RAS Techniques and Instruments

RASTAI **2,** 78–90 (2023) https://doi.org/10.1093/rasti/rzad004
Advance Access publication 2023 February 8

The entropy of galaxy spectra: how much information is encoded?

Ignacio Ferreras [1,2,3★] Ofer Lahav,[1] Rachel S. Somerville[4,5] and Joseph Silk[6,7,8]

[1]*Department of Physics and Astronomy, University College London, London WC1E 6BT, UK*
[2]*Instituto de Astrofísica de Canarias, Calle Vía Láctea s/n, E-38205 La Laguna, Tenerife, Spain*
[3]*Departamento de Astrofísica, Universidad de La Laguna, E-38206 La Laguna, Tenerife, Spain*
[4]*Center for Computational Astrophysics, Flatiron Institute, New York, NY 10010, USA*
[5]*Department of Physics and Astronomy, Rutgers University, 136 Frelinghuysen Rd, Piscataway, NY 08854, USA*
[6]*Institut d'Astrophysique de Paris, UMR7095:CNRS & UPMC, Sorbonne University, F-75014, Paris, France*
[7]*Department of Physics and Astronomy, The Johns Hopkins University Homewood Campus, Baltimore, MD 21218, USA*
[8]*Beecroft Institute of Particle Astrophysics and Cosmology, Department of Physics, University of Oxford, Keble Road, Oxford OX1 3RH, UK*

Accepted 2023 January 20. Received 2022 November 11; in original form 2022 May 27

ABSTRACT

The inverse problem of extracting the stellar population content of galaxy spectra is analysed here from a basic standpoint based on information theory. By interpreting spectra as probability distribution functions, we find that galaxy spectra have high entropy, thus leading to a rather low effective information content. The highest variation in entropy is unsurprisingly found in regions that have been well studied for decades with the conventional approach. We target a set of six spectral regions that show the highest variation in entropy – the 4000 Å break being the most informative one. As a test case with real data, we measure the entropy of a set of high-quality spectra from the Sloan Digital Sky Survey, and contrast entropy-based results with the traditional method based on line strengths. The data are classified into star-forming (SF), quiescent (Q), and active galactic nucleus (AGN) galaxies, and show – independently of any physical model – that AGN spectra can be interpreted as a transition between SF and Q galaxies, with SF galaxies featuring a more diverse variation in entropy. The high level of entanglement complicates the determination of population parameters in a robust, unbiased way, and affects traditional methods that compare models with observations, as well as machine learning (especially deep learning) algorithms that rely on the statistical properties of the data to assess the variations among spectra. Entropy provides a new avenue to improve population synthesis models so that they give a more faithful representation of real galaxy spectra.

Key words: galaxies: stellar content – methods: data analysis – methods: statistical – techniques: spectroscopic.

1. INTRODUCTION

Stars and gas constitute the most fundamental observables in extragalactic astrophysics. While the gaseous component, in its different guises, mostly reflects the ongoing physical state, the stellar populations encode valuable information about the past formation history, both through their collisionless nature – that keep track of the past dynamical history – and their age and chemical composition – that trace the past star-formation history. The most relevant spectral window to probe the stellar content is the interval between near ultraviolet and near infrared, including the optical range, where most of the light from stellar photospheres is emitted. The standard way to study the underlying stellar population focuses on comparisons of photospectroscopic observables with population synthesis models [to name a few, Bruzual & Charlot 1993, Bruzual & Charlot 2003 (hereafter BC03); Worthey 1994; Fioc & Rocca-Volmerange 1997; Maraston 2005; Conroy & Gunn 2010; Vazdekis et al. 2010)], that combine our understanding of stellar formation and evolution, along with a determination (empirical, theoretical, or mixed) of the stellar atmospheres (see e.g. Walcher et al. 2011; Conroy 2013, for a general view of these models).

While this paper focuses on an alternative analysis of galaxy spectra, there is a large number of papers developing the traditional methods of extracting information from the observations focusing on the determination of stellar age, metallicity, and targeted abundance ratios (see e.g. González 1993; Worthey et al. 1994; Trager et al. 2000; Gallazzi et al. 2005; Thomas et al. 2005; Graves, Faber & Schiavon 2009; La Barbera et al. 2013). These models range from methods based on the concept of a simple stellar population (SSP; uniquely defined by a stellar initial mass function and age, along with a fixed chemical composition) to complex mixtures spanning a range of those parameters. The most detailed models invoke comparisons of targeted line strengths (LSs; such as the Lick system, and variations thereof) or full spectral fitting (e.g. Cid Fernandes et al. 2005; Ocvirk et al. 2006; Cappellari 2012). However, all these methods are hampered by the so-called age-metallicity-dust degeneracy whereby changes in age can mimic the effects on the photometric and spectroscopic information of a change in chemical composition (Worthey 1994). Most notably, this degeneracy is present at all scales regarding spectral resolution, with similar behaviour in broad-band photometry and spectroscopy (Ferreras, Charlot & Silk 1999). It is

★ E-mail: iferreras@iac.es

in fact a major paradox that the amount of information that can be gathered either from a few broad-band colours, or from thousands of fluxes in a spectrum is rather comparable, reflecting the strong correlation of spectral features due to the underlying astrophysics.

Alternatively to model comparison methods, other works adopt data-driven multivariate techniques aimed at disentangling the information via principal component analysis (PCA; Ronen, Aragon-Salamanca & Lahav 1999; Ferreras et al. 2006), factor analysis (Nolan, Raychaudhury & Kabán 2007), independent component analysis (Kabán, Nolan & Raychaudhury 2005), Fisher information matrices (Heavens, Jimenez & Lahav 2000), Bayesian latent variable analysis (Nolan et al. 2006), or the information bottleneck (Slonim et al. 2001; Ferreras 2012). The central tenet of these methods is that the spectra of galaxies represent superpositions of more basic units – eventually individual stellar spectra – so that a statistical analysis would reveal the basis spectra that provide a more fundamental set to decipher the formation and evolution of galaxies. For instance, it was found that a basis set comprising just a few components could accurately describe sets of thousands of galaxies (Madgwick et al. 2002), and these components were mostly dependent on age, metallicity, or recent star formation (Madgwick et al. 2003; Rogers et al. 2007; Wild et al. 2014). This drastic dimensional reduction can be exploited, for instance, to speed up the retrieval of star-formation histories from spectra (e.g. Panter, Heavens & Jimenez 2003; Tojeiro, et al. 2007). However, while these methods only depend on the input data and are not affected by potential systematics of the synthesis models, the results are rather similar in the sense that only very generic constraints can be produced regarding the star-formation histories. More recently, similar results have been obtained with deep learning methods based on convolutional neural networks (see e.g. Lovell et al. 2019; Portillo et al. 2020; Liew-Cain et al. 2021; Teimoorinia et al. 2022), where the algorithms are only capable of robustly deriving a few, general properties from galaxy spectra.

This paper performs a novel analysis based on the interpretation of galaxy spectra as an information carrier, by redefining a spectrum as a probability distribution function, with its associated entropy, following the standard approach of information theory. Entropy encodes the amount of 'surprise' in the outcome of a given system. In spectra, this can be naturally defined when the acquisition of a spectrum is interpreted as a photon counting experiment. In this case, the simplest scenarios are represented by a laser beam and a white light source. In the former, the entropy will be very low, given its highly monochromatic nature. In contrast, 'white' light will produce maximum entropy. Throughout this paper we often link entropy to 'information content', as defined in Shannon (1953). This is the most fundamental way of assessing information, as in, e.g. determining the number of bits needed for a full representation of the data. Subsequent definitions of information content focus instead on cross-entropy (see e.g. Slonim et al. 2001; MacKay 2003) involving an additional set of data or a classification scheme.

Note that the concept of information as number of bits needed to encode the spectra is intimately linked to the extraction of stellar population parameters, although in a non-trivial way: Low information would imply either a highly low-dimensional representation of the problem, or deep entanglement that produces strong degeneracies that makes the extraction of the underlying parameters very challenging. As an example, the spectrum of a blackbody radiator depends on a single parameter (temperature), thus allowing for a minimal representation. In stellar populations, one can envision a large number of fundamental parameters (age and chemical composition distribution of the stars, initial mass function, etc.). This paper explores whether the information content somehow reflects this level of complexity or whether the degeneracies result in a much simpler scenario regarding spectra as probability distributions.

Extraction of population parameters from galaxy spectra depends on diverse practical issues – spectral resolution, modelling line widths, etc. – that vary between data sets and plague model comparisons. Our proposal to use entropy content is meant to provide a fundamental procedure to assess the information content of galaxy spectra and the ultimate limits attainable for inferring the parametric nature of the stellar populations. We compare the results from models of population synthesis with actual galaxies from the Sloan Digital Sky Survey (SDSS). We stress that our aim here is not to suggest new spectral regions or to supersede the traditional methodology based on model comparisons, but to explore the inverse problem of extracting stellar population parameters from galaxy spectra from a standpoint based on information theory. As we will see, the analysis illustrates the inherent entanglement that explains why highly sophisticated algorithms, such as those based on deep learning, can only produce generic descriptions of the stellar content. We emphasize that any study aimed at the derivation of detailed properties about the stellar populations from galaxy spectra will be subject to this fundamental limit concerning information content.

The paper has a concise structure, and is organized as follows: The basic methodology based on entropy is presented in Section 2, motivated by the use of models of stellar population synthesis. The models produce interesting results regarding the restricted region where information (negentropy) is substantial, and motivates the selection of a set of targeted spectral windows. The analysis then turns to the spectra of galaxies from the SDSS to confirm whether the suggestions from the models are applicable to real data, presented in Section 3. Finally, our conclusions are given in Section 4 followed by an epilogue that proposes a similar outlook with a more standard treatment based on covariance.

2. THE ENTROPY OF GALAXY SPECTRA

The best way to illustrate the connection between spectroscopy and information theory is to consider the measurement of a spectrum as a photon counting process. The specific flux density can then be interpreted as a conditional probability distribution. For a galaxy g, the flux at wavelength λ is thus

$$\Phi_g(\lambda) \quad \Longrightarrow \quad p(\lambda|g). \tag{1}$$

This formalism was introduced by Slonim et al. (2001) to motivate a classification algorithm based on the so-called 'information bottleneck', whereby a sample of galaxy spectra is subject to an agglomerative binning procedure based on the evolution of entropy, as spectra are progressively binned into classes. The mutual information contained in the class representation of the original set serves as a target to achieve a description of the data with the maximum amount of information encoded into the smallest number of classes.

In this paper, we do not follow this approach. However, this interpretation allows us to exploit the concept of entropy in galaxy spectra as a fundamental way to quantify the way information is stored across the spectral window. The Shannon definition of entropy for a probability distribution is (Shannon & Weaver 1975)

$$H(\Phi) \equiv -\sum_i p(\lambda_i) \log p(\lambda_i). \tag{2}$$

Hereafter, we denote the entropy with the letter H, following standard notation. This methodology allows us, for instance, to quickly identify the language of a given text, by use of the frequency of the different letters of the alphabet, leading to the definition of the

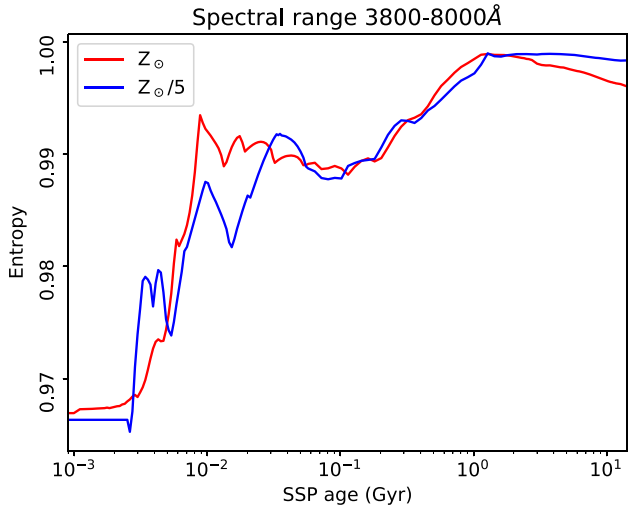

Figure 1 Evolution of entropy with respect to age for a set of SSPs from the models of BC03, corresponding to two values of metallicity, as labelled. The entropy, defined with the standard definition, equation (2), is measured within the 3500–8000 Å spectral range, and is defined such that $H = 1$ corresponds to a maximally uninformative case (i.e. $p(\lambda) =$ constant).

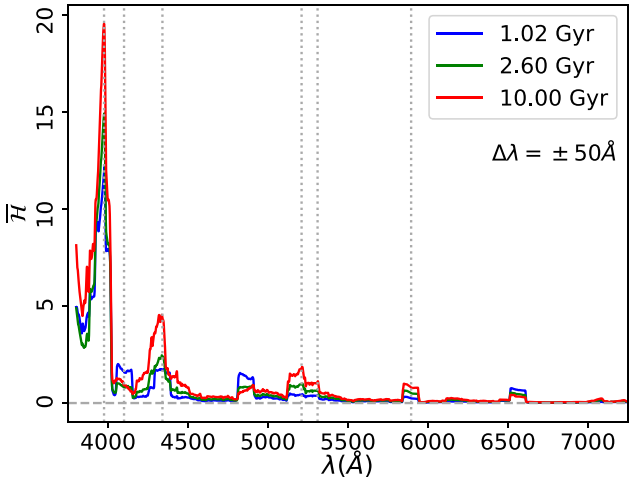

Figure 2 Negentropy spectra ($\overline{\mathcal{H}}(\lambda)$, see equation 3), corresponding to three SSPs from the models of BC03 at solar metallicity, with age as labelled. They have been computed by measuring the entropy in windows of $\delta\lambda =$ 100 Å. The vertical dotted lines mark the wavelengths chosen for the reduced dimensional analysis of SDSS data based on maximum entropy variations.

entropy of a language (Shannon 1951). In our case, the goal is to identify the most likely star-formation history corresponding to a given spectrum. In contrast with the traditional methods based on spectral fitting or targeted LSs (see e.g. Walcher et al. 2011), we focus here on the information content as quantified by the entropy.

We begin with a simple test determining H in a set of well-known and tested population synthesis models. These models incorporate our knowledge of the formation and evolution of stars, as well as the radiative properties of stellar atmospheres that ultimately give rise to spectra (Conroy 2013). These models are typically produced for a reduced set of variables, mainly stellar initial mass function, age, total chemical composition (metallicity), and possibly a number of non-solar abundance ratios, most notably [Mg/Fe]. At present, state-of-the-art models rely on a reduced set of libraries of stellar spectra (see e.g. Le Borgne et al. 2003; Coelho et al. 2005; Sánchez-Blázquez et al. 2006), so that – from the pure point of view of information theory – any synthetic model, no matter how complex, is always defined by a linear combination of a relatively reduced set of stellar spectra, of order 10^3. This is the reason why information-based models can easily discriminate between synthetic and 'real' spectra, as shown in Slonim et al. (2001) where the information content from observed 2dFGRS spectra was found to be much higher than the one derived from theoretical galaxy formation models that incorporated population synthesis[1] (see fig. 3 in that paper).

Fig. 1 shows the entropy defined within a relatively wide spectral window, between 3500 and 8000 Å, normalized such that a totally uninformative scenario – i.e. a constant $p(\lambda)$ in this interval – corresponds to $H = 1$, and the extreme case of a perfectly monochromatic signal, i.e. $p(\lambda) = \delta(\lambda - \lambda_0)$ has $H = 0$. The SSP models of BC03 have been used here to explore the variation with respect to stellar age (horizontal axis), with two choices of metallicity, as labelled. We adopt the Chabrier (2003) stellar initial mass function in this analysis, but the differences are minimal for other reasonable choices. From an information theory point of view, the first salient feature of galaxy spectra is the closeness to a fully uninformative scenario: The deviation from the maximum entropy case is less than one part in ∼30. However, this is caused by the large spectral interval chosen and by the fact that, aside from absorption features (that usually amount to less than a few Angstrom in equivalent width), the continuum is a fairly smooth function of wavelength, with a relatively shallow gradient. Moreover, the trend found with respect to age stems from this: The younger populations are dominated by hotter stars with steeper, and blue, continua, whereas older populations have a much shallower wavelength dependence, except for the prominent 4000 Å break.

In order to be able to discriminate better among the models, we measure the entropy within a narrower spectral window. This is the equivalent of defining individual probability distributions within prescribed wavelength intervals: $p(\lambda|[\lambda_1, \lambda_2])$, where only photons with wavelength between λ_1 and λ_2 are considered in the definition of the probability distribution function. We can also do this experiment adopting a running interval, of width $\Delta\lambda$, as we traverse the optical window, defining an 'entropy spectrum', $H_{\Delta\lambda}(\lambda)$. The choice of $\Delta\lambda$ is important: too wide and we wash out all information, too narrow and the result becomes prohibitively dependent on the signal-to-noise ratio (S/N) or velocity dispersion. At the resolution of SDSS galaxy spectra, around $R \sim 2000$, Rogers et al. (2010) found that $\Delta\lambda \sim$ 50–100 Å provides an optimal window, a result that is equally optimized for stellar spectra in SDSS (Hawkins et al. 2014). Fig. 2 shows the 'entropy spectrum' of a few synthetic populations from the models of BC03, using a running interval of width $\Delta\lambda = 100$ Å. As comparison, we show the corresponding entropy spectrum for a sample of real galaxy spectra from SDSS in Fig. 3 – showing the median of subsamples classified according to their nebular emission properties. We will explore this sample in more detail in Section 3. Hereafter, we quote estimates of entropy with the following definition:

$$\overline{\mathcal{H}}(\lambda) \equiv [1 - H_{\Delta\lambda=100A}(\lambda)] \times 1000, \quad (3)$$

which can be interpreted as negentropy, or information content. The maximally uninformative case corresponds to $\overline{\mathcal{H}} = 0$. Note that in Fig. 1 the entropy variations are, at most of a few per 100. The number 1000 in equation (3) is just a scaling factor so that the values of negentropy are given in numbers of order ∼10.

[1] Also noting that noise – inherent to any observational data – will affect the information content.

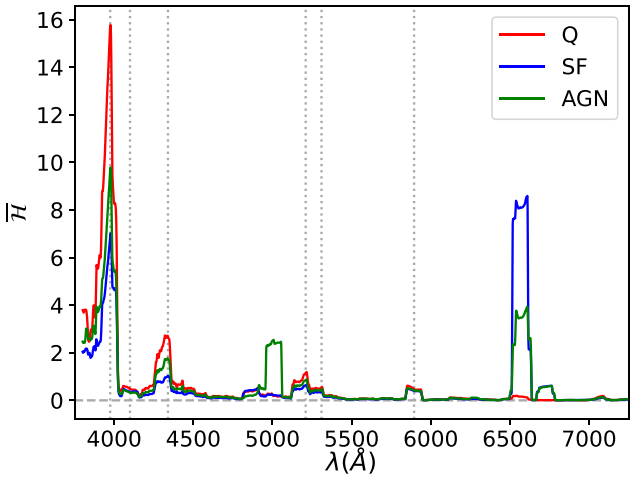

Figure 3 Negentropy spectra of SDSS galaxies, equivalent to the one shown for synthetic spectra in Fig. 2. Here, we plot separately quiescent (Q, red), star-forming (SF, blue), and active galactic nucleus (AGN, green) galaxies based on their nebular emission properties, following the standard BPT classification (see text for details). The vertical dotted lines mark the wavelengths chosen for the reduced dimensional analysis. Note the substantial contribution to the entropy from emission lines in SF and AGN spectra. These lines correspond to the diffuse gas component, not the stellar populations, and thus are omitted in our analysis.

Fig. 2 provides insight into the information encoded in galaxy spectra. At redder wavelengths, the spectrum becomes less informative, as it is dominated by a continuum with relatively weak absorption features, except for some prominent regions, as labelled. The strongest contribution comes from the region around 4000 Å, where the pile-up of spectral lines produces a prominent break, with the entropy reaching its lowest value. While minute, targeted LSs can be exploited to constrain stellar population properties (see e.g. La Barbera et al. 2013), other methods based on full spectral fitting, or machine learning algorithms – where the full spectra are directly input into the algorithm – will be subject to this problem.

Note how this blind approach to the information content of spectra produces the standard features targeted for the analysis of stellar populations: The Balmer regions at 4100 Å (Hδ), 4340 Å (Hγ), 4862 Å (Hβ), and 6563 Å (Hα) are more prominent in younger populations around 1 Gyr, a well-known result caused by the dominance of A-type stars to the net luminosity budget in populations with that age, i.e. when that stellar type becomes the main-sequence turn-off. Older populations produce conspicuous features around 4300 Å (*G* band) and 5100–5400 Å, where a number of Mg and Fe absorption lines are very prominent. The optical Na absorption at 5890 Å is also present in the entropy spectra. These features have been intensively explored in galaxy spectra for decades, so their presence is not surprising. However, our analysis confirms that these regions are the ones that carry most of the information – as in negentropy – from a pure information theory viewpoint. This result is especially relevant to deep learning algorithms, where the complex 'machines' critically depend on the information content of the spectra to classify or learn about the underlying components. In essence, machine learning methods rely on a quantitative assessment, based on some figure of merit, whose maximization leads either to a decision (in a classification algorithm), or to a set of best-fitting parameters (in a regression). The success of any Machine Learning (ML) method – either supervised or unsupervised – will ultimately depend on how this figure of merit can discriminate among the basis of sources. These sources correspond here to the generic base set of SSPs that describe any star-formation history. Blind source separation methods also operate on the same basis (see e.g. Hyvärinen, Oja & Karhunen 1997), namely that the original sources that produce the final data carry enough information to be told apart from the observed mixtures. Whilst we do not optimally fit specific LSs given a set of models, the analysis of entropy marks an intrinsic limit regarding the ability of any such method to unambiguously extract the details of the underlying stellar populations from spectroscopy.

This paper is focused on how information is stored, and can be retrieved, from galaxy spectra in the most fundamental way based on entropy. The results from the models justify a simplification of the methodology by selecting a small set of targeted regions that correspond to the highest values of negentropy. Please note we avoid on purpose Hβ and Hα, as they are strongly affected by emission from the gas component, and therefore bias any analysis focused on stellar populations. Traditional studies typically remove these lines from the analysis or perform detailed subtraction of the emission from the gas, but this introduces a substantial dependence on the models, and leads to systematics. Moreover, the emission lines trace 'instantaneous' properties of the galaxy, namely the ionization of the diffuse gas by young stars, AGN activity, or shocks. The analysis of populations focuses, rather, on the so-called 'fossil record', i.e. the star-formation activity throughout the whole formation history of the galaxy. Entropy in each of these six intervals is measured within the $\Delta\lambda = 100$ Å window. The vertical dotted lines mark those regions – adopted below for the analysis of real galaxy spectra from SDSS (see Section 3). While Fig. 2 is shown for a few choices of stellar age and solar metallicity, the models feature similar behaviour at different ages and chemical composition. Therefore, we propose a simplified description of each galaxy spectrum with only six 'coordinates' that represent the regions that carry the maximum amount of variation in their entropy.

It is worth noting that the interpretation of a galaxy spectrum as a probability distribution of the energy of the incoming photons is affected by two observational pitfalls: the S/N and the effective spectral resolution. The former depends on the efficiency of the observing apparatus and the exposure time, whereas the latter depends both on the spectrograph as well as the galaxy under scrutiny, as the distribution of stellar orbits impose an effective 'kinematic kernel' that broadens – via Doppler shifts – all spectral features. Fig. 4 illustrates the effect on the entropy, with the noiseless case at the fiducial resolution of the stellar models shown as a red line. A reasonable S/N still produces manageable results, with the shaded region spanning the range of entropy produced by a bootstrap that adds Gaussian noise to the model spectra, corresponding to S/N = 5 per Å. A lower spectral resolution – here shown by producing the equivalent observation of a (massive) galaxy with velocity dispersion of 300 km s^{-1} (dark red) degrades the entropy quite dramatically. As expected, lower resolution (or higher velocity dispersion) produces smoother spectra, with less-defined absorption features, approaching a homogeneous probability distribution, and therefore tending towards maximum entropy, i.e. minimum information. For those reasons, the analysis presented in the next section – comprising actual galaxy spectra – will constrain the data to high S/N and relatively low velocity dispersion.

Fig. 5 shows how entropy changes with respect to spectral resolution within the six features defined here. The results are shown for two population synthesis models at two different stellar ages (both with solar chemical abundance), as labelled. The BC03 models are shown as solid lines, and the MIUSCAT models (Vazdekis et al. 2012) are plotted with dashed lines. We emphasize that these models

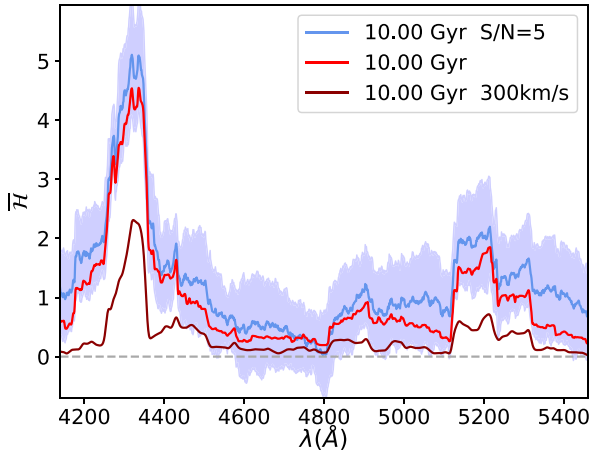

Figure 4 Variation of the negentropy spectrum with respect to velocity dispersion and noise. A test case is shown for a 10 Gyr SSP at solar metallicity, within a relatively narrow spectral window for ease of visualization. The red line represents the noiseless reference at 'zero velocity dispersion'.

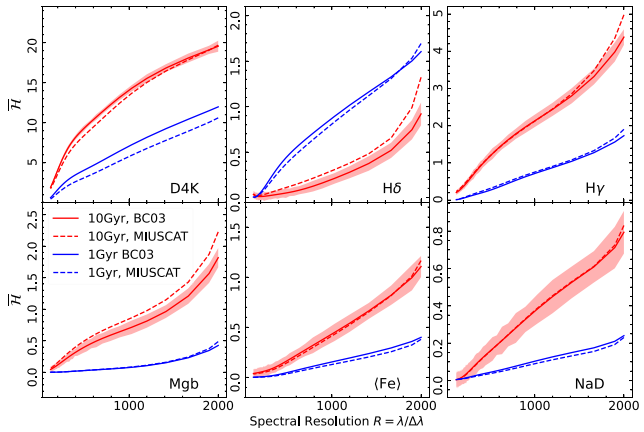

Figure 5 Variation of the information content of galaxy spectra with respect to the effective spectral resolution (horizontal axes), shown in the six spectral windows targeted in this paper. The blue (red) lines represent a young (old) population at solar metallicity and solar abundance ratios from the BC03 (solid) and MIUSCAT (Vazdekis et al. 2012, dashed) population synthesis models, as labelled. The red shaded region shows, for reference, the expected 1σ uncertainty at an S/N of 20 Å^{-1}.

are independent, and even adopt different stellar libraries. Note the trends are equivalent in both sets of models, although there are some systematic variations caused by the different choice of model prescriptions and stellar libraries. For reference, we include an estimate of the 1σ uncertainties if the spectra have an S/N of 20 Å^{-1}, as a shaded region for the 10 Gyr BC03 model. Noting that $R \sim 2000$ is an optimal resolution for galaxy spectra given the typical velocity dispersion of the stellar component, we find that the information content – defined as the entropy – gets halved at $R \sim 500$–1000 (varying with the spectral index). The left part of the horizontal axis corresponds to the resolution expected in slitless grism spectroscopy or medium band photometry, where the discriminatory power of spectral indices is reduced – in this case, the continuum is used instead, to constrain the population parameters (see e.g. Ferreras et al. 2009; Pérez-González et al. 2013; Díaz-García et al. 2015).

3. THE ENTROPY CONTENT OF SDSS GALAXY SPECTRA

The trends presented in the previous section have been obtained with synthetic models of stellar populations. In this section, we turn the entropy analysis to a set of high quality galaxy spectra from the SDSS (York et al. 2000). We use the original, legacy, survey that consists of single fibre spectroscopy at resolution $\mathcal{R} \sim 2000$. From the original data set (we retrieve the data from their Data Release 16; Ahumada et al. 2020), we select a sample of 76 569 spectra with a high S/N ($\gtrsim$10 Å^{-1}, measured in the r band), constrained in velocity dispersion between $\sigma = 100$ and 150 km s^{-1}. It is a well-known fact that the stellar velocity dispersion – a tracer of the gravitational potential – correlates strongly with population properties (e.g. Bernardi et al. 2003; Ferreras et al. 2019). By choosing a relatively reduced range in σ, we aim at simplifying the working sample. Our goal in this paper is to explore how well entropy can disentangle subtle differences in stellar populations, so a focused data set is preferred. The median redshift is $z = 0.079$ with a 95 per cent interval $\Delta z = [0.054, 0.099]$. The spectra are corrected for foreground dust contamination using the dust extinction law of Fitzpatrick (1999), are brought to the rest-frame, and are resampled to a common wavelength grid with 1 Å spacing. In each spectrum, bad pixels flagged as problematic by the SDSS team – that represent a very small fraction in these high S/N data – were replaced by their best model fits, also provided within the data structure of the SDSS spectra.

Fig. 3 shows the negentropy spectrum of SDSS data, i.e. the measurement of $\overline{\mathcal{H}}(\lambda)$ (equation 3) with a running $\Delta\lambda = 100$ Å interval, as in Fig. 2, which was presented for a few SSPs from the synthetic models of BC03. The spectra are median-stacked into three classes, defined according to nebular emission as Q (red), SF (blue), and AGN (green). These three classes represent a fundamental transition during the evolution of galaxies, from an initial SF stage (in the Blue Cloud) to passive evolution in a Q phase (in the Red Sequence). AGN activity appears to dominate the transitioning stage (mostly populating the intermediate region, termed the Green Valley). These regions stem from the inherent bimodality of galaxy properties (Strateva et al. 2001; Salim 2014). Therefore, by splitting the sample according to this classification scheme, we can assess how such a transition affects the entropy of galaxy spectra. The classification follows the standard BPT diagram (Baldwin, Phillips & Terlevich 1981), and is taken from the official galspecExtra SDSS catalogue (Brinchmann et al. 2004), where we include an additional constraint on the equivalent width of Hα emission to select Q galaxies (e.g. Cid Fernandes et al. 2011) – as the bpt $= -1$ flag only refers to galaxies whose spectra cannot be mapped on the standard BPT diagram. The real spectra show a similar behaviour to the models, except at the location of prominent emission lines such as Hβ, [O III], Hα+[N II], and [S II].

We focus on the most relevant features of the negentropy spectra, as found in the models, namely 4000 Å break strength (termed D4K), two Balmer absorption lines (Hγ and Hδ), and the strongly metallicity dependent indices Mgb, ⟨Fe⟩, and NaD (see e.g. Worthey & Ottaviani 1997; Trager et al. 1998; Balogh et al. 1999). We will compare the results from those six spectral windows in three complementary ways: (1) LSs, measuring the equivalent width in those windows, following the methodology of Rogers et al. 2010); (2) negentropy, as defined earlier; and (3) relative entropy, also termed Kullback–Leibler divergence ($D_{\rm KL}$; Kullback & Leibler 1951, bottom panel). The colour coding separates subsamples based on nebular emission (Q: red, SF: blue, and AGN: green). The definition of relative entropy rests on the idea of the information content of a probability distribution, $p(\lambda)$, with respect to another

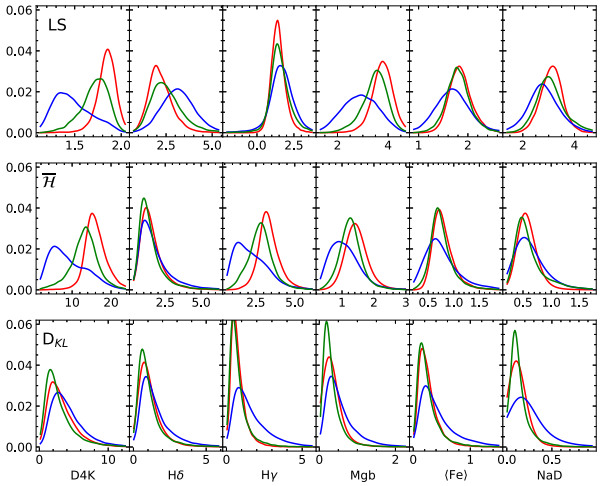

Figure 6 Distribution of LS equivalent widths (top panel), information content based on negentropy ($\overline{\mathcal{H}}$, middle panel), and $D_{\rm KL}$ (bottom panel) for a set of galaxy spectra from SDSS. The sample is subdivided into SF (blue), Q (red), and AGN (green) galaxies based on the nebular emission properties from the standard BPT classification scheme (see text for details). The histograms can be interpreted as probability distributions, and are plotted with a Gaussian kernel density estimator.

one that acts as baseline, $q(\lambda)$. In this case, entropy refers to the amount of information relative to the baseline:

$$D_{\rm KL} \equiv -\sum_i p(\lambda_i) \log\left[p(\lambda_i)/q(\lambda_i)\right]. \quad (4)$$

Note that when $q(\lambda)$ is constant (i.e. a totally uninformative baseline), we recover Shannon entropy. Here we use as reference the median of all spectra, taking the galaxies from the whole sample, regardless of nebular emission (i.e. including Q, SF, and AGN), to define the median $q(\lambda)$. Therefore, relative entropy in this paper is a measure of departure from the overall behaviour of the general sample. Fig. 6 shows the distribution of the six spectral windows in the three different representations of the spectral information: LS (top panel), negentropy ($\overline{\mathcal{H}}$, middle panel), and relative entropy ($D_{\rm KL}$, bottom panel).

The most significant difference between SF, AGN, and Q subsamples in Fig. 6 is obtained for D4K, and the SF subset is the one with the widest range of values in all cases. Comparing the standard LSs with negentropy, we find some analogous results (D4K, Mgb, and ⟨Fe⟩), whereas Balmer absorption behaves differently: While the LS analysis finds a larger variation regarding nebular activity in Hδ, negentropy discriminates better with Hγ. In contrast, relative entropy does not produce such clear segregation, with only a more extended tail in the distribution of SF galaxies. The figure shows that there is no evident trend when looking at any of the six spectral windows. While the separation of spectra based on nebular emission is a fundamental discriminator of activity, the distributions do not reveal substantial differences, neither in LS nor in the two definitions of entropy. Therefore, in order to explore in more detail the potential drivers of variance/information, we resort to further reduce the dimensionality of parameter space via PCA. PCA provides a simple way to visualize the variance of the data by diagonalizing the corresponding covariance (represented here by a 6×6 matrix). The eigenvectors, or principal components, are ranked in decreasing order of their associated variance, and the original data are projected on to these eigenvectors to produce an alternative representation that is ranked as a function of variance.

Table 1. Normalized PCA eigenvalues.

Rank	Eigenvalue ratio		
	LS	Entropy	$D_{\rm KL}$
1	0.500 22	0.874 99	0.754 79
2	0.254 49	0.069 83	0.092 25
3	0.116 22	0.020 57	0.059 72
4	0.096 91	0.016 87	0.046 88
5	0.026 16	0.014 78	0.039 86
6	0.006 00	0.002 96	0.006 50

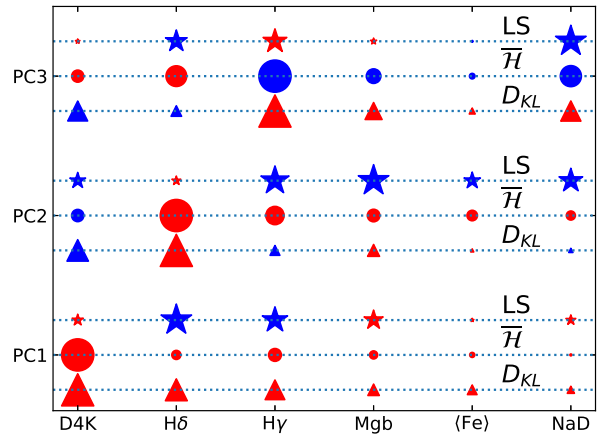

Figure 7 Weights of the first three principal components (PC1, PC2, and PC3) using as input either LSs (stars, LS), negentropy (circles, $\overline{\mathcal{H}}$), or relative entropy (triangles, $D_{\rm KL}$), shown for the six features adopted to describe the spectra (horizontal axis). The symbol size corresponds to the value of the weight, and the colour gives the sign as positive (red), or negative (blue). Note negentropy gives the simplest interpretation of the variations in terms of 4000 Å break strength (PC1) and Balmer absorption Hδ (PC2).

We apply PCA independently to the data from the three methods described above, finding the eigenvalues shown in Table 1. They are quoted as a ratio of total variance, and show that the standard LS analysis produces the most mixed distribution, needing up to four components to account for 90 per cent of total variance, whereas the first principal component for standard entropy already carries ∼87 per cent of total variance, with the second one accounting for around 7 per cent and the subsequent components featuring a lower share. Relative entropy also shows a high weight of the first principal component, accounting for ∼75 per cent of total variance.

Fig. 7 is a graphic representation of the weights of the first three principal components, where symbol size denotes the weight of a given spectral feature, as labelled in the horizontal axis, and the colour means positive (red) or negative (blue) sign for each weight. For reference, these weights are quantified in Table 2. The principal components derived from the standard LS data correspond to mixtures of Balmer absorption (mostly in PC1) and metallicity dependent strengths (PC2). In contrast, the components corresponding to Shannon (neg)entropy ($\overline{\mathcal{H}}$) produce a simpler scheme, where PC1 is mostly represented by D4K and PC2 is dominated by the information in the Balmer absorption lines (Hδ and Hγ). The third component – whose associated variance is only 2 per cent – depends mostly on Hγ, with residual dependence on the other features, most notably Hδ and NaD. The relative entropy representation ($D_{\rm KL}$) produces a mixed set of weights, with PC1 mostly having a strong dependence on D4K and Balmer absorption (both Hδ and Hγ being represented). The second component is dominated by

Table 2. First three PCs (normalized).

	LS			$\overline{\mathcal{H}}$			$D_{\rm KL}$		
Region	PC1	PC2	PC3	PC1	PC2	PC3	PC1	PC2	PC3
D4K	0.105 81	−0.178 78	0.017 60	0.981 89	−0.135 20	0.119 92	0.849 49	−0.403 48	−0.325 69
Hδ	−0.778 59	0.057 07	−0.393 72	0.079 52	0.925 08	0.348 29	0.395 36	0.902 03	−0.092 39
Hγ	−0.532 60	−0.565 43	0.475 13	0.157 91	0.300 10	−0.838 79	0.319 89	−0.086 16	0.846 11
Mgb	0.303 26	−0.668 13	0.032 45	0.063 21	0.142 42	−0.173 15	0.108 98	0.125 65	0.238 71
⟨Fe⟩	0.009 90	−0.194 17	−0.002 98	0.025 05	0.099 90	−0.025 60	0.077 78	0.008 82	0.029 85
NaD	0.082 90	−0.401 21	−0.786 04	0.003 90	0.075 00	−0.360 71	0.042 51	−0.016 25	0.334 09

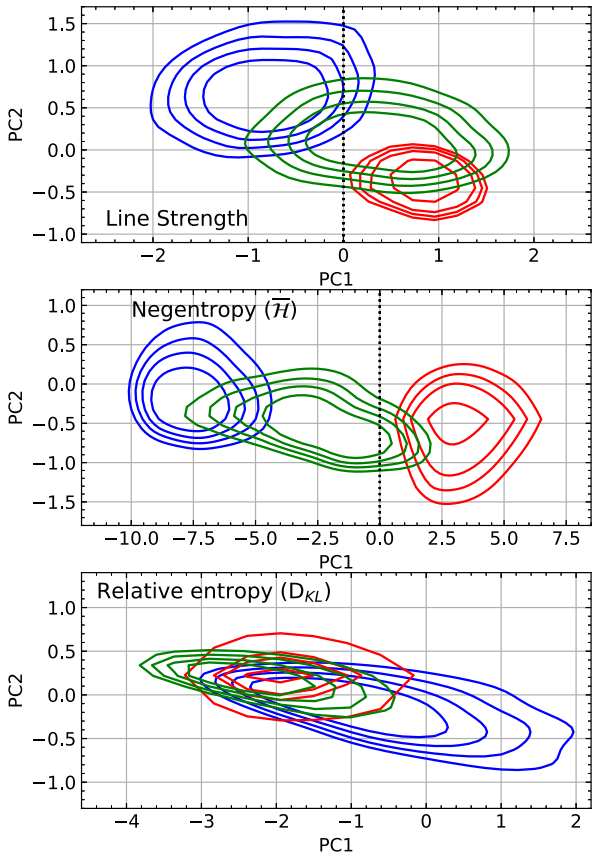

Figure 8 The sample of SDSS spectra is projected on to the first two principal components, with the distributions shown as contours (corresponding to their number density, from 50 per cent to 80 per cent of each sample, in steps of 10 per cent). The sample is shown separately for SF (blue), Q (red), and AGN (green) spectra. The projection for LS, negentropy, and relative entropy, is shown in the top , middle, and bottom panels, respectively. Negentropy produces the largest discrimination among these with respect to a single component (PC1).

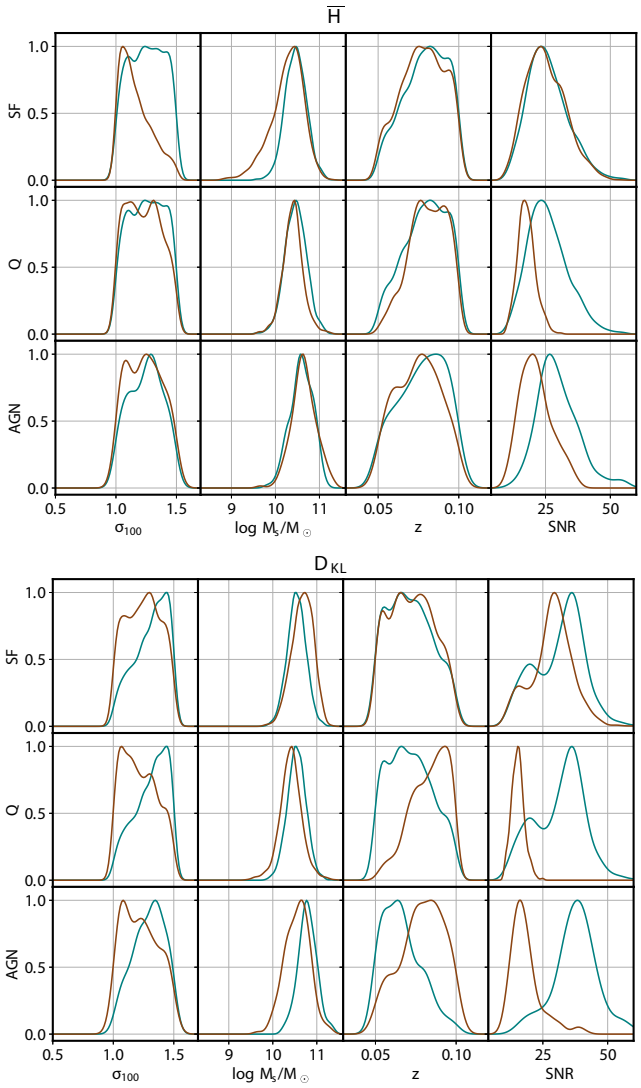

Figure 9 Distribution of some of the physical parameters of galaxies at the extreme ranges of PC1 (5th and 95th percentiles in teal blue and brown, respectively), according to negentropy ($\overline{\mathcal{H}}$, top panel) and relative entropy ($D_{\rm KL}$, bottom panel). From top to bottom in each figure, the rows correspond to SF, Q, and AGN spectra, respectively. From left to right, the measurements are stellar velocity dispersion in units of 100 km s^{-1} (σ_{100}), (log) stellar mass in solar units, redshift, and the average S/N in the SDSS r band. The histograms are plotted with a Gaussian kernel density estimator.

Hδ, with a substantial dependence on D4K (although with negative sign, in contrast to PC1). The third component mostly relates to Hγ with some dependence on D4K and NaD. Surprisingly, the metallicity dependent features (Mgb and ⟨Fe⟩) play a minor role in these components, which may be caused by the strong correlation among spectral indices (see epilogue below). Overall, it is worth emphasizing that the bulk of the entropy variation found in both representations refers to the 4000 Å break strength and Hδ. This type of scheme has been extensively used in standard analyses of stellar populations in galaxies (e.g. Kauffmann et al. 2003) as the 4000 Å break indicator is overall sensitive to the average age and metallicity,

$$\mathbb{C}_{\rm BC03} = \begin{array}{c} \begin{array}{cccccc} D4K & H\delta & H\gamma & \mathrm{Mgb} & \langle\mathrm{Fe}\rangle & \mathrm{NaD} \end{array} \\ \begin{pmatrix} 1. & -0.8705 & -0.8557 & 0.7263 & 0.7313 & 0.6483 \\ -0.8705 & 1. & 0.9410 & -0.6839 & -0.6458 & -0.6114 \\ -0.8557 & 0.9410 & 1. & -0.6772 & -0.6303 & -0.6011 \\ 0.7263 & -0.6839 & -0.6772 & 1. & 0.5505 & 0.5012 \\ 0.7313 & -0.6458 & -0.6303 & 0.5505 & 1. & 0.5085 \\ 0.6483 & -0.6114 & -0.6011 & 0.5012 & 0.5085 & 1. \end{pmatrix} \end{array} \begin{array}{c} D4K \\ H\delta \\ H\gamma \\ \mathrm{Mgb} \\ \langle\mathrm{Fe}\rangle \\ \mathrm{NaD} \end{array} \quad (5)$$

Figure 10 Covariance matrix of the six LSs targeted in the analysis of SDSS spectra, derived from the BC03 population synthesis models, renormalized to unit variance (i.e. defined as a correlation). The spectra have been convolved to an equivalent velocity dispersion of 100 km s^{-1}, with Gaussian noise corresponding to S/N = 10 Å^{-1} in the SDSS *r* band, mimicking the observational data used in this paper. Note the high degree of correlatedness between LSs (see text for details).

whereas Hδ, or Hγ, trace recent episodes of star formation (see e.g. Kauffmann 2014; Wu et al. 2018). In this paper, we retrieve this result purely from the information point of view, without any reference to the modelling of stellar populations. In Appendix A, we show the actual distribution of LSs for the six targeted features, confirming the importance of D4K and Balmer absorption, with a strong relation to nebular activity, which acts as a rough proxy of evolution – i.e. from SF to AGN to quiescence.

Fig. 8 shows the distribution of projections of the observed spectra on the first two principal components for LS data and both definitions of entropy, as labelled. The contours map the density from 50 per cent to 80 per cent of total number of data points (in steps of 10 per cent), and are plotted separately with respect to nebular activity, as above, into SF (blue), Q (red), and AGN (green). We emphasize that the PCA output depends only on data in spectral regions where emission line activity is absent or weak, so that most of the variance cannot be directly ascribed to their classification into SF, Q, and AGN. The standard LS data (top panel) separate the subsamples along a diagonal line in PC1–PC2 space, with AGN galaxies located between the SF and Q subsets. Similarly, Shannon entropy (middle panel) shows substantial differences between the SF and the Q subsets, with the AGN group populating an intermediate area, something that hints at nuclear activity representing a transition from star formation to quiescence, as already proposed in the literature (see e.g. Schawinski et al. 2007; Salim 2014). This diagram suggests that the 'horizontal direction', i.e. PC1 (mainly dependent on D4K for $\overline{\mathcal{H}}$), is the main discriminator, supporting the definition of the green valley, i.e. transitioning galaxies from star formation to quiescence, using this spectral feature (Angthopo, Ferreras & Silk 2019, 2020). In contrast, relative entropy (bottom panel) does not separate the spectra so well, and only produces an extended tail of SF galaxies towards high, positive values of PC1.

Fig. 9 shows the distribution of some observables when taking subsets corresponding to the 5th (teal colour) and 95th (brown) percentiles of the distribution of PC1 projections from negentropy ($\overline{\mathcal{H}}$) or relative entropy ($D_{\rm KL}$), shown in both cases for the SF (top), Q (middle), and AGN (bottom) samples. We should stress that the parent sample is restricted to a relatively narrow range of stellar velocity dispersion (100–150 km s^{-1}), so that the distributions are as 'homogeneous' as possible, concerning variations derived from the general mass–age and mass–metallicity trends (see e.g. Gallazzi et al. 2005). It is worth pointing out that this analysis produces a noticeable segregation in the distribution of S/N – a property only dependent on the technical aspects of the observations – in many instances, except for SF galaxies. This is an important trend that all methods based on deep/shallow learning should take into account. Taken at face value, it would imply that an important part of the variance/information/signal one can extract from spectra could be due to technical details of the data acquisition process. Note that such segregation is not present in velocity dispersion, stellar mass, or redshift, for $\overline{\mathcal{H}}$, although relative entropy also shows a redshift and velocity dispersion dependence. Appendix B performs the same analysis on a smaller subset of galaxy spectra restricted at a higher S/N threshold, showing that the overall trends are consistent. This result illustrates the need to select clean, well-defined samples in any statistical, data-driven study of the underlying star-formation history of galaxies. In addition, these methods can be applied to separate out the observational effects from the physical ones – as, in, e.g. the removal of night sky lines (Wild & Hewett 2005).

4. AFTERTHOUGHT: THE COVARIANCE OF STELLAR POPULATION SPECTRA

The entropy-based analysis presented here illustrates the challenge of extracting a unique solution to this particular inverse problem: Starting from a set of observed galaxy spectra, we aim at determining the fundamental components, i.e. the base stellar populations that lead to the underlying star-formation histories. While signal extraction from superpositions is a relatively easier task in, say, audio files or superposition of independent time series, galaxy spectra reveal a high level of entanglement regarding the information content as a function of wavelength, that prevents us from extracting details of the age distribution or chemical composition even in spectra with the highest quality. In this epilogue, we show a simple exercise that makes use of the optimal set of six spectral windows presented in this paper, defined by the entropy content (see Fig. 2). Fig. 10 shows the covariance matrix of a set of 16 386 synthetic spectra representing SSPs that cover a wide range of age ($0.5 < t_{\rm SSP} < 14$ Gyr) and metallicity ($-1.5 < \log Z/Z_\odot < +0.3$), from the models of BC03. Other sets of models produce very similar results. Each of the six indices is renormalized to unit variance, to be able to compare all indices in the same way. To produce data comparable to the SDSS spectra used in this paper, we convolve the models to an equivalent velocity dispersion of 100 km s^{-1}, and add Gaussian noise corresponding to an S/N in the SDSS *r* band of 10 Å^{-1}. Note that all indices are correlated or anticorrelated at a rather high level,

amounting to $\gtrsim$70 per cent of the intrinsic variance, being in some cases as high as ∼90 per cent. Also note that these indices have been chosen to yield most of the information content, or variance, across the optical window, and are typically adopted in most analyses of stellar populations in the literature. For instance, the 4000 Å break strength is often associated to an overall stellar age, and Balmer absorption reveals episodes of star formation within the past ∼1 Gyr. However, the otherwise strong metallicity sensitive indices Mgb, ⟨Fe⟩, and NaD are also strongly correlated with the former, as expected from the well-known age–metallicity degeneracy (see e.g. Worthey et al. 1994; Ferreras et al. 1999).

Observational errors and spectral resolution obviously affects the outcome: The original models with no velocity dispersion of noise have an even higher covariance. This result should be taken as a note of caution in the interpretation of any statistic that produces a figure of merit based on the comparison of observations with population synthesis models. The most commonly used one, χ^2, is expected to produce a minimum at the best fit, with a value roughly similar to the number of degrees of freedom (Andrae, Schulze-Hartung & Melchior 2010). If we use the more standard definition of the χ^2 statistic for a generic covariance $\mathbb{C}$

$$\chi^2(\pi) \equiv \Big(\mathbf{x} - \mathbf{m}(\pi)\Big)^T \cdot \mathbb{C}^{-1} \cdot \Big(\mathbf{x} - \mathbf{m}(\pi)\Big), \tag{6}$$

where $\mathbf{x}$ represents the measurements and $\mathbf{m}$ the model output for a specific set of parameters (π). Perfectly uncorrelated data – the assumption most often invoked – implies a covariance $\mathbb{C} = \mathbb{I}$, and typically yields a minimum χ^2 around 6, whereas the covariance estimated from the model grid produces a minimum χ^2 of ∼20, clearly inconsistent with the assumption that the six LSs are independent units of information. At face value, the substantially higher χ^2 when accounting for covariance suggests that the fitting procedure effectively involves fewer degrees of freedom than the six (or more) spectral features. Therefore, even though we may have a large number of observable constraints from galaxy spectra, we can only produce a rough estimate of stellar age, metallicity, and possibly some non-solar abundance ratios. This is not only applicable to LS analysis, but to full spectral fitting, where each 'pixel', i.e. a flux unit within a relatively narrow spectral window ($\Delta\lambda \sim 1$ Å), is highly correlated with most of the other pixels in the spectrum within the typical NUV-optical-NIR spectral interval used in such analyses. Highly targeted studies of LSs are perhaps the only effective way to constrain the stellar population content of galaxies, whereas methods based on full spectral fitting should take into consideration this important caveat about covariance, and should not make the assumption that the large number of flux measurements are independent, a result that would falsely lead to a high constraining power.

Equally interesting to the result obtained from the models is a comparison with the covariance of the real galaxy spectra from the SDSS sample (Fig. 11). While substantial, the covariance is smaller than the one found in synthetic data. This should come as no surprise. More realistic data regarding noise and the 'kinematic kernel' will lower the covariance. However, this effect should be subdominant: For instance, the covariance between D4k and Hδ weakens from -0.9013 in the noiseless case to -0.8705 at S/N = 10, a change that is not large enough to explain the lower covariance of the observations. A similar variation is obtained if we include the effect of the stellar velocity dispersion. It is important to note that the synthetic models boil down to linear superpositions of ∼1000 stellar spectra (e.g. MILES; Sánchez-Blázquez et al. 2006). This result was, for instance, evident when comparing the mutual information in a set of 2dFGRS spectra to data from galaxy formation models (Slonim et al. 2001). A comparison of these two covariance matrices should convince the reader that models still lack the complexity needed for more detailed constraints of the star-formation histories of galaxies.

5. CONCLUSIONS

This paper focuses on the fundamental limitations of extracting information from galaxy spectra to derive star-formation histories. This work is complementary to the standard procedure based on comparisons of stellar population synthesis models with the observational data via full spectral fitting or targeted LSs. There is a vast literature devoted to the traditional methodology that perform careful analyses to mitigate the inherent degeneracies among the population properties (to name a few: Panter et al. 2003; Cid Fernandes et al. 2005; Ocvirk et al. 2006; Tojeiro et al. 2007; Conroy & Gunn 2010; Koleva et al. 2011; Cappellari 2012; La Barbera et al. 2013; Wilkinson et al. 2017). However, it is not the goal of this paper to assess these methods, but to explore the limitations at a more fundamental level, based on the information content (here defined as negentropy, alternatively defined as variance, as in PCA-based work, e.g. Rogers et al. 2007). The essence of the problem lies in the information content of the spectra, that we base in this paper on entropy, either taking the fundamental definition of Shannon & Weaver (1975) or relative entropy as defined by the $D_{\rm KL}$ (Kullback & Leibler 1951). While alternative methods are adopted in the literature to define information content, entropy represents the 'building block' as it directly relates to, e.g. how many bits of information are needed for a full representation of the data (Shannon 1953). Entropy lies at

$$\mathbb{C}_{\rm SDSS} = \begin{array}{c} \begin{array}{cccccc} D4K & H\delta & H\gamma & \mathrm{Mgb} & \langle\mathrm{Fe}\rangle & \mathrm{NaD} \end{array} \\ \begin{pmatrix} 1. & -0.5193 & -0.0420 & 0.7573 & 0.3901 & 0.3722 \\ -0.5193 & 1. & 0.5249 & -0.3881 & -0.0042 & -0.1074 \\ -0.0420 & 0.5249 & 1. & -0.0290 & 0.1793 & 0.0551 \\ 0.7573 & -0.3881 & -0.0290 & 1. & 0.4445 & 0.3099 \\ 0.3901 & -0.0042 & 0.1793 & 0.4445 & 1. & 0.1841 \\ 0.3722 & -0.1074 & 0.0551 & 0.3099 & 0.1841 & 1. \end{pmatrix} \end{array} \begin{array}{c} D4K \\ H\delta \\ H\gamma \\ \mathrm{Mgb} \\ \langle\mathrm{Fe}\rangle \\ \mathrm{NaD} \end{array} \tag{7}$$

Figure 11 Equivalent covariance matrix as in Fig. 10, but derived from the observed SDSS spectra. Note the lower correlatedness of real galaxy spectra with respect to the models' covariance.

the core of all methods aimed at blind classification of galaxies, such as PCA, independent component analysis, cross-entropy methods, factor analysis, convolutional neural networks, etc. We explore both synthetic models and real galaxy spectra from SDSS, to find that a reduced set of spectral regions encode most of the information, unsurprisingly tracing the traditional Lick system (e.g. Trager et al. 1998) commonly used in most studies.

The definition of entropy is initially determined for the full spectral range of interest, say the optical window, and then narrower intervals can be explored. Fig. 2 suggests that stellar populations feature a reduced set of well-defined regions where the information content is highest. Overall, the entropy of galaxy spectra is rather high – as in uninformative – reflecting the difficult task of constraining the details of the underlying populations, regardless of the apparently high number of data points in each spectrum. Applying this method to SDSS galaxy spectra confirms this high entropy (Fig. 3), with a clear, and fully expected, dependence on nebular emission, with AGN systems representing an intermediate stage between SF and Q galaxies. A detailed analysis based on PCA applied to the entropy estimates suggests that the 4000 Å break strength and Balmer absorption are the most important sources of entropy variation, with metal-dependent indicators being subdominant. This trend confirms previous work that focuses on, e.g. the $D_n(4000)$ versus Hδ bivariate diagram as a fundamental tool in the analysis of star-formation histories (see e.g. Kauffmann et al. 2003). Intriguingly, Fig. 9 shows that S/N may also affect the analysis (even at S/N $\gtrsim 10$ Å^{-1}), which implies that blind methods are highly susceptible to the quality of the data.

Our results suggest that detailed estimates of star-formation histories are hampered by the sizeable covariance of the spectral elements, which result in an effectively low number of 'degrees of freedom', regardless of the large number of data units in a spectrum. Somehow, this addresses the long-known fact that a reduced set of high-quality colours based on broad-band photometry can produce constraints on the stellar populations that are comparable with the analogous study at high spectral resolution. For instance, the derivation of stellar masses from photometry is comparable with the equivalent analysis making use of spectra – once the redshift is well known (e.g. Santini et al. 2015). Moreover, non-solar abundance ratios produce variations over large spectral windows, especially at shorter wavelengths (e.g. Vazdekis et al. 2015), so that broad- and medium-band photometry, in principle, carry such detailed information. This would make surveys such as J-PAS (Benitez et al. 2014, featuring 56 filters with 150 Å bandwidth), or slitless grism spectroscopy (with e.g. ACS or WFC3 at the *Hubble Space Telescope*, NIRISS at the *JWST*, or NISP at *Euclid*) as informative, from the point of view of stellar populations, as the more expensive spectral surveys. Studies of entropy variations *within* spectral data of the same galaxy – from integral field unit observations – can also be exploited to understand radial variations of the underlying stellar populations, a concept that will be explored in future work.

ACKNOWLEDGEMENTS

IF acknowledges support from the Spanish Research Agency of the Ministry of Science and Innovation (AEI-MICINN) under the grant with reference PID2019-104788GB-I00. OL acknowledges STFC Consolidated Grant ST/R000476/1 and a Visiting Fellowship at All Souls College, Oxford. This paper was meant to be produced during a visit to the Flatiron Institute (FI), but it was not possible due to the SARS-CoV-2 pandemic. Nevertheless, IF warmly thanks the FI for offering to host his visit. Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is http://www.sdss3.org/.

DATA AVAILABILITY

This work has been fully based on publicly available data: Galaxy spectra were retrieved from the SDSS DR16 archive (https://www.sdss.org/dr16/) and stellar population synthesis models can be obtained from the respective authors.

REFERENCES

Ahumada R. et al., 2020, ApJS, 249, 3
Andrae R., Schulze-Hartung T., Melchior P., 2010, preprint (arXiv:1012.3754)
Angthopo J., Ferreras I., Silk J., 2019, MNRAS, 488, L99
Angthopo J., Ferreras I., Silk J., 2020, MNRAS, 495, 2720
Baldwin J. A., Phillips M. M., Terlevich R., 1981, PASP, 93, 5
Balogh M. L., Morris S. L., Yee H. K. C., Carlberg R. G., Ellingson E., 1999, ApJ, 527, 54
Benitez N. et al., 2014, preprint (arXiv:1403.5237)
Bernardi M. et al., 2003, AJ, 125, 1849
Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, MNRAS, 351, 1151
Bruzual A. G., Charlot S., 1993, ApJ, 405, 538
Bruzual G., Charlot S., 2003, MNRAS, 344, 1000
Cappellari M., 2012, Astrophysics Source Code Library, record ascl:1210.002
Chabrier G., 2003, PASP, 115, 763
Cid Fernandes R., Mateus A., Sodré L., Stasińska G., Gomes J. M., 2005, MNRAS, 358, 363
Cid Fernandes R., Stasińska G., Mateus A., Vale Asari N., 2011, MNRAS, 413, 1687
Coelho P., Barbuy B., Meléndez J., Schiavon R. P., Castilho B. V., 2005, A&A, 443, 735
Conroy C., 2013, ARA&A, 51, 393
Conroy C., Gunn J. E., 2010, ApJ, 712, 833
Díaz-García L. A. et al., 2015, A&A, 582, A14
Ferreras I., 2012, in Tuffs R. J., Popescu C. C., eds, Proc. IAU Symp. 284, The Spectral Energy Distribution of Galaxies. Cambridge Univ. Press, Cambridge, p. 38
Ferreras I., Charlot S., Silk J., 1999, ApJ, 521, 81
Ferreras I., Pasquali A., de Carvalho R. R., de la Rosa I. G., Lahav O., 2006, MNRAS, 370, 828
Ferreras I. et al., 2009, ApJ, 706, 158
Ferreras I. et al., 2019, MNRAS, 489, 608
Fioc M., Rocca-Volmerange B., 1997, A&A, 500, 507
Fitzpatrick E. L., 1999, PASP, 111, 63
Gallazzi A., Charlot S., Brinchmann J., White S. D. M., Tremonti C. A., 2005, MNRAS, 362, 41
González J. J., 1993, PhD thesis, Univ. California, Santa Cruz
Graves G. J., Faber S. M., Schiavon R. P., 2009, ApJ, 698, 1590
Hawkins K., Jofré P., Gilmore G., Masseron T., 2014, MNRAS, 445, 2575
Heavens A. F., Jimenez R., Lahav O., 2000, MNRAS, 317, 965
Hyvärinen A., Oja E., Karhunen J., 1997, Independent Component Analysis. Wiley-Interscience, New York
Kabán A., Nolan L. A., Raychaudhury S., 2005, in Kargupta H., Srivastava J., Kamath C., Goodman A., eds, Proc. SIAM 2005 Int. Conf. Data Min., Finding Young Stellar Populations in Elliptical Galaxies from Independent Components of Optical Spectra. SIAM, USA, p. 183
Kauffmann G., 2014, MNRAS, 441, 2717
Kauffmann G. et al., 2003, MNRAS, 341, 33
Koleva M., Prugniel P., Bouchard A., Wu Y., 2011, Astrophysics Source Code Library, record ascl:1104.007
Kullback S., Leibler R. A., 1951, Ann. Math. Statist., 22, 79

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
Le Borgne J.-F. et al., 2003, A&A, 402, 433
Liew-Cain C. L., Kawata D., Sánchez-Blázquez P., Ferreras I., Symeonidis M., 2021, MNRAS, 502, 1355
Lovell C. C., Acquaviva V., Thomas P. A., Iyer K. G., Gawiser E., Wilkins S. M., 2019, MNRAS, 490, 5503
MacKay D., 2003, Information Theory, Inference and Learning Algorithms. Cambridge Univ. Press, Cambridge
Madgwick D. S. et al., 2002, MNRAS, 333, 133
Madgwick D. S., Somerville R., Lahav O., Ellis R., 2003, MNRAS, 343, 871
Maraston C., 2005, MNRAS, 362, 799
Nolan L. A., Harva M. O., Kabán A., Raychaudhury S., 2006, MNRAS, 366, 321
Nolan L. A., Raychaudhury S., Kabán A., 2007, MNRAS, 375, 381
Ocvirk P., Pichon C., Lançon A., Thiébaut E., 2006, MNRAS, 365, 46
Panter B., Heavens A. F., Jimenez R., 2003, MNRAS, 343, 1145
Pérez-González P. G. et al., 2013, ApJ, 762, 46
Portillo S. K. N., Parejko J. K., Vergara J. R., Connolly A. J., 2020, AJ, 160, 45
Rogers B., Ferreras I., Lahav O., Bernardi M., Kaviraj S., Yi S. K., 2007, MNRAS, 382, 750
Rogers B., Ferreras I., Peletier R., Silk J., 2010, MNRAS, 402, 447
Ronen S., Aragon-Salamanca A., Lahav O., 1999, MNRAS, 303, 284
Salim S., 2014, SerAJ, 189, 1
Sánchez-Blázquez P. et al., 2006, MNRAS, 371, 703
Santini P. et al., 2015, ApJ, 801, 97
Schawinski K., Thomas D., Sarzi M., Maraston C., Kaviraj S., Joo S.-J., Yi S. K., Silk J., 2007, MNRAS, 382, 1415
Shannon C. E., 1951, Bell Syst. Tech. J., 30, 50
Shannon C., 1953, IEEE, 1, 44
Shannon C. E., Weaver W., 1975, The Mathematical Theory of Communication. Univ. Illinois Press, Urbana
Slonim N., Somerville R., Tishby N., Lahav O., 2001, MNRAS, 323, 270
Strateva I. et al., 2001, AJ, 122, 1861
Teimoorinia H., Archinuk F., Woo J., Shishehchi S., Bluck A. F. L., 2022, AJ, 163, 71
Thomas D., Maraston C., Bender R., Mendes de Oliveira C., 2005, ApJ, 621, 673
Tojeiro R., Heavens A. F., Jimenez R., Panter B., 2007, MNRAS, 381, 1252
Trager S. C., Worthey G., Faber S. M., Burstein D., González J. J., 1998, ApJS, 116, 1
Trager S. C., Faber S. M., Worthey G., González J. J., 2000, AJ, 119, 1645
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., Ricciardelli E., Cenarro A. J., Rivero-González J. G., Díaz-García L. A., Falcón-Barroso J., 2012, MNRAS, 424, 157
Vazdekis A. et al., 2015, MNRAS, 449, 1177
Walcher J., Groves B., Budavári T., Dale D., 2011, Ap&SS, 331, 1
Wild V., Hewett P. C., 2005, MNRAS, 358, 1083
Wild V. et al., 2014, MNRAS, 440, 1880
Wilkinson D. M., Maraston C., Goddard D., Thomas D., Parikh T., 2017, MNRAS, 472, 4297
Worthey G., 1994, ApJS, 95, 107
Worthey G., Ottaviani D. L., 1997, ApJS, 111, 377
Worthey G., Faber S. M., Gonzalez J. J., Burstein D., 1994, ApJS, 94, 687
Wu P.-F. et al., 2018, ApJ, 855, 85
York D. G. et al., 2000, AJ, 120, 1579

APPENDIX A: LS COVARIANCE

Fig. A1 shows a comparison of the distribution of LSs of galaxy spectra from SDSS in the six spectral regions targeted in this paper. The data are colour-coded according to nebular emission into SF (blue), AGN (green), and Q (red), see main text for details. Note the strong correlation between the 4000 Å break strength and either Hδ or Mgb, as can also be seen in the covariance matrix of the SDSS spectra (Fig. 11).

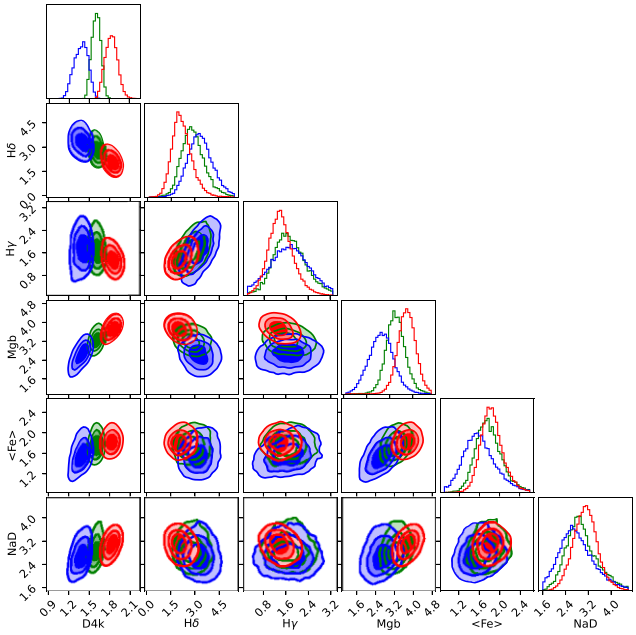

Figure A1 Distribution of observed LSs in our working sample of SDSS galaxies. The contours engulf 25, 50, and 75 per cent of each subsample of SF (blue), AGN (green), and Q (red) galaxies.

APPENDIX B: SELECTING A SAMPLE AT HIGHER S/N

In Fig. 9, we saw that the distribution of galaxy spectra at extreme ends of the projection on to PC1 feature a significant correlation with S/N. While this could be an indirect trend caused by the sample selection, we reanalyse here a subset of spectra with a more stringent cut on S/N (>30, in contrast with S/N > 10 for the original sample). The sample decreases by a factor ∼7, from the original 76 569 spectra to 10 574, still a reasonable number for our statistical analysis. The results are presented in Figs B1, B2, and B3. The

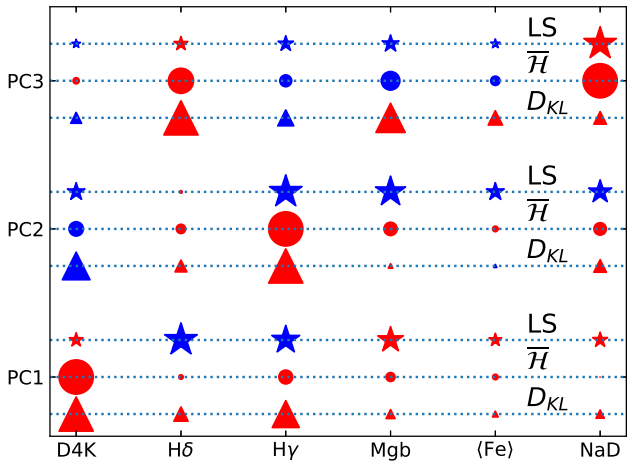

Figure B1 Equivalent of Fig. 7 for a subsample of spectra restricted to S/N > 30 (instead of 10 in the original sample).

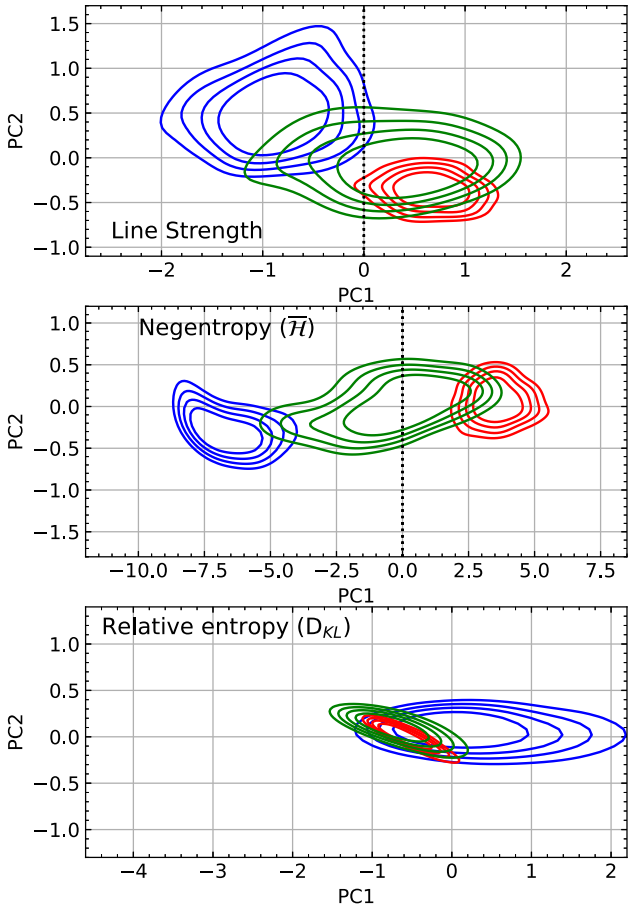

Figure B2 Equivalent of Fig. 8 for a subsample of spectra restricted to S/N > 30 (instead of 10 in the original sample).

eigenvalues of the PCA in this subset are given in Table B1, and the principal components are quantified in Table B2. The results are very similar to those presented in the main text for the original sample. The eigenvalues show that the variance is, once more, dominated by ∼2 principal components when using entropy estimates, and appears more extended with the traditional LSs – requiring, for instance, four components with the LS inputs to engulf 98 per cent of variance, whereas entropy (relative entropy) only require two (three) components. This subset also depends on S/N in a similar way, although the sample appears culled at low velocity dispersion – understandably from the fact that lower σ galaxies are overall fainter and thus will correspond to lower S/N in the SDSS spectra. It is also worth noting the redshift trend in Q and AGN spectra, as the higher S/N produces an effectively shallower survey, and is thus more prone to Malmquist bias. Consistently to the original sample, the principal components also feature the 4000 Å break strength as the dominant one (PC1), but the second component is now mostly dependent on Hγ (for both definitions of entropy). The contribution from Hδ is now relegated to PC3, in combination with NaD. In any case, the distribution of principal component projections (Fig. B2) is similar to the original case, suggesting that while S/N should be taken into account, it does not dominate the variations found in the data.

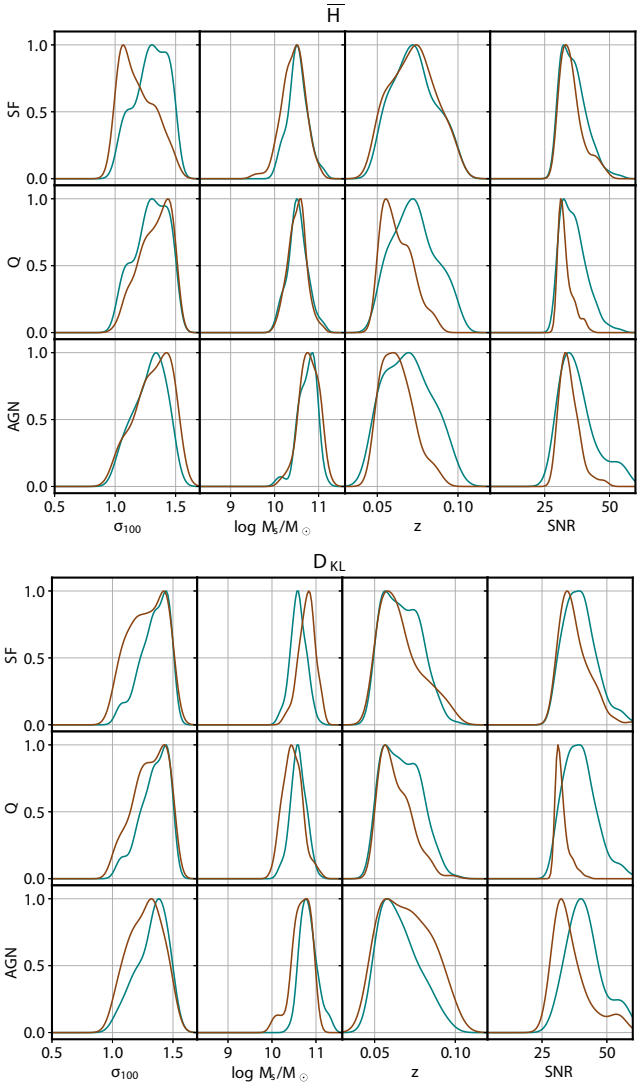

Figure B3 Equivalent of Fig. 9 for a subsample of spectra restricted to S/N > 30 (instead of 10 in the original sample).

Table B1. Normalized PCA eigenvalues (high S/N).

Rank	Eigenvalue ratio		
	LS	Entropy	D_{KL}
1	0.527 06	0.976 71	0.894 43
2	0.290 45	0.013 64	0.071 26
3	0.128 32	0.003 79	0.013 37
4	0.039 83	0.002 78	0.010 67
5	0.010 74	0.002 23	0.005 56
6	0.003 60	0.000 85	0.004 72

Table B2. The first three PCs of the high S/N subsample (normalized).

	LS			$\overline{\mathcal{H}}$			$D_{\rm KL}$	
Region	PC1	PC2	PC3	PC1	PC2	PC3	PC1	PC2	PC3
D4K	0.127 86	− 0.198 73	− 0.064 82	0.983 95	− 0.170 63	0.027 82	0.836 37	− 0.541 42	− 0.079 80
Hδ	− 0.703 98	0.005 36	0.169 15	0.016 93	0.071 90	0.455 32	0.145 86	0.100 89	0.779 44
Hγ	− 0.517 32	− 0.682 09	− 0.207 33	0.162 01	0.962 80	− 0.102 10	0.522 13	0.826 80	− 0.171 14
Mgb	0.431 58	− 0.580 63	− 0.251 18	0.068 14	0.146 75	− 0.250 22	0.059 06	0.015 03	0.571 33
⟨Fe⟩	0.110 80	− 0.202 73	− 0.075 46	0.025 70	0.026 06	− 0.063 14	0.023 99	− 0.007 19	0.137 97
NaD	0.147 99	− 0.342 04	0.924 88	0.000 16	0.128 54	0.845 51	0.050 27	0.113 18	0.106 59

This paper has been typeset from a TEX/LATEX file prepared by the author.