stars (geometry); the enhanced attenuation of young stars in their birth clouds relative to older stars; and disc inclination.⁵ The mean MW, LMC and SMC extinction curves often adopted in galaxy spectral analyses do not account for any of these features. The (angle-averaged) attenuation curve of Calzetti et al. (1994),

⁴ A special case in which the use of mean extinction curves is physically motivated is the analysis of occulting galaxy pairs. In this case, dust in the foreground galaxy attenuates the light coming from the background one, hence motivating the adoption of a screen geometry to describe the effect of dust on the light emitted by the background galaxy (e.g. Holwerda, Keel & Bolton 2007; Holwerda & Keel 2013).

⁵ We do not consider explicitly here changes in the size distribution and optical properties of dust grains, as these are implicitly included in Table 1, either through the difference between the MW, LMC and SMC extinction curves or the possibility to change the slope of the attenuation curve in the CF00 and C13 prescriptions. Such changes are expected to have less effect on integrated galaxy SEDs than geometry, age dependence and inclination (e.g. Granato et al. 2000; Fontanot et al. 2009, C13), except perhaps around the 2175 Å absorption feature (e.g. Conroy, Schiminovich & Blanton 2010b).

Table 1. Different prescriptions for dust attenuation implemented in the BEAGLE tool (see Section 2.6 for details).

Prescription for dust attenuation	Effects accounted for			Adjustable parameters
	Geometry ^a	Age dependence ^b	Inclination ^c	
Mean extinction curve (MW, LMC, SMC)	no	no	no	$\hat{\tau}_V$
Calzetti et al. (1994)	yes	yes (implicitly)	no	$\hat{\tau}_V$
Charlot & Fall (2000)	yes	yes	no	$\hat{\tau}_V, \mu$
Chevallard et al. (2013) ‘quasi-universal’ relation	yes	yes	yes (implicitly)	$\hat{\tau}_V, \mu$
Chevallard et al. (2013) ‘full model’	yes	yes	yes	$\hat{\tau}_V, \mu, \theta, t_{\text{thin}}, t_{\text{thick}}, t_{\text{bulge}}$

Notes. ^aInfluence of the relative distributions of stars and dust on attenuation (a ‘no’ refers to a uniform foreground screen model).

^bInfluence of stellar age on attenuation (enhanced attenuation of young stars in their birth clouds relative to older stars).

^cInfluence of disc inclination on attenuation (when some stars and dust components are distributed in discs).

empirically determined from the spectra of 39 nearby starburst and blue-compact galaxies, implicitly incorporates the effects of geometry (as indicated by the shallower slope of this curve relative to standard extinction curves) and of the dependence of attenuation on stellar age (as indicated by the factor of ~ 2 difference in the attenuation affecting line and continuum photons in this sample). The main limitation of this prescription is that it was derived by neglecting the strong dependence of the slope of the attenuation curve on dust optical depth (see section 3.1.2 of C13). In Table 1, only the C13 prescriptions include the dependence of attenuation on galaxy inclination.

We now describe in slightly more detail the implementation in the BEAGLE tool of the CF00 and C13 dust prescriptions. The CF00 prescription is based on an angle-averaged, two-component model, which accounts for the fact that stars are born in dense molecular clouds, which dissipate on a time-scale of typically 10 Myr. The dust attenuation optical depth in this model can be expressed as

$$\hat{\tau}_\lambda(t') = \begin{cases} \hat{\tau}_\lambda^{\text{BC}} + \hat{\tau}_\lambda^{\text{ISM}} & \text{for } t' \leq 10 \text{ Myr,} \\ \hat{\tau}_\lambda^{\text{ISM}} & \text{for } t' > 10 \text{ Myr,} \end{cases} \quad (8)$$

where $\hat{\tau}_\lambda^{\text{BC}}$ and $\hat{\tau}_\lambda^{\text{ISM}}$ are the dust optical depths in stellar birth clouds and the ambient (diffuse) ISM, respectively. The attenuation curves in these two components are parametrized as

$$\hat{\tau}_\lambda^{\text{BC}} = \hat{\tau}_V^{\text{BC}} \left(\frac{\lambda}{0.55 \mu\text{m}} \right)^{-n_V^{\text{BC}}}, \quad (9)$$

$$\hat{\tau}_\lambda^{\text{ISM}} = \hat{\tau}_V^{\text{ISM}} \left(\frac{\lambda}{0.55 \mu\text{m}} \right)^{-n_V^{\text{ISM}}}, \quad (10)$$

where $\hat{\tau}_V^{\text{BC}}$ and $\hat{\tau}_V^{\text{ISM}}$ are the V -band attenuation optical depths in stellar birth clouds and in the ambient ISM and, following Wild et al. (2007, see also da Cunha et al. 2008), we adopt $n_V^{\text{BC}} = 1.3$ and $n_V^{\text{ISM}} = 0.7$. We parametrize this model in terms of the total V -band attenuation optical depth, $\hat{\tau}_V = \hat{\tau}_V^{\text{BC}} + \hat{\tau}_V^{\text{ISM}}$, and the fraction of this arising from dust in the ambient ISM, $\mu = \hat{\tau}_V^{\text{ISM}} / \hat{\tau}_V$, which is typically constrained in the range [0.3, 0.5] (e.g. Wild et al. 2007).

Recently, C13 proposed a new approach to account for the effects of inclination and dust/star geometry on the attenuation of galaxy SEDs. C13 incorporate the generic predictions of different types of sophisticated models of radiative transfer in dusty media into the two-component dust model of CF00. C13 show that these predictions can be subsumed in a quasi-universal relation between V -band attenuation optical depth in the diffuse ISM and shape of the attenuation curve. This relation, which accounts for the effects of dust/star geometry (including ISM clumpiness) and galaxy inclination, exhibits a steepening of the attenuation curve (from more starburst-like to more MW-like) at increasing dust optical depth (as

a consequence of either a rise in the amount of dust or a higher inclination). A fit to a wide range of models yields (see fig. 4 of C13)

$$n_V^{\text{ISM}} = \frac{2.8}{1 + 3\sqrt{\hat{\tau}_V^{\text{ISM}}}} \quad (\pm 25 \text{ per cent}), \quad (11)$$

where the typical scatter is indicated in parentheses. The above expression was derived at optical wavelengths. C13 show that, by adopting a wavelength-dependent exponent of the power law in equation (10), one can reproduce the generic predictions of radiative transfer models over the entire wavelength range from the near ultraviolet to the near-infrared, neglecting the 2175 Å bump. This can be achieved by rewriting equation (10) as

$$\hat{\tau}_\lambda^{\text{ISM}}(\theta) = \hat{\tau}_V^{\text{ISM}} \left(\frac{\lambda}{0.55 \mu\text{m}} \right)^{-n_\lambda^{\text{ISM}}}, \quad (12)$$

where the exponent of the power law is a linear function of wavelength,

$$n_\lambda^{\text{ISM}}(\hat{\tau}_V^{\text{ISM}}) = n_V^{\text{ISM}} + b(\lambda/\mu\text{m} - 0.55), \quad (13)$$

valid over the range $0.1 \leq \lambda \leq 2.5 \mu\text{m}$, and the coefficient b is given by

$$b = 0.3 - 0.05 \hat{\tau}_V^{\text{ISM}} \quad (\pm 10 \text{ per cent}). \quad (14)$$

We note that the above implementation of the C13 prescription does not require more parameters than the original CF00 model, since equations (11)–(14) depend only on the V -band attenuation optical depth in the diffuse ISM, $\hat{\tau}_V^{\text{ISM}}$ (see column ‘Adjustable parameters’ of Table 1).

Finally, we also implement in the BEAGLE tool the more sophisticated dust prescription proposed by C13 (referred to as ‘full model’ in Table 1), which allows one to explicitly express the dependence of attenuation in the ambient ISM on the viewing angle θ . This is achieved by associating stars in different age ranges with the thin-disc, thick-disc and bulge components of the versatile model of radiative transfer of Tuffs et al. (2004). In this case, we rewrite equation (10) as

$$\hat{\tau}_\lambda^{\text{ISM}}(\theta, t') = \begin{cases} \hat{\tau}_\lambda^{\text{thin}}(\theta) & \text{for } t' \leq t_{\text{thin}}, \\ \hat{\tau}_\lambda^{\text{thick}}(\theta) & \text{for } t_{\text{thin}} < t' \leq t_{\text{thick}}, \\ \hat{\tau}_\lambda^{\text{bulge}}(\theta) & \text{for } t_{\text{thick}} < t' \leq t_{\text{bulge}}, \\ 0 & \text{for } t' > t_{\text{bulge}}, \end{cases} \quad (15)$$

where $\hat{\tau}_\lambda^{\text{thin}}(\theta)$, $\hat{\tau}_\lambda^{\text{thick}}(\theta)$ and $\hat{\tau}_\lambda^{\text{bulge}}(\theta)$ are the attenuation curves at inclination θ for a thin stellar disc, thick stellar disc and bulge in the Tuffs et al. (2004) model. We stress that equations (11)–(15) affect only the dust attenuation arising from the diffuse ISM (the term $\hat{\tau}_\lambda^{\text{ISM}}$ in equation 8), while the birth-cloud component




Figure 1. Predictions of the IGM absorption models of Madau (1995, dotted lines) and Inoue et al. (2014, solid lines) for sources at different redshifts.

($\hat{\tau}_\lambda^{\text{BC}}$ in equation 8) is treated in C13 as in the original CF00 model. We also emphasize that the flexible modular structure of the BEAGLE tool enables one to easily substitute the predictions of any other model of radiative transfer for those of the Tuffs et al. (2004) one.

2.7 IGM absorption

We now consider the absorption of photons emerging from galaxies by gas in the IGM along the line of sight. This is modelled through the transmission function $T_{\lambda,z}^{\text{IGM}}$ of the IGM in equation (4). The IGM is composed primarily of hydrogen and helium and contains three main cloud components: the Ly α forest, Lyman-limit systems (LLS) and damped Ly α systems (DLA), in order of increasing H I column density and decreasing number density (e.g. Rauch 1998; Péroux et al. 2003; Wolfe, Gawiser & Prochaska 2005). The Ly α forest consists mainly of primordial gas, while LLS and DLA, which are thought to be associated with haloes and galaxies, are enriched in metals. Neutral hydrogen is the primary contributor to $T_{\lambda,z}^{\text{IGM}}$, which can thus be estimated from the H I column density distributions and number density evolutions of the different cloud components.

Madau (1995) proposed a first, simple analytic model to compute the contributions to $T_{\lambda,z}^{\text{IGM}}$ by H I in the Ly α forest and LSS, neglecting the much rarer DLA. This model includes photoelectric absorption of Lyman-continuum photons and blanketing by Lyman-series lines (from the combined absorption in many intervening clouds) of the background galaxy SED, which produces a characteristic ‘staircase’ profile of $T_{\lambda,z}^{\text{IGM}}$ as a function of observed wavelength (Fig. 1). Recently, Inoue et al. (2014) updated this widely used model by revising the H I column density distributions and number density evolutions of the Ly α forest and LSS and by adding the contribution to $T_{\lambda,z}^{\text{IGM}}$ by H I in DLA. In Fig. 1, we compare the IGM transmission functions predicted by this model (solid lines) and the Madau (1995) model (dotted lines), for background sources at different redshifts. For sources at $z \lesssim 5$, the model of Madau (1995) predicts a lower transmission than that of Inoue et al. (2014) at all rest wavelengths $\lambda_{\text{obs}}/(1+z) < \lambda_{\text{Ly}\beta}$, while at $z \gtrsim 6$, the trend is partially reversed. For sources at $z = 7$, the Inoue et al. (2014) model predicts that nearly all photons emitted at rest wavelengths $\lambda_{\text{obs}}/(1+z) < \lambda_{\text{Ly}\alpha}$ are absorbed by the IGM, while 1–7 per cent of the emission is transmitted at $\lambda_{\text{obs}}/(1+z) \in [\lambda_{\text{Ly}\beta}, \lambda_{\text{Ly}\alpha}]$ in the

Madau (1995) model. It is important to note that, because of the steepness of the IGM transmission curves in Fig. 1, these differences between the two models can translate into differences of up to ~ 1 mag in the observed colours of high-redshift galaxies (see fig. 8 of Inoue et al. 2014).

The analytic prescriptions of Madau (1995) and Inoue et al. (2014) for $T_{\lambda,z}^{\text{IGM}}$ are both limited by the fact that they pertain to averages over infinite numbers of sight lines, while individual galaxies probe single lines of sight through the IGM. In the future, we plan to account for variations in $T_{\lambda,z}^{\text{IGM}}$ along different lines of sight in the BEAGLE tool by appealing to the prescription of Harrison, Meiksin & Stock (2011).

2.8 Line-of-sight velocity distribution

The precise fitting of spectroscopic galaxy observations requires one to also account for the effects of stars and gas kinematics on the emergent SED. We implement this feature in the BEAGLE tool by introducing a flexible description of the line-of-sight velocity distribution (LOSVD) taken from van der Marel & Franx (1993, see also Gerhard 1993). This consists in decomposing the LOSVD into orthogonal functions via a Gauss–Hermite series, which enables the clean modelling of deviations from pure Gaussian line profiles (the orthogonality of the Hermite polynomials minimizing correlations among the adjustable coefficients). In terms of the standardized variable $x = (v - v_{\text{sys}})/\sigma$, we therefore express the LOSVD as

$$\mathcal{L}_x(v_{\text{sys}}, \sigma, h_3, h_4) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) \times [1 + h_3 H_3(x) + h_4 H_4(x)], \quad (16)$$

where v_{sys} is the galaxy systemic velocity, σ the velocity dispersion, and $H_3(x)$ and $H_4(x)$ the Hermite polynomials of order 3 and 4,

$$H_3(x) = \frac{1}{\sqrt{3}} (2\sqrt{2}x^3 - 3\sqrt{2}x), \quad (17)$$

$$H_4(x) = \frac{1}{\sqrt{24}} (4x^4 - 12x^2 + 3). \quad (18)$$

The adjustable coefficients h_3 and h_4 measure, respectively, asymmetric and symmetric deviations from pure Gaussian LOSVD. They can be determined separately for stars and gas.

2.9 Instrumental effects

So far, we have discussed the production of starlight in galaxies and its transfer through the interstellar and the intergalactic media. In this section, we address important instrumental effects altering the SED of a galaxy observed through a telescope: the spectral response and the spectroscopic flux calibration. In a Bayesian framework, accounting for instrumental effects is straightforward so long as these can be parametrized, as this amounts to adding ‘nuisance’ parameters.⁶ In the next subsections, we describe two instrumental effects incorporated in the BEAGLE tool: the instrumental spectral response (i.e. the line spread function) and the spectroscopic flux calibration. Other effects potentially biasing the interpretation of galaxy SEDs, such as the calibration of photometric zero-points,

⁶ A nuisance parameter is a parameter of no direct interest, but which must be included in the analysis to obtain reliable inference about the parameters of interest.

require a different approach. For example, a rigorous treatment of zero-point offsets can be performed by adopting a hierarchical Bayesian modelling approach. This and other instrumental effects will be investigated in future work.

2.9.1 Line spread function

The line spread function, noted $\mathcal{L}_{\lambda_{\text{obs}}}^{\text{LSF}}$, describes the instrumental spectral response, which relates the spectral flux density $F_{\lambda_{\text{obs}}}$ reaching the telescope at wavelength λ_{obs} (equation 4) to that effectively measured by the observer, $F'_{\lambda_{\text{obs}}}$, in such a way that

$$F'_{\lambda_{\text{obs}}} = F_{\lambda_{\text{obs}}} \otimes \mathcal{L}_{\lambda_{\text{obs}}}^{\text{LSF}}, \quad (19)$$

where the symbol \otimes indicates convolution. A common first-order approximation is to model the line spread function as a Gaussian function,

$$\mathcal{L}_{\lambda_{\text{obs}}}^{\text{LSF}}(\mu_{\text{LSF}}, \sigma_{\text{LSF}}) = \frac{1}{\sigma_{\text{LSF}}\sqrt{2\pi}} \times \exp\left[-\frac{1}{2}\left(\frac{\lambda_{\text{obs}} - \mu_{\text{LSF}}}{\sigma_{\text{LSF}}}\right)^2\right], \quad (20)$$

with mean $\mu_{\text{LSF}} = 0$ and dispersion $\sigma_{\text{LSF}}(\lambda_{\text{obs}}) = \lambda_{\text{obs}}/R$, where $R = \lambda_{\text{obs}}/\Delta\lambda_{\text{obs}}$ is the instrumental spectral resolution. In practice, this may provide only a poor approximation of the true line spread function, which depends on, for example, the intrinsic light profile of the source, the point spread function and the width of the spectroscopic aperture (slit or fibre). Ideally, therefore, the line spread function should be determined on an object-by-object basis. Another complication is that, for spectral analyses involving models based on observed stellar spectral libraries, the line spread function affecting the original stellar spectra should also be accounted for to perform meaningful comparisons with galaxy observations (see e.g. Koleva et al. 2009).

In the BEAGLE framework, we adopt a flexible parametrization of the line spread function, in which the parameters μ_{LSF} and σ_{LSF} in equation (20) are taken to be polynomials of adjustable degree in λ_{obs} , i.e.

$$\mu_{\text{LSF}}(\lambda_{\text{obs}}) = \sum_{i=0}^n c_i \mathcal{P}_i(\lambda_{\text{obs}} - \lambda_{\text{obs}}^c), \quad (21)$$

$$\sigma_{\text{LSF}}(\lambda_{\text{obs}}) = \sum_{i=0}^n d_i \mathcal{P}_i(\lambda_{\text{obs}} - \lambda_{\text{obs}}^c), \quad (22)$$

where n indicates the degree of the polynomial, c_i and d_i the coefficients of the polynomial, and \mathcal{P}_i the i th term of the polynomial expanded around the central wavelength λ_{obs}^c . The coefficients of the polynomials c_i and d_i are treated as nuisance parameters and marginalized out when computing the posterior probability distributions of the model parameters of interest. This ensures that uncertainties arising from an inaccurate knowledge of the line spread function are correctly propagated to the statistical constraints on interesting model parameters. Moreover, the possibility for $\mu_{\text{LSF}}(\lambda_{\text{obs}})$ to differ from zero enables one to straightforwardly account for any inaccurate wavelength calibration.

2.9.2 Spectroscopic flux calibration

A major challenge in the reduction of spectroscopic data is to achieve a reliable flux calibration, both absolute and relative. We refer here by ‘absolute’ to the calibration of an observed SED on

an absolute flux scale using observations of standard stars (and accounting for any required aperture correction). By ‘relative’, we refer to the calibration of the flux at any wavelength with respect another. The quality of the absolute calibration determines how well galaxy properties depending on total flux, such as stellar mass and star formation rate, can be evaluated. The quality of the relative flux calibration can potentially affect all galaxy properties, as it alters the shape of the SED. Obtaining an accurate relative flux calibration is challenging because of the contamination by several factors, such as the wavelength dependences of the point spread function and the galaxy light profile.

Two main approaches are generally considered to deal with inaccurate relative flux calibrations: one is to continuum-normalize the observed and model spectra before comparing them (e.g. Wolf et al. 2007; Spiniello et al. 2014); the other is to introduce a smooth correction function by which to multiply the model continuum, in order to bring it in agreement with the observed one (e.g. Kelson et al. 2000; Koleva et al. 2009). We implement both strategies in the BEAGLE tool to account for inaccurate relative flux calibrations. We allow either the observed and model spectra to be continuum-normalized, or a series of Legendre polynomials of adjustable order to be used to correct the model continuum shape, the coefficients of this series being treated as nuisance parameters (as in the case of the LOSVD in Section 2.8, we choose orthogonal polynomials to minimize correlation between adjustable parameters). This second, more flexible approach presents several advantages over a continuum normalization, which requires a fixed determination of the continuum, is highly sensitive to noise and erases valuable information contained in the spectral continuum shape. The subtlety is to select the smallest possible order of the series of Legendre polynomials able to account for the continuum-shape mismatch between model and observed spectra, while preserving informative spectral features such as the 4000 Å break and broad molecular absorption lines. This order is typically around 2–3.

3 MAIN CHARACTERISTICS OF BEAGLE

In Section 2, we presented the astrophysical ingredients used to model galaxy SEDs in the framework of the BEAGLE tool. These include, at the present time: the consistent modelling of the emission from stars and ionized gas by means of combined stellar population synthesis and photoionization codes; a prescription to account for the effects of changes in the α element-to-iron abundance ratio on stellar population spectra; different prescriptions for the star formation and chemical histories of galaxies, dust attenuation and IGM absorption; and simple analytic models of the LOSVD and instrumental effects. In this section, we describe the statistical approach at the basis of the BEAGLE tool, the combined implementation of the different model ingredients and the main output products.

3.1 Statistical approach

Our main goal in this paper is to develop a new-generation tool for the analysis of galaxy SEDs. In addition to enabling the production of mock catalogues of any spectroscopic and photometric galaxy properties, this tool must allow one to derive statistical constraints on a wide range of galaxy physical parameters from observed SEDs. To achieve this, we appeal to the modelling approach outlined in Section 2 to perform statistical inference on physical parameters from observed galaxy samples. A plethora of approaches have been proposed in the literature to perform such inference on the basis of

various statistical methods (e.g. minimum χ^2 , maximum likelihood, Bayesian; see Section 1). To select the approach optimal for our purpose, we require the constraints on galaxy physical parameters derived through an inference process to allow:

- (i) a rigorous propagation of observational errors into statistical constraints on model parameters, to compute realistic uncertainties in these parameters;
- (ii) a full characterization of correlations among model parameters, to deal with parameter degeneracies and multimodal solutions;
- (iii) a proper accounting of instrumental effects, to minimize the impact of instrumental systematics on inference products.

In addition to these requirements, the full exploitation of high-quality data will require complex models with multiple parameters. The choice of the optimal statistical approach therefore also depends on the computational and memory requirements of such multidimensional models.

In the framework of the BEAGLE tool, we adopt a Bayesian approach to perform inference on galaxy physical parameters from observed SEDs. This approach satisfies the above requirements through the characterization of prior and posterior probability distributions of model parameters, the consideration of nuisance parameters, and a precise framework to perform model comparisons and hierarchical analyses of multilevel observational constraints using multiparameter models. This kind of approach is now routinely employed to interpret large astrophysical data sets in the context of parameter spaces of very high dimensions ($>10^7$ parameters, e.g. Jasche, Leclercq & Wandelt 2015) by appealing to efficient computational techniques (e.g. Markov chain Monte Carlo, Hamiltonian Monte Carlo, Nested Sampling).

In practice, Bayes' theorem allows one to perform inference on a set of model parameters by combining information obtained from an experiment (through the 'likelihood function') with any prior knowledge about the parameters (through the 'prior probability distribution'). Past and current knowledge is therefore combined in the posterior probability distribution of the parameters, which can be expressed as (e.g. Jeffreys 1961)

$$\text{posterior} = \frac{\text{prior} \cdot \text{likelihood}}{\text{evidence}}, \quad (23)$$

and, mathematically, as

$$P(\Theta \mid \mathbf{D}, H) = \frac{P(\Theta \mid H) P(\mathbf{D} \mid \Theta, H)}{\int P(\Theta \mid H) P(\mathbf{D} \mid \Theta, H) d\Theta}, \quad (24)$$

where Θ refers to a set of parameters of a model (hypothesis) H , \mathbf{D} to a data set, and the denominator (i.e. the evidence, or marginal likelihood) is often written simply as $P(\mathbf{D} \mid H)$. For simplicity, in the remainder of this paper, we refer to the prior distribution of a model H with parameters Θ as $\pi(\Theta) \equiv P(\Theta \mid H)$, and to the likelihood function of a data set \mathbf{D} given a model H with parameters Θ as $\mathcal{L}(\Theta) \equiv P(\mathbf{D} \mid \Theta, H)$. This function depends on both the physical model and the statistical description of the noise in the data (e.g. Gaussian, Poisson). Given a prior probability distribution, which reflects our belief about the model parameters before considering the data set \mathbf{D} , equation (24) enables rigorous statistical constraints on these parameters from observations at any S/N ratio. These constraints do not require any assumption about the shape of the posterior probability distribution (unlike, e.g. the Gaussian shape implicitly assumed when estimating confidence intervals with the criterion $\Delta\chi^2 < 1$). We note that, as described in Section 2.9, instrumental systematics can also be incorporated in a Bayesian approach, by means of nuisance parameters, which can be

marginalized out when computing the posterior probability distributions of the model parameters of interest (see, e.g. the treatment of nuisance parameters in Planck Collaboration 2015).

Another main interest of equation (24) is the possibility to fully characterize complex, non-linear correlations among model parameters, an achievement often ignored in statistical approaches focusing on simple point-wise (e.g. minimum χ^2) estimates of best-fitting model parameters. Such correlations can reveal parameter degeneracies and lead to multimodal solutions, i.e. to different parameter combinations providing similarly good fits to the data (which one may then try to break by appealing to new observables; see, e.g. fig. 19 of Planck Collaboration 2015). The problem is particularly acute, for example, in the context of galaxy redshift determinations from deep photometric observations (see Section 4.3). A Bayesian analysis provides a rigorous solution to this problem, through the methodical comparison of models populating the different modes of a posterior probability distribution (Section 5.2). Such statistically driven comparison allows the selection of the best subset of models able to account for a data set, beyond the selection of simply the best set of model parameters.

It is worth discussing briefly the role of the prior probability distribution entering equation (24). In most situations, the data entering equation (24) through the likelihood function will be informative enough for the prior probability distribution to have a negligible impact on inference results. In the case of poorly informative data, the influence of the prior can be studied by testing different choices and by comparing the prior and posterior probability distributions of model parameters.⁷ As an example, the analysis of low-S/N photometric observations does not allow one to put strong constraints on all model parameters, but the presence of a prior term in equation (24) allows one to properly incorporate our ignorance about the value of these parameters into their posterior probability distribution (see Section 4.4 below). The current version of the BEAGLE tool allows the user to choose, for any adjustable parameter of the model, between uniform (linear or logarithmic), Gaussian (linear or logarithmic) or Cauchy (to allow for broader tails) prior distributions. Additional prior distributions will be made available in the future. Finally, we note that prior probability distributions can also be used in a hierarchical way in Bayesian analyses, by adopting 'hyper-priors' with their own 'hyper-parameters' to describe the parameter priors. This enables a multilevel analysis, which reduces the final uncertainties in model parameters by allowing data to share information with one another (see, e.g. Sonnenfeld et al. 2015).

3.2 Model implementation

We now describe our strategy to implement the various astrophysical ingredients of BEAGLE (Section 2) into the Bayesian statistical framework outlined in Section 3.1, with the aim of creating a flexible, physically motivated tool for the analysis of galaxy SEDs. The flexibility of this tool is crucial to overcome a major challenge faced by any SED modelling approach: to adapt to a wide variety of data (from photometric to spectroscopic; from high- to low-S/N; at diverse spectral resolution over different wavelength ranges) without neither sacrificing model completeness nor over-fitting data with too many uncontrolled adjustable parameters.

⁷ We refer the interested user to Loredó (2012) for an interesting discussion about common misconceptions regarding Bayesian methods, including the role of prior distributions.

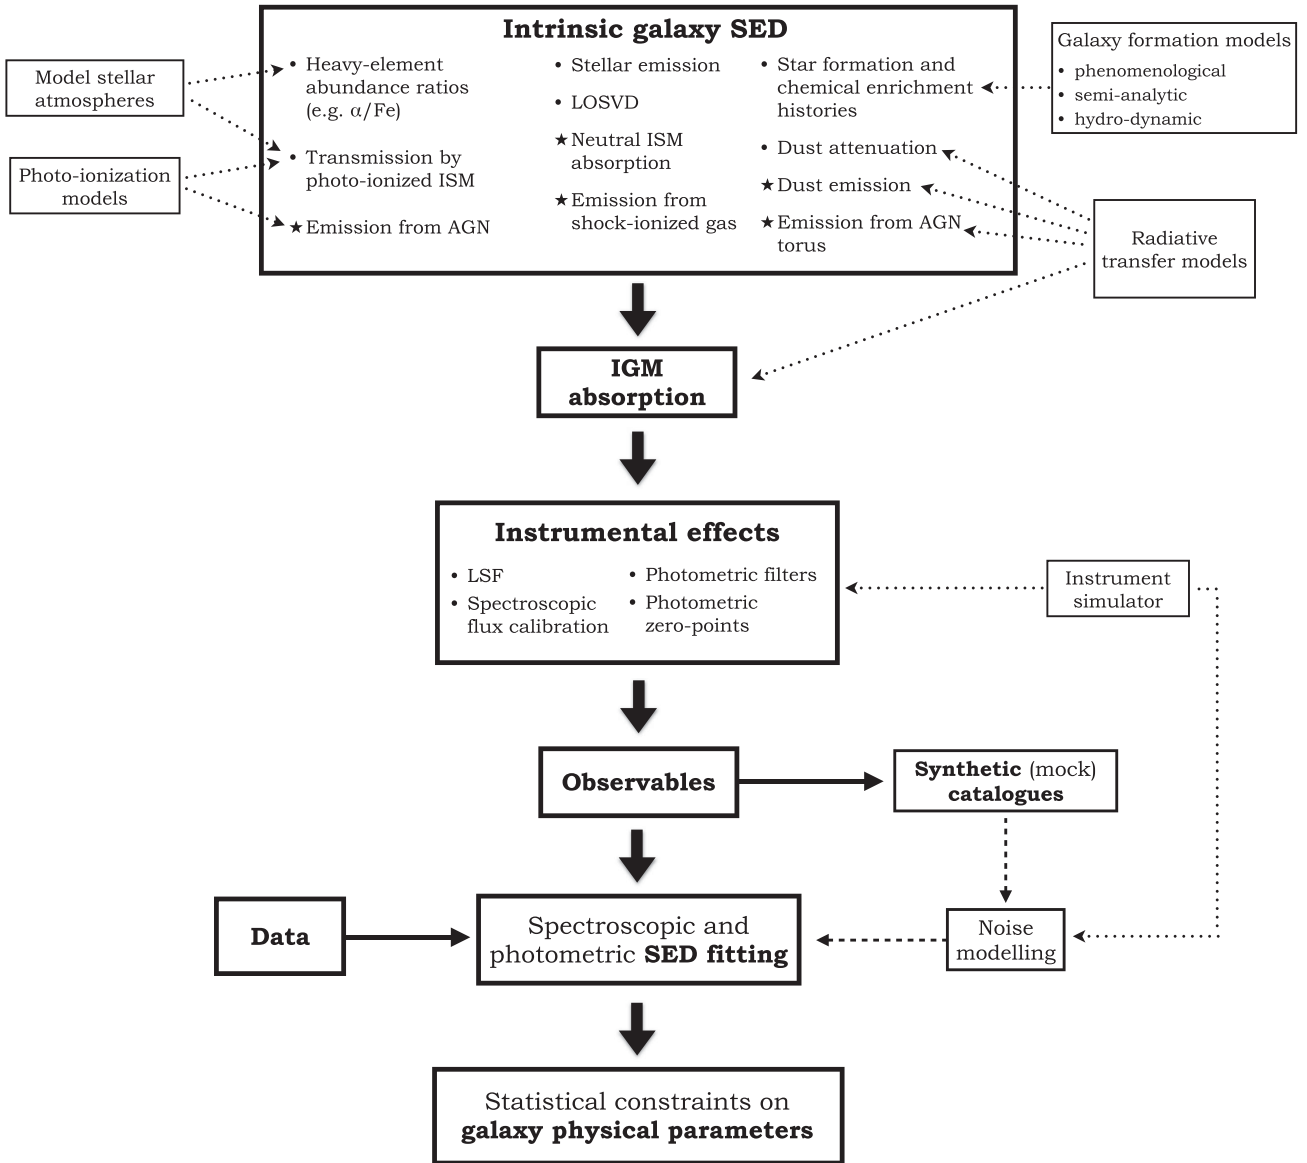


Figure 2. Workflow diagram showing the different building blocks of the BEAGLE tool. The astrophysical ingredients indicated with stars are currently being implemented and will soon be available. Dotted arrows indicate how external models can be incorporated into BEAGLE to inform the various ingredients. See Section 2 for details.

Fig. 2 shows our implementation in a fully integrated tool of the different astrophysical ingredients described in Section 2. This workflow diagram highlights the different steps leading to either the production of synthetic catalogues of galaxy SEDs or the quantitative interpretation of observed SEDs with model ones. The ingredients used to define an intrinsic galaxy SED are shown at the top, and the external inputs used to inform these ingredients on the sides (e.g. galaxy formation and spectral evolution models). The light emerging from the model galaxy is then processed through IGM absorption and instrumental effects to produce observables comparable with data.

As a complement to Fig. 2, Table 2 summarizes the adjustable parameters used to compute galaxy observables in the BEAGLE tool. To achieve the flexibility required to handle different types of (both true and pseudo) observations, we allow each adjustable parameter to be either ‘free’, ‘fixed’ or ‘dependent’. Free parameters are

drawn from prior probability distributions (e.g. uniform, normal, log-normal), both when building synthetic catalogues of galaxy SEDs and when fitting an observation, in which case they enter equation (24) as elements of array Θ . Alternatively, an adjustable parameter can be fixed to a default (or standard) value, both when producing synthetic observations and when fitting an observed SED. This is useful to preserve the full coherence of a multiparameter model when dealing with observations with low constraining power, such as limited-band photometry at moderate S/N. Finally, the value of an adjustable parameter can also be set to depend on other adjustable parameters, for example, through an analytic relation.

We note that, from a Bayesian point view, the consideration of an adjustable parameters as free, fixed or dependent corresponds simply to different choices of prior probability distributions. The prior probability distribution of a fixed parameter is a Dirac δ

Table 2. Summary of the adjustable parameters available to build the intrinsic SED of a galaxy in the BEAGLE tool. This table does not include the adjustable parameters used to describe instrumental effects (see Section 2.9).

Physical module	Adjustable parameters	Default range	Physical meaning
Star formation and chemical enrichment histories	$t_{\text{SSP}}/\text{Gyr}$	$[0.01, t_{\text{U},z}^a]$	Age of stars when approximating a galaxy by a simple stellar population ^b
	$t_{\text{start}}/\text{Gyr}, t_{\text{end}}/\text{Gyr}$	$[0, t_{\text{U},z}^a]$	Start/end look-back times of an arbitrary number of star formation periods ^c
	M/M_{\odot}	$[0, 10^{12}]$	Mass of stars formed during a star formation period
	$\tau_{\text{SFR}}/\text{Gyr}$	$[0.01, 2 \times t_{\text{U},z}^a]$	Star formation time-scale during a star formation period (for analytic star formation histories, such as exponentially declining, delayed, rising, etc.)
	$\log(Z/Z_{\odot})$	$[-2, 0.25]$	Metallicity of stars formed during a star formation period
	$[\alpha/\text{Fe}]$	$[0, 0.4]$	α/Fe ratio of stars formed during a star formation period (relative to Solar) ^d
	$\Delta t_{\text{SFR}}/\text{yr}$	$[10^7, 10^8]$	Duration of the current episode of star formation
	$\log(\psi_{\text{S}}/\text{yr}^{-1})$	$[-12, -7]$	Specific star formation rate (averaged over Δt_{SFR})
	$\log(Z_{\text{young}}/Z_{\odot})$	$[-2, 0.25]$	Metallicity of stars younger than 10 Myr
	z	$[0, 15]$	Redshift of observation
Dust attenuation	$z_{\text{form}} (t_{\text{form}})$	$[z, 50]$	Formation redshift (look-back time) of the first stellar generation
	$\hat{\tau}_V$	$[0, 5]$	Vband attenuation optical depth ^{e,f}
	μ	$[0, 1]$	Fraction of $\hat{\tau}_V$ arising from the dust in the diffuse ISM ^{e,f}
	θ/deg	$[0, 90]$	Galaxy inclination ^f
	$t_{\text{thin}}/\text{Gyr}$	$[0.1, t_{\text{form}}]$	Stars with ages $t \leq t_{\text{thin}}$ pertain to thin stellar disc ^f
	$t_{\text{thick}}/\text{Gyr}$	$[t_{\text{thin}}, t_{\text{form}}]$	Stars with ages $t_{\text{thin}} \leq t \leq t_{\text{thick}}$ pertain to thick stellar disc ^f
	$t_{\text{bulge}}/\text{Gyr}$	$[t_{\text{thick}}, t_{\text{form}}]$	Stars with ages $t_{\text{thick}} \leq t \leq t_{\text{bulge}}$ pertain to stellar bulge ^f
Nebular emission	$\log U$	$[-4, -1]$	Effective galaxy-wide ionization parameter ^g
	$\log(Z_{\text{ISM}}/Z_{\odot})$	$[-2, 0.25]$	Effective galaxy-wide interstellar metallicity ^g
	ξ_{d}	$[0.1, 0.5]$	Effective galaxy-wide dust-to-metal mass ratio ^g
Kinematics ^h	$v_{\text{sys}}/\text{km s}^{-1}$	$[0, 10^4]$	Systemic velocity
	$\sigma/\text{km s}^{-1}$	$[0, 400]$	Velocity dispersion
	h_3	$[-1, 1]$	Coefficient of the third-order Hermite polynomial ⁱ
	h_4	$[-1, 1]$	Coefficient of the fourth-order Hermite polynomial ⁱ

Notes. ^aThis indicates the age of the Universe at redshift z , assuming a cosmological model and a fixed set of cosmological parameters.

^bAlthough we do not favour the adoption of SSPs to describe galaxy star formation and chemical enrichment histories, we have implemented this model to allow an easier comparison with previous analysis tools adopting SSPs (and for spectral analyses of individual star clusters).

^cAs defined in Section 2.5.

^dCoelho et al. (2007); Coelho (2014); Walcher et al. (2009).

^eCF00.

^fC13.

^gCharlot & Longhetti (2001); Pacifici et al. (2012); Gutkin et al. (2016).

^hDifferent kinematic parameters can be used to describe the LOSVD of stellar and nebular emission.

ⁱvan der Marel & Franx (1993); Gerhard (1993).

function, as is that of a dependent parameter, only conditional in this case on the values of other parameters. In the current version of the BEAGLE tool, the prior probability distribution of each model parameter is independent of the others, which makes the generalized Bayesian formulation of the above three parameter classes not easy to implement. In the future, we will expand the flexibility of prior probability distributions to allow the choice of any type of joint and conditional probability distributions. This will enable the incorporation in the inference process of a priori relations between different galaxy physical parameters (e.g. the mass–metallicity relation), accounting at the same time for the scatter about such relations.

The number of adjustable parameters of the BEAGLE tool in Table 2 is fairly large. Since many of these parameters cannot be easily set to standard values nor related to other parameters, in typical situations, many will have to be considered as free parameters. In practice, the actual number of free parameters will be chosen on a case by case basis as a function of the available data. It is worth briefly pausing here to comment on a common misconception regarding the influence of the number of free parameters on a statistical analysis. It is often stated that, when fitting a model to a data set, the number of free parameters should not exceed that of (independent) data

points. In reality, this statement is true only for linear models, i.e. models depending linearly on free parameters (e.g. polynomials, linear least-squares). In the more general case of non-linear models, such as in the BEAGLE tool (Fig. 2), the rule does not apply. While simple models should generally be preferred over more complex ones at equal predictive power, in some situations, having more free parameters than data points may be recommended to account for uncertainties about these parameters.⁸

Another potential source of ambiguity relates to the parametrization of galaxy physical properties in SED fitting tools. In standard spectral analyses, galaxies are often characterized simply in terms of stellar mass, age and metallicity. The stellar mass of a galaxy reflects the integral of the star formation history.⁹ Thus, a same galaxy stellar mass could result from an infinite number of different star

⁸See the interesting discussion about the relative numbers of free parameters versus data points at <https://jakevdp.github.io/blog/2015/07/06/model-complexity-myth/>

⁹A fraction of the stellar mass is actually returned to the ISM through stellar winds and SNe explosions during the evolution of the stellar population. This fraction, and hence the mass currently locked into stars, is computed using stellar population synthesis models and is recorded in BEAGLE as the quantity

formation histories. This implies that, in turn, the distributions of stellar age and metallicity will depend on the specific star formation and chemical enrichment histories of a galaxy. The definition of global galaxy age and metallicity at fixed stellar mass is therefore ambiguous. In this context, it is customary to define light- and mass-weighted ages and metallicities, which are also integral quantities computed from the star formation and chemical enrichment histories of a galaxy. Finally, we note that, in many spectral analyses, the age of a galaxy is defined as the age of the oldest stars in that galaxy. This quantity is most relevant to studies of young galaxies near the reionization epoch, as it is otherwise difficult to constrain, the oldest stars tending to be out-shined by younger ones.

3.3 Output products

In this section, we briefly describe the output products of the `BEAGLE` tool (we refer the reader to the code manual for a more detailed description of these products).¹⁰ These are organized in FITS files with multiple extensions, each extension containing information about a physical module (e.g. star formation and chemical enrichment histories, dust). We use a standard output format to produce synthetic SED catalogues, in the sense that the properties of different simulated galaxies are written on different rows. We describe in more detail here the output products of a model fit to an observed galaxy SED, which are more specific to the statistical approach inherent in the `BEAGLE` tool.

In the output FITS extension produced for a given physical module by a spectral fit, the entries on each row are the properties predicted by a model with free parameters drawn from the posterior probability distribution of equation (24) using a dedicated Bayesian algorithm (e.g. `MULTINEST`; see Section 4.2). The posterior probability distributions of the parameters themselves are reported in a separate extension of the output FITS file. This generalized format allows one to interpret in a probabilistic way the constraints on not only the model free parameters, but also all other physical quantities included in the output files, which we refer to as ‘derived quantities’. For example, the uncertainties in the mass-weighted age and metallicity of a galaxy can be computed from the values of these quantities across all rows of the corresponding output file, weighted by the posterior probability value associated with each row. This approach allows one to easily compute the probability distribution of any theoretical or observable quantity predicted by the model (e.g. mass-weighted age, rate of ionizing photons, ultraviolet spectral slope, broad-band magnitude, emission-line luminosity, absorption-line strength), as resulting from the posterior probability distribution of the adjustable parameters.

Another advantage of the above approach is that it simplifies the adoption of posterior predictive checks to quantify how well model predictions match the data in a given observable (see Section 5.1 for an example). This can be readily extended to the study of residuals, which in this context are no more a point-wise estimate of the difference between model predictions and observations, but rather a probabilistic distribution of such differences. We believe

M_* , which differs from the model parameter M reflecting the integral of the star formation history.

¹⁰ The `BEAGLE` tool will be released in the near future as an open-source project. To be informed about the code release, please visit and register at <http://www.jacopochevallard.org/beagle>. In the meantime, interested users should contact the corresponding author of this paper.

that this will be a powerful means of identifying model failures and driving the development of better models with the `BEAGLE` tool.

4 PHOTOMETRIC SED ANALYSIS WITH `BEAGLE`

In this section, we present a first application of the `BEAGLE` tool to interpret the broad-band SEDs of a published sample of about 10^4 galaxies at redshifts $0.1 \lesssim z \lesssim 8$. The observational challenge of gathering large spectroscopic samples of galaxies at high redshifts is forcing much progress in the field of galaxy formation and evolution to rely on photometric surveys. Interpreting broad-band SEDs is therefore a common task in galaxy evolution studies. Historically, the derivation of galaxy physical parameters from broad-band galaxy SEDs has often been decoupled from that of redshift. In fact, redshift is considered to be the most robust quantity that can be constrained from broad-band photometric data, while determinations of galaxy physical parameters are expected to be more model-dependent (e.g. Dahlen et al. 2013; Mobasher et al. 2015). Dedicated photometric-redshift codes generally rely on a small library of ‘representative’ SEDs of different types of galaxies, either built using spectral evolution models (e.g. `EAZY` Brammer, van Dokkum & Coppi 2008) or consisting of a small number of observed galaxy SEDs (e.g. `BPZ` Benítez 2000; `LEPHARE`, Arnouts et al. 1999; Ilbert et al. 2006). Such libraries enable the extraction of only limited information about physical parameters from photometric SEDs of galaxies.¹¹ In contrast, most codes designed to constrain galaxy physical parameters from photometric data require an independent determination of redshift, often estimated using one of the above photometric-redshift codes (e.g. `FAST`, Kriek et al. 2009; `CIGALE`, Burgarella, Buat & Iglesias-Páramo 2005; Noll et al. 2009).

In the `BEAGLE` tool, we follow an alternative approach, similar that adopted by Acquaviva, Raichoor & Gawiser (2015, see also `HYPERZ`, Bolzonella et al. 2000; Pozzetti et al. 2007), and consider redshift just as an additional model parameter to be constrained along with the other physical parameters when fitting a galaxy SED. As we shall see below, a major advantage of this approach is to allow a rigorous propagation of the uncertainty on photometric redshift to the statistical constraints on other galaxy physical parameters, accounting at the same time for any potential correlation between redshift and other parameters. In addition, our statistical approach, based on the `MULTINEST` algorithm (a Bayesian analysis tool based on the Nested Sampling algorithm of Skilling et al. 2006; see Appendix A and Feroz, Hobson & Bridges 2009), can naturally deal with multimodal solutions, i.e. with the occurrence of different combinations of parameters yielding similar posterior probability distributions (Section 3.1). This can often be the case in determinations of photometric redshifts of faint sources (e.g. Edmondson, Miller & Wolf 2006). We now describe the photometric sample of distant galaxies we appeal to for our study (Section 4.1) and the results of SED fitting of this sample using the `BEAGLE` tool (Section 4.2 for redshift and Section 4.4 for other physical parameters).

¹¹ Other approaches to estimate photometric redshifts also exist: ‘data-driven’ methods are based on the application of machine-learning techniques to determine the relation between galaxy colours and redshift (e.g. Collister & Lahav 2004; Hogan, Fairbairn & Seeburn 2015), while ‘clustering-based’ methods are based on the redshift evolution of spatial correlations between galaxies (e.g. Schmidt et al. 2013; Rahman, Ménard & Scranton 2016). We do not discuss such approaches here, as they typically do not provide any information about galaxy physical parameters.




Figure 3. Relation between total signal-to-noise ratio, computed by summing in quadrature the S/N of all bands with available measurements, and observed ACS/WFC *F850LP* magnitude, along with their marginal distributions, for all galaxies in the UVUDF catalogue.

4.1 Galaxy sample

To illustrate our approach to broad-band SED fitting, we appeal to the photometric galaxy catalogue assembled by Rafelski et al. (2015, hereafter ‘UVUDF’ catalogue). This consists of (up to) 11 bands at ultraviolet to near-infrared wavelengths for 9927 galaxies in the *Hubble Ultra Deep Field* (Beckwith et al. 2006). The ultraviolet data come from observations in three *HST*/WFC3 filters, *F225W*, *F275W* and *F336W* (observations in Teplitz et al. 2013), reanalysed by Rafelski et al. (2015) to improve the photometric and astrometric calibrations. The optical data come from observations in four ACS filters, *F435W*, *F606W*, *F775W* and *F850LP* (imaging in Beckwith et al. 2006), and the near-infrared ones in four WFC3 bands, *F105W*, *F125W*, *F140W* and *F160W*, from the UDF09, UDF12 (Oesch et al. 2010a,b; Bouwens et al. 2011; Ellis et al. 2013; Koekemoer et al. 2013) and CANDELS GOODS-S programs (Grogin et al. 2011; Koekemoer et al. 2011). To maximize the depth of their catalogue, Rafelski et al. (2015) combine, pixel by pixel, the four optical and four near-infrared images to create a ‘detection image’. They identify sources in this image by running a standard extraction algorithm with different settings, varying the detection and deblending threshold parameters, then merging in a single catalogue the sources identified with the different settings. Finally, to produce a homogeneous source catalogue, they extract PSF-corrected aperture-matched photometry in all bands, and compute the total (Kron 1980) flux in each band by applying an aperture correction to the isophotal flux.

We show in Fig. 3 the relation between total signal-to-noise ratio $(S/N)_{\text{tot}}$, computed by summing in quadrature the S/N of all bands with available measurements, and observed ACS/*F850LP* magnitude, along with the marginal distributions of these quantities, for all UVUDF galaxies. We select the *F850LP* band as it is the one with the largest fraction (9919/9927) of detected objects (defined as entries in the catalogue with positive measurements of both flux and flux error). The catalogue peaks around $m_{\text{AB}} = 29.5$, with a typical $(S/N)_{\text{tot}}$ in the range $\sim[10, 30]$ and a long tail of bright objects with $(S/N)_{\text{tot}} \gtrsim 100$. For a small fraction (about 1.7 per cent) of the galaxies in the UVUDF catalogue, reliable spectroscopic redshifts are available from the literature. In Fig. 4, we compare the frequency distribution of redshift for this spectroscopic subsample (in blue)

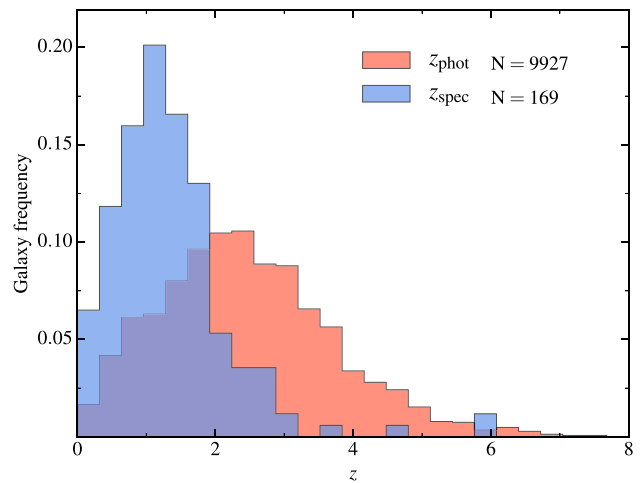


Figure 4. Frequency distribution of redshift for the full UVUDF photometric catalogue (in red) and for galaxies with spectroscopic redshift measurements (in blue). The photometric redshifts are computed with the BEAGLE tool (see Section 4.3).

to that for the full photometric sample (in red), using photo- z estimates obtained with the BEAGLE tool (see Section 4.2 below). Aside from the low number of spectroscopic detections, Fig. 4 highlights the different redshift distributions of the two samples. The spectroscopic sample peaks around $z \sim 1$ and has very few galaxies at $z \gtrsim 3$, while the photometric sample displays a broader distribution, which extends out to $z \gtrsim 6$.

4.2 Modelling approach

To analyse the photometric SEDs of galaxies in the UVUDF catalogue with the BEAGLE tool, we adopt the following prescriptions for the adjustable parameters listed in Table 2. For simplicity, we describe the star formation histories of model galaxies as delayed exponential functions over a single star formation period, $\psi(t') \propto t' \exp(-t'/\tau_{\text{SFR}})$ for $t' \leq t_{\text{U},z} - t_{\text{form}}$, and assume that all stars in a given galaxy have the same metallicity Z . We account for the stochastic nature of star formation by adding a ‘current’ burst of star formation to describe stars assembled in the last $\Delta t_{\text{SFR}} = 10$ Myr of a galaxy star formation history. The strength of the burst component is parametrized in terms of the specific star formation rate ψ_s . We compute the photoionization of interstellar gas by young stars in the burst as described in Section 2.4. For simplicity, in the absence of spectroscopic constraints, we adopt the same metallicity for the interstellar gas as for the stars ($Z_{\text{ISM}} = Z$), a fixed (intermediate) dust-to-metal mass ratio ($\xi_d = 0.3$), and the following relation between Z_{ISM} and $\log U_s$, derived from the analysis of SDSS galaxies by Brinchmann et al. (2004, and Carton et al. 2016, in preparation):

$$\log U_s = \begin{cases} -3.638 + 0.055x & \text{for } x \leq -0.04, \\ -3.640 & \text{for } x > -0.04, \end{cases} \quad (25)$$

where $x = \log(Z_{\text{ISM}}/Z_{\odot})$. We use the prescription of C13 to describe dust attenuation by dust, fixing the fraction of the dust optical depth arising from the diffuse ISM at $\mu = 0.4$ (Wild et al. 2011). Together with the mass M of stars formed, which provides the absolute scaling of the luminosity, and the redshifts of observation and formation, z and z_{form} , these represent seven free parameters, in the sense defined in Section 3.2. We also include absorption by the

Table 3. Prior distributions of the seven free parameters adopted to interpret the photometric SEDs of UVUDF galaxies in Sections 4.2 and 4.4. A galaxy is assumed to form over a single star formation period extending from z_{form} to z (see Table 2 for a description of the different parameters).

Parameter	Prior range
$\log(M/M_{\odot})$	[5, 12]
$\log(\tau_{\text{SFR}}/\text{yr})$	[7, 10.5]
$\log(Z/Z_{\odot})$	[-2.2, 0.25]
$\log(\psi_{\text{S}}/\text{yr}^{-1})$	[-14, -7]
$\hat{\tau}_V$	[0.001, 5]
z	[0, 15]
z_{form}	[z , 50]

IGM using the prescription of Inoue et al. (2014, see Section 2.7 above).

To compute the posterior probability distribution $P(\Theta | \mathbf{D}, H)$ of the free model parameters favoured by the observations of a given UVUDF galaxy, we must specify the likelihood function $\mathcal{L}(\Theta)$ and prior distribution $\pi(\Theta)$ entering the right-hand side of equation (24). We adopt flat linear prior distributions for z , z_{form} , $\log M$, $\log \tau_{\text{SFR}}$, $\log Z$ and $\log \psi_{\text{S}}$, and a flat logarithmic prior distribution for $\hat{\tau}_V$, within the ranges reported in Table 3. We do not introduce any luminosity function-based prior in this analysis, since we wish to investigate the presence of multimodal solutions in the posterior distribution of model parameters in the absence of any external information. To build the likelihood function, we model the observed fluxes \mathbf{y} of the UVUDF galaxy as a multivariate Gaussian random variable, with mean given by the prediction $\hat{\mathbf{y}}(\Theta^k)$ of our model for a set of parameters $\Theta^k = [M, \tau_{\text{SFR}}, Z, \psi_{\text{S}}, \hat{\tau}_V, z, z_{\text{form}}]^k$, and noise described by a diagonal covariance matrix Σ .¹² We therefore write the likelihood function of that galaxy as

$$-2 \ln \mathcal{L}(\Theta^k) = \sum_i \left[\frac{y_i - \hat{y}_i(\Theta^k)}{\sigma_i} \right]^2, \quad (26)$$

where the summation index i runs over all bands observed, even in the absence of detection (i.e. with negative flux after an uncertain background subtraction), and the σ_i s are the diagonal elements of matrix Σ . Bands with no detection are of crucial importance to track the absorption of radiation blueward of hydrogen Ly α by the IGM (and blueward of the Lyman limit by the ISM) at high redshift.

The parameter σ_i in equation (26) is not purely the observational error. In fact, when fitting broad-band photometry with spectral evolution models, it is customary to introduce an additional error term to account for uncertainties potentially unaccounted for in the observed fluxes (linked to, e.g. background subtraction, flux calibration, aperture effects) and for (unquantified) systematic uncertainties in model predictions (e.g. Brammer et al. 2008; Dahlen et al. 2013; Acquaviva et al. 2015). We account for this effect by

¹² We note that the use of a Gaussian likelihood function to describe data originating from the difference between two Poisson processes (i.e. the difference between source+background and background counts on a detector) is justified only in the limit of large counts, in which case the Poisson distributions can be approximated by Gaussian ones. The presence of objects with negative fluxes in the UVUDF catalogue suggests that this approximation may not be valid for the faintest sources. In those cases, a better approach would be to introduce a background term in the likelihood function and use a Poisson distribution to describe the combined source+background signal (e.g. Thompson 1999). Unfortunately, the information in the UVUDF catalogue is not sufficient to allow us to perform such an analysis.

adding a 2 per cent relative error in quadrature to the fluxes in all photometric bands, i.e. we write $\sigma_i = \sqrt{\sigma_{\text{obs}}^2 + (\sigma_0 y_i)^2}$, where σ_{obs} is the observational error and $\sigma_0 = 0.02$.¹³ Although this additional error term does not correct for potential biases originating from systematic uncertainties, it reduces their impact on the results by widening the posterior probability distribution of model parameters. We note that adopting $\sigma_0 = 0.02$ translates into increasing flux errors in each band by factors of 1.08, 1.16 and 1.28 for S/N = 20, 30 and 40, respectively. We can estimate the typical S/N of a UVUDF band starting from Fig. 3, by considering the quantity $(\text{S/N}) \sim (\text{S/N})_{\text{tot}}/\sqrt{N_{\text{bands}}}$, where N_{bands} varies from 8 to 11 depending on the galaxy. This leads to $\langle \text{S/N} \rangle \lesssim 10$ for most galaxies, implying that only the few brightest UVUDF galaxies might be significantly affected by the addition of this error term. These also tend to be the lowest redshift galaxies with spectroscopic redshifts (Fig. 4). Without the additional error term, therefore, the results for the high- and low-S/N subsamples could be dominated by different error sources: (uncontrolled) systematics originating from both data and models for the former, and S/N ratio of the data for the latter. In this case, the conclusions drawn from, for example, the comparison between spectroscopic and photometric redshifts for bright galaxies, would be irrelevant to the fainter sample.

To compute the posterior probability distribution of the model parameters in Table 3 (and of other derived quantities; see Section 3.3) with the BEAGLE tool, we adopt MULTINEST, an efficient Bayesian inference tool based on the Nested Sampling algorithm of Skilling et al. (2006, see Appendix A for detail). We start by focusing on photometric redshift estimates, with the purpose of comparing these with estimates derived using the dedicated photometric redshift codes BPZ and EAZY in the UVUDF catalogue. To this goal, we adopt as a measurement of z the mean of the marginal posterior probability distribution provided by the BEAGLE tool. In the case of multiple-redshift (i.e. multimodal) solutions, we identify the mode with highest local evidence, as computed using MULTINEST, and take the posterior mean within that mode to be a measurement of z (see Section 5.2 for a more detailed discussion of multimodal solutions).

4.3 Photometric redshifts

In their original study, Rafelski et al. (2015) estimate the photometric redshifts of all UVUDF sources by appealing to two standard codes, BPZ (Benítez 2000) and EAZY (Brammer et al. 2008). Both codes rely on a small set of template galaxy SEDs computed using the PEGASE population synthesis code (Fioc & Rocca-Volmerange 1997), although with different prescriptions for the contamination of broad-band fluxes by nebular emission lines. The BPZ code relies on the original prescription of PEGASE for nebular emission, while the EAZY code incorporates a simplified model relating the Ly α , H α , H β , H γ , [O II] $\lambda 3727$ and [O III] $\lambda \lambda 4959, 5007$ line luminosities to the star formation rate (see Brammer et al. 2011). The BPZ code also includes a Bayesian prior, based on previously measured galaxy luminosity functions, to help constrain redshift estimates in cases of multiple solutions.

Rafelski et al. (2015) use the subsample of 169 galaxies with spectroscopic redshifts in the UVUDF catalogue to assess the quality of photometric redshifts derived using the BPZ and EAZY codes, although this quality check is limited by construction to redshifts

¹³ After experimenting with values in the range $0.01 \leq \sigma_0 \leq 0.04$, we find that fixing $\sigma_0 = 0.02$ minimizes the dispersion between z_{BEAGLE} and z_{SPEC} , as measured by the quantity σ_{NMAD} (equation 27).

Table 4. Bias, normalized median absolute deviation and outlier fraction for the comparisons between BEAGLE-derived photometric redshifts and spectroscopic (169 galaxies) and BPZ- and on EAZY-derived photometric redshifts (9927 galaxies) from the UVUDF catalogue.

Comparison	Bias ^a	σ_{NMAD}^b	Outlier fraction ^c
$z_{\text{BPZ}} - z_{\text{spec}}$	0.005	0.028	2.4 per cent
$z_{\text{EAZY}} - z_{\text{spec}}$	-0.013	0.030	5.9 per cent
$z_{\text{BEAGLE}} - z_{\text{spec}}$	0.007	0.047	7.1 per cent
$z_{\text{BEAGLE}} - z_{\text{BPZ}}$	0.005	0.042	8.3 per cent
$z_{\text{BEAGLE}} - z_{\text{EAZY}}$	0.039	0.058	7.3 per cent

Notes. ^aMedian of the distribution of $(z - z_{\text{ref}})/(1 + z_{\text{ref}})$, where z_{ref} stands for z_{spec} , z_{BPZ} or z_{EAZY} , depending on the comparison.

^bComputed using equation (27) by substituting z_{spec} , z_{BPZ} or z_{EAZY} for z_{ref} , depending on the comparison.

^cFraction of galaxies with $|\Delta z|/(1 + z_{\text{ref}}) > 0.15$ for the comparisons $z_{\text{BPZ}} - z_{\text{spec}}$, $z_{\text{EAZY}} - z_{\text{spec}}$ and $z_{\text{BEAGLE}} - z_{\text{spec}}$, and with $|\Delta z|/(1 + z_{\text{ref}}) > 0.25$ in the other two cases.

$z \lesssim 3$ (Fig. 4). They compute for both BPZ and EAZY the scatter of the difference Δz between photometric redshift and spectroscopic redshift, which they quantify through the normalized median absolute deviation

$$\sigma_{\text{NMAD}} = 1.48 \times \text{median} \left| \frac{\Delta z - \text{median}(\Delta z)}{1 + z_{\text{ref}}} \right|, \quad (27)$$

where the reference redshift is the spectroscopic one ($z_{\text{ref}} = z_{\text{spec}}$), and the factor 1.48 ensures that σ_{NMAD} be equal to the standard deviation for a Gaussian distribution. Rafelski et al. (2015) find $\sigma_{\text{NMAD}} = 0.028$ with the BPZ code and $\sigma_{\text{NMAD}} = 0.030$ with the EAZY code. Following Brammer et al. (2008), they define the fraction of outliers (OLF) as the fraction of galaxies with $|\Delta z|/(1 + z_{\text{spec}}) > 5\sigma_{\text{NMAD}}$, i.e. $|\Delta z|/(1 + z_{\text{spec}}) > 0.15$. This yields $\text{OLF} \in [1.2 \text{ per cent}, 4.2 \text{ per cent}]$ (4 outliers) with the BPZ code, and $\text{OLF} \in [4.1 \text{ per cent}, 8.4 \text{ per cent}]$ (10 outliers) with the EAZY code, where the interval in brackets indicates the 68 per cent confidence range computed assuming a Poisson distribution.¹⁴ We follow Rafelski et al. (2015) and define the OLF in the comparison between BEAGLE-derived and spectroscopic redshifts as the fraction of galaxies with $|\Delta z|/(1 + z_{\text{spec}}) > 0.15$. We increase the threshold from 0.15 to 0.25 when comparing BEAGLE-derived photometric redshifts with BPZ- and EAZY-derived ones, as σ_{NMAD} in those cases is typically around 0.05 (see Table 4).

Fig. 5 shows that, for the 169 galaxies with spectroscopic detections in the UVUDF catalogue, the photometric redshifts computed as described in Section 4.2 above using the BEAGLE tool agree well with the spectroscopic ones. The corresponding normalized median absolute deviation is $\sigma_{\text{NMAD}} = 0.047$, and the fraction of outliers $\text{OLF} \in [5.1 \text{ per cent}, 9.8 \text{ per cent}]$ (12 outliers; see Section 4.1). The value of σ_{NMAD} is larger than that obtained using both the BPZ and EAZY codes. This is likely because the two standard codes BPZ and EAZY rely on restricted sets of spectral templates optimized for the determination of photometric redshifts, while we consider a full model spanning a broad parameter space to describe the emission from a galaxy. Thus, in our approach, a large number of templates corresponding to different sets of parameters can potentially be consistent with the observed fluxes within the errors, which tends

¹⁴ Given k observations originating from a Poisson process with mean m , the confidence level $1 - \alpha$ on the mean can be computed as $\frac{1}{2}\chi^2(\alpha/2; 2k) \leq m \leq \frac{1}{2}\chi^2(1 - \alpha/2; 2k + 2)$, where $\chi^2(p; n)$ indicates the quantile function of the chi-squared distribution with n degrees of freedom.




Figure 5. Photometric redshift z_{BEAGLE} derived with the BEAGLE tool plotted against spectroscopic redshift z_{spec} , for the 169 galaxies with spectroscopic detection in the UVUDF catalogue. The green line shows the identity relation.

to increase the dispersion in the photometric redshifts derived for a galaxy at a given spectroscopic redshift. In return, the BEAGLE tool has the advantage of providing valuable constraints on galaxy physical properties other than redshift (Figs 8 and 9 below), as well as unique insight into the fundamentals of photometric redshift determinations (Section 5). Regarding this last point, for example, we can identify those galaxies with multimodal posterior probability distributions and quantify the integrated probability in each mode (Section 5.2). We note that, when considering only those galaxies with a single significant mode (i.e. 107 out of 169 galaxies exhibiting $2 \ln K > 10$, where K is the ‘Bayes factor’ defined by equation 31), we obtain $\sigma_{\text{NMAD}} = 0.037$ and a single outlier, object no. 4721.

It is instructive to investigate the origin of the 12 outliers in the comparison between z_{BEAGLE} and z_{spec} . For five galaxies (no. 1990, no. 8292, no. 21130, no. 22245, no. 50714), the BEAGLE tool identifies two redshift solutions of comparable probability, for which the second solution matches the spectroscopic redshift (see Section 5.2 for an in-depth discussion of multimodal solutions). To gain insight into the origin of the remaining seven outliers, we run a query at the location of each galaxy to obtain its classification in the Simbad¹⁵ data base. We also visually examine the SED and image of each outlier from the UVUDF website.¹⁶ Two outliers (no. 7024 and no. 10157) appear to be contaminated by nearby objects. One (no. 4721), classified in Simbad as an AGN, exhibits a strong ultraviolet upturn, which cannot be reproduced by the AGN-free model adopted here (but see Section 7). The SEDs of two galaxies (no. 4658 and no. 31320) with $z_{\text{spec}} \sim 0.2$ lack ultraviolet observations, and therefore do not sample any strong spectral continuum break useful to constrain photometric redshifts (Lyman, Balmer; these galaxies are outliers in the BPZ-based analysis too). For one outlier (no. 4562, $z_{\text{spec}} = 2.15$), classified as a ‘near-IR’ source in Simbad and with no ultraviolet observation, we obtain the same photometric redshift (≈ 2.71) as that derived by Rafelski et al. (2015) using BPZ,

¹⁵ <http://simbad.u-strasbg.fr/simbad/>

¹⁶ <http://asd.gsfc.nasa.gov/UVUDF/catalogs.html>




Figure 6. Comparison of photometric redshifts estimated using different approaches (z_{BEAGLE} : this work; z_{BPZ} and z_{EAZY} : Rafelski et al. 2015), for the 9927 galaxies in the UVUDF catalogue. The distributions are those of $(z_{\text{BEAGLE}} - z_{\text{BPZ}})/(1 + z_{\text{BPZ}})$ (blue filled histogram), $(z_{\text{BEAGLE}} - z_{\text{EAZY}})/(1 + z_{\text{EAZY}})$ (red filled histogram) and $(z_{\text{BPZ}} - z_{\text{EAZY}})/(1 + z_{\text{EAZY}})$ (black open histogram). In each case, a triangle indicates the median of the distribution.

indicating that the information contained in the SED is insufficient to provide an accurate photometric redshift estimate. Another galaxy (no. 10496, $z_{\text{spec}} = 1.10$) exhibits a similarly featureless continuum, with no ultraviolet observation.

In Fig. 6, we compare the photometric redshifts computed using the BEAGLE tool with those obtained by Rafelski et al. (2015) using the BPZ and EAZY codes. The corresponding normalized median absolute deviations and outlier fractions are reported in Table 4. Together, Fig. 6 and Table 4 show that the photometric redshifts computed with our approach and those derived by Rafelski et al. (2015) are globally consistent with each other, as σ_{NMAD} and OLF are typical of comparisons between different photometric redshift codes (e.g. see section 4.2 of Dahlen et al. 2013). Fig. 6 reveals a difference between the distribution of $(z_{\text{BEAGLE}} - z_{\text{BPZ}})/(1 + z_{\text{BPZ}})$ (blue filled histogram), which is roughly centred around zero, and that of $(z_{\text{BEAGLE}} - z_{\text{EAZY}})/(1 + z_{\text{EAZY}})$ (red filled histogram), which is highly skewed towards positive values. This is mainly because of the presence of a second peak including ~ 30 per cent of the objects around $z_{\text{BEAGLE}} - z_{\text{EAZY}} \approx 0.08(1 + z_{\text{EAZY}})$ in the latter distribution. The similarity between BEAGLE- and BPZ-derived photometric redshifts implies that the distribution of $(z_{\text{BPZ}} - z_{\text{EAZY}})/(1 + z_{\text{EAZY}})$ is also skewed towards positive values, although in a less severe way (black open histogram in Fig. 6).

Fig. 7 shows a more detailed comparison of the photometric redshifts derived using the BEAGLE tool with those derived by Rafelski et al. (2015) using the EAZY (top panel) and BPZ (bottom panel) codes. In each panel, different grey levels correspond to different logarithmically spaced galaxy densities, while red circles mark galaxies with multiple redshift solutions of comparable probability (see Section 5.2 for an extended discussion). In Fig. 7a, the vast majority of galaxies lie around the identity relation (green line) at all redshifts. A few galaxies with small z_{BEAGLE} ($\lesssim 1$) have large associated z_{BPZ} , while a few with high z_{BEAGLE} have low z_{BPZ} ($\lesssim 1$). The presence of red circles in these outlying regions suggests that multiple redshift solutions may be related to large discrepancies between photometric redshift estimates. In Fig. 7b, most galaxies also lie close to the

identity relation, but, as expected from Fig. 6, the estimates of z_{EAZY} for a substantial fraction of galaxies (at redshifts $2 \lesssim z_{\text{EAZY}} \lesssim 5$) are systematically smaller than those of z_{BEAGLE} . As in the case of Fig. 7a, the outlying regions occupied by galaxies with discrepant z_{BEAGLE} and z_{EAZY} are also populated by galaxies with multiple redshift solutions (see Section 5.2). We conclude from Figs 6 and 7 that the photometric redshifts estimated using the BEAGLE tool are in good general agreement with those estimated by Rafelski et al. (2015) using the BPZ and EAZY codes. The agreement between z_{BPZ} and z_{BEAGLE} is good at all redshifts, while at $2 \lesssim z_{\text{EAZY}} \lesssim 5$, the redshifts estimated with the EAZY code are systematically lower, by $\sim 0.08(1 + z_{\text{EAZY}})$, than those estimated using the BEAGLE and BPZ tools.

4.4 Posterior probability distribution of model parameters

The advantage of the BEAGLE tool over dedicated photometric redshift codes, such as BPZ and EAZY, is that it also allows the derivation of rigorous statistical constraints on galaxy physical parameters. We compute the one-dimensional (i.e. marginal) and two-dimensional (joint) probability distributions of the model parameters in Table 3 (and of other derived quantities) using the GETDIST Python package, which we integrate into PYP-BEAGLE, our package for post-processing BEAGLE results obtained with the MULTINEST algorithm (Appendix A). The GETDIST package¹⁷ has been developed within COSMOMC,¹⁸ a powerful Bayesian framework for the analysis of cosmological data originally presented in Lewis & Bridle (2002) and extensively exploited to interpret *Planck* data. Both MULTINEST and GETDIST are well documented and can be used as stand-alone packages. In particular, GETDIST allows one to compute continuous posterior probability distributions from the samples obtained through MULTINEST by means of one-dimensional and two-dimensional kernel density estimates. This presents several advantages over density estimation through standard histograms (e.g. continuity, well defined derivatives, no requirement of a bin width). The major difficulty associated with the use of kernel density estimates lies in the handling of boundary effects, which are, however, rigorously treated in GETDIST, following the approach outlined in Lewis (2015).

We compute in this way the posterior probability distribution of model parameters and derived quantities for all galaxies in the UVUDF sample. For the sake of illustration, we focus here on two particularly instructive cases: a faint, high-redshift galaxy (no. 1021); and a galaxy exhibiting multiple probability modes (no. 5866), i.e. different regions of high posterior probability in the multidimensional parameter space of the model. The main panel of Fig. 8 shows, on the diagonal, the marginal posterior probability distributions of the derived quantity M_* (the current stellar mass, which accounts for the fraction of mass returned to the ISM during stellar evolution; see footnote 9) and the model parameters τ_{SFR} , Z , ψ_{S} , $\hat{\tau}_{\text{V}}$, z and z_{form} for the galaxy no. 1021 (F850LP = 27.6 AB mag), and off diagonal, the joint posterior probability distribution of every combination of these parameters. The small inset panel shows the observed SED of the galaxy (blue diamonds), along with the distribution of predicted fluxes (orange ‘violins’) resulting from the posterior probability distribution of the model parameters. This was computed by considering the posterior probability and predicted broad-band fluxes corresponding to each set of parameters

¹⁷ Available at <https://github.com/cmbant/getdist>

¹⁸ Downloadable from <http://cosmologist.info/cosmomoc/>




Figure 7. (a) Photometric redshift z_{BPZ} derived by Rafelski et al. (2015) using the BPZ code plotted against that derived in this work using the BEAGLE tool, z_{BEAGLE} , for the 9927 galaxies in the UVUDF catalogue. (b) Same as (a), but comparing the photometric redshift z_{EAZY} derived by Rafelski et al. (2015) using the EAZY code to z_{BEAGLE} . In both panels, different grey levels correspond to different logarithmically spaced galaxy densities (indicated on the right), red circles mark galaxies with multiple redshift solutions of comparable probability, and the green line shows the identity relation. See Section 4.2 for detail.

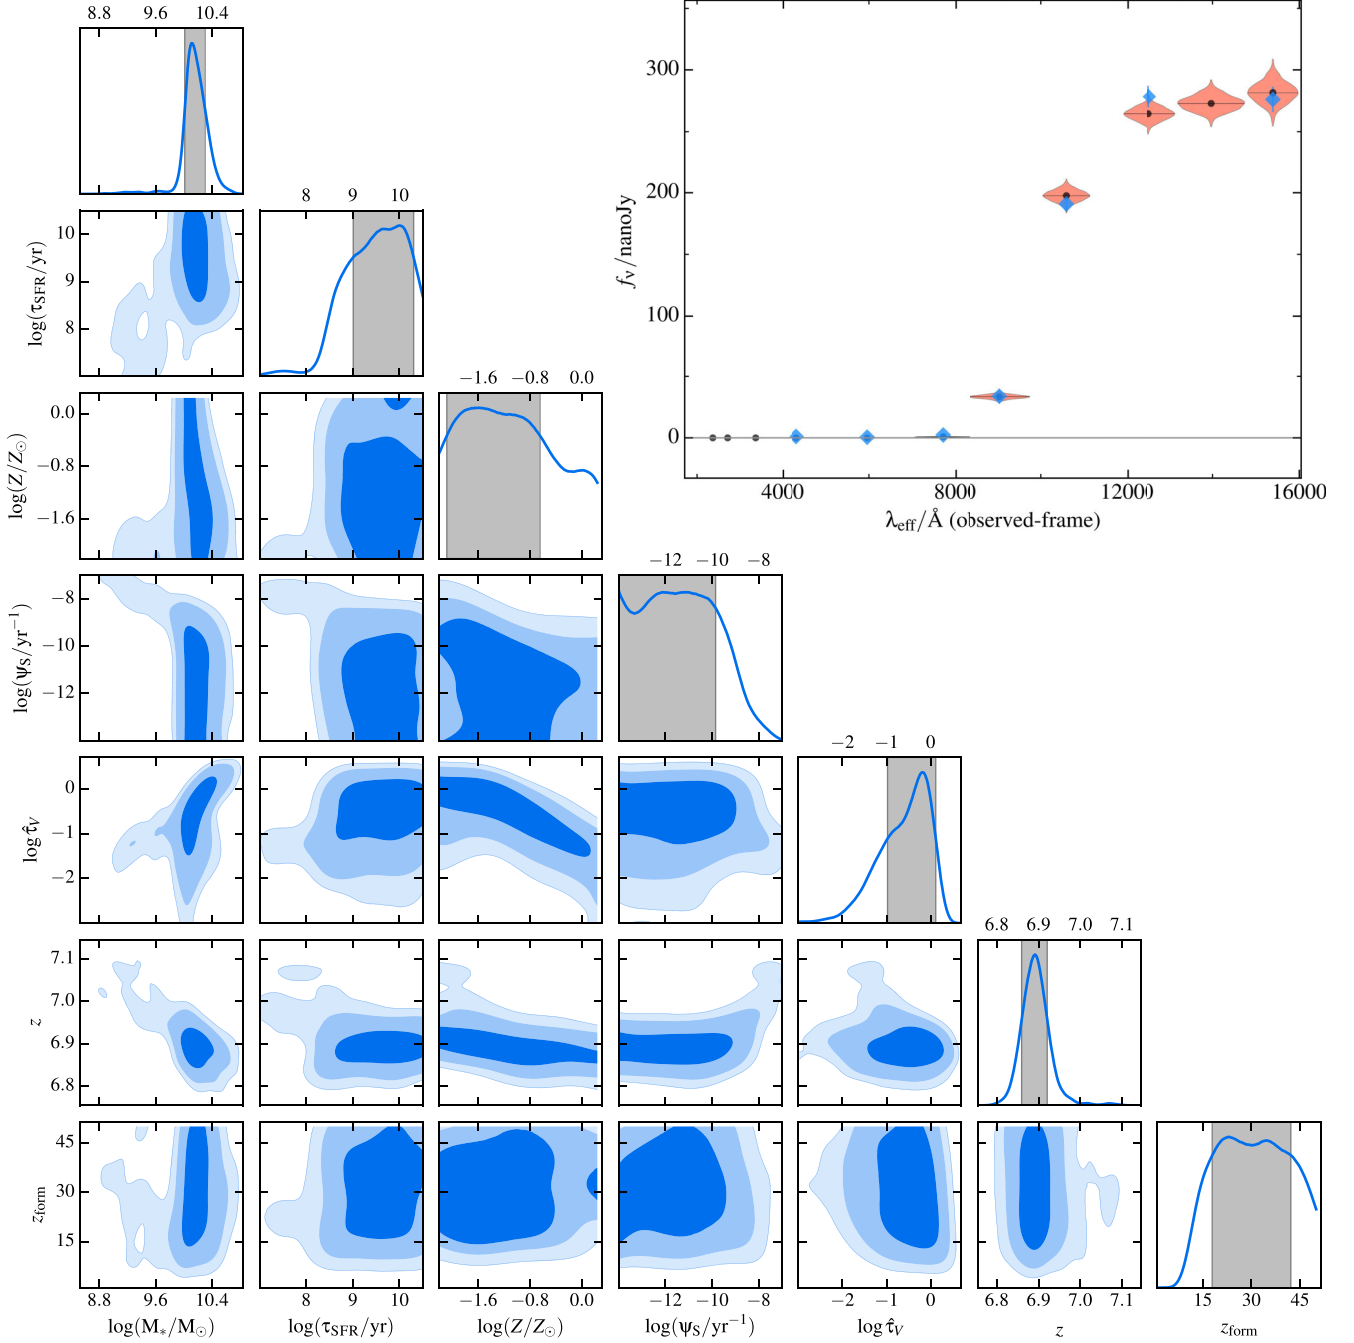


Figure 8. Posterior probability distribution of the model parameters in Table 3 for the UVUDF galaxy #1021 ($F850LP = 27.6$ AB mag, p -value = 0.13). The diagonal panels show the marginal distributions of M_* , τ_{SFR} , Z , ψ_S , $\hat{\tau}_V$, z and z_{form} , and the off-diagonal panels the joint distribution of every pair of these parameters. The inset panel on the top right shows the observed SED of the galaxy (blue diamonds), along with the distribution of predicted fluxes (orange ‘violins’) resulting from the posterior probability distribution of the model parameters (see Section 4.4 for details). This object does not have any measurement in the $F225W$, $F275W$ and $F336W$ ultraviolet bands nor in the $F140W$ near-infrared band.

in the posterior probability distribution sampled by MULTINEST. A kernel density estimate was then performed separately for each band to obtain a smooth flux distribution visualized by the violin (see Section 5.1 for details about violin plots). The SED of object no. 1021 reveals strong IGM absorption at $\lambda \lesssim 8000$ Å, suggesting a high redshift for this galaxy.

The diagonal panels of Fig. 8, in which we show for each parameter the 68 per cent central credible interval as a grey shaded

area, indicate that object no. 1021 is a moderately massive galaxy [$\log(M_*/M_\odot) \sim 10.2$] at redshift $z \sim 6.9$, with a moderate dust content ($\hat{\tau}_V \sim 0.4$). The model favours a long star formation time-scale [$\log(\tau_{\text{SFR}}/\text{yr}) \gtrsim 9$], high formation redshift [$z_{\text{form}} \gtrsim 20$], low metallicity [$\log(Z/Z_\odot) \lesssim -0.7$] and low specific star formation rate [$\log(\psi_S/\text{yr}^{-1}) \lesssim -10$], although the widths of the marginal posterior probability distributions of these parameters indicate that they are barely constrained by the observations. The off-diagonal

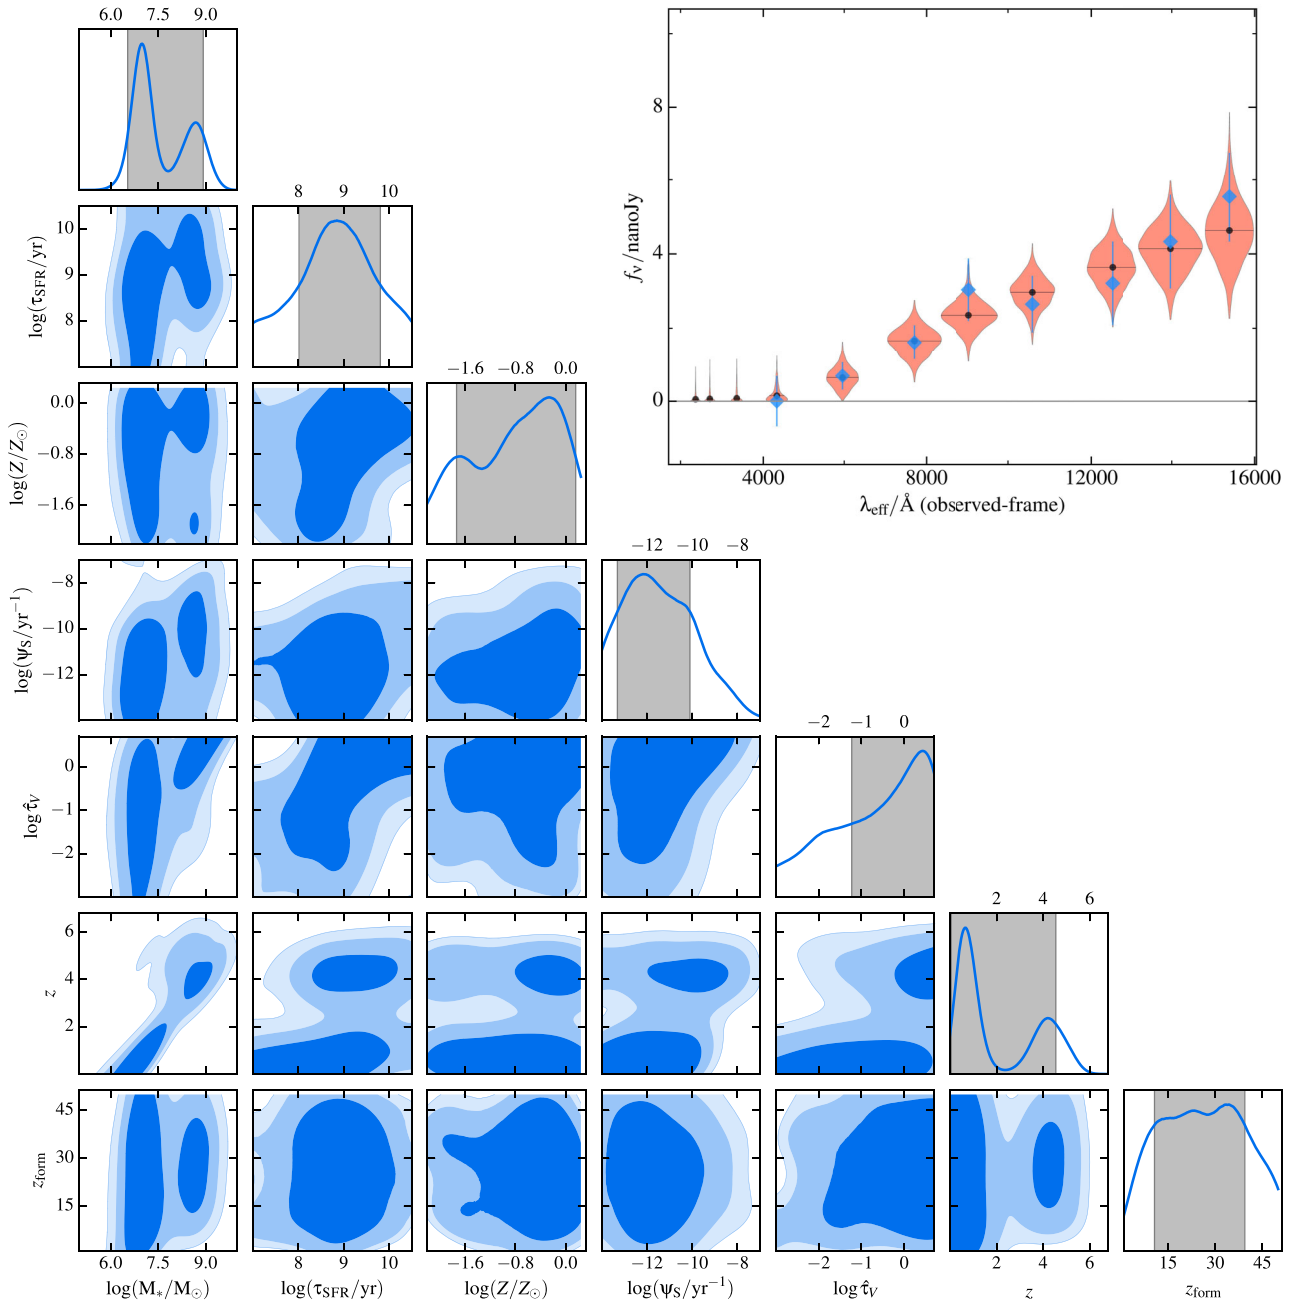


Figure 9. Same as Fig. 8, but for the galaxy no. 5866 ($F850LP = 30.2$ AB mag, p -value = 0.78). This object does not have any measurement in the $F225W$, $F275W$ and $F336W$ ultraviolet bands.

panels illustrate the potential of a Bayesian approach to study degeneracies between model parameters: the three contour levels, showing the 68, 95 and 99 per cent central credible regions, reveal a correlation between stellar mass and dust attenuation optical depth, in the sense that larger M_* implies larger $\hat{\tau}_V$. This is because the enhanced flux produced by a more massive galaxy must be attenuated by more dust to produce the same observed flux, when all other parameters are fixed. The figure also shows a mild anti-correlation between dust attenuation optical depth and metallicity, a well known degeneracy resulting from the similar effects of an increase in Z and $\hat{\tau}_V$ on galaxy colours.

Fig. 9 shows the analogue of Fig. 8 for the much fainter galaxy no. 5866 ($F850LP = 30.2$ AB mag). The inset panel il-

lustrates how the larger observational errors in this case allow much more extended distributions of the predicted fluxes. The marginal posterior probability distributions of τ_{SFR} , Z , ψ_S , $\hat{\tau}_V$ and z_{form} in the diagonal panels of Fig. 9 show that these parameters are only weakly constrained by the observations. Also, in contrast to Fig. 8, the marginal posterior probability distributions of M_* and z show two peaks, indicating the presence of two solutions of comparable probability. The joint posterior probability distribution of M_* and z further shows that these multiple solutions are correlated (i.e. the parameters are degenerate), since the low-redshift solution at $z \sim 0.8$ favours a lower mass [$\log(M_*/M_\odot) \sim 7$] than the high-redshift one at $z \sim 4$ [$\log(M_*/M_\odot) \sim 8.7$]. Such a correlation is expected because, at

comparable mass-to-light ratio, a larger mass is required to produce the same apparent luminosity at high relative to low redshift. We note that the existence of this multimodal solution arises primarily from the faintness of the galaxy no. 5866: the low observational S/N ratio of this galaxy causes both the Balmer break at $\lambda \sim 3600 \text{ \AA}$ and the Lyman break at $\lambda \sim 1216 \text{ \AA}$ to be compatible with the observed drop in flux between the *F435W* and *F606W* bands. Fig. 9 also shows correlations between dust attenuation optical depth and stellar mass, for the same reason as outlined above, and dust attenuation optical depth and redshift, for which solutions implying high values of $\hat{\tau}_V$ correspond to higher z . Finally, the tighter constraint on z_{form} in the case of the high-redshift solution ($z \sim 4$) results from the smaller age spread (and hence luminosity range) of stars in the galaxy in this case, implying less uncertainty on the age of the oldest stellar generation.

It is also important to note that, while dedicated photometric redshift codes, such as those adopted in Rafelski et al. (2015), can warn against the presence of multiple redshift solutions, the Bayesian approach implemented in the BEAGLE tool allows one to accurately characterize these solutions, for instance by providing their respective integrated probability. We return to this point in Section 5.2 below, where we illustrate the power of the BEAGLE tool in such situations by performing a Bayesian model comparison of the different modes of a posterior probability distribution.

5 ADVANCED STATISTICAL ANALYSIS WITH BEAGLE

In the previous section, we have seen that the BEAGLE tool enables one to fit broad-band galaxy SEDs to derive constraints not only on redshift, but also on other model parameters (e.g. stellar mass, star formation history, dust content) and derived quantities. We now describe another main feature of the BEAGLE tool, which is to allow advanced statistical analyses, such as the rigorous quantification of goodness of fit and the detailed study of potential correlations (i.e. degeneracies) between model parameters, to identify and characterize multimodal solutions.

5.1 Posterior predictive checks

The outcome of any Bayesian analysis is the posterior probability distribution of a set of parameters *conditional* to a set of observations. Still, the posterior probability distribution alone does not allow one to determine whether the assumed model is a good description of the data, or whether it needs to be changed and improved. In a ‘frequentist’ statistical framework, the goodness of a fit can be evaluated by comparing the value of a test statistics, such as χ^2 , to a reference distribution. This is often achieved via the computation of the p -value, i.e. the integrated tail probability. A major challenge in this case is to define the appropriate reference distribution for the adopted test statistics. For example, the reference distribution for a χ^2 statistics is tied to the number of degrees of freedom of the assumed model. While this number can be known a priori in a few special cases, such as linear models, it is difficult to estimate in many practical situations. In contrast, in a Bayesian approach, one can design goodness-of-fit tests in which the reference distribution of the adopted test statistics is estimated straightforwardly using the model itself. Such tests, called posterior predictive checks (e.g. Guttman 1967; Rubin et al. 1984; Gelman, Meng &

Stern 1996), enable the probabilistic assessment of whether a model is a reasonable description of a data set.

The general idea motivating posterior predictive checks is the following: if the assumed model is a good description of a set of observations, the model should be able to produce ‘replicated observations’ statistically indistinguishable from the true ones. These replicated data can be thought of as observations which *could* have been measured, assuming that the variability in the data is entirely captured by the adopted statistical model. The ‘posterior predictive probability distribution’ of replicated data can be written as (e.g. Gelman et al. 1996)

$$P(\mathbf{D}^{\text{rep}} | \mathbf{D}, H) = \int P(\mathbf{D}^{\text{rep}} | \Theta, H) P(\Theta | \mathbf{D}, H) d\Theta, \quad (28)$$

where H , Θ and \mathbf{D} have the same meaning as in equation (24), and \mathbf{D}^{rep} is the replicated data set. The first factor in the integral on the right-hand side of equation (28) is the probability distribution of replicated data conditional to the model parameters, while the second is the posterior probability distribution obtained by applying Bayes’ theorem (equation 24). An advantage of a Bayesian goodness-of-fit test based on equation (28) is that, once a set of samples drawn from the posterior probability distribution is available, no further heavy computation is required.

We consider here two types of posterior predictive checks: a numerical one, based on the χ^2 test statistics, and a graphical one. In both cases, for each galaxy in the UVUDF catalogue, we need a set of replicated data drawn from the probability distribution defined in equation (28). For this, we start from the output of MULTINEST, which consists in an ensemble of N_{out} sets of parameters weighted by the posterior probability distribution (see Appendix A). Then, we draw $N_{\text{rep}} = 2000$ replicated data sets as follows (we have checked that adopting a larger N_{rep} has a negligible influence on the results):

- (1) we draw N_{rep} sets of model parameters Θ^k , with $1 \leq k \leq N_{\text{rep}}$, from the posterior probability distribution obtained with MULTINEST, using a ‘weighted sampling with replacement’ scheme (see Appendix B for detail);
- (2) for each set of parameters Θ^k , we draw the replicated data \mathbf{y}^{rep} from a Gaussian distribution $\mathcal{N}[\hat{\mathbf{y}}(\Theta^k), \Sigma]$, where $\hat{\mathbf{y}}(\Theta^k)$ indicates the fluxes predicted by the model given the set of parameters Θ^k , and Σ is the covariance matrix of the data, i.e. in our case a diagonal matrix with elements $\Sigma_{i,i} = \sigma_i^2$ (as in Section 4.2).

To perform the numerical posterior predictive check, we adopt the χ^2 test statistics as a measure of deviance between model predictions and data. For each galaxy in the UVUDF catalogue, we compute the χ^2 deviance as

$$\chi^2(\Theta^k) = \sum_i \left[\frac{y_i - \hat{y}_i(\Theta^k)}{\sigma_i} \right]^2, \quad (29)$$

where, as in equation (26), y_i indicates the observed flux in the i th band, $\hat{y}_i(\Theta^k)$ the flux predicted by the model in the same band for a set of parameters Θ^k , and σ_i the error (which, as in equation 26, includes both the observational error and the 2 per cent relative error added in quadrature). We compute $\chi^2(\Theta^k)$ for all N_{rep} sets of model parameters Θ^k . Then, for each set of parameters, we substitute y_i^{rep} for y_i in equation (29) and compute the corresponding $\chi_{\text{rep}}^2(\Theta^k)$. In this way, we obtain two distributions of the χ^2 statistics, one pertaining to the true data, and one to replicated data. To compare the two distributions, we adopt as p -value the fraction of replicated




Figure 10. Distribution of p -values, computed according to equation (30), for the 9927 galaxies in the UVUDF catalogue fitted with the model presented in Section 4.2.

data with χ^2 larger than the corresponding one obtained with the true data (Gelman et al. 1996), i.e.

$$p\text{-value} = \frac{N \left[\chi_{\text{rep}}^2(\Theta) > \chi^2(\Theta) \right]}{N_{\text{rep}}}. \quad (30)$$

Fig. 10 shows the distribution of p -values computed in this way for all galaxies in the UVUDF catalogue. This reveals that most galaxies have p -values in the range 0.05–0.95, indicating a satisfactory model fit, while 7 (13) per cent have p -values less than 0.01 (0.05). Such very low p -values suggests that there may be a problem with either the data or the model, or both, in the fit of these galaxies. We note that a way to discriminate between a data or model origin of a bad fit is to study the fit residuals in the observer and rest frames: while data-driven mismatches arise at the instrument (or reduction) level, hence acting in the observer frame, model-driven ones are caused by an inaccurate physical representation of galaxy SEDs, hence acting in the galaxy rest frame.

While Fig. 10 summarizes the global quality of photometric fits of UVUDF galaxies obtained with the model of Section 4.2, it does not allow us to characterize this performance in detail, i.e. band by band. For this reason, we also perform a graphical posterior predictive check. We randomly select 12 well-fitted galaxies with p -values greater than 0.1, and 12 badly fitted galaxies with p -values less than 0.01. For each galaxy and each set of replicated data of that galaxy, we compute the residual between replicated and true data, $(y_i^{\text{rep}} - y_i)/\sigma_i$, in each photometric band (where the symbols have the same meaning as in equation 29). Fig. 11 shows a ‘violin’ plot of the resulting distribution of residuals for the 12 galaxies with p -values greater than 0.1. Each violin was computed by performing, for each band separately, a kernel density estimate to obtain a continuous distribution of the N_{rep} residuals and then plotting the 0.997 central credible region of this distribution. As is customary in violin plots, the distribution was mirrored with respect to an axis parallel to the ordinate axis and the maximum width of each violin adapted to avoid overlap. Each violin of Fig. 11 therefore reflects the probability of obtaining a given residual in a given photometric band: the more extended the violin in the ordinate direction, the broader the distribution of the residual.

As a complement to Fig. 11, Table 5 lists for each galaxy the p -value computed using equation (30). We also report the p -value computed in the same way for each photometric band separately,

by considering the contribution of only that band to χ^2 in equation (29). In the case of (random) noise-driven residuals, and assuming that σ (where we have dropped the band index) accounts for all possible sources of noise, we expect residual distributions centred around zero with a dispersion comparable to σ . This is what most violins reflect in Fig. 11, although some bands for some galaxies display larger residuals ($\sim 2\sigma$; e.g. no. 22363: band *F336W* and *F140W*; no. 1811: band *F105W*). Table 5 quantifies this information by providing the significance of residual deviations by means of the p -value: band *F336W* and *F140W* for object no. 22363 and band *F105W* for object no. 1811 all have p -values in the range $\sim [0.07, 0.03]$.

Fig. 12 and Table 6 show the analogues of Fig. 11 and Table 5 for the 12 randomly selected galaxies with p -values less than 0.01. In contrast to Fig. 11, Fig. 12 shows significant residual deviations between replicated and true data. These deviations differ from galaxy to galaxy, as Table 6 also highlights. Although a comprehensive analysis of fitting residuals from posterior predictive checks goes beyond the scope of the present paper, we stress that studying the distributions of these residuals in the observer and galaxy rest frames can help discriminate between a data or model origin (see above).

Hence, by means of posterior predictive checks, we have shown that the relatively simple model of Section 4.2 provides a satisfactory fit to the photometry of a vast majority of UVUDF galaxies. We have illustrated how the combination of graphical and numerical posterior predictive checks can provide valuable insight into the origin of discrepancies between model and data. Distributions of residual deviance between model predictions and data, of the type shown in Figs 11 and 12 for individual galaxies, can be straightforwardly extended to combine residuals from different objects, enabling the identification of data- and model-driven discrepancies too subtle to be detected in single galaxies. In this context, we believe that the BEAGLE tool will be valuable both to identify subtle systematics in observed data sets, and to characterize current limitations and drive future developments of spectral models.

5.2 Multimodal solutions

So far, we have shown that: (i) multiband photometric fitting of UVUDF galaxies with the BEAGLE tool leads to photometric redshift estimates in good agreement with those derived from standard dedicated codes (Section 4.3); (ii) the accurate Bayesian characterization of the posterior probability distribution of model parameters with the BEAGLE tool allows a rigorous study of multiple solutions and degeneracies between model parameters (Section 4.4); and (iii) the relatively simple model of Section 4.2 provides satisfactory fits to the photometry of most UVUDF galaxies (Section 5.1). In this section, we focus on the study of multimodal solutions, which as noted above could be a major cause of discrepancy between photometric redshifts derived using different approaches (Fig. 7).

We start by noting that the BPZ and EAZY codes adopted by Rafelski et al. (2015) to estimate photometric redshifts provide a quantity, the ‘odds’, which is sensitive to the occurrence of multiple redshift solutions. Both codes compute the marginal probability distribution of redshift and integrate this over some fixed range. In the BPZ code, this range is an interval of width $0.06 \times (1 + z)$ around the peak, while in the EAZY code, the interval is $0.2 \times (1 + z)$ wide. Such a fixed interval of integration can provide some indication about the concentration of the marginal probability distribution of redshift




Figure 11. Graphical posterior predictive check for 12 randomly selected galaxies with p -value > 0.1 (as computed using equation 30 when fitting with the model presented in Section 4.2) in the UVUDF catalogue (the galaxy ID is indicated in the top right corner of each panel). In each panel, each ‘violin’ shows the probability distribution of the residual between replicated and true data in a given photometric band, with the median marked by a black horizontal line. Crosses indicate bands with no measurement. The solid and dashed grey lines indicate zero and $\pm 3\sigma$ residuals, respectively (see Section 5.1 for details).

Table 5. Posterior predictive p -values in the different photometric bands, and global p -value computed using equation (30) (rightmost column), for the 12 galaxies in Fig. 11 (see Section 5.1 for details).

Galaxy ID	Photometric band										All	
	F225W	F275W	F336W	F435W	F606W	F775W	F850LP	F105W	F125W	F140W		F160W
22774				0.53	0.38	0.30	0.47	0.54	0.57	0.53	0.50	0.41
1730				0.48	0.53	0.56	0.55	0.45	0.53		0.60	0.59
8871				0.47	0.44	0.21	0.32				0.63	0.24
6223				0.56	0.58	0.51	0.64	0.86	0.69		0.14	0.58
4888	0.49	0.75	0.51	0.38	0.59	0.22	0.51	0.55	0.54	0.38	0.47	0.53
438				0.43	0.58	0.23	0.50	0.52				0.44
3788				0.97	0.56	0.55	0.52	0.76	0.38		0.60	0.76
22363		0.72	0.063	0.65	0.58	0.71	0.54	0.39	0.67	0.072	0.29	0.21
280				0.46	0.16	0.21	0.76	0.76	0.47		0.78	0.33
1811				0.86	0.55	0.51	0.42	0.031	0.73		0.58	0.22
10630				0.51	0.52	0.59	0.58	0.79	0.28		0.14	0.47
3270				0.36	0.46	0.25	0.21	0.44	0.61		0.68	0.28

around the peak, but not a measure of the relative probabilities of different potential solutions. For reference, Rafelski et al. (2015) consider a redshift estimate to be reliable is the corresponding odds are greater than 0.9.

As emphasized in Sections 4.3 and 4.4, our approach in the BEAGLE tool differs from that of most template-based photometric redshift codes, in that we consider redshift as just one of several model parameters influencing the predicted observables (Table 2). Thus, each value of photometric redshift is explored along with a set of other galaxy parameters. This implies that different redshift solutions also correspond to different solutions of other physical parameters, via the multidimensional posterior probability distribution of equation (24). We note in passing that this may also potentially provide useful information to exclude some redshift solutions, for example, because of unlikely combinations of parameters (e.g. massive galaxy with very low metallicity). In the case of multimodal solutions, the BEAGLE tool allows one to compute the probability associated with

each mode and perform a Bayesian model comparison to interpret the results in a probabilistic way. In practice, we identify different modes (higher probability contours in the multidimensional parameter space separated by lower probability valleys) in the posterior probability distribution of model parameters provided by MULTINEST (Appendix A) and compute the local evidence within each mode. We then consider the two modes with largest local evidence, which we label H_1 and H_2 in order of decreasing local evidence. To assess the plausibility of these two solutions, we compute the ‘Bayes factor’ K given by the ratio of local evidences

$$K = \frac{P(\mathbf{D} | H_1)}{P(\mathbf{D} | H_2)} = \frac{\int P(\boldsymbol{\Theta}_1 | H_1) P(\mathbf{D} | \boldsymbol{\Theta}_1, H_1) d\boldsymbol{\Theta}_1}{\int P(\boldsymbol{\Theta}_2 | H_2) P(\mathbf{D} | \boldsymbol{\Theta}_2, H_2) d\boldsymbol{\Theta}_2}. \quad (31)$$

In this expression, the local evidence within each mode is the integral of the prior distribution times the likelihood (equation 24) in the subset of the entire parameter space occupied by that mode. We stress that this Bayesian framework enables the straightforward




Figure 12. Same as Fig. 11, but for 12 randomly selected galaxies with p -value ≤ 0.01 .

Table 6. Same as Table 5, but for the 12 galaxies in Fig. 12.

Galaxy ID	Photometric band											
	<i>F225W</i>	<i>F275W</i>	<i>F336W</i>	<i>F435W</i>	<i>F606W</i>	<i>F775W</i>	<i>F850LP</i>	<i>F105W</i>	<i>F125W</i>	<i>F140W</i>	<i>F160W</i>	All
5024	0.046	0.065	0.57	0.062	0.42	0.042	0.43	0.56	0.61	0.35	0.43	0.0025
6640					0.28	0.30	0.13	0.67	0.046	0.035	0.55	0.0095
3458					0.0035	0.13	0.076	0.19	0.48		0.38	0.0030
5725				0.0070	0.66	0.54	0.47	0.023	0.038		0.52	0.0050
10404			0.19	0.54	0.49	0.46	0.031	0.0090	0.53	0.51	0.52	0.0090
4897				0.018	0.40	0.24	0.0005	0.35	0.58		0.30	0.0020
4863				0.35	0.13	0.39	0.069	0.082	0.020		0.48	0.0030
3422				0.51	0.55	0.52	0.069	0.088	0.020		0.15	0.0090
1493				0.43	0.032	0.37	0.0035	0.81	0.13		0.88	0.0035
9724				0.23	0.052	0.19	0.085	0.037	0.093		0.49	0.0040
2077					0.57	0.016	0.0020	0.61	0.85	0.27	0.0055	
5606				0.072	0.046	0.042	0.60	0.082	0.40	0.60	0.0025	

implementation of any type of prior distribution of model parameters (through equation 31), such as those based on galaxy luminosity functions often adopted in photometric redshift codes (e.g. Benítez 2000; Brammer et al. 2008).

We analyse in this way the posterior probability distributions of all galaxies in the UVUDF catalogue and label as multimodal those with redshift solutions separated by $|\Delta z| > 0.5$. This minimum Δz threshold, which corresponds roughly to the typical posterior standard deviation of redshift for a faint source (with single solution) in the catalogue, allows us to remove multiple redshift solutions with statistically non-significant separation. To interpret the results of our analysis, we consider the quantity $2 \ln K$, which, according to Kass & Raftery (1995), is better suited than the Bayes factor K itself to drive the choice between different modes.¹⁹ Fig. 13 shows

the distribution of $2 \ln K$ for the 3243 UVUDF galaxies identified to have multiple redshift solutions separated by $|\Delta z| > 0.5$. Dashed vertical lines divide the $2 \ln K$ axis into the four ‘belief’ categories defined by Kass & Raftery (1995, regions A, B, C and D) and reported in Table 7. As the table indicates, about 18 per cent (1800) of all galaxies in the UVUDF catalogue exhibit multiple redshift solutions separated by $|\Delta z| > 0.5$, for which the Bayes factor does not allow a highly confident redshift selection. This is when adopting a high-confidence threshold, corresponding to $K = \exp(5) \approx 150$ (i.e. $2 \ln K = 10$), as recommended by Kass & Raftery (1995). Adopting instead a lower threshold, corresponding to $K = \exp(3) \approx 20$ (i.e. $2 \ln K = 6$), would lower the fraction of galaxies with ambiguous redshift solutions to 13 per cent (Table 7).

The distribution of Bayes factors in Fig. 13 and Table 7 may also be used to gain insight into the origin of the widely different redshift solutions found for some objects using different photometric redshift codes (illustrated by the outliers in Figs 8 and 9). In fact, as the rightmost columns of Table 7 show, about half of the outliers in the z_{BEAGLE} versus z_{BPZ} and z_{BEAGLE} versus z_{EAZY} comparisons, defined as galaxies with $|\Delta z|/(1+z) > 0.25$ (Section 4.2 and Table 4), are objects with multiple, ambiguous redshift solutions

¹⁹ The choice of an optimal scale for the interpretation of the Bayes factor is a classical challenge in statistics, with no unique solution. The scale adopted here, first proposed by Kass & Raftery (1995), is similar to the original scale of Jeffreys (1961), hence switching between the two scales will not significantly alter our conclusions.




Figure 13. Distribution of the quantity $2 \ln K$, where K is the Bayes factor (equation 31), for the 3243 UVUDF galaxies identified to have multiple redshift solutions separated by $|\Delta z| > 0.5$. The vertical lines mark the thresholds proposed by Kass & Raftery (1995) to separate different categories of ‘belief’ (regions A, B, C and D), as specified in Table 7. Note that region D includes all galaxies with $2 \ln K > 10$, but for clarity we only show the plot up to $2 \ln K = 20$.

($2 \ln K \leq 10$; see section 5 of Ilbert et al. 2006 for a similar conclusion based on a comparison between spectroscopic and photometric redshifts, and the discussion of the ‘reliability parameter’ in section 4.2 of Brammer et al. 2008). Moreover, for about 75 per cent of the outliers with ambiguous redshift solutions, the second redshift solution identified with the BEAGLE tool matches the redshift estimated by Rafelski et al. (2015) with the BPZ or EAZY codes.

It is important to stress that insights of the kind provided by Fig. 13 and Table 7, and the associated ability with the BEAGLE tool to reduce by over an order of magnitude (from 12 to 1) the number of outliers in the z_{BEAGLE} versus z_{spec} comparison through the exclusion of objects identified to have ambiguous redshifts (modulo a reduction of the sample size; see Section 4.3), represent a new, promising way to study the origin of photometric redshift outliers, from which future large photometric surveys, such as those planned with *Euclid* and the Large Synoptic Survey Telescope (LSST), can greatly benefit.

6 RELATION TO EXISTING SED FITTING CODES

The interpretation of galaxy SEDs at ultraviolet, optical and infrared wavelengths in terms of physical parameters has been the subject of many studies, leading to the development of several public codes. In practice, most codes designed to fit galaxy spectra (e.g. PPF, Cappellari & Emsellem 2004; STARLIGHT, Cid Fernandes et al. 2005; STECKMAP, Ocvirk et al. 2006; ULYSS, Koleva et al. 2009) suffer from intrinsic limitations. For example, to select the best-fitting models, most codes rely on simple χ^2 minimization techniques, which do not allow the computation of realistic uncertainties in the derived physical parameters. Classical codes also tend to be programmed in ‘interpreted’ (e.g. IDL, Yorick, MATLAB) rather than ‘compiled’ (e.g. FORTRAN, C, C++) languages, at the expense of performance.²⁰ In addition, standard SED interpretation tools are usually tied to

specific choices of built-in physical ingredients (e.g. stellar evolution and dust attenuation prescriptions), which cannot easily be changed, nor tested (e.g. by means of posterior predictive checks; see Section 5.1), nor extended to include new physical ingredients (e.g. emission from an AGN, neutral ISM absorption). An additional specificity of most current tools focused on the interpretation of photometric (rather than spectroscopic) galaxy SEDs is that these tend to be optimized for either redshift estimation (e.g. HYPERZ: Bolzonella et al. 2000; LEPHARE: Arnouts et al. 1999; Ilbert et al. 2006; KCORRECT: Blanton & Roweis 2007; EAZY: Brammer et al. 2008) or the determination of galaxy physical parameters (e.g. CIGALE: Burgarella et al. 2005; Noll et al. 2009; MAGPHYS: da Cunha et al. 2008, 2015; FAST: Kriek et al. 2009), but not both simultaneously.

In this context, and given the large number of existing SED fitting codes, we restrict our discussion in this section to the comparison of the BEAGLE tool with publicly available codes relying on an analogous Bayesian approach: CIGALE, MAGPHYS, GALMC and BAYESED, of which we report the main characteristics (and references) in Table 8. As indicated, none of these codes is designed to interpret spectroscopic galaxy observations. Also, among them, only the CIGALE code includes at the same time nebular, dust and AGN emission, although in an approximate way. In this code, for example, nebular emission is incorporated using only two fixed emission-line templates to represent the emission from gas heated by young stars, ignoring the contributions by recombination continuum radiation and any AGN component. The MAGPHYS code has been designed to consistently interpret ultraviolet to far-infrared SEDs, by adopting an ‘energy budget’ approach to account for dust emission (da Cunha et al. 2008). This code does not include a full model of nebular emission, accounting for only H α and H β emission, and it is based on a wide (albeit predefined) library of galaxy star formation histories. Finally, the BAYESED code does not include models for nebular and AGN emission, and it is based on a rigid model (exponentially declining function) to describe a galaxy star formation history.

The limitations of existing SED fitting codes mentioned above and in Table 8 are the main motivation for our development of the BEAGLE tool presented in this paper. This tool incorporates the most recent prescriptions for stellar and nebular emission (and the dependence of these components on chemical composition), attenuation by dust, IGM absorption, etc., in a physically consistent and highly flexible way (Section 2): the modular design of the BEAGLE tool allows any of these prescriptions to be easily replaced by an alternative one. It also enables the straightforward implementation of additional physical ingredients, such as described in Section 7 for the BEAGLE 2.0 version. The Bayesian approach adopted in the tool allows the user to rigorously quantify the uncertainties and degeneracies affecting model parameters. Finally, the use of MCMC techniques (implemented in the MULTINEST algorithm; see Appendix A) and a compiled language (FORTRAN 2003) makes the exploration of complex, multidimensional parameter spaces much more efficient, and much less memory-demanding, than in conventional, grid-based approaches.

7 SUMMARY AND FUTURE DEVELOPMENTS

We have introduced a novel tool, named BEAGLE, to model and interpret any combination of photometric and spectroscopic galaxy observation at ultraviolet to infrared wavelengths. Currently, the BEAGLE tool allows one to model the emission from stellar populations, over wide ranges of age, metallicity and α -element to iron abundance ratio, the emission from gas photoionized by young stars and the absorption of the light by interstellar dust and the

²⁰ An exception is the public software STARLIGHT (Cid Fernandes et al. 2005) which is written in FORTRAN, but the code is not open-source.

Table 7. Fractional distribution of the quantity $2 \ln K$, where K is the Bayes factor (equation 31), among the four ‘belief’ categories defined by Kass & Raftery (1995, regions A, B, C and D), of the total sample of 9927 UVUDF galaxies. The two rightmost columns indicate the fractional distribution among the same categories (with actual numbers in parentheses) of the outliers in the z_{BEAGLE} versus z_{BPZ} and z_{BEAGLE} versus z_{EAZY} comparisons (from Table 4).

Label	$2 \ln K$	Kass and Raftery	Total sample	BPZ outliers	EAZY outliers
A	[0, 2]	‘Barely worth mentioning’	0.040	0.19 (160)	0.21 (155)
B	(2, 6]	‘Positive’	0.093	0.19 (159)	0.29 (210)
C	(6, 10]	‘Strong’	0.049	0.043 (35)	0.053 (38)
A, B, C	≤ 10	Not ‘very strong’	0.18	0.43 (354)	0.56 (403)
D	> 10	‘Very strong’	0.15	0.089 (73)	0.13 (92)

Table 8. Main characteristics of different publicly available Bayesian SED fitting codes. In the case of photometric SED fitting, the ability to estimate photometric redshifts is indicated in parentheses.

Name and reference	Non-stellar emission component			Type of SED to be interpreted		
	Nebular	Dust	AGN	Photometry	Spectroscopy	Mix
CIGALE ^a	yes (approximate)	yes	partial (accretion disc, dust torus)	yes	no	no
MAGPHYS ^b	H α and H β only	yes	no	yes	no	no
GALMC ^c	yes (approximate)	no	no	yes (+ photo- z)	no	no
BAYESED ^d	no	no	partial (dust torus)	yes (+ photo- z)	no	no
BEAGLE (this work)	yes	no	no	yes (+ photo- z)	yes	yes
BEAGLE 2.0 (in preparation)	yes	yes	yes	yes (+ photo- z)	yes	yes

Notes. ^a(Burgarella et al. 2005; Noll et al. 2009; Ciesla et al. 2015).

^b(da Cunha et al. 2008, 2015).

^c(Acquaviva et al. 2011, 2015).

^d(Han & Han 2012, 2014).

intergalactic medium. The tool also includes a flexible parametrization of galaxy star formation and chemical enrichment histories, which can be drawn from analytic functions or from different flavours of galaxy formation models, such as phenomenological and semi-analytic models and hydro-dynamic simulations. A main strength of the BEAGLE tool with respect to other existing spectral analysis tools is the flexible, modular implementation of sophisticated prescriptions for the production of light and its transfer through the interstellar and the intergalactic media, in a physically consistent way. This enables one to adapt model complexity (i.e. the number of adjustable parameters let to vary freely; Section 3.2) to the available observational constraints, without having to sacrifice the physical coherence of the model. Statistical inference on galaxy physical parameters from observations is achieved by means of a Bayesian approach. Unlike widely used statistical techniques focusing on simple point-wise estimates of best-fitting model parameters (e.g. minimum χ^2) and confidence intervals (e.g. $\Delta\chi^2 < 1$), this approach allows the rigorous propagation of observational uncertainties into the output statistical constraints on model parameters. It can also reveal correlations (i.e. degeneracies) among model parameters, at the origin of multimodal solutions, and provides a well-defined framework to account for parameter interdependency, to explore the properties of galaxy populations beyond those of single galaxies (through hierarchical modelling; Section 3.1) and to incorporate instrumental effects (Section 2.9).

We presented a first application of the BEAGLE tool to interpret the photometric SEDs of 9927 galaxies in the redshift range $0.1 \lesssim z \lesssim 8$ from the UVUDF sample of Rafelski et al. (2015). Adopting a relatively simple model with seven free parameters (stellar mass, star formation time-scale, metallicity, specific star formation rate, attenuation by dust, redshift of observation and formation redshift of the oldest stars; see Table 3), we find that the photometric

redshifts derived using the BEAGLE tool are globally consistent with the spectroscopic redshifts available for a small sub-sample of UVUDF galaxies and with the redshifts derived by Rafelski et al. (2015) for the full sample using two standard, dedicated photometric-redshift codes, BPZ and EAZY. The statistical sophistication of the BEAGLE tool allows us to gain unique quantitative insight into the origin of occasional discrepancies between photometric and spectroscopic redshifts, and between photometric redshifts estimated using different codes. Such outliers appear to arise mainly from the presence of multiple modes of comparable probability in the posterior probability distribution of model parameters, corresponding to different redshifts. In fact, the accurate Bayesian characterization of the posterior probability distribution of model parameters with the BEAGLE tool allows a rigorous study of multiple solutions and degeneracies between model parameters (Section 4.4). We have illustrated the strength of posterior predictive checks in the framework of the BEAGLE tool to identify and interpret systematic offsets between models and data, pointing either to limitations in the data or necessary improvements of the models (Section 5.1). In the case of the UVUDF sample, a global (Bayesian) goodness-of-fit test indicates that the simplified model mentioned above reproduces well the photometric SEDs of 93 per cent of all galaxies in the catalogue. A complementary, graphical posterior predictive check further shows the potential of this approach to characterize systematic errors in the data, and limitations in the adopted physical model (Section 5.1).

The flexible BEAGLE tool is designed to evolve as more modules are incorporated to account for new physical ingredients (Fig. 2). The next version of the tool, currently in development (BEAGLE 2.0), will include several novelties: (i) an enlarged grid of photoionization models describing the emission from gas in wider ranges of C/N/O abundance ratios, gas densities and IMF upper mass cutoffs

than considered here (Gutkin et al. 2016). This is important to reproduce and interpret the emission from chemically pristine galaxies (e.g. Erb et al. 2010; Stark et al. 2014); (ii) a model to describe the emission from AGN narrow-line regions, fully consistent with the models of nebular emission from stellar populations (Feltre, Charlot & Gutkin 2016, see in particular their section 2). This will allow the reliable exploitation of emission-line diagnostic diagrams at rest-frame ultraviolet and optical wavelengths to interpret the emission from active and inactive galaxies (e.g. Baldwin, Phillips & Terlevich 1981; Feltre, Charlot & Gutkin 2016), hence opening a new window on studies of the co-evolution of black holes and galaxies; and (iii) a model to describe the ultraviolet and optical absorption features from stars and the neutral ISM in and around galaxies (Vidal-García et al., in preparation). Finally, to extend the capabilities of the BEAGLE tool at mid- and far-infrared wavelengths, we also plan to include models to describe the emission from dust heated by stars (da Cunha et al. 2008) and an AGN (Fritz, Franceschini & Hatziminaoglou 2006; Feltre et al. 2012).

With the addition of these and other future modules, the BEAGLE tool will incorporate a panchromatic, physically consistent description of galaxy SEDs. Together with the highly flexible implementation of star formation and chemical enrichment histories of galaxies, dust attenuation and IGM absorption, this will allow the coherent modelling and interpretation of any combination of photometric and spectroscopic galaxy observation, such as those gathered by modern ground-based (e.g. Atacama Large Millimeter Array, Extremely Large Telescopes) and space-based (e.g. *James Webb Space Telescope*) observatories, in terms of powerful constraints on galaxy formation models.

ACKNOWLEDGEMENTS

We thank the anonymous referee for a thorough reading of the manuscript, which helped us improve this paper. We are also grateful to Emma Curtis-Lake and David Stenning for valuable comments on an early draft of the paper, which helped greatly improve our work. JC acknowledges support from FAPESP (project no. 2013/11837-9). JC and SC acknowledge support from the European Research Council via an Advanced Grant under grant agreement no. 321323 (NEOGAL).

REFERENCES

- Acquaviva V., Gawiser E., Guaita L., 2011, *ApJ*, 737, 47
 Acquaviva V., Raichoor A., Gawiser E., 2015, *ApJ*, 804, 8
 Arimoto N., Yoshii Y., 1987, *A&A*, 173, 23
 Arnouts S., Cristiani S., Moscardini L., Matarrese S., Lucchin F., Fontana A., Giallongo E., 1999, *MNRAS*, 310, 540
 Baldwin J. A., Phillips M. M., Terlevich R., 1981, *PASP*, 93, 5
 Bauer A. E. et al., 2013, *MNRAS*, 434, 209
 Beckwith S. V. W. et al., 2006, *AJ*, 132, 1729
 Benítez N., 2000, *ApJ*, 536, 571
 Blanton M. R., Roweis S., 2007, *AJ*, 133, 734
 Bolzonella M., Miralles J.-M., Pelló R., 2000, *A&A*, 363, 476
 Bolzonella M. et al., 2010, *A&A*, 524, A76
 Bouchet P., Lequeux J., Maurice E., Prevot L., Prevot-Burnichon M. L., 1985, *A&A*, 149, 330
 Bouwens R. J. et al., 2011, *ApJ*, 737, 90
 Brammer G. B., van Dokkum P. G., Coppi P., 2008, *ApJ*, 686, 1503
 Brammer G. B. et al., 2011, *ApJ*, 739, 24
 Bressan A., Chiosi C., Fagotto F., 1994, *ApJS*, 94, 63
 Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, *MNRAS*, 427, 127
 Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, 351, 1151
 Bruzual G., 1983, *ApJ*, 273, 105
 Bruzual G., Charlot S., 2003, *MNRAS*, 344, 1000
 Burgarella D., Buat V., Iglesias-Páramo J., 2005, *MNRAS*, 360, 1413
 Buzzoni A., 1989, *ApJS*, 71, 817
 Calzetti D., Kinney A. L., Storchi-Bergmann T., 1994, *ApJ*, 429, 582
 Cappellari M., Emsellem E., 2004, *PASP*, 116, 138
 Cappellari M. et al., 2012, *Nature*, 484, 485
 Cardelli J. A., Clayton G. C., Mathis J. S., 1988, *ApJ*, 329, L33
 Cardelli J. A., Clayton G. C., Mathis J. S., 1989, *ApJ*, 345, 245
 Cerviño M., Luridiana V., Castander F. J., 2000, *A&A*, 360, L5
 Chabrier G., 2003, *PASP*, 115, 763
 Chabrier G., Hennebelle P., Charlot S., 2014, *ApJ*, 796, 75
 Charlot S., Bruzual A. G., 1991, *ApJ*, 367, 126
 Charlot S., Fall S. M., 2000, *ApJ*, 539, 718 (CF00)
 Charlot S., Longhetti M., 2001, *MNRAS*, 323, 887
 Charlot S., Worthey G., Bressan A., 1996, *ApJ*, 457, 625
 Chernetoff I., 2010, in Leitherer C., Bennett P., Morris P., van Loon J., eds, *ASP Conf. Ser. Vol. 425, Hot and Cool: Bridging Gaps in Massive Star Evolution*. Astron. Soc. Pac., San Francisco, p. 237
 Chevallard J., Charlot S., Wandelt B., Wild V., 2013, *MNRAS*, 432, 2061 (C13)
 Cid Fernandes R., Gu Q., Melnick J., Terlevich E., Terlevich R., Kunth D., Rodrigues Lacerda R., Joguet B., 2004, *MNRAS*, 355, 273
 Cid Fernandes R., Mateus A., Sodré L., Stasińska G., Gomes J. M., 2005, *MNRAS*, 358, 363
 Cid Fernandes R., Asari N. V., Sodré L., Stasińska G., Mateus A., Torres-Papaqui J. P., Schoenell W., 2007, *MNRAS*, 375, L16
 Ciesla L. et al., 2015, *A&A*, 576, A10
 Citro A., Pozzetti L., Moresco M., Cimatti A., 2016, *A&A*, 592, A19
 Clayton G. C., Martin P. G., 1985, *ApJ*, 288, 558
 Codis S., Pichon C., Devriendt J., Slyz A., Pogosyan D., Dubois Y., Sousbie T., 2012, *MNRAS*, 427, 3320
 Coelho P. R. T., 2014, *MNRAS*, 440, 1027
 Coelho P., Bruzual G., Charlot S., Weiss A., Barbuy B., Ferguson J. W., 2007, *MNRAS*, 382, 498
 Collister A. A., Lahav O., 2004, *PASP*, 116, 345
 Conroy C., Gunn J. E., 2010, *ApJ*, 712, 833
 Conroy C., van Dokkum P., 2012a, *ApJ*, 747, 69
 Conroy C., van Dokkum P. G., 2012b, *ApJ*, 760, 71
 Conroy C., Gunn J. E., White M., 2009, *ApJ*, 699, 486
 Conroy C., White M., Gunn J. E., 2010a, *ApJ*, 708, 58
 Conroy C., Schiminovich D., Blanton M. R., 2010b, *ApJ*, 718, 184
 Crain R. A. et al., 2015, *MNRAS*, 450, 1937
 da Cunha E., Charlot S., Elbaz D., 2008, *MNRAS*, 388, 1595
 da Cunha E., Eminian C., Charlot S., Blaizot J., 2010, *MNRAS*, 403, 1894
 da Cunha E. et al., 2015, *ApJ*, 806, 110
 Dahlen T. et al., 2013, *ApJ*, 775, 93
 Edmondson E. M., Miller L., Wolf C., 2006, *MNRAS*, 371, 1693
 Ellis R. S. et al., 2013, *ApJ*, 763, L7
 Erb D. K., Pettini M., Shapley A. E., Steidel C. C., Law D. R., Reddy N. A., 2010, *ApJ*, 719, 1168
 Feltre A., Hatziminaoglou E., Fritz J., Franceschini A., 2012, *MNRAS*, 426, 120
 Feltre A., Charlot S., Gutkin J., 2016, *MNRAS*, 456, 3354
 Ferland G. J. et al., 2013, *Rev. Mex. Astron. Astrofis.*, 49, 137
 Feroz F., Hobson M. P., 2008, *MNRAS*, 384, 449
 Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
 Fioc M., Rocca-Volmerange B., 1997, *A&A*, 326, 950
 Fontanot F., Somerville R. S., Silva L., Monaco P., Skibba R., 2009, *MNRAS*, 392, 553
 Fritz J., Franceschini A., Hatziminaoglou E., 2006, *MNRAS*, 366, 767
 Gallazzi A., Bell E. F., 2009, *ApJS*, 185, 253
 Gallazzi A., Charlot S., Brinchmann J., White S. D. M., Tremonti C. A., 2005, *MNRAS*, 362, 41
 Gelman A., Meng X.-L., Stern H., 1996, *Stat. Sin.*, 6, 733

- Gerhard O. E., 1993, *MNRAS*, 265, 213
- Gonzalez-Perez V., Lacey C. G., Baugh C. M., Lagos C. D. P., Helly J., Campbell D. J. R., Mitchell P. D., 2014, *MNRAS*, 439, 264
- Granato G. L., Lacey C. G., Silva L., Bressan A., Baugh C. M., Cole S., Frenk C. S., 2000, *ApJ*, 542, 710
- Grogin N. A. et al., 2011, *ApJS*, 197, 35
- Guideroni B., Rocca-Volmerange B., 1987, *A&A*, 186, 1
- Guo Y. et al., 2011, *ApJ*, 735, 18
- Gutkin J., Charlot S., Bruzual G., 2016, *MNRAS*, preprint ([arXiv:1607.06086](https://arxiv.org/abs/1607.06086))
- Guttman I., 1967, *J. R. Stat. Soc. B*, 29, 83
- Haas M. R., Schaye J., Booth C. M., Dalla Vecchia C., Springel V., Theuns T., Wiersma R. P. C., 2013a, *MNRAS*, 435, 2931
- Haas M. R., Schaye J., Booth C. M., Dalla Vecchia C., Springel V., Theuns T., Wiersma R. P. C., 2013b, *MNRAS*, 435, 2955
- Hahn O., Teyssier R., Carollo C. M., 2010, *MNRAS*, 405, 274
- Han Y., Han Z., 2012, *ApJ*, 749, 123
- Han Y., Han Z., 2014, *ApJS*, 215, 2
- Harrison C. M., Meiksin A., Stock D., 2011, preprint ([arXiv:1105.6208](https://arxiv.org/abs/1105.6208))
- Heavens A. F., Jimenez R., Lahav O., 2000, *MNRAS*, 317, 965
- Heavens A., Panter B., Jimenez R., Dunlop J., 2004, *Nature*, 428, 625
- Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R., Guo Q., Lemson G., Springel V., Overzier R., 2015, *MNRAS*, 451, 2663
- Hernán-Caballero A. et al., 2013, *MNRAS*, 434, 2136
- Höfner S., 2009, in Henning T., Grün E., Steinacker J., eds, *ASP Conf. Ser. Vol. 414, Cosmic Dust - Near and Far*. Astron. Soc. Pac., San Francisco, p. 3
- Hogan R., Fairbairn M., Seeburn N., 2015, *MNRAS*, 449, 2040
- Holwerda B. W., Keel W. C., 2013, *A&A*, 556, A42
- Holwerda B. W., Keel W. C., Bolton A., 2007, *AJ*, 134, 2385
- Ilbert O. et al., 2006, *A&A*, 457, 841
- Ilbert O. et al., 2013, *A&A*, 556, A55
- Inoue A. K., Shimizu I., Iwata I., Tanaka M., 2014, *MNRAS*, 442, 1805
- Jasche J., Leclercq F., Wandelt B. D., 2015, *J. Cosmol. Astropart. Phys.*, 1, 36
- Jeffreys H., 1961, *Theory of Probability*. Oxford Univ. Press, Cambridge
- Jones A. P., 2004, in Witt A. N., Clayton G. C., Draine B. T., eds, *ASP Conf. Ser. Vol. 309, Astrophysics of Dust*. Astron. Soc. Pac., San Francisco, p. 347
- Jones A. P., Nuth J. A., 2011, *A&A*, 530, A44
- Kass R. E., Raftery A. E., 1995, *J. Am. Stat. Assoc.*, 90, 773
- Kauffmann G. et al., 2003, *MNRAS*, 341, 33
- Kawinwanichakij L. et al., 2016, *ApJ*, 817, 9
- Kelson D. D., Illingworth G. D., van Dokkum P. G., Franx M., 2000, *ApJ*, 531, 159
- Kochiashvili I. et al., 2015, *A&A*, 580, A42
- Koekemoer A. M. et al., 2011, *ApJS*, 197, 36
- Koekemoer A. M. et al., 2013, *ApJS*, 209, 3
- Koleva M., Prugniel P., Bouchard A., Wu Y., 2009, *A&A*, 501, 1269
- Kriek M., van Dokkum P. G., Labbé I., Franx M., Illingworth G. D., Marchesini D., Quadri R. F., 2009, *ApJ*, 700, 221
- Kron R. G., 1980, *ApJS*, 43, 305
- Kroupa P., 2001, *MNRAS*, 322, 231
- La Barbera F., Ferreras I., Vazdekis A., de la Rosa I. G., de Carvalho R. R., Trevisan M., Falcón-Barroso J., Ricciardelli E., 2013, *MNRAS*, 433, 3017
- Lackner C. N., Cen R., Ostriker J. P., Joung M. R., 2012, *MNRAS*, 425, 641
- Lagos C. d. P., Davis T. A., Lacey C. G., Zwaan M. A., Baugh C. M., Gonzalez-Perez V., Padilla N. D., 2014, *MNRAS*, 443, 1002
- Le Borgne D., Rocca-Volmerange B., Prugniel P., Lançon A., Fioc M., Soubiran C., 2004, *A&A*, 425, 881
- Le Borgne J.-F. et al., 2003, *A&A*, 402, 433
- Leitherer C., Heckman T. M., 1995, *ApJS*, 96, 9
- Lewis A., 2015, *GetDist: Kernel Density Estimation*, available at <http://cosmologist.info/notes/GetDist.pdf>.
- Lewis A., Bridle S., 2002, *Phys. Rev. D*, 66, 103511
- López Fernández R. et al., 2016, *MNRAS*, 458, 184
- Loredo T. J., 2012, preprint ([arXiv:1208.3036](https://arxiv.org/abs/1208.3036))
- Lu Y. et al., 2014, *ApJ*, 795, 123
- Lundgren B. F. et al., 2014, *ApJ*, 780, 34
- McConnell N. J., Lu J. R., Mann A. W., 2016, *ApJ*, 821, 39
- McDermid R. M. et al., 2015, *MNRAS*, 448, 3484
- Madau P., 1995, *ApJ*, 441, 18
- Maraston C., 2005, *MNRAS*, 362, 799
- Maraston C., Strömbäck G., 2011, *MNRAS*, 418, 2785
- Marigo P., Bressan A., Nanni A., Girardi L., Pumo M. L., 2013, *MNRAS*, 434, 488
- Martins L. P., Coelho P., 2007, *MNRAS*, 381, 1329
- Milone A., Barbuy B., Schiavon R. P., 2000, *AJ*, 120, 131
- Mobasher B. et al., 2015, *ApJ*, 808, 101
- Mortlock A. et al., 2015, *MNRAS*, 447, 2
- Muzzin A. et al., 2013, *ApJ*, 777, 18
- Noll S., Burgarella D., Giovannoli E., Buat V., Marcellac D., Muñoz-Mateos J. C., 2009, *A&A*, 507, 1793
- Ocvirk P., Pichon C., Lançon A., Thiébaud E., 2006, *MNRAS*, 365, 74
- Oesch P. A. et al., 2010a, *ApJ*, 709, L16
- Oesch P. A. et al., 2010b, *ApJ*, 709, L21
- Pacifici C., Charlot S., Blaizot J., Brinchmann J., 2012, *MNRAS*, 421, 2002
- Pacifici C., Kassin S. A., Weiner B., Charlot S., Gardner J. P., 2013, *ApJ*, 762, L15
- Pacifici C. et al., 2015, *MNRAS*, 447, 786
- Panther B., Jimenez R., Heavens A. F., Charlot S., 2007, *MNRAS*, 378, 1550
- Papovich C., Dickinson M., Ferguson H. C., 2001, *ApJ*, 559, 620
- Percival S. M., Salaris M., 2009, *ApJ*, 703, 1123
- Péroux C., McMahon R. G., Storrie-Lombardi L. J., Irwin M. J., 2003, *MNRAS*, 346, 1103
- Planck Collaboration, 2015, preprint ([arXiv:1502.01589](https://arxiv.org/abs/1502.01589))
- Pozzetti L. et al., 2007, *A&A*, 474, 443
- Prevot M. L., Lequeux J., Prevot L., Maurice E., Rocca-Volmerange B., 1984, *A&A*, 132, 389
- Prugniel P., Soubiran C., 2001, *A&A*, 369, 1048
- Prugniel P., Soubiran C., Koleva M., Le Borgne D., 2007, preprint ([arXiv:e-prints](https://arxiv.org/abs/e-prints))
- Rafelski M. et al., 2015, *AJ*, 150, 31
- Rafieferantsoa M., Davé R., Anglés-Alcázar D., Katz N., Kollmeier J. A., Oppenheimer B. D., 2015, *MNRAS*, 453, 3980
- Rahman M., Ménard B., Scranton R., 2016, *MNRAS*, 457, 3912
- Rauch M., 1998, *ARA&A*, 36, 267
- Rocha-Pinto H. J., Maciel W. J., Scalo J., Flynn C., 2000a, *A&A*, 358, 850
- Rocha-Pinto H. J., Scalo J., Maciel W. J., Flynn C., 2000b, *A&A*, 358, 869
- Rubin D. B. et al., 1984, *Ann. Stat.*, 12, 1151
- Salpeter E. E., 1955, *ApJ*, 121, 161
- Sánchez-Blázquez P. et al., 2006, *MNRAS*, 371, 703
- Scannapieco C. et al., 2012, *MNRAS*, 423, 1726
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schmidt S. J., Ménard B., Scranton R., Morrison C., McBride C. K., 2013, *MNRAS*, 431, 3307
- Serra P., Amblard A., Temi P., Burgarella D., Giovannoli E., Buat V., Noll S., Im S., 2011, *ApJ*, 740, 22
- Silva L., Granato G. L., Bressan A., Danese L., 1998, *ApJ*, 509, 103
- Skilling J. et al., 2006, *Bayesian Anal.*, 1, 833
- Smith R. J., 2014, *MNRAS*, 443, L69
- Sonnenfeld A., Treu T., Marshall P. J., Suyu S. H., Gavazzi R., Auger M. W., Nipoti C., 2015, *ApJ*, 800, 94
- Spiniello C., Trager S., Koopmans L. V. E., Conroy C., 2014, *MNRAS*, 438, 1483
- Stark D. P. et al., 2014, *MNRAS*, 445, 3200
- Tang B., Worthey G., 2015, *MNRAS*, 453, 4431
- Teplitz H. I. et al., 2013, *AJ*, 146, 159
- Thompson W. J., 1999, *Comput. Sci. Eng.*, 1, 84
- Tinsley B. M., 1978, *ApJ*, 222, 14
- Tojeiro R., Heavens A. F., Jimenez R., Panter B., 2007, *MNRAS*, 381, 1252
- Tojeiro R. et al., 2013, *MNRAS*, 432, 359
- Torrey P., Vogelsberger M., Genel S., Sijacki D., Springel V., Hernquist L., 2014, *MNRAS*, 438, 1985
- Tuffs R. J., Popescu C. C., Völk H. J., Klyafis N. D., Dopita M. A., 2004, *A&A*, 419, 821

- Valdes F., Gupta R., Rose J. A., Singh H. P., Bell D. J., 2004, *ApJS*, 152, 251
- van der Marel R. P., Franx M., 1993, *ApJ*, 407, 525
- van Dokkum P. G., Conroy C., 2010, *Nature*, 468, 940
- Vazdekis A., Sánchez-Blázquez P., Falcón-Barroso J., Cenarro A. J., Beasley M. A., Cardiel N., Gorgas J., Peletier R. F., 2010, *MNRAS*, 404, 1639
- Vazdekis A. et al., 2015, *MNRAS*, 449, 1177
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
- Walcher C. J., Coelho P., Gallazzi A., Charlot S., 2009, *MNRAS*, 398, L44
- Wild V., Kauffmann G., Heckman T., Charlot S., Lemson G., Brinchmann J., Reichard T., Pasquali A., 2007, *MNRAS*, 381, 543
- Wild V. et al., 2011, *MNRAS*, 410, 1593
- Wise M. W., Silva D. R., 1996, *ApJ*, 461, 155
- Wolf M. J., Drory N., Gebhardt K., Hill G. J., 2007, *ApJ*, 655, 179
- Wolfe A. M., Gawiser E., Prochaska J. X., 2005, *ARA&A*, 43, 861
- Worthey G., 1994, *ApJS*, 95, 107
- Worthey G., Tang B., Serven J., 2014, *ApJ*, 783, 20
- Wuyts S. et al., 2011, *ApJ*, 742, 96
- Zieleniewski S., Houghton R. C. W., Thatte N., Davies R. L., 2015, *MNRAS*, 452, 597

APPENDIX A: NESTED SAMPLING WITH MULTINEST

The MULTINEST package (Feroz & Hobson 2008; Feroz et al. 2009) allows one to explore complex, multidimensional posterior probability distributions by appealing to the nested sampling algorithm of Skilling et al. (2006). This algorithm was initially developed to calculate the Bayesian evidence (see equations 24 and A1), a quantity which requires long computational times when appealing to ‘standard’ MCMC-based methods. The nested sampling algorithm transforms the problem of computing the multidimensional evidence integral

$$\mathcal{Z} = \int \pi(\Theta) \mathcal{L}(\Theta) d\Theta, \quad (\text{A1})$$

where $\pi(\Theta)$ and $\mathcal{L}(\Theta)$ refer to, respectively, the prior distribution and likelihood function for a model described by the parameters Θ , into the computation of a one-dimensional integral. This is based on the property that Bayesian evidence can be expressed as

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X) dX, \quad (\text{A2})$$

where the function $\mathcal{L}(X) = F^{-1}[X(l)]$ is the inverse of the prior hyper-volume X over iso-likelihood hyper-surfaces $\mathcal{L}(\Theta) = l$, expressed by

$$X(l) = \int_{\mathcal{L}(\Theta) > l} \pi(\Theta) d\Theta. \quad (\text{A3})$$

The domain of this integral is the volume enclosed by the iso-likelihood surface defined by the parameter l . Using the above transformation, the evidence can be computed in a straightforward way by numerically integrating equation (A2) after calculating the function $X(l)$ for increasing likelihood thresholds l .

The difficulty of applying equation (A2) to compute the evidence \mathcal{Z} lies in the evaluation of the function $X(l)$ for increasing values of l , which correspond to iso-likelihood surfaces encompassing smaller and smaller regions of the prior volume. For this, MULTINEST employs the ‘simultaneous ellipsoidal nested sampling method’: at iteration $i = 1$, the algorithm starts by computing the likelihood function $\mathcal{L}(\Theta^k)$ for N sets of parameters Θ^k (called ‘active points’), with $1 \leq k \leq N$, drawn randomly from the prior

probability distribution. The point with lowest likelihood l_{\min}^i is removed, and a new point Θ^i is drawn from the prior with the requirement that $\mathcal{L}(\Theta^i) > l_{\min}^i$. This is achieved by decomposing the prior volume into ellipsoids bounded by iso-likelihood surfaces corresponding to $\mathcal{L}(\Theta) = l_{\min}^i$, whose number and shapes are optimized via an ‘expectation-maximization’ method. At each iteration, the algorithm enables in this way the computation of $X(l_{\min}^i)$ using equation (A3), while the prior volume shrinks to regions of higher likelihood. The algorithm stops when the product of the shrunk prior volume $X(l_{\min}^i)$ and the likelihood $\mathcal{L}(\Theta^i)$, where Θ^i is the active point with largest likelihood, falls below an adopted ‘evidence tolerance factor’, i.e. $\mathcal{L}(\Theta^i) X(l_{\min}^i) < \mathcal{Z}_{\text{tol}}$.

As a side product of the evidence calculation, the MULTINEST algorithm computes the posterior probability $P(\Theta | D, H)$ of any set of parameters Θ ever drawn during an iteration, providing an ensemble of $N_{\text{out}}^{\text{tot}}$ sets of parameters weighted by the posterior probability distribution. This ensemble can be used to perform inference on model parameters, such as the computation of posterior means and marginal and joint distributions.

The interest of MULTINEST relative to other algorithms (such as MCMC) in the context of the current study is the possibility to identify multiple modes in the posterior distribution of model parameters, and to evaluate the ‘local’ evidence in each mode (see Section 5.2). In practice, to analyse the SEDs of UVUDF galaxies with the model described in Section 4.2 using MULTINEST, we must specify three parameters: the number of active points, the evidence tolerance factor and the ‘sampling efficiency’. The number of active points must be large enough to probe all potential modes of the posterior probability distribution, which we expect to be around 2–3 at most, mainly caused by degeneracies between redshift, mass and dust attenuation. We fix $N = 300$, similar to the value adopted by Feroz et al. (2009) to estimate cosmological parameters. We also fix $\mathcal{Z}_{\text{tol}} = 0.1$, lower than the value of 0.5 suggested by Feroz et al. (2009). This is to accurately evaluate the evidence when computing the relative probabilities of different modes, for those galaxies exhibiting multimodal solutions. The sampling efficiency is an additional factor introduced to account for the potentially inaccurate characterization of the iso-likelihood surfaces by the ellipsoidal decomposition mentioned above. This parameter is the inverse of the factor by which the ellipsoids are ‘inflated’ at each iteration, before a new candidate active point is drawn from these ellipsoids. A low sampling efficiency guarantees a more accurate evidence evaluation, but with a higher chance of drawing candidate active points that do not satisfy the criterion $\mathcal{L}(\Theta^i) > l_{\min}^i$, and hence must be rejected and redrawn. We adopt a sampling efficiency of 0.3, as recommended by Feroz et al. (2009) for an accurate estimate of the evidence. Finally, another adjustment of MULTINEST pertains to the ‘clustering algorithm’ employed to search for different modes in the multidimensional posterior probability distribution of model parameters. Since this type of algorithm loses accuracy with increasing dimensionality, we restrict the clustering analysis here to those three parameters most correlated with one another, namely redshift, stellar mass and attenuation optical depth.

APPENDIX B: WEIGHTED SAMPLING WITH REPLACEMENT

We outline here the procedure used to produce ‘replicated observations’ to perform posterior predictive checks (Section 5.1). For each observed galaxy analysed with the BEAGLE tool, MULTINEST outputs an ensemble of $N_{\text{out}}^{\text{tot}}$ sets of parameters weighted by

the posterior probability distribution (Appendix A). We consider the subset N_{out} of these corresponding to posterior probabilities greater than 10^{-4} times that of the parameter set with largest posterior probability, noted $\tilde{\Theta}$, i.e. the N_{out} sets of parameters with $P(\Theta | \mathbf{D}, H) > 10^{-4} P(\tilde{\Theta} | \mathbf{D}, H)$. From this ensemble of typically $N_{\text{out}} \sim 10^3$ sets of model parameters, we draw N_{rep} replicated parameter sets Θ^k , with $1 \leq k \leq N_{\text{rep}}$, by appealing to a ‘weighted sampling with replacement algorithm’. This means that the probability of drawing the parameter set Θ^i (with $1 \leq i \leq N_{\text{out}}$) is equal to

the posterior probability $P(\Theta^i | \mathbf{D}, H)$ (the ‘weight’). The drawn sets of parameters Θ^k are not removed from the ensemble of N_{out} points (hence the term ‘with replacement’), implying that any Θ^i can be drawn multiple times (in practice, this is the case only for those sets of parameters Θ^i with the largest weights).

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