

Probing cool giants in unresolved galaxies using fluctuation eigenspectra: A demonstration using high-resolution MUSE observations of NGC 5128

Russell J. Smith  

Centre for Extragalactic Astronomy and Department of Physics, University of Durham, Durham DH1 3LE, United Kingdom

Accepted 2021 November 18. Received 2021 November 12; in original form 2021 September 6

ABSTRACT

I describe and demonstrate a new approach to using spectroscopic data to exploit Poisson sampling fluctuations in unresolved stellar populations. The method is introduced using spectra predicted for independent samples of stars from a 10 Gyr population using a simple stochastic spectral synthesis model. A principal components analysis shows that >99 per cent of the spectral variation in the red-optical can be attributed to just three ‘fluctuation eigenspectra’, which can be related to the number of giant stars present in each sample, and their distribution along the isochrone. The first eigenspectrum effectively encodes the spectrum of the coolest giant branch stars, and is equivalent to the ratio between high- and low-flux pixels discussed in previous literature. The second and third eigenspectra carry higher-order information from which the giant-star spectral sequence can in principle be reconstructed. I demonstrate the method in practice using observations of part of NGC 5128, obtained with the MUSE narrow-field adaptive optics mode. The expected first eigenspectrum is easily recovered from the data, and closely matches the model results except for small differences around the Ca II triplet. The second eigenspectrum is below the noise level of the present observations. A future application of the method would be to the cores of giant ellipticals to probe the spectra of cool giant stars at high metallicity and with element abundance patterns not accessible in the Milky Way.

Key words: methods: data analysis – galaxies: individual: NGC 5128 – galaxies: stellar content.

1 INTRODUCTION

Integrated-light spectroscopy is the pre-eminent method for studying the stellar content of passive galaxies that are too distant to be resolved into individual stars. The total observed spectrum mixes signal from (at least) cool dwarfs, warm turn-off stars, and evolved giants, in varying proportions. Yet by comparison to the equivalent spectrum predicted from population synthesis models, a rich variety of properties can be constrained, including age or star-formation history, the initial mass function (IMF), and metallicity or individual element abundances (see, e.g. Conroy 2013). Such studies perform the twin roles of (a) probing the formation histories of galaxies beyond the reach of resolved photometry, and (b) testing the model inputs and stellar physics under conditions which differ from those pertaining in the Milky Way. The cores of elliptical galaxies are especially interesting from both perspectives: they are thought to have formed at the earliest epochs in violent starbursts, and they exhibit peculiarities, such as super-solar metallicity, enhancements in α elements and in sodium, and strong dwarf-star spectral features, attributed to a bottom heavy IMF (e.g. Conroy, Graves & van Dokkum 2014; van Dokkum et al. 2017; La Barbera et al. 2019).

Although truly resolved studies of individual stars in the cores of ellipticals are impossible, additional information about the brightest stars can in principle be retrieved from high-resolution observations, via the measurable statistical fluctuations in the small number of giants contributing to each resolution element. In the imaging

domain, this effect has long been employed in the form of surface brightness fluctuations (SBFs) as introduced by Tonry & Schneider (1988). The fluctuations in broad-band flux, on the scale of the image point spread function (PSF), reflect the magnitude of the brightest stars present, which dominate the Poisson variance. The SBFs can thus be used as a distance indicator, by calibrating the fluctuation magnitude as a standard candle (e.g. Tonry et al. 2001) and as a probe of the stellar populations, where they carry information which is distinct from the integrated colour or spectrum (e.g. Liu, Graham & Charlot 2002).

Several works have considered the exploitation of Poisson fluctuations in the spectroscopic domain. One approach is simply to apply the standard SBF concept to narrow wavelength channels. The ‘SBF spectrum’ can be straightforwardly calculated from stellar population synthesis models (e.g. Mitzkus et al. 2018; Vazdekis et al. 2020), as the ratio of the variance to the mean flux at each wavelength. In practice, the SBF spectrum can be measured in datacubes from integral field unit (IFU) spectrographs, by applying the classic imaging-processing technique on a channel-by-channel basis. Mitzkus et al. (2018) have pioneered this approach, using MUSE¹ observations of the post-starburst S0 galaxy NGC 5102. They confirmed that the observed SBF spectrum is dominated by the spectral features of M-type giant stars, and showed that using a combination of the mean and SBF spectrum could tighten the inferred

¹The Multi-Unit Spectroscopic Explorer (Bacon et al. 2010) on the ESO 8.2m Very Large Telescope.

* E-mail: russell.smith@durham.ac.uk




Figure 1. Ten samples from the 10 Gyr stochastic spectral synthesis model, showing the variation caused by Poisson fluctuations in the giant stars contributing to independent pixels. The plotting colour is ordered according to the first eigenvalue, P_1 , derived in the principal components analysis. The black line shows the mean spectrum over all 3500 samples computed, not just those shown. A linear continuum has been divided out of each spectrum.

constraints on the age and metallicity of the stellar population, compared to the mean spectrum alone.

An alternative approach to the Poisson fluctuation regime was discussed by van Dokkum & Conroy (2014). Noting that the SBFs in imaging data are driven by variation in the number of luminous stars, they used synthesis models to compute the expected ratio of spectra from high- and low-brightness pixels. In contrast to SBF spectra, which treat each wavelength channel independently, this ‘fluctuation-ratio spectrum’ explicitly captures the correlated spectral structure in the pixel-to-pixel variation. In van Dokkum & Conroy (2014), the method was demonstrated not on true spectroscopic data but instead by using the spectra to predict how narrow-band pixel colours should correlate with brightness fluctuation in images from the *Hubble Space Telescope* (*HST*).

In this paper, I introduce an alternative approach to spectroscopy in the fluctuation regime, based on an eigenspectrum decomposition of the variation among independent samples from the population. Section 2 first describes the construction of synthetic data by Poisson sampling from an underlying simple stellar population model. A principal components analysis is then applied to retrieve the leading eigenspectra, and the relationship of these to the underlying stellar content is explored. Section 3 presents a first observational application of this method, using adaptive optics observations of a stellar field in the peculiar early-type galaxy NGC 5128. The results are evaluated through comparison to the model eigenspectra, and by inspection of the corresponding eigenvalue maps. Section 4 discusses the merits of and challenges for the method, and prospects for its further application to more distant, but less disturbed, elliptical galaxies. The main conclusions are summarized in Section 5.

2 FLUCTUATION EIGENSPECTRA

In this section, I describe the process of building spectra from stochastically-sampled population synthesis models, and decomposing them into eigenspectra which describe the major contributions to the variation among independent samples. Some properties of the synthetic data are chosen to match the characteristics of the MUSE data analysed in Section 3, but most of the principles are quite general.

The calculations are based on the usual ingredients and processes used for spectral synthesis; the only unusual feature is the inclusion of stochastic sampling from the mass function. The process begins with an isochrone describing the locus of the stellar parameters (temperature, surface gravity, luminosity, etc), as a function of initial mass. All of the models in this paper are computed using the 10 Gyr isochrone with solar metallicity from the PARSEC² stellar tracks³ (Bressan et al. 2012). Appropriate stellar spectra are then associated to the isochrone points, using the SPI approximator⁴ to interpolate between empirical spectra from the MILES⁵ (Sánchez-Blázquez et al. 2006) and Extended IRTF⁶ (Villaume et al. 2017) libraries. For ease of comparison to the data for NGC 5128, all spectra are smoothed to its characteristic velocity dispersion of 150 km s^{-1} (Wilkinson et al. 1986). Given the IMF (here a single power law with the Salpeter slope), the mean number of stars present on each segment of the isochrone can be defined up to a normalizing constant, which depends on the surface brightness, the distance, and the angular area of the spatial sample. The isochrone is then populated by drawing Poisson-distributed star counts according to these means, and the stellar contributions are summed with the corresponding number weights and luminosity weights, to generate the integrated spectrum $S_i(\lambda)$ for the i ’th individual sample. For the illustrative purposes of this paper, the samples are assumed to be fully independent, but the effects of smoothing by a PSF could be included in more realistic calculations.

Fig. 1 shows the red-optical spectra computed for ten samples from the population, compared to the mean synthesized spectrum. The calculations are matched to the observational case described in Section 3. In particular, the adopted surface brightness ($\mu_l = 18.3 \text{ mag arcsec}^{-2}$), and distance modulus (27.9 mag) are

²Padova and TRIeste Stellar Evolution Code.

³These tracks include evolution on the asymptotic giant branch up to the first thermal pulse; including the subsequent evolution using the tracks of Marigo et al. (2013) did not affect any of the results for the old populations used here.

⁴Spectrum Polynomial Interpolator (Villaume et al. 2017).

⁵Medium-resolution Isaac Newton Telescope Library of Empirical Spectra.

⁶Infrared Telescope Facility.

appropriate to the observed field in NGC 5128, the spectra are rebinned at 1.25 Å intervals (as for MUSE) and limited to 6750–9200 Å (where the observed PSF is sharpest), and the spatial samples are 0.125×0.125 arcsec² in area (5×5 MUSE pixels in the narrow-field mode). With this configuration, the mean luminosity per sample is $\sim 4000 L_{\odot}$, of which ~ 60 per cent is generated by ~ 130 giant stars. In Fig. 1, the structured effect of the Poisson fluctuations in these giants is clearly visible: samples that are fainter at $\sim 0.715 \mu\text{m}$ are typically brighter at $\sim 0.755 \mu\text{m}$ and $\sim 0.815 \mu\text{m}$, but fainter at $\sim 0.770 \mu\text{m}$. Note that a linear continuum shape has been divided out of all the spectra before plotting, to emphasize the narrower spectral features, and to mimic the treatment of the models and data below.

The variations in Fig. 1 arise from the most luminous giants, which are present in small number and hence subject to the largest statistical fluctuations. The spectra of these cool stars are dominated by molecular band heads from the TiO γ , δ , and ϵ systems. As described in Section 1, the Poisson sampling variation has sometimes been described by computing the SBF spectrum, defined as $\text{var}[S(\lambda)]/\langle S(\lambda) \rangle$ (Mitzkus et al. 2018). This method recovers a strong signal in the Ti O bands because M-giants dominate the fluctuations, but since it treats each spectral channel independently, the SBF spectrum does not explicitly encode the *correlated* spectral variations seen in Fig. 1. The van Dokkum & Conroy (2014) approach, based on a ratio of spectra between samples with higher- versus lower-than-average flux, i.e. more or fewer giants, does capture the correlated variations, but requires a level of relative photometric calibration that may be challenging to achieve with IFU data.

The new direction taken in this paper is to decompose the spectral variations between samples into a small set of eigenspectra, using principal components analysis (PCA) to identify the optimal parametrization. In preparation for the PCA, each synthetic spectrum is divided by a low-order polynomial continuum, to remove sample-to-sample colour variations. This ensures that the extracted information is independent of luminosity or broadband colour, which may be affected by calibration uncertainties. The spectra are also divided by the mean synthesized spectrum, to isolate the variations: $\Delta S_i(\lambda) = S_i(\lambda)/\langle S(\lambda) \rangle - 1$.

The PCA algorithm then decomposes the correlation structure of the $\Delta S_i(\lambda)$ into eigenspectra $E_j(\lambda)$, ordered by their contribution to the total variance, such that

$$\Delta S_i(\lambda) \approx P_1 E_1(\lambda) + P_2 E_2(\lambda) + \dots + P_m E_m(\lambda)$$

(with different values of P_1 , P_2 , etc for each i) provides an acceptable approximation, for some small number m of retained components. In the absence of observational error, the eigenspectra are uniquely determined by the inputs chosen for the synthesis model (e.g. the isochrone, stellar library), and are independent of the number of stars in each sample. The eigenspectra can be considered to have the same ‘status’ as an output of the model as does the mean spectrum, or the SBF spectrum. From the discussion above, we can expect that the first fluctuation eigenspectrum, E_1 , should be driven by the varying number of cool luminous giants present in the samples, similar to the van Dokkum & Conroy (2014) fluctuation spectrum. Subsequent E_j can be expected to reflect the difference in distribution of the giants in each sample, and hence probe the systematic trend in stellar spectra along the giant sequence.

Fig. 2 shows the mean spectrum $\langle S \rangle$, and the eigenspectra $E_1 - E_3$, computed for the 10 Gyr age, solar metallicity, Salpeter-IMF population model, derived from 3500 samples similar to those in Fig. 1. The sign of the eigenspectra is arbitrary in the PCA process, so they have been plotted here such that the strong band at $0.71 \mu\text{m}$ is shown in ‘absorption’ in each spectrum.

The first principal component accounts for 77.6 per cent of the total variance among samples. The corresponding eigenspectrum, E_1 , is dominated by the Ti O bands as foreseen from Fig. 1, while no trace of the Ca II triplet lines is present. E_1 closely resembles the van Dokkum & Conroy (2014) fluctuation-ratio spectrum, as derived from the same model predictions. The second principal component accounts for a further 16.2 per cent of the total variance. Its eigenspectrum, E_2 , again shows the strong Ti O band at $0.71 \mu\text{m}$, but the band at $0.76 \mu\text{m}$ is barely present, and those at $0.82 \mu\text{m}$ and $0.88 \mu\text{m}$ are reversed in sign compared to E_1 . The Ca II triplet lines are also more noticeable. Thus the second component encodes a ‘correction term’, which strengthens Ca II, and the Ti O bands in the blue, but weakens those in the red. The third principal component accounts for a further 5.6 per cent of the total variance; its eigenspectrum, E_3 shows weaker molecular features, and the Ca II lines are more prominent. All subsequent components collectively contribute only ~ 0.6 per cent of the variance. Thus although the stars can be distributed along the model isochrone in infinite combinations, almost all of the emergent spectral variation can be compressed into just three linear terms.

Fig. 3 illustrates how the eigenvalues P_1 and P_2 relate to the distribution of stars on the giant branches in individual samples from the population. As expected from the above discussion, a larger value for P_1 (e.g. sample ‘B’ in the figure) indeed corresponds to a larger contribution from the coolest giants, while the lowest values are found for samples having effectively no giants cooler than ~ 4000 K (e.g. sample ‘A’). Comparing samples with similar P_1 values, it can be seen that higher P_2 corresponds to more continuously populated giant branches (e.g. sample ‘D’), while lower P_2 is associated with a wider gap below the brightest star present (e.g. sample ‘C’).

Guided by Fig. 3, some physical interpretation of the first two principal components can be inferred. The first component evidently encodes the total cool-giant content. Since the PCA was applied to the ratio between sample spectra and the mean spectrum, E_1 should be simply the ratio between a characteristic giant star and the mean. Comparison to each stellar spectrum along the isochrone reveals a close match for stars at the very tip of the giant branch, with $T_{\text{eff}} = 3136$ K (spectral type $\sim M7$); the match to any star warmer than 3500 K ($\sim M4$) is notably poor. Thus E_1 effectively isolates the spectrum of the coolest giant stars present in the population, which even collectively contribute only a small fraction of the total integrated light (e.g. ~ 7 per cent for giants with $T_{\text{eff}} < 3500$ K).

The second principal component, E_2 , seems to encode the presence or absence of a ‘gap’ on the giant branch, and so should be related to the spectrum of warmer giants than those traced by E_1 . If the giant sequence were ‘linear’, in the sense that any star could be represented as a linear combination of the M7 star with a K0 giant, for example, then the gap would not need to be described with a separate component. It follows that E_2 should represent the *departure* of the intermediate giants from a linear spectral sequence, expressed again as a ratio with the mean spectrum. Minimizing a least-squares statistic over the possible combinations, a very close match is found for $E_2 \approx [M2 - (0.70 K0 + 0.30 M7)]/\langle S \rangle$, where the spectral types are shorthand for unity-normalized stellar spectra with temperatures 3671 K, 4547 K, and 3136 K. Fig. 4 shows how closely this approximation, and the equivalent $E_1 \approx M7/\langle S \rangle$, reproduces the first two fluctuation eigenspectra. A comparably simple interpretation for the third principal component has not been identified. The eigenvalue P_3 is somewhat correlated with the number of red clump stars present, and the features in E_3 certainly suggest a sensitivity to higher temperatures, where the molecular bands are weaker. E_3 will not be discussed further here.




Figure 2. Fluctuation eigenspectra computed from principal component analysis of 3500 samples from the 10 Gyr stochastic synthesis model (blue), compared to the mean spectrum (black). The first three components account for >99 per cent of the spectral variation caused by Poisson fluctuations. For comparison, two alternative indicators of the variations are shown, computed from the same model samples: the wavelength-dependent SBFs as discussed, e.g. by Mitzkus et al. (2018) and Vazdekis et al. (2020) (purple), and the fluctuation-ratio spectrum introduced by van Dokkum & Conroy (2014) (red). The wavelengths of some prominent Ti O bandheads and the Ca II triplet lines are indicated.

To summarize this section, the Poisson variation among independent samples from an old stellar population can be efficiently described using a linear combination of two or three eigenspectra. The first two components, at least, have an intuitive meaning in terms of the spectra of the most luminous stars present in the population. The PCA approach describes the spectral fluctuations with more generality than related methods such as the SBF spectrum, and is well suited to IFU data with imperfect flatfielding and relative flux calibration.

3 OBSERVATIONAL DEMONSTRATION

In this section, I present a practical demonstration of extracting the fluctuation eigenspectra from IFU data, using new high-resolution adaptive optics (AO) observations of a field in NGC 5128.

The data were acquired using MUSE at the ESO Very Large Telescope, operated in its narrow-field mode, which implements Laser Tomographic Adaptive Optics using four laser beacons (Ströbele et al. 2012). The instrument provides a ~ 7 arcsec field of view, sampled with 0.025 arcsec spatial pixels. NGC 5128 was chosen as the target for this observation mainly for its proximity (3.8 ± 0.1 Mpc; Harris, Rejkuba & Harris 2010), since the Poisson fluctuations are larger for closer galaxies, at a given surface brightness. A bright on-axis guide star was required for the tip/tilt correction in the narrow-field mode when these observations were proposed, and so a field was selected at ~ 80 arcsec (~ 1.5 kpc) from the galaxy nucleus, centred on a $H = 11.0$ foreground star (see top left-hand panel of Fig. 5). *HST* imaging shows that this field is largely free from dust and from bright young stars (top right-hand panel). The surface brightness at this radius is $\mu_J = 17.4$ mag arcsec $^{-2}$ (Jarrett et al. 2003), which converts to $\mu_I \approx 18.3$ mag arcsec $^{-2}$ for an old stellar population.

The observations were made between 2020 February and 2021 April, and comprise four on-target exposures with total integration

time 7122 s. Each individual ~ 1800 s exposure was bracketed by two ~ 330 s offset sky frames. (One further exposure was taken but not used in the analysis as the image quality was much poorer.) Three of the frames have approximately the same pointing, while the fourth was unintentionally offset ~ 2 arcsec south; each exposure was acquired at a different position angle, with relative rotations of 90 degrees.

The custom analysis began from the four observatory-processed pipeline-reduced, and sky-subtracted data cubes. The frames were registered with nearest-pixel shifts determined from the image centroid of the AO star, and the cubes were combined with a median, to clean residual cosmetic artifacts. Only the area covered by at least three pointings was retained for further analysis. Fig. 6 shows the profile of the PSF core in the resulting combined frame, for two faint stars located ~ 3.5 arcsec from the reference star, and not used in the registration process. These profiles can be described by a Gaussian with 0.06 arcsec full width at half maximum (FWHM), combined with a second component of 0.25 arcsec FWHM. Beyond this well-corrected core, the PSF has a diffuse halo which can be traced to ~ 2 arcsec in the bright star. A square 5×5 (or 9×9) pixel block centred on the reference star encloses 38.2 (or 50.4) per cent of the total energy in the PSF.

The halo of the reference star must be subtracted from the data to prevent its spatially-varying light contribution from affecting the PCA fit. It proved too difficult to model out the star based on its distinct spectrum compared to the galaxy, so the method adopted instead assumes a circular spatial profile. All pixels within 1 arcsec of the star were first masked out, and a polynomial approximation to the wings of the radial profile was then fit and subtracted independently for each spectral channel. Forming a collapsed image over the red wavelength region of the subtracted cube confirms detection of numerous faint stars, a few of which can be matched to sources in the *HST* image (lower panels of Fig. 5). The brightest stars seen in the

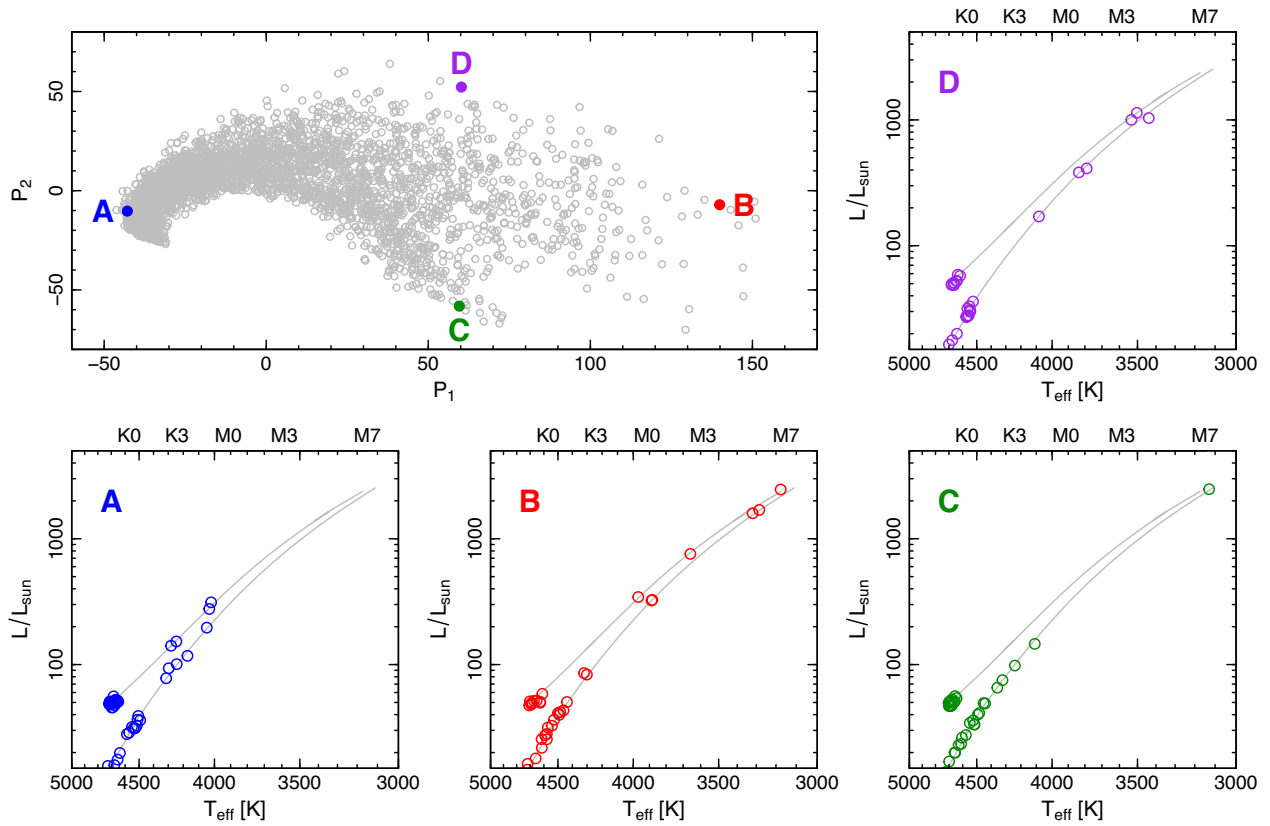


Figure 3. A visual guide to interpreting the spectral fluctuations. The upper left-hand panel shows the first two eigenvalues, P_1 and P_2 , for all of the stochastic samples computed from the 10 Gyr stellar population. The remaining panels show the distribution of stars on the model isochrone for four samples chosen from the extremes of the eigenvalue distribution. (Small random dithers have been applied to help separate multiple stars at the same isochrone point.) Increasing P_1 corresponds to a larger contribution from cool giants (compare samples ‘A’ and ‘B’); increasing P_2 corresponds to a less pronounced gap in the temperature distribution (compare samples ‘C’ and ‘D’). The temperature versus spectral type relation for giants is taken from van Belle et al. (1999).




Figure 4. Relationship of the first two fluctuation eigenspectra to the spectra of individual giant stars in the model. The first eigenspectrum, E_1 , is essentially the ratio between an M7 giant and the mean spectrum of the population. The second, E_2 , describes the deviation of an M2 giant from the closest approximation that can be derived from a linear combination of K0 and M7 giants, also expressed as a ratio with the mean spectrum.

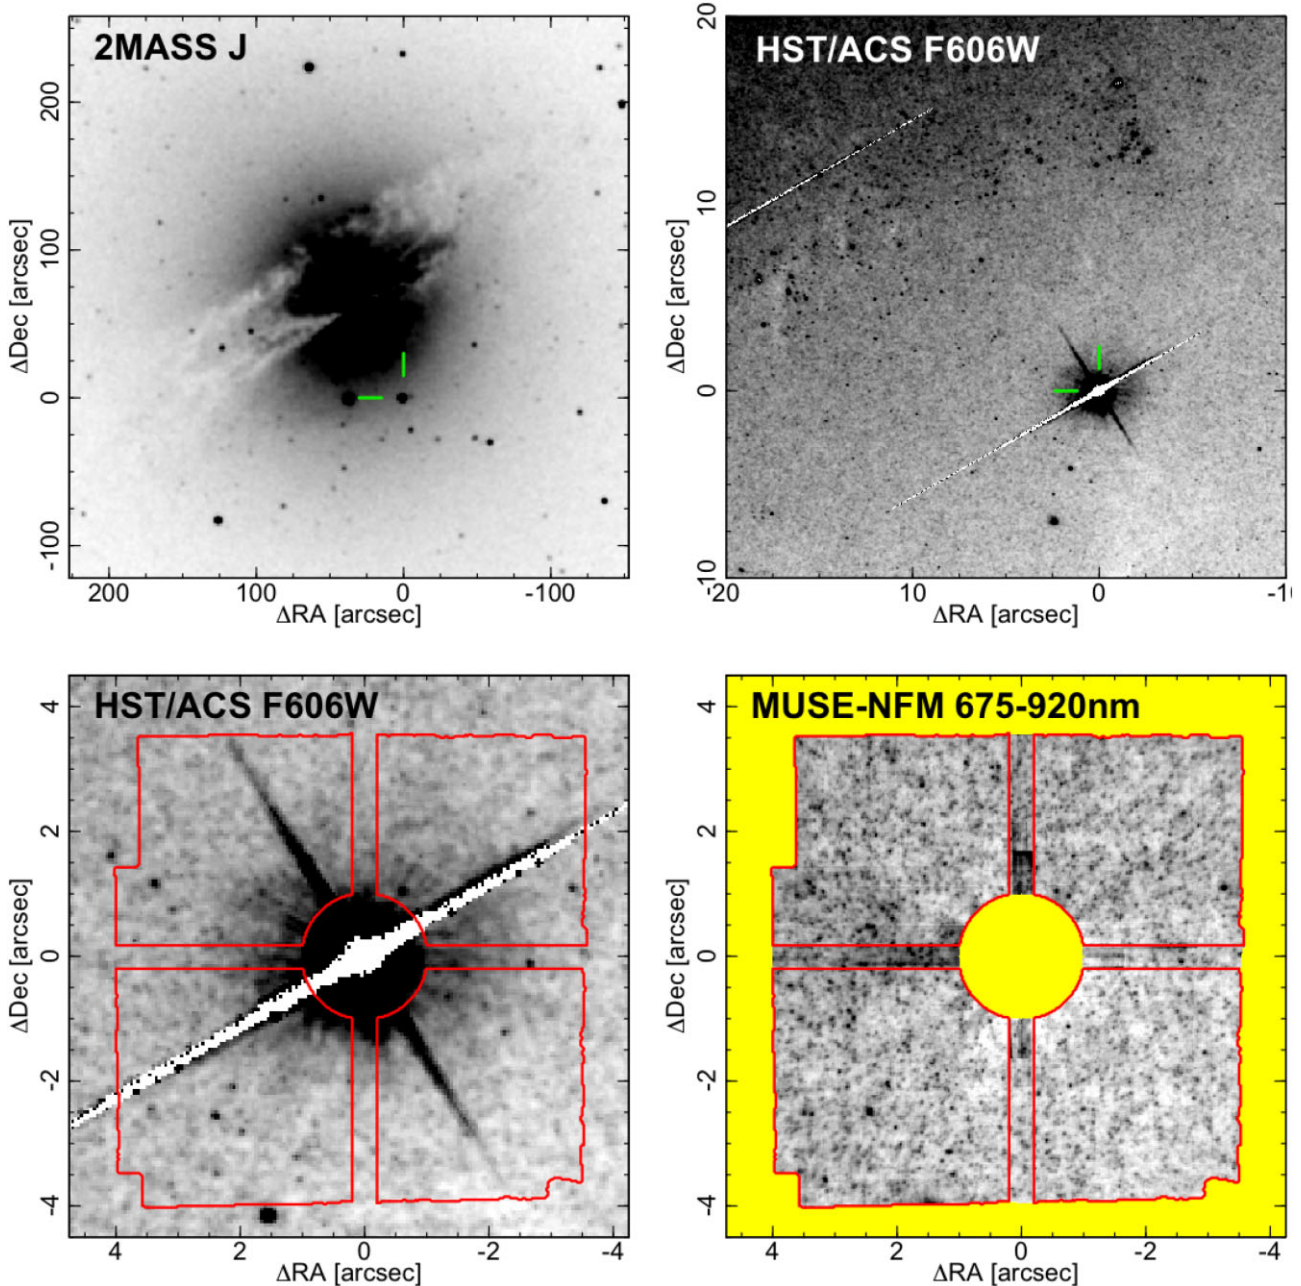


Figure 5. The NGC 5128 MUSE AO field in context. The upper-left-hand panel shows a J-band image from the 2MASS Large Galaxy Atlas (Jarrett et al. 2003), highlighting the foreground reference star used for the low-order wavefront correction. The upper-right-hand panel shows an extract from *HST* imaging in F606W from Harris et al. (2006). The targeted field lies away from the region of recent star formation seen to the north, and is free from prominent dust features. The lower-left-hand panel shows a zoom of the *HST* image, while the lower-right-hand shows the collapsed red-optical image from MUSE on the same scale. The red outline shows the default mask region; pixels in the ‘cross’ are affected by scattered light from the AO star, but are retained here for comparison to the *HST* image.

MUSE image have $I \approx 22$, while the typical fainter stars have $I \approx 24$, consistent with being red giants in NGC 5128. Note however, that the image combines signal over nearly 2000 wavelength channels; very few stars are detectable in the individual channels.⁷

The collapsed image reveals some residual scattered light in a cross pattern around the bright star, which traces the MUSE field-splitter

⁷Only the star at (+ 2.5 arcsec, -2.5 arcsec), with $I \approx 21.8$, is bright enough to extract an informative spectrum. The radial velocity confirms this star as a member of NGC 5128, two magnitudes above the tip of the red giant branch.

boundaries. These regions are masked out from the PCA analysis. A region of diffuse nebular emission ($H\beta$, $H\alpha$, [N II], etc) is present in the north–east corner of the field of view, i.e. towards the disk of NGC 5128; this region was not masked for the default analysis, since the fluctuation analysis is restricted to redder wavelengths, where no strong emission lines are present.

The unmasked pixels were combined into individual spatial samples by median binning over 5×5 pixel (0.125×0.125 arcsec²) blocks, chosen to roughly match the resolution of the PSF core. (Repeating the analysis with 9×9 pixel blocks did not qualitatively alter the results.) For the default masking scheme described above,

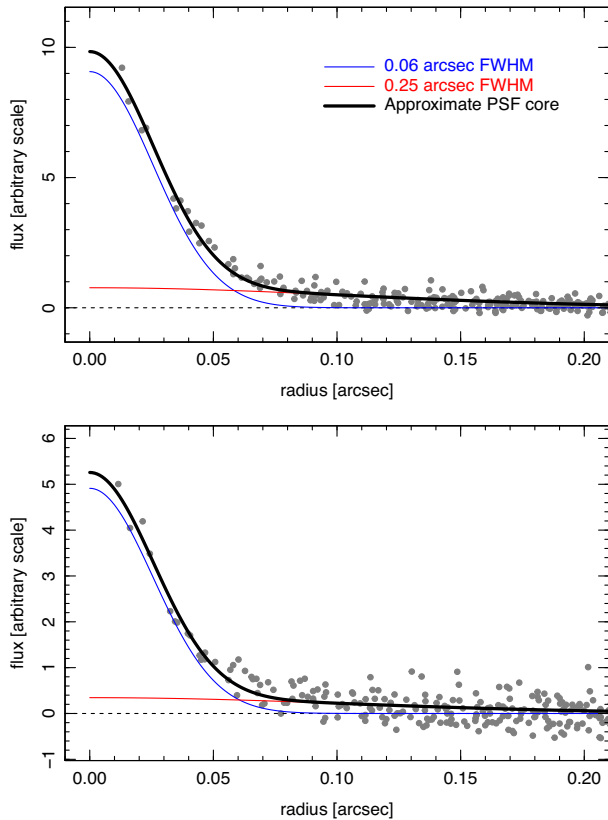


Figure 6. The core of the AO PSF as traced by two faint stars that were not used in the frame registration process. The upper panel shows the $I \approx 21.8$ star seen at (+ 2.5 arcsec, -2.5 arcsec) in Fig. 5, while the lower panel is for the $I \approx 22.4$ star at (-3.3 arcsec, + 1.1 arcsec). The inner profile can be adequately described with a double-Gaussian fit as shown, though a significant outer halo is also present.

this yields 2490 spectra, each of which is individually extremely noisy, with no recognizable features. Due to the method adopted to remove the AO star contamination, the bin spectra do not have a correct continuum level, since a mean was already subtracted at each radius for each channel. Thus for the observed data, the PCA was applied not to $\Delta S_i(\lambda)$ as defined in Section 2, but simply to $S_i(\lambda) - \langle S_i(\lambda) \rangle$. The wavelength range was restricted to 6750–9200 Å, since the AO correction is more effective in the red. The particular choice of blue limit was made to exclude the [S II] lines in the area with nebular emission, while the red limit avoids any contamination from faint second-order light from the AO star. Channels in the interval 7590–7670 Å were also excised, to avoid systematic residuals from imperfect correction of the atmospheric A-band in the pipeline reduction. As a final step before applying the PCA, a low-order polynomial fit was subtracted from each spectral sample, to remove any sensitivity to broadband flux response. Again, this treatment differs from that applied to the models in Sections 2, where the continuum fit was divided out.

The first two fluctuation eigenspectra recovered from the NGC 5128 data are shown in Fig. 7. For comparison to these results, the model predictions were regenerated to match the slightly different treatment used for the data, i.e. with subtraction of the mean and the continuum from each sample. Additionally, the PCA was run on a version of the model samples to which uncorrelated random noise was added to match the wavelength-dependent signal-to-noise ratio of the data. Recall, however, that the model is not a fit to the data,

and has not been tuned to the specific population characteristics of NGC 5128.

The first eigenspectrum is extracted cleanly from the observed spectra, and is remarkably free from sky subtraction artifacts. The strong molecular band features seen in the predicted E_1 are immediately recognizable in the observational equivalent, and the detailed structures of the bands are very well reproduced, after applying a smoothing to $\sigma = 150 \text{ km s}^{-1}$ to match the intrinsic velocity dispersion. On the other hand, the overall relative strengths of the features differ between the observations and the model, with TiO bands in the blue appearing weaker in the data compared to those in the red. This might be astrophysical, but could also arise from the wavelength-dependent AO PSF, which leads to the spectral samples being less fully independent in the blue than in the red, which was not accounted for in the model calculations. A more notable difference between the observed and model E_1 is apparent in the Ca II triplet lines, which are more prominent in the data than in the model. Intuitively, this suggests the presence of warmer stars among the luminous giants which dominate the fluctuation signal. A future exploration of the model parameter space should help to interpret this result.

While the fluctuation eigenspectra E_1 , E_2 , etc, carry meaningful information about the spectra of the giant stars, the values of P_1 , P_2 etc measured for each spatial sample should be essentially random, if the samples are truly drawn from the same parent population. To test whether this is the case, the right-hand panels of Fig. 7 show the derived eigenvalues mapped across the observed field of view. Any coherent structure here would suggest either residual systematic errors in the data (e.g. a radial trend might reveal incomplete subtraction of the AO star light) or true astrophysical variation across the field of view (e.g. if the fluctuations were driven by areas with younger populations). The map for the first eigenvalue, P_1 , in fact does show a uniform distribution within the field, as expected for stochastic sampling variation. As further confirmation of this behaviour, applying the PCA analysis independently for each ‘quadrant’ of the observed field yields E_1 spectra consistent with the whole-field results.

The observed second eigenspectrum and its corresponding eigenvalue map indicate that this component is *not* recovering the variation in giant-star distribution encoded in the model E_2 . Strong features are visible in the spectrum around the major OH sky line complexes, and a distinct boundary is present in the P_2 map, dividing regions observed in four exposures from those with only three. Evidently E_2 is being driven by variations in the sky-subtraction residuals between the four frames. Fig. 8 shows the equivalent results after applying a more conservative masking of the data, which retains only pixels covered by all four exposures, and also excises the region where extended ionized gas was detected. While E_1 is essentially unchanged by this more restrictive mask, the second component no longer shows any spatial structure in the eigenvalue map, or obvious sky features in the eigenspectrum. However, except for a possible weak dip at $0.715 \mu\text{m}$, the recovered E_2 does not show any clear signal from the stellar features seen in the models.

Applying the PCA algorithm to the noisy observational data (typical $S/N \approx 2$ per sample), the leading components naturally account for a very much smaller fraction of the total variance than in the noise-free simulations. For the default masking scheme, E_1 contributes 0.51 per cent, while E_2 carries 0.26 per cent and many subsequent components then contribute at an almost constant level of ~ 0.19 per cent. (The hundred-fold reduction in variance for the lead component is reproduced in the model PCA, when noise is added to match the data quality.) With the more conservative




Figure 7. First and second fluctuation eigenspectra (left) and corresponding eigenvalue maps (right), derived from the NGC 5128 data using the default masking scheme. In the spectrum plots, the grey trace is the recovered eigenspectrum, while the black shows the same after smoothing by $\sigma = 150 \text{ km s}^{-1}$ to suppress noise below the effective resolution. The red trace for E_1 is the equivalent computed from the noise-free 10 Gyr solar-metallicity stellar population model; the blue trace shows results from the same model after adding noise to match the data quality, and smoothing as for the data. (The model PCA has here been recomputed to match the slightly different treatment necessary for the data; see text for details.) The eigenvalue map shows that the variations associated with E_1 are smoothly distributed across the field of view, as expected if P_1 traces stochastic fluctuations in the giant-star content per spatial sample. The second eigenspectrum shows sharp features due to variations in the sky-subtraction residuals, and the coherent structure in its eigenvalue map confirms that this signal is *not* astrophysical in origin.




Figure 8. First and second fluctuation eigenspectra retrieved using a more conservative masking scheme. The first component is almost unchanged from the default result. The spatial structure in P_2 is removed, but the corresponding eigenspectrum is now simply dominated by random noise.

mask, E_1 contributes 0.63 per cent of the variance, while E_2 and subsequent components each carry ~ 0.25 per cent. Fig. 9 compares these variance contributions between the model and the observations, normalizing to E_1 in each case, to assess whether a reliable measurement of E_2 is feasible. Apart from the sky-contaminated component in the default-mask results, the observational curves flatten at a level

around twice as high as the expected intrinsic variation from the second eigenspectrum, showing that substantially deeper data would be necessary to retrieve E_2 from this field.

To summarize this section, the MUSE observations of NGC 5128 provide enough information to extract the spectrum of the most luminous giants, as characterized by E_1 , but the second-order

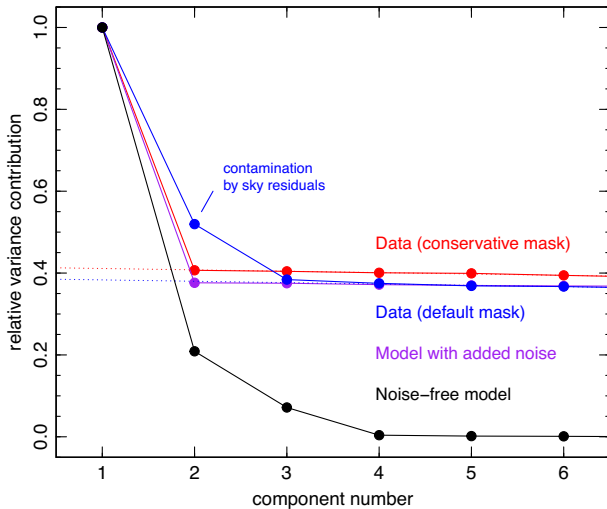


Figure 9. Total variance associated with each principal component, relative to that of the leading component. The noise floor of the NGC 5128 data (where the variance contributions flatten) is a factor of ~ 2 higher than the expected level of the second component. The dotted lines here show a linear fit to components 4–10 for the observations.

information carried by E_2 lies below the noise level of the present data.

4 DISCUSSION

The method outlined in this paper provides a general framework to extract the spectroscopic properties of the brightest stars in an unresolved stellar population. Applied to the red-optical spectrum of an old population, the fluctuation eigenspectra probe cool giants on the first and early asymptotic giant branch (AGB).⁸ However in principle, the method could be applied to younger populations, where thermally-pulsing AGB stars or red supergiants might drive the fluctuations in the red or infrared. If applied in the blue (notwithstanding the difficulty of AO observations at such wavelengths), hot horizontal-branch stars might contribute significant fluctuations in the Balmer lines, if present.

Compared to other methods proposed for the Poisson fluctuation regime (van Dokkum & Conroy 2014; Mitzkus et al. 2018), the PCA approach has the advantage of capturing more than just a single spectral component, and hence reconstructing the stellar content in greater detail, if data quality allows. From a practical perspective, it is also advantageous in bypassing any dependence on broadband throughput and wavelength response. An inherent disadvantage of the method, however, is that the eigenspectra, and their ordering, are not guaranteed to be stable as the configuration of the data change. For example, extending the wavelength range might bring in a new spectral feature such that a different type of star begins to contribute to, or even dominate, the variance. Interpreting the observations thus requires new predictions to be generated for each new observational configuration.

In this proof-of-concept paper, I have compared the observational results only to the predictions from a single fiducial stellar population model, without exploring the dependence on age and metallicity, composite populations, or the specific choice of isochrones etc;

⁸Resolved spectroscopy of such stars beyond the local group, and in NGC 5128 in particular, has been advanced as a key science case for instruments on 30m-class telescopes (Gonzalez & Battaglia 2018).

all such questions are deferred to future work. For NGC 5128 in particular, it will be possible to test whether the difference in E_1 around the Ca II triplet can be resolved by including a broad metallicity distribution, as indicated by resolved photometry (Soria et al. 1996; Harris, Harris & Poole 1999), or AGB stars from a younger component (Rejkuba et al. 2021). More realistic modelling should also take into account the wavelength-dependent PSF, which will introduce a systematic variation in the fluctuation strength across the spectral range.

The ultimate application of the method would be to determine the properties of cool giant stars not in a system like NGC 5128, where they probably closely resemble those in the Milky Way, but in the centres of very massive elliptical galaxies. The stellar populations in such regions have super-solar metallicities and enhanced $[\alpha/\text{Fe}]$ and $[\text{Na}/\text{Fe}]$ ratios (Conroy et al. 2014), as well as other spectral signatures that seem to indicate an IMF enriched in low-mass stars (Smith 2020). The integrated spectra of elliptical galaxies are analysed through sophisticated models built usually from empirical libraries of Milky Way stars, and augmented with theoretical stellar atmosphere calculations to account for non-solar abundance patterns (e.g. Conroy & van Dokkum 2012; La Barbera et al. 2017). Such calculations are especially complex for cool atmospheres with strong molecular features, and the evolutionary tracks for AGB stars also remain challenging to compute (Marigo et al. 2013). Empirically isolating the spectral contribution of giants would thus provide a novel test of some of the uncertain ingredients in stellar population models for elliptical galaxies, which cannot be performed directly in the Milky Way.

Perhaps counterintuitively, the spectral fluctuation method is not limited to the situation seen in the NGC 5128 observations, where the number of luminous giants per spectral sample is of order unity. At a given distance, working at higher surface brightness reduces the relative amplitude of the fluctuations, but suppresses the photon-noise contribution to the error by the same factor, so the signal-to-noise ratio is independent of surface brightness in this regime. In practice, for the very small spatial pixels used in the narrow-field mode, the read noise is generally significant or dominant (as it is for the NGC 5128 data), so there is substantial scope to improve precision by working in brighter regions.⁹ The choice of field (and consequent low surface brightness) in Section 3 was driven by the need for a bright AO reference star in early observations with the MUSE narrow-field mode. However, after recent upgrades to the wavefront sensor system,¹⁰ the low-order correction can be now achieved with point sources as faint as $J \approx 19$ or even using the core of the target galaxy itself. With these advances, retrieving at least the first fluctuation eigenspectrum in the cores of ellipticals in the Virgo or Fornax clusters might be feasible with deep MUSE observations.

5 SUMMARY

I have described a new approach to characterizing the spatial variation in spectra caused by Poisson fluctuations in the population of bright giants sampled by each resolution element in IFU observations of galaxies. The key results are as follows:

⁹For a population of identical stars each contributing f counts, with a mean of n stars per element, the intrinsic fluctuations are simply $\sqrt{n}f$, and the error on each measurement is $\sqrt{nf + \sigma^2}$ where σ is the read noise. Hence the signal-to-noise ratio increases with surface brightness at low n and then flattens when the photon-noise regime is reached.

¹⁰<https://www.eso.org/sci/publications/announcements/sciann17429.html>

(i) For old stellar population models, >99 per cent of the variation in the red-optical region can be characterized by a linear combination of just three eigenspectra that can be extracted through a standard principal components analysis.

(ii) The first two principal components are driven by the variation in the number of cool giant stars present, and their distribution along the isochrone, in each spatial sample.

(iii) For an old stellar population, the first fluctuation eigenspectrum, E_1 , describes the coolest (late-M) giants, while the second, E_2 , can be used to reconstruct the spectrum of warmer (early-M) stars.

(iv) Narrow-field mode observations of NGC 5128 with MUSE are able to retrieve E_1 with high precision; the observed features are a good, but not perfect, match to predictions from a 10 Gyr solar-metallicity model.

(v) Differences in E_1 around the Ca II triplet lines suggest the presence of warmer giant stars in NGC 5128 than in the fiducial model.

(vi) The higher order information encoded in E_2 lies below the noise level of the present NGC 5128 data.

Future application of this technique to the centres of the nearest massive ellipticals should be feasible, with recent upgrades to the MUSE AO system. Looking further ahead, future facilities such as the IFU mode of MAVIS¹¹ (Rigault et al. 2020) at the VLT, and HARMONI¹² (Thatte et al. 2021) at the ESO Extremely Large Telescope, offer great promise for further such studies in the optical and infrared respectively. Such observations would directly probe the properties of cool giant stars with element abundance patterns not found in the Milky Way, and hence test key ingredients of current spectral synthesis models for old stellar populations.

ACKNOWLEDGEMENTS

I am grateful for helpful comments and suggestions from the anonymous reviewer. This work was supported by the Science and Technology Facilities Council through the Durham Astronomy Consolidated Grant 2020–2023 (ST/T000244/1). This research is based on observations collected at the European Organisation for Astronomical Research in the Southern Hemisphere under ESO programme 0104.B-0174(A).

DATA AVAILABILITY

The observational data used in this paper are publicly available in the ESO archive (archive.eso.org). Other products can be provided upon reasonable request to the author.

¹¹Multi-conjugate adaptive optics Assisted Visible Imager and Spectrograph.

¹²High Angular Resolution Monolithic Optical and Near-infrared Integral field spectrograph.

REFERENCES

- Bacon R. et al., in I. S. McLean, S. K. Ramsay, Takami H., eds, 2010, in Proc. SPIE Conf. Ser. Vol. 7735, Ground-Based and Airborne Instrumentation for Astronomy III. SPIE, Bellingham, p. 773508
- Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, *MNRAS*, 427, 127
- Conroy C., 2013, *ARA&A*, 51, 393
- Conroy C., van Dokkum P., 2012, *ApJ*, 747, 69
- Conroy C., Graves G. J., van Dokkum P. G., 2014, *ApJ*, 780, 33
- Gonzalez O. A., Battaglia G., 2018, Proc. IAUS347, Early Science with ELTs, preprint ([arXiv:1810.04422](https://arxiv.org/abs/1810.04422))
- Harris G. L. H., Harris W. E., Poole G. B., 1999, *AJ*, 117, 855
- Harris W. E., Harris G. L. H., Barmby P., McLaughlin D. E., Forbes D. A., 2006, *AJ*, 132, 2187
- Harris G. L. H., Rejkuba M., Harris W. E., 2010, *PASA*, 27, 457
- Jarrett T. H., Chester T., Cutri R., Schneider S. E., Huchra J. P., 2003, *AJ*, 125, 525
- La Barbera F. et al., 2019, *MNRAS*, 489, 4090
- La Barbera F., Vazdekis A., Ferreras I., Pasquali A., Allende Prieto C., Röck B., Aguado D. S., Peletier R. F., 2017, *MNRAS*, 464, 3597
- Liu M. C., Graham J. R., Charlot S., 2002, *ApJ*, 564, 216
- Marigo P., Bressan A., Nanni A., Girardi L., Pumo M. L., 2013, *MNRAS*, 434, 488
- Mitzkus M., Walcher C. J., Roth M. M., Coelho P. R. T., Cioni M.-R. L., Raimondo G., Rejkuba M., 2018, *MNRAS*, 480, 629
- Rejkuba M., Harris W. E., Greggio L., Crnojević D., Harris G. L. H., 2021, *A&A*, in press, preprint ([arXiv:2111.02945](https://arxiv.org/abs/2111.02945))
- Rigault F. et al., 2020, in Proc. SPIE Conf. Ser. Vol. 11447, Ground-based and Airborne Instrumentation for Astronomy VIII. SPIE, Bellingham, p. 114471
- Sánchez-Blázquez P. et al., 2006, *MNRAS*, 371, 703
- Smith R. J., 2020, *ARA&A*, 58, 577
- Soria R. et al., 1996, *ApJ*, 465, 79
- Ströbele S. et al., 2012, in Ellerbroek B. L., Marchetti E., Véran J.-P., eds, Proc. SPIE Conf. Ser. Vol. 8447, Adaptive Optics Systems III. SPIE, Bellingham, p. 844737
- Thatte N. et al., 2021, *The Messenger*, 182, 7
- Tonry J., Schneider D. P., 1988, *AJ*, 96, 807
- Tonry J. L., Dressler A., Blakeslee J. P., Ajhar E. A., Fletcher A. B., Luppino G. A., Metzger M. R., Moore C. B., 2001, *ApJ*, 546, 681
- van Belle G. T. et al., 1999, *AJ*, 117, 521
- van Dokkum P. G., Conroy C., 2014, *ApJ*, 797, 56
- van Dokkum P., Conroy C., Villaume A., Brodie J., Romanowsky A. J., 2017, *ApJ*, 841, 68
- Vazdekis A., Cerviño M., Montes M., Martín-Navarro I., Beasley M. A., 2020, *MNRAS*, 493, 5131
- Villaume A., Conroy C., Johnson B., Rayner J., Mann A. W., van Dokkum P., 2017, *ApJS*, 230, 23
- Wilkinson A., Sharples R. M., Fosbury R. A. E., Wallace P. T., 1986, *MNRAS*, 218, 297

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