

Resolving galaxies in time and space: Applying STARLIGHT to CALIFA data cubes

11 June 2012

ABSTRACT

Fossil methods based on spectral synthesis techniques have matured during the past decade, and their application to integrated galaxy spectra fostered substantial advances on the understanding of galaxies and their evolution. Yet, because of the lack of spatial resolution, these studies are limited to a global view, providing no information about the internal physics of galaxies. Motivated by the CALIFA survey, which is gathering Integral Field Spectroscopy (IFS) over the full optical extent of 600 galaxies, we have developed an end-to-end pipeline which: (i) partitions the observed data cube into Voronoi zones in order to, when necessary and taking due account of correlated errors, increase the signal-to-noise ratio, (ii) extracts rest-framed and resampled spectra, including propagated errors and bad-pixel flags, (iii) feeds the spectra into the STARLIGHT spectral synthesis code, (iv) packs the results for all galaxy zones into a single fits file, (v) performs a series of post-processing operations, including zone-to-pixel image reconstruction and unpacking the spectral and stellar population properties derived by STARLIGHT into multi-dimensional time, metallicity, and spatial coordinates.

This paper provides an illustrated description of this whole pipeline—dubbed Python CALifa Starlight Synthesis Organizer (PYCASSO)—and its many products. Using data for the nearby spiral NGC ??? as a show case, we go through each of the steps involved, and present a series of ways of visualizing and analyzing this manifold. These include 2D maps of properties such as the velocity and velocity dispersion fields, stellar extinction, mean ages and metallicities, surface densities of light, mass, star formation rates on different time scales and normalized in different ways, plus 1D averages in the temporal and spatial dimensions, leading to evolutionary curves and radial profiles of physical properties. Projections of the stellar light and mass growth onto radius-age diagrams are introduced as a means of visualizing galaxy evolution in time and space simultaneously, something which can also be achieved in 3D with snapshot cuts through the (x, y, t) cubes. The results provide a vivid illustration of the richness of the combination of IFS data with spectral synthesis, of the insights on galaxy physics provided by the variety of diagnostics and semi-empirical constraints obtained, as well as a glimpse of what is to come from CALIFA and future IFS surveys.

Key words: galaxies: evolution – galaxies: stellar populations – galaxies: stellar content – galaxies: metallicity– techniques: spectroscopic

1 INTRO... JUST A DEPOSIT!


This paper presents an end-to-end methodology to process CALIFA data cubes through the STARLIGHT spectral synthesis code, leading to spatially resolved SFHs and a detailed mapping of series of other physical and observed properties.

, including emission lines.... STARLIGHT spectral synthesis code, leading to spatially resolved star formation histories as retrieved from the fossil record encoded in their spectra ... shifting the focus from mapping galaxies as a whole to within galaxies

The purpose of this paper is twofold: (a) Firstly, to present a detailed account of all processing steps employed from the reduced CALIFA $xy\lambda$ datacube to the multidimensional products derived from the STARLIGHT analysis. (b) Secondly, to illustrate the richness of the multidimensional products of this analysis.

The technical and methodological aspects discussed here are The main motivation for this paper is to use it

Barring updates in the modeling ingredients and the reduction pipeline,
other publications

Figure 1.  RGB things INCLUDING S/N image and S/N radial profile, with pixels as points, and average as a line... As in other images throughout this paper, the circles mark $R = 1, 2$ and $3 R_e$

Both these items are in fact Both a and b are of general interest

The main purpose of this paper is to present a detailed account of all steps employed in the processing of CALIFA $xy\lambda$ datacubes to the multidimensional products derived from the STARLIGHT analysis.

with CALIFA. This is important because we will use it a lot in future papers. (b) Secondly, it serves as a general guide to the analysis of IFS data with STARLIGHT.

... we work several ways to present and analyse the data, some are more convenient/useful than others ... the main steps involve spatial binning to ensure spectral quality, treatment of errors and flags, spectra ... spatial masks, starlight input and run, putting together the results in a coherent multi-layered structure containing spectra and physical props... .. construction of diagnostic of fit quality and spatially resolved stellar population tracers ... visualization. ...and on some general considerations on issues related to the application of STARLIGHT to data cubes.

2 DATA

few technical words about CALIFA . REF to main paper and others

3 SPECTRAL EXTRACTION: FROM DATACUBES TO STARLIGHT INPUT

In order for STARLIGHT to do a precise job, it is needed to pre-process each spectrum in a particular way. Moreover, STARLIGHT was designed as a 1D spectrum fitting tool, and therefore it only stores the spectral information related to a spectrum, losing all the spatial information. Thus, we developed a 3D datacube pipeline to process all spectra, store the 2D spatial information and save a number of quality control data to evaluate the goodness of the datacube and potential errors that might affect the spectral synthesis fitting.

3.1 Spatial Masks

Foreground stars and artefacts can affect both the fitting and the subsequent analysis of the properties of the galaxy. Thus, a previous step of the pre-processing consist on creating spatial mask to avoid this non physical related objects. As an starting point, we used SDSS r-band mask obtained by (SEXTRACTOR or whatever: Nadine, Jakob, need some info here). These masks subtend a size of $\sim 13' \times 10'$. We first have to cut this images to the size of the CALIFA Field of View (FoV). We use the WCS information available in the headers to reproject the SDSS masks to the scale, orientation and pixel size of the CALIFA datacubes.

After checking carefully a hundred of cases, we decided to broaden the SDSS mask applying a gaussian filter with σ

$= 0.15$, since a few unmasked pixels remained in the borders of each masked object, probably due to a combination of the slightly better PSF SDSS images and the accuracy of the CALIFA WCS information.

Once we have each reprojected and broadened mask, we perform a one-by-one visual inspection correction of the SDSS masks. We looked at each masked object spectra to check if it did not belong to the galaxy (i.e. foreground star or artefact). In a number of cases, we found that some bright HII regions were banned, so we unmasked this regions. Some galaxies with characteristic dusty regions were also unmasked. In other cases, the masked regions are too big, losing spatial information, so we unmasked completely these particular cases and build new masks for them. In a few cases, very dim point-like sources can be distinguished in the SDSS image but, due to the SN, we do not find counterpart in the CALIFA spectra. In these cases we decided to unmask these points since they had no imprint in the local galaxy spectra.

3.2 3D2Starlight

As previously stated, STARLIGHT needs some pre-processing of the original spectra. Moreover, an accurate handling of the errors and flags can produce better quality fit in terms of adev of χ^2 .

This first pre-processing step is packed in a fully automatic pipeline package, named QBICK. Although was built to work for CALIFA data, it is sufficiently general to handle any FITS format to process the datacube for a 3D spectral fitting with STARLIGHT.

The CALIFA Pipeline 1.3 (Husemann, in preparation) produces carefully derived errors and flags. Nevertheless, there are a number of cases, not included in the flags, that can affect the quality of the spectral fitting, being these issues mainly related with the errors. We identified four cases and assigned a ID label to each one:

- *Pipeline flags*: Original flags of the pipeline
- *Sky lines flags*: Flags assigned to certain sky lines bands
- *Big Error flags*: Flags assigned to pixels that follow the condition $\epsilon_\lambda > \alpha_b |F_\lambda|$
- *Small Error flags*: Flags assigned to pixels that follow the condition $\epsilon_{Z_\lambda} < \alpha_s \tilde{\epsilon}_Z$

where ϵ_λ is the error of one particular λ associated to the flux in that wavelength (F_λ) and $\tilde{\epsilon}_Z$ is the median value of the errors for one particular zone (see below). All the variables along this pre-processing step can be adjusted at convenience. For this particular case, we set $\alpha_b = 5$ and $\alpha_a = 0.1$.

As mentioned above, we have flagged a number of wavelength windows containing the strongest sky lines seen in the CALIFA data. These flags include the Hg I 4358.33 Å, Hg I 5460.74 Å, O I 5577.34 Å, Na D (around 5890 Å), O I - OH (6300 Å), O I 6363.78 Å, and the B-Band atmospheric absorption band.

Some values have to be given to the pipeline, either as header FITS entries or in a external file in order to perform some corrections of changes in the output files names. Galactic extinction is applied if needed, provided the E(B-V) or A_V is given. For CALIFA data this is not a necessary step, since the data cubes are already corrected for this effect.

The recessional velocity is used to perform a K-correction to the flux and errors and then, shift all the spectra to their rest-frame.

The next step consist in applying Signal-to-Noise (SN) and spatial masks. We used a small (rest wavelength) window between 5590-5680 Å to estimate certain quality control measurements. This window was chosen because its almost absence of sky lines (for the CALIFA redshift range) or spectral lines. We estimate the signal, noise (detrended), and derived SN, rejecting flagged pixels in these calculations. We apply the spatial mask described in subsection 3.1 and afterwards apply a gaussian or median filter to smooth the image and estimate the iso-contour for a SN clipping value. This threshold value was set to 5, except for a few objects where we reduced this value because of the low surface-brightness. Once we get this iso-contour, we mask all the pixels below this value (outside the iso-contour). After this step, we have three 2D images with the signal, noise and SN, with all pixels below a defined SN threshold and foreground stars and other artefacts not related the the object masked.

All CALIFA observations center the object in the FoV. We divide the FoV into 9 quadrants (each axis divided in three equal parts) and define the center of the object, nucleus hereafter, as the maximum value in the center quadrant of the signal masked image. Also in this image, we estimate the the Half Light Radius (HLR) of the galaxy.

The following step consist on applying a segmentation structure (spatial binning) to the data if a certain minimum SN is required, grouping pixels into zones. A external segmentation map FITS file can be injected into the pipeline, where all the pixels belonging to the same zone shared a common ID integer label, which is unique for each zone. In this work, we have made use of the Voronoi segmentation code by Cappellari & Copin (2003), setting the minimum SN to 20 and the first zone the nucleus of the object. Most of the zones inside 1 HLR are just individual pixels, while the pixels of the outer parts are grouped zones containing several dozens of pixels, depending on the SN of the object. Since the CALIFA datacubes suffer from correlated noise, we apply our own correction method to the error spectra when performing spatial binning (see section).

Once the spatial binning map is obtained, we start summing the flux, error, and flag spectra for each spatial zone. The three output zone spectra are resampled to a uniformly wavelength step $d\lambda$ (set to 2 Å). After and before the resampling process the signal, noise and SN are estimated in th output spectra. For the flux (and error) array, we sum following the equation:

$$F_{Z\lambda} = \frac{N}{\Phi_\lambda} \sum_{i=1}^N F_{\lambda i} \phi_{\lambda i} \quad (1)$$

$F_{Z\lambda}$ is the total flux for a particular Z zone in wavelength λ , N is the total number of spectra contained in that zone, $F_{\lambda i}$ is the flux at λ for spectra i , $\phi_{\lambda i}$ is 1 if $F_{\lambda i}$ is not flagged and 0 otherwise, and Φ_λ is the total number of not flagged $F_{\lambda i}$. If all $F_{\lambda i}$ are flagged, Φ_λ is set to N and $\phi_{\lambda i}$ to 1. Errors are also accommodated to this scheme, but added in quadratic form.

For each zone Z, we flagged a particular λ if $\frac{N}{\Theta} > \Phi_\lambda$, where Φ_λ is the total number of “good” values (not flagged).

We have followed a “half-full bottle” (or “optimistic”) rule, setting $\Theta = 2$.

Finally, once all the spectra zones are calculated, the total integrated spectra is obtained summing all zones, following the same criteria described above for the flag rule.

The zone spectra are stored in ASCII files (and FITS files if needed), with its flux, error and flag column. All the quality control images produced along the pipeline are packed in a single FITS file along with all the parameters used in the different steps (SN clipping, segmentation SN, ...).


The pre-processing pipeline can take advantage of multi-core systems to speed-up the process.

...zone-to-yx array puzzle-solver tensor *zoneMask_{xy}*

3.3 Spatial binning and error covariance

The input data cube to PyCASSO has a spatial sampling of 1 arcsec/spaxel. Because of the three-fold dithering pattern used for the CALIFA survey observations (Sánchez et al. 2012), each spaxel has a varying contribution from a number of fibers; typically three fibers contribute partially with flux to a spaxel. Thus because one fiber can contribute to more than one spaxel, the noise of the spectrum in each 1 arcsec spaxel is partially correlated with the noise of the spectrum in adjacent spaxels. This implies that when spectral fluxes from adjacent spaxels are added together the noise in the resulting flux spectrum has to be calculated taking spatial covariances into account. As an illustrative extreme example, assume that we add spectra (with their associated noise spectrum) which are completely correlated, e.g. if we add a spectrum to itself N times, the signal keeps increasing linearly with N , and because it is always the same spectrum the actual noise also scales linearly, thus if we propagate errors quadratically we would compute the error in the N times summed spectrum as $\epsilon_{sum} = \sqrt{\sum \epsilon_i^2} = \epsilon_i \sqrt{N}$, and we would compute a signal-to-noise of the summed spectrum as \sqrt{N} times the original signal-to-noise, while the actual value of the signal-to-noise of the summed spectrum would not have changed in fact. Thus it is necessary to take into account a correction in the computation of the resulting signal-to-noise when neighbouring spatially correlated spectra are added together.

In our processing of the CALIFA datacubes we perform a Voronoi spatial binning (Cappellari & Copin 2003) in order for the spectra to have a minimum S/N of 20 in the continuum at 5635 Å (cf. section 3.2), so that the stellar population synthesis fits provide a sensible result. The computation of the signal-to-noise in the Cappellari & Copin algorithm assumes that uncorrelated spectra are summed, so we have computed an empirical correction algorithm that takes into account the global behaviour of the spatial correlation of the signal in CALIFA spectra. Notice that the three-fold dithering pattern is not completely uniform, and thus the amount of correlation is not exactly uniform across the face of the fiber bundle, so our correction is a statistical one, that takes into account globally the number of spaxels summed (see figure 2 below). Once the correction function is computed, it is implemented in the signal-to-noise computation within the Voronoi binning scheme so that Voronoi zones are computed with the correct signal-to-noise values

Figure 2.  β ...

and output binned spectra have the correct errors associated. The way we proceed is as follows: first we compute the Voronoi zones with the original signal-to-noise scheme which assumes uncorrelated input. We then measure the error in the Voronoi summed spectrum and compare with the error computed in the Voronoi algorithm. The ratio of these two numbers, which we call β , is represented versus the number of pixels in each Voronoi zone in figure ??, where β values for the individual zones are plotted as small grey dots, the large red dots represent the median value computed at each zone area, while the shaded region represents the ± 1 standard deviation of the mean. The red line represents the fit of the following function

$$\beta = \left(\frac{15 \text{area}}{15 + \text{area} - 1} \right)^{1/2} \quad (2)$$

The implementation of this correction factor within the Voronoi zoning algorithm is then given by

$$\epsilon_{sum} = \sqrt{\sum_i (\beta \epsilon_i)^2}$$

—spatial binning also helps alleviating the stochastic business no region has mass smaller than??? Mo...

4 STARLIGHT RUNS AND PYCASSO

STARLIGHT is a spectral synthesis method which combines the spectra from a base, in order to match an observed spectrum O_λ . The search for an optimal model M_λ and the coefficients of the population vector \vec{x} associated with the N_* base elements also allows for reddening (parametrized by the V-band extinction, A_V), a velocity off-set v_* and the stellar velocity dispersion σ_* . The code was first introduced in Cid Fernandes et al (2004) and substantially optimized by Cid Fernandes et al (2005). The code, didactic introductions and a user manual are available at www.starlight.ufsc.br.

In this paper and associated ones, a new version of STARLIGHT was used. The major changes are related to the possibility of including observational constraints related to the dust reprocessed (far-IR) luminosity, photometric measurements, and recombination emission line luminosities (tracers of the ionizing flux). Besides these extensions, STARLIGHT now uses a smarter overall normalization scheme, a friendlier output of mass and luminosity related information, more flexible mask and clipping options, as well as a number of formatting related bug fixes.

The STARLIGHT results reported in the next section rely on a spectral base of Simple Stellar Populations (SSPs) comprising 4 metallicities, $Z = \dots, 1, 1, 5Z_\odot$, and 39 ages between $t = 10^6$ and 1.4×10^{10} yr. These SSPs combine the models of González Delgado et al (2005) for ages $t < 63$ Myr and those of Vazdekis et al (2010), as updated by Falcón-Barroso et al (2011). They are based on the Salpeter IMF and the evolutionary tracks by ???. Details on this set of SSPs as well as comparison of results obtained other bases

are discussed in detail in González Delgado et al (in preparation), to which the reader is also referred for examples of the spectral fits.

The virtues and caveats of full spectral fits are discussed in a number of articles (REFS...McArthur et al 200?...). A general consensus in the field is that uncertainties in the results for individual objects average out for large statistical samples (Panter ???). With CALIFA, each galaxy becomes a statistical sample per se! Hence, even if the results for single spaxels are not iron-clad, the overall trends should be robust. It is also worth pointing out that despite the diversity of spectral synthesis methods currently employed, substantial changes on the results are more likely to come from revisions in the input data (as happened, for instance, with the recalibration of SDSS spectra between data releases 5 and 6) and from updates in the base models (see González Delgado & Cid Fernandes 2010), the single most important ingredient in any spectral synthesis analysis.

As documented in the user manual, STARLIGHT outputs a large array of quantities. In broad terms, these can be categorized as

(i) Input data: File names, configuration options, and other information provided either by the user explicitly or derived from the user-provided spectrum. Base-related data are also reported, including ages and metallicities of each component, as well as the corresponding light-to-mass ratios (at the chosen normalization wavelength λ_N) and a returned-mass correction factor.

(ii) Fit results: Figures of merit (χ^2 , $\overline{\Delta}$), kinematic parameters (v_* , σ_*), the V-band extinction (A_V), total stellar masses and luminosities, and population vectors, expressed in terms of light (\vec{x}) and mass ($\vec{\mu}$) fractions.

(iii) Spectral data: The input (O_λ), best model fit (M_λ), and weight spectra (w_λ , which reflects the adopted error spectrum as well as masked, flagged and clipped pixels).

All these data (and more) are stored in a plain ascii file, one for each fitted spectrum. The population vectors can be post-processed in a number of ways to obtain star-formation rates on different time-scales, to compute the mass and metallicity growth as a function of time, to condense the results into first moments like the (light or mass weighted) mean t and Z , etc. This scheme works fine for studies of independent objects, like SDSS galaxies (e.g., Asari et al 2007) or globular clusters (Gonzalez Delgado & Cid Fernandes 2010).

IFS data, however, require a more structured level of organization, as it is cumbersome and inefficient to handle so many files separately. To handle data cubes we developed PycASSO, the Python Califa Starlight Synthesis Organizer. PycASSO comprises three main parts:

(i) A *writer* module, which packs the output of all STARLIGHT fits of the individual zones of a same galaxy into a single FITS file, which also contains all information generated during the pre-STARLIGHT processing steps (Section 3) as well as information propagated from the original reduced data cube.

(ii) A *reader* module, which reads this "super-FITS" file and structures all the data in an easy to access and manipulate format.

(iii) A *post-processing* module which performs a series of




Figure 3. Maps of spectral-fit quality indicators. (a) Mean relative deviation $\bar{\Delta}$. (b) χ^2 per fitted flux. (c) Number of fluxes used in the fit. (d) Number of clipped fluxes.

common operations such as mapping any property from zones to pixels, resampling and smoothing the population vectors, computing growth functions, averaging in spatial, time or metallicity dimensions, etc. All these functionalities are conveniently wrapped as python packages, easily imported into the users own analysis code.

PycASSO will be the main work-horse behind a series of articles by our collaboration, so we dedicate the remainder of this paper to illustrating some of its main products. Despite the focus on CALIFA and PycASSO, the issues addressed below are of interest to STARLIGHT applications to data cubes in general.

5 RESULTS

NGC 5947 (CALIFA 938) was chosen as a show-case. This SBbc galaxy (the last on the survey’s mother sample) is $d = 83$ Mpc away, such that 1 arcsec (the sampling adopted for all CALIFA reduced data cubes) corresponds to 400 pc. It has ??? and ??? (??REFS??) **!⊙j⊙!** **SOME COMMENTS ABOUT K0938! CHECK LITERATURE** Its bar is ??pc long, has a ring here, arms there, a compact bulge.... Our Voronoi binning of the CALIFA 938 data cube produces $N_z = 1117$ zones. The corresponding $N_z + 1$ spectra (the +1 corresponding to the spatially integrated spectrum, which we also analyse for comparison) were fitted with STARLIGHT and the results feed into PycASSO.

5.1 Fit quality assessment

As with 1D spectra, the quality of the fits should be examined before proceeding to interpretation of the results. This subsection describes some standard methods to do this, the caveats involved and strategies to circumvent them.

Fig. 3a shows the map of the fit quality indicator

$$\bar{\Delta} = \frac{1}{N_\lambda^{eff}} \sum_\lambda \frac{|O_\lambda - M_\lambda|}{M_\lambda} \quad (3)$$

where the sum is carried over the N_λ^{eff} wavelengths actually used in the fit, ie, discarding masked, flagged and clipped pixels.¹ For this galaxy $\bar{\Delta}$ spans the range between 1.1 and 9.1%, with a median value of 4.3%. $\bar{\Delta}$ increases towards the outer regions, where the S/N is smaller (Fig. 1). The χ^2 map (panel b) shows the opposite behavior, since the errors also increase outwards (see Section 3). Because of its non-explicit reliance on the (often hard to compute, if not altogether unavailable) ϵ_λ spectrum, as well as because of its easily grasped meaning, $\bar{\Delta}$ is a more useful figure of merit to assess fit quality.

One can use such maps to define spatial masks over which the spectral fits satisfy some quality threshold. Inspection of the highest $\bar{\Delta}$ spectra often reveal artifacts, like imperfect masking of foreground sources, slightly misplaced bad pixel flags, or O_λ values which should have been clipped by STARLIGHT but were not because of large ϵ_λ . We emphasize that such bugs are *very rare* in CALIFA. The median $\bar{\Delta}$ for the 110k zone spectra analysed so far is just 4% (corresponding to an equivalent S/N of 25), and in less than 3% of the cases $\bar{\Delta}$ exceeds 10%. This high rate of success is mainly due to the carefully derived errors and bad pixel flags in the reduction pipeline and further refined in our pre-processing steps (Section 3).

Experiments with the whole data set suggest that a reasonable quality-control limit should be in the $\bar{\Delta} = 8$ to 10% range. Only 7 zones of our example galaxy fail $\bar{\Delta} < 8\%$, so we will not apply any cut here. For the benefit of STARLIGHT users, it is nevertheless worth making some further remarks on quality control.

!⊙j⊙! UP TO 2 Re OK...AFTER, FEW and NOISE ...

5.1.1 Subtleties and caveats on quality control

In full spectral fitting methods, seemingly trivial operations like imposing a fit-quality threshold often hide subtleties which can easily go unnoticed, particularly when processing tons of data. The paragraphs below (focused on STARLIGHT but extensible to other codes) review some of these.

First, one should distinguish cases where $\bar{\Delta}$ is large because of unreliable data from those where large residuals arise because of the user’s failure to properly inform, through masks or bad pixel flags, spectral regions which should be ignored in the fit. For instance, one can have an excellent spectrum of an HII region with several strong emission lines besides those included in a generic emission line mask feed into STARLIGHT, causing an artificially large $\bar{\Delta}$. Large velocity offsets can have a similar effect, as emission lines get

¹ Notice that this is nearly, but not exactly identical to what is called "adev" in STARLIGHT’s output, the difference being the use of M_λ instead of O_λ in the denominator. This minor modification avoids faulty (but not flagged) pixels with very low flux affecting the statistics. This very rarely occurs in CALIFA data, and when it does the pixels invariably have large errors, and so are irrelevant for the fit, but can still affect the value of $\bar{\Delta}$.

shifted out of the fixed m_λ windows. Similarly, sky residuals and other artifacts missed out by the b_λ flags lead to large $\bar{\Delta}$ even when the overall spectrum is good. In short, not everything which fails a blind quality control is actually bad.

The clipping options implemented in STARLIGHT help spotting pixels which are too hard to fit and thus probably represent non-stellar or spurious features. Our fits use the "NSIGMA" clipping method and a conservative 4σ threshold, meaning that we only clip pixels when $|O_\lambda - M_\lambda| > 4\epsilon_\lambda$. Fig. 3c maps the number of clipped points in CALIFA 938. One sees a peak in the central pixels, where the S/N is so high that our 4σ clipping does not forgive even small $|O_\lambda - M_\lambda|$ deviations (most often associated with problems with the base models rather than with the data). Elsewhere, very few pixels were clipped. The point to highlight here is clipping only worked because we have a reliable ϵ_λ . Other clipping methods can be used when this is not the case, but in all cases the user should always check carefully the output, as it may easily happen that too many points are clipped. In case of doubt, or if ϵ_λ , b_λ and m_λ are all fully trustworthy, do not clip.

Custom-made spectral masks also help. For instance, since Mateus et al (2006), the STARLIGHT analysis of SDSS spectra employs individual spectral masks constructed by searching for emission lines in the $R_\lambda = O_\lambda - M_\lambda$ residual spectrum obtained from a first spectral fit, and taking into consideration the local noise level. O_λ is then refitted with this tailored m_λ , circumventing some of the issues above. This procedure has not yet been implemented in our CALIFA pipeline, which uses the same general emission line mask used by Cid Fernandes et al (2005) for all spectra.

Bad pixel flags, errors and spectral masks are of course secondary actors in any spectral synthesis analysis, but these general remarks illustrate that it pays off to dedicate them due attention.

5.2 2D maps: Stellar light, mass and extinction

One of the main products of spectral synthesis is to convert light to mass. Fig. 4 illustrates the results for CALIFA 938. Its top-left panel shows the surface density of the luminosity at $\lambda_N = 5635 \text{ \AA}$, while the top-right panel show the derived extinction map. The dust corrected image is show in the bottom left, while the mass surface density is in the bottom right.

The effects of the Voronoi zoning are present in all panels, but are much more salient in the A_V map. This is because the light and mass images were "dezonified" by scaling the value of each xy pixel by its fractional contribution to the total flux in a zone (z). For instance, for the mass surface density this operation reads

$$\mathcal{M}_{xy} = \frac{M_{xy}}{A_{xy}} = \frac{M_z}{A_z} \times w_{xyz} \quad (4)$$

where

$$w_{xyz} = \frac{I_{xy}}{\sum_{xy|z} I_{xy}} \quad (5)$$

and I is the mean observed flux in the 5590–5680 \AA region,




Figure 4. (a) Observed surface brightness at 5635 \AA (b) stellar extinction map (A_V), (c) dereddened $\mathcal{L}_{\lambda 5635}$ image, (d) stellar mass surface density.

the same used in the Voronoi zoning (Section 3). This operation, which produces somewhat smoother images than obtained with $w_{xyz} = 1$, was applied only to luminosity and mass related quantities. Intensive properties like A_V , mean ages, σ_* and etc. cannot (and are not) be dezonified.

The extinction map shows an $A_V \sim 0.8$ in the nucleus, surrounded by a ring of low extinction in the radial region coinciding with the bar extent, a rise to ~ 0.5 mag in a around the bar edge and a slow decline from there (see also Fig. 10). Small scale fluctuations at the edges (beyond $2R_e$ in this case) are probably a result of the lower S/N . As with other properties, the general behaviour of A_V in the outer regions may be better mapped with radial averaging.

The light and mass images have a similar structure, both showing the bulge, bar and disc components. The spiral arms are less prominent in terms of mass than in light because of the higher L/M of stars in the arms.

5.3 2D maps: Kinematics

STARLIGHT provides estimates v_* and σ_* (obtained by shifting and broadening the model spectrum in its search for a better spectral match). Fig. 5 shows v_* and σ_* fields, as well as a position velocity diagram and the radially averaged $\sigma_*(R)$ profile. The velocity dispersion was corrected by the difference in spectral resolution between the base and observed spectra. The σ_* field shows a mild increase towards the center in the inner R_e and a noisier behaviour outwards, while the v_* field suggests a projected rotation velocity of ~ 200 km/s.

Both the v_* and σ_* maps show small scale variations which can be traced to either noisy data or to young stellar populations, whose weak absorption lines prevent a robust determination of kinematical parameters. One also sees several regions where $\sigma_* = 0$, which means that the broadening fitted by STARLIGHT is smaller than expected due to the 139




Figure 5. Kinematical products of the spectral fitting.

km/s difference between the spectral resolution of the data and base.

While these kinematical maps are clearly useful, CALIFA can do better. Better spatial resolution on the outer regions can be obtained lowering the target S/N in the Voronoi binning, while better velocity resolution can be obtained using the V1200 data alone. Also, there are codes better suited than STARLIGHT to derive velocity fields (e.g. ??ASK JESUS??). In short, the kinematical information derived from our analysis will be superseded by these more kinematics-gearred studies. Eventually, one can envisage feeding the parameters derived from these more precise analysis back into the STARLIGHT fits (using its fixed-kinematics mode). In fact, this feedback may well turn out to improve our stellar population analysis, given the potential degeneracy between σ_* and Z_* (decreasing the former while increasing the latter have the same global effect of making metal lines deeper; Koleva et al 2009).

5.4 2D maps: Mean ages and metallicities

The crudest way to quantify the SFH of a system is to compress the age and metallicity distributions encoded in the population vectors to their first moments. For this purpose we will use the following definitions:

$$\langle \log t \rangle_L = \sum_{t,Z} x_{tZ} \log t \quad (6)$$




Figure 6. Luminosity (top) and mass (bottom) weighted mean age (left) and metallicity (right) of CALIFA 938.

$$\langle Z \rangle_L = \sum_{t,Z} x_{tZ} Z \quad (7)$$

where x_{tZ} is the fraction of light at λ_N due to the base population with age t and metallicity Z . The mass weighted versions of these indicators, $\langle \log t \rangle_M$ and $\langle Z \rangle_M$, are obtained replacing x_{tZ} by the mass fraction μ_{tZ} in these equations.

Fig. 6 shows the light and mass weighted mean age and metallicity maps. The $\langle \log t \rangle_L$ map shows a general increase towards the center, with a marked steepening at $R \sim 0.7R_e$ and a slight drop in the nucleus, which must therefore have had some recent star formation, perhaps fueled by gas flowing along the bar. These structures are hardly noticeable in the $\langle \log t \rangle_M$ map, which, as usual, spans a smaller dynamical range because of the large weight of old populations.

The metallicity maps is essentially saturated at the maximal Z of the base models ($1.5Z_\odot$) in the inner R_e , declining outwards, but with a large scatter. As is well known, Z is the harder property to derive in mixed stellar populations (??REF??). Besides these intrinsic difficulties, it is clear that our upper limit is affecting the analysis.

Fig. 7 plots age ($\langle \log t \rangle_L$) versus metallicity ($\langle Z \rangle_L$) versus extinction (A_V), the main properties involved in spectral synthesis. Points are color and symbols coded by their distance from the nucleus. !ojo! One sees thatold populations with large A_V are confined to the nucleus blablabla ROSA-HELP !ojo! The saturation in Z is even more evident in this plot. The data seems to be calling for larger Z 's than allowed by our base. BC03 models reach $2.5 Z_\odot$, but using a stellar library which is does not reach such high Z (see Koleva et al 2008 for a thorough discussion on this issue). ... This is a reminder on how much full spectral synthesis relies on SSP models...




Figure 7. Age versus metallicity versus extinction, color coding points by their distance to the nucleus. Green: $R/R_e < 0.5$. Red: $0.5 < R/R_e < 1$. Blue: $1 < R/R_e < 2$. Black: $R/R_e > 2$.




Figure 8. Maps of the percentage contribution of "Young" ($t < 0.14$ Gyr), "Intermediate age" (0.14–1.4 Gyr), and "Old" (> 1.4 Gyr) stars to the observed light at 5635 Å.

5.5 2D maps: x_Y , x_I , x_O

Population synthesis studies in the past found that a useful way to summarize the SFH is to condense the age distribution encoded in the population vector into age ranges. This strategy, comes from a time when the analysis was based on equivalent widths and/or colors (Bica 1988; Schmitt et al ???; Cid Fernandes et al 2001, 2003), but which was also successfully applied to full spectral fits (Gonzalez Delgado et al 2004; Cid Fernandes et al 2004).

Fig. 8 presents images of the light fraction due to Young, Intermediate and Old populations (x_Y , x_I and x_O), defined as those with $t \leq 0.14$, $0.14 < t \leq 1.4$, and $t > 1.4$ Gyr, respectively. As usual, the choice of borderlines is somewhat subjective, constrained only by underlying idea of grouping several base elements covering relatively wide age ranges.




Figure 9. Spatially resolved SFR surface density (eq. 8) on the last $t_{SF} = 14$ (top), 140 (middle) and 1400 (bottom) Myr, in units of different reference values. The left panels compares $\overline{SFR}_{xy}(t_{SF})$ to $\overline{SFR}_{xy}(t_{\infty})$, i.e., the life-time average in pixel xy , thus providing a local version of Scalo's b parameter (eq. 9). Middle panels compares the local SFR surface density to the life-time average over the whole galaxy (eq. 10). Panels on the right column compare $\overline{SFR}_{xy}(t_{SF})$ to the SFR per unit area of the galaxy as a whole over the same time scale (eq. 11), thus highlighting regions which are forming ...!ojo! FIND WORDS! All images are on logarithmic scale from log 1 to log 5, such that only $\overline{SFR}_{xy}(t_{SF})$ values above the corresponding reference value are visible.

For the base used in our STARLIGHT runs the Y, I and O bins contain 4×20 , 4×10 and 4×9 SSPs, respectively (where the $4 \times$ comes from the 4 metallicities).

!ojo! NEED HELP FROM IMAGINATIVE MINDS ...The plots show that the spiral arms stand out more clearly in the x_Y map, while the bar seems mixed population of I and O stars, cutting a predominantly old central bulge... BLABLABLA

5.6 2D maps: Star formation rates and IFS-based variations over Scalo's b parameter

The base used in our fits comprises SSPs, i.e., instantaneous bursts, so, despite the large number of ages considered, our SFHs are not continuous and hence not derivable. There are, nonetheless, ways of defining rates.

The simplest way to estimate an SFR is to accumulate all the stellar mass formed since a lookback time of t_{SF} :

$$\overline{SFR}_{xy}(t_{SF}) = \frac{1}{t_{SF}} \sum_{t < t_{SF}} \sum_Z \mathcal{M}_{tZxy} \quad (8)$$

where \mathcal{M}_{tZxy} is the mass (at pixel xy and per unit area) of stars formed at lookback time t and with metallicity Z . Eq. 8 gives the mean SFR surface density ($M_{\odot} yr^{-1} pc^{-2}$) since $t = \tau_{SF}$, and one can tune t_{SF} to reach different depths in the past. Of course, as t_{SF} increases this estimator becomes

increasingly useless, converging to the mass density divided by $t_\infty \equiv 14$ Gyr (the largest age in the base).

It is often more useful to consider SFRs in relation to some fiducial value, instead of absolute units. The classical example is Scalo's b parameter, which measures the SFR in the recent past (t_{SF}) with respect to its average over the whole lifetime of the system.² This is trivially obtained dividing $\overline{\mathcal{SFR}}_{xy}(t_{SF})$ by its asymptotic value

$$b_{xy}^{local}(t_{SF}) = \frac{\overline{\mathcal{SFR}}_{xy}(t_{SF})}{\overline{\mathcal{SFR}}_{xy}(t_\infty)} = \frac{t_\infty}{t_{SF}} \frac{\sum_{t < t_{SF}} \sum_Z \mathcal{M}_{tZxy}}{\mathcal{M}_{xy}} \quad (9)$$

The superscript "local" is to emphasize that the reference lifetime average SFR is that of same location xy . One should recall that while young stars in a pixel were probably born there, old ones may have migrated from different locations. Thus, despite the identical definitions, the physical meaning of b_{xy}^{local} is not exactly the same as for galaxies as a whole.

IFS data allow for other definitions of b . For instance, one might prefer to compare $\overline{\mathcal{SFR}}_{xy}(t_{SF})$ to the $\overline{\mathcal{SFR}}(t_\infty)$ of the galaxy as a whole, the bulge, the disc or any general region. This variation measures the "present" ($t < t_{SF}$) "here" (at xy) to the "past" ($t < t_\infty$) in a spatial region \mathcal{R} :

$$b_{xy}^{\mathcal{R}}(t_{SF}) = \frac{\overline{\mathcal{SFR}}_{xy}(t_{SF})}{\overline{\mathcal{SFR}}_{\mathcal{R}}(t_\infty)} \quad (10)$$

A formally similar but conceptually different definition is obtained using for the reference value in the denominator the SFR surface density within the *same time-scale* but evaluated at a *different region*:

$$c_{xy}^{\mathcal{R}}(t_{SF}) = \frac{\overline{\mathcal{SFR}}_{xy}(t_{SF})}{\overline{\mathcal{SFR}}_{\mathcal{R}}(t_{SF})} \quad (11)$$

ie, to compare not "now and then", but "now here and now there".

Equations 9–11 cover the different combinations of time and space afforded by the application of fossil methods to IFS data. Despite some degree of redundancy (the numerator is always the same), these definitions offer different views of the recent star formation in a galaxy. Fig. 9 shows maps of b_{xy}^{local} (left panels), $b_{xy}^{\mathcal{R}}$ (middle), and $c_{xy}^{\mathcal{R}}$ (right) for three values of t_{SF} : 14 (top), 140 (middle) and 1400 Myr (bottom).

!⊙j⊙! It is seen that ... a shows ... BLABLABLA Notice that the color scale is deliberately saturated to highlight regions where $\overline{\mathcal{SFR}}_{xy}(t_{SF})$ is larger than the chosen reference value Conclusion ... Depending what one wants to see, one definition is more convenient than the other ... highlight some feature in space and/or time.

5.7 1D spatial maps: radial profiles

The information contained in 2D maps like the ones shown above can sometimes be hard to process. Azimuthal averaging is a useful way to compress the 2D images into radial profiles. PyCASSO provides both circular and elliptical xy to

R conversion tensors. For simplicity, and also because our examples galaxy is nearly face on, the examples below are based on a circular mapping.

For quantities like \mathcal{M} , one can think of two types of radial averaging. The first is to add up all the stellar mass in xy pixels within a given R bin and then dividing by the bin area. The second is to average the surface density values for each pixels directly. The same applies to, for instance, $\langle \log t \rangle_L$. One can either add up the value of $\sum L_{t,xy} \log t$ for all pixels and then divide by the $\sum L_{xy}$, or else average the $\langle \log t \rangle_{L,xy}$ values in the R -bin directly. In principle the first method seems more reasonable, as it effectively collapses galaxy to 1D, but it cannot be applied to quantities which do not involve light or mass, like A_V , σ_* and $\overline{\Delta}$. We therefore chose the second type of radial averaging, noting that, the two methods almost always give nearly identical results. The exceptions occur in cases where the geometry is not favorable, which is not the case of CALIFA 938.

Fig. 10 shows ... !⊙j⊙! The individual points represent individual pixels (the horizontal stripes correspond to pixels in the same zone). The solid line is the mean value of the plotted quantity among the pixels in the same R bin, and the error bars map the corresponding dispersion. Several of the effects briefly discusses above while commenting the 2D images are clearly seen in these plots. BLABLABLA !⊙j⊙!

5.8 1D temporal maps: evolutionary curves

In all plots so far the time dimension was compressed either to its mean or to two or three intervals. The power of fossil methods, however is precisely is to open the time axis. The caveat is that it does so with the typical logarithmic resolution inherent to stellar evolution, but the studies of ??? Panter... Tojeiro ... Cid Fernandes et al 2007, Asari et al (2007) show that much can be learned about galaxy evolution despite this fundamental constraint. One way to obtain more robust produce SFHs is to smooth the over-detailed (39 ages) population vectors produced with our base. As discussed by Asari et al (2007), who introduced this method, age resolutions between 0.5 and 1 dex are reasonable. We chose to smooth x_t and μ_t with a gaussian of FWHM = 0.5 dex in $\log t$.

Fig. 11 plots a number of our synthesis products against time. To keep an at least partial representation of the spatial information, we plot evolutionary curves derived for different radial regions: The nucleus, defined as the central pixel (plotted in gray dashed lines), $R < 0.5R_e$ (orange dashed), $R < R_e$ (red solid), $R > R_e$ blue. Two other lines are also shown, one derived by adding the results obtained for all galaxy zones (black solid), and the other for the results obtained from the STARLIGHT analysis of the integrated galaxy spectrum (magenta dashed). The comparison of these last two is particularly interesting, as it illustrates the difference between properties inferred from the analysis of the whole and the more realistic properties inferred from the sum of the parts.

The top panels in Fig. 11 show the growth of luminosity (panel a) and mass (b). Among other things, this plot illustrates the difference between light and mass, in the sense that, regions inside and outside R_e do not have the same mass, despite having the same light. The difference in the

² A related index is the so called specific SFR, which differs from b just by the t_∞ factor in equation 9.

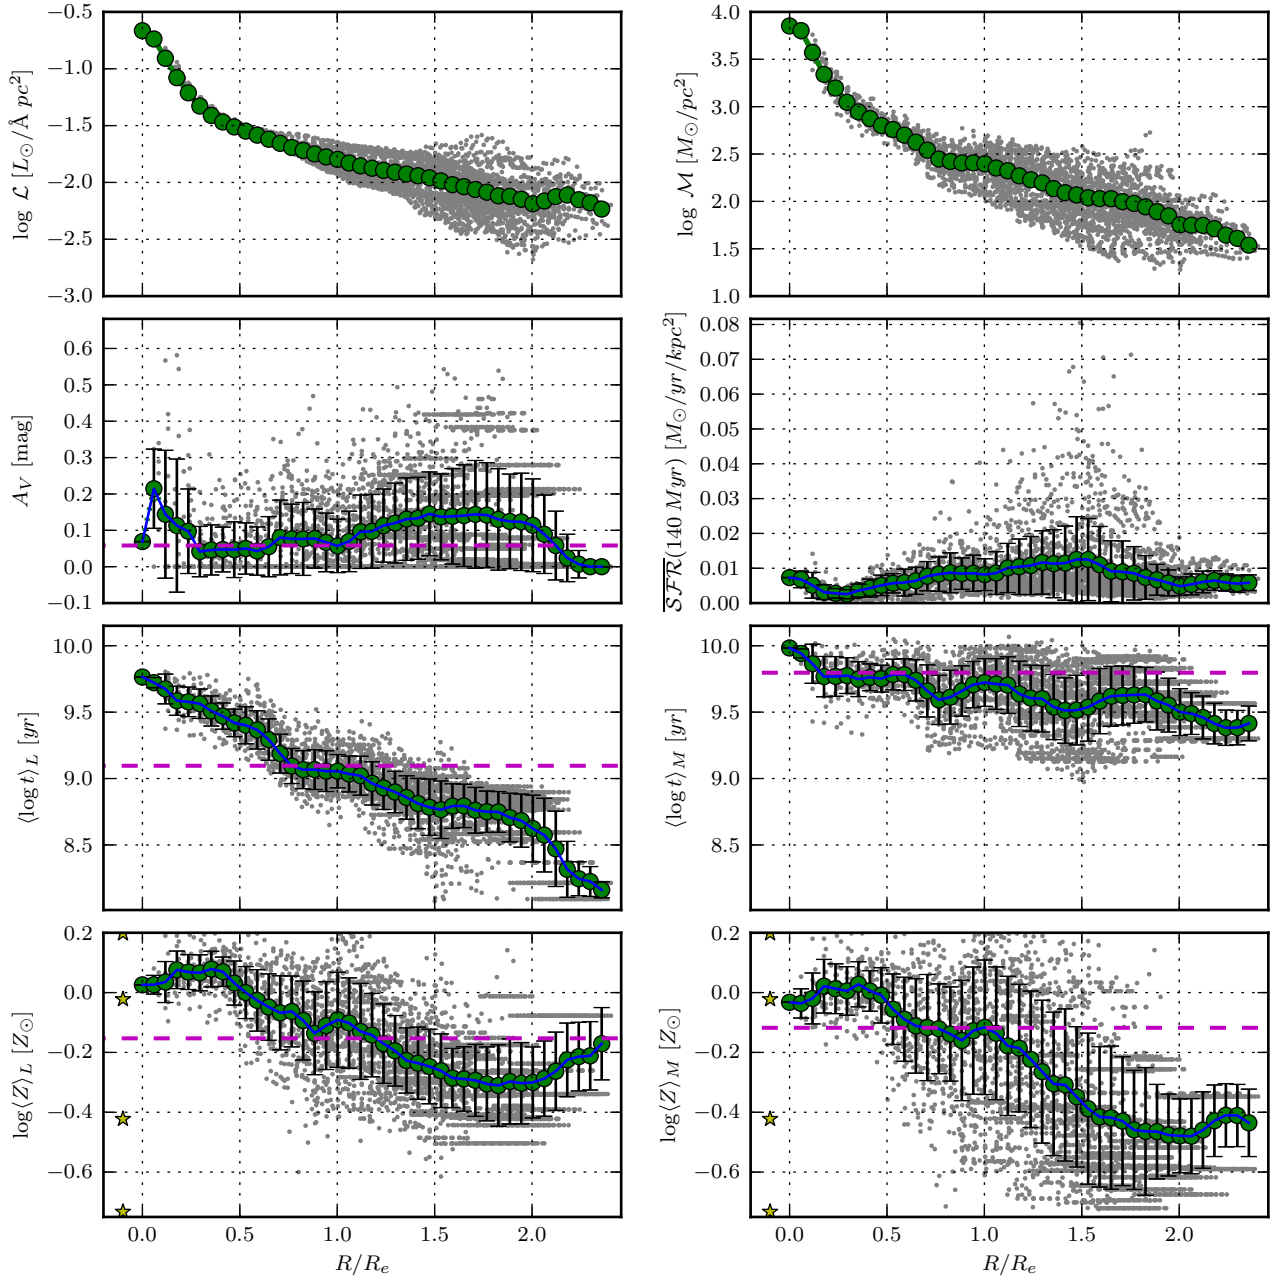


Figure 10. Radial profiles of several properties. Points in the background correspond to values in individual xy pixels (the horizontal strips of dots towards large R correspond to Voronoi zones). Solid lines represent the average of the plotted quantity, while the error bars represent the dispersion in the R -bin. The dashed horizontal line marks the value derived from the STARLIGHT analysis of the integrated spectrum (i.e. collapsing the xy dimensions of the data-cube).

expected direction, with inner producing less light per unit mass than outer ones. It is also seen that the mass growth at different speeds for different regions, something which is better appreciated in the second row of panels, where each growth curve is plotted on a 0 to 1 scale, with 1 represent the present values (which tantamounts to cumulating the equivalent x_t and μ_t vectors for each region). The $R > R_e$ regions reaches 80% of its mass at $t = ??$, while the inner regions do so earlier: at $t = ??$ Gyr for the central pixel, $??$ Gyr for $R < R_e$ and $??$ Gyr for the whole $< R_e$ region.

The differential version of these curves are show in the 3rd row of plots. Notice that, because of the smoothing and resampling, x and μ become continuous functions, so what is actually plotted is $dx/d \log t$ (left) and $d\mu/d \log t$ (right). The bottom left panel shows $SFR(t)$ the SFR surface density for each of the regions, $SFR(t)$. The overall level of SFR per unit area decrease from the nucleus outwards, even though, because of the area factor, the inner regions contribute little to the total SFR of the galaxy. Notice that, again because of the smoothing, $SFR(t)$ is now a continuous function of

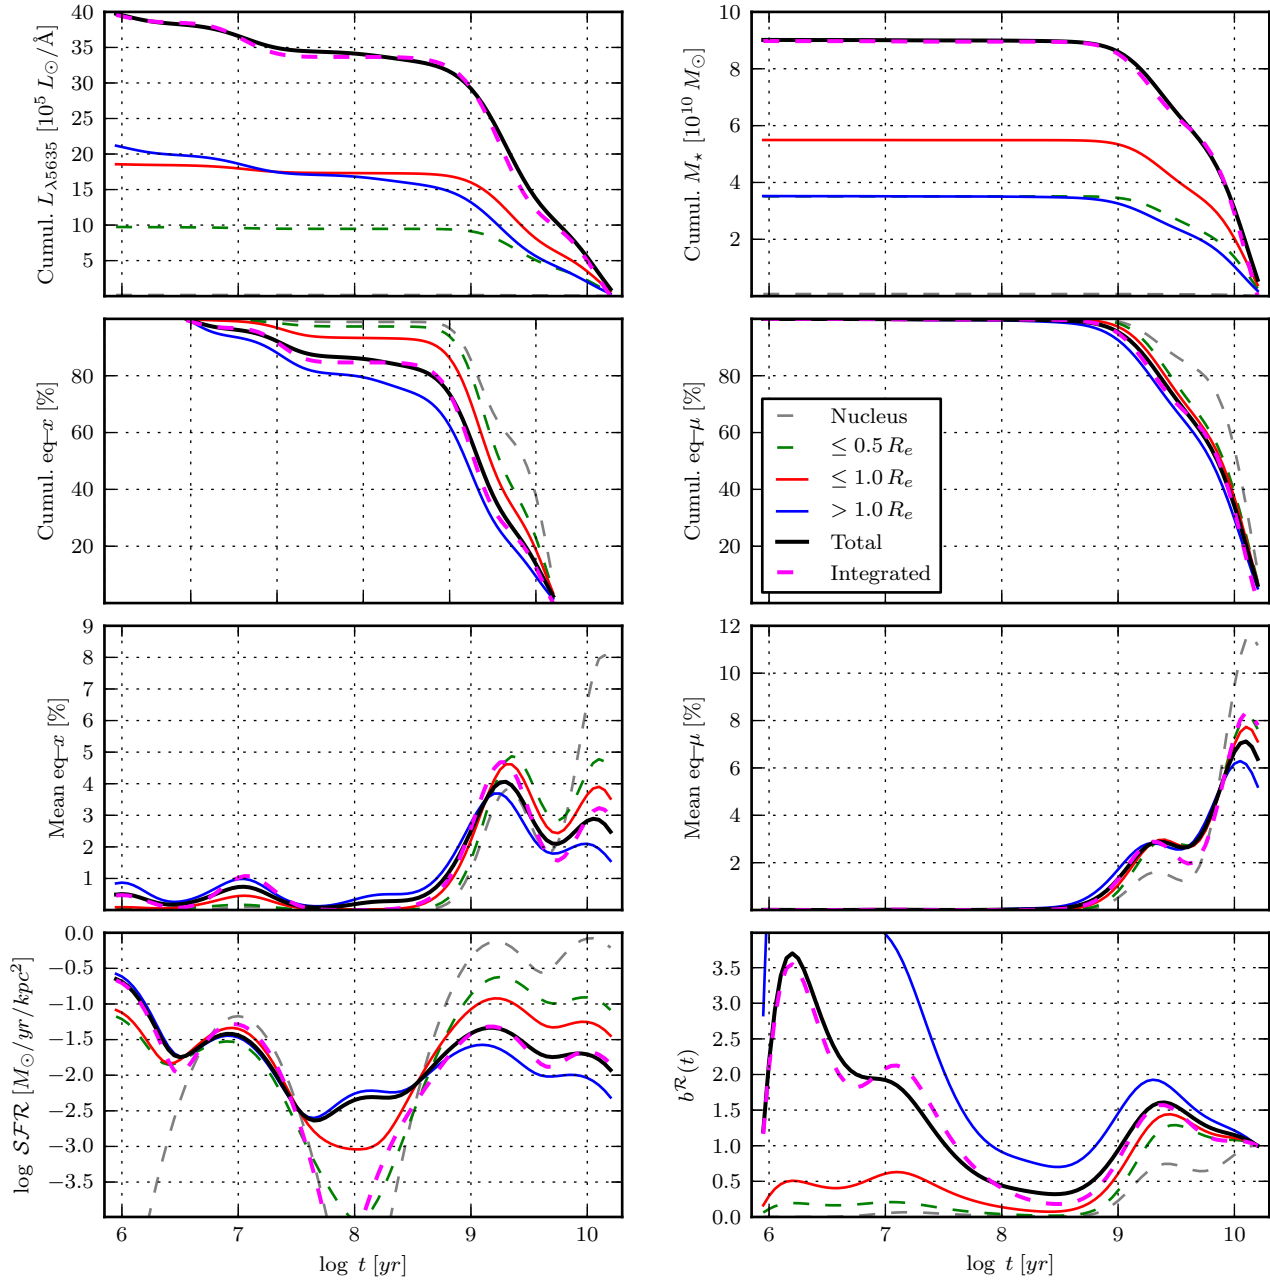


Figure 11. Time evolution of light and mass related properties, as derived from the STARLIGHT fits. The discrete x_{tZxy} and μ_{tZxy} "population vectors" involved in these curves were smoothed by 0.5 dex in $\log t$, marginalized in Z and integrated over the spatial regions indicated. Dotted gray lines indicate the nucleus (central pixel). Orange, red, and blue lines correspond to $R < 0.5R_e$, $R < 1R_e$ and $R > 1R_e$ regions, respectively. The black line represents the whole galaxy evolution as reconstructed from the sum of its parts, while (as in Fig. 10) the magenta dashed line is used to represent the results obtained for the spatially integrated spectrum. The top panels show the luminosity (left) and mass (right) growth curves, which give the total light (mass) due to populations older than t .

time, so that this panel represent the instantaneous SFH. The right panel shows Scalos's b parameter for each region, which does use the running mean $\overline{SFR}(t)$ of equation 8. The comparison with the previous panel illustrates the care one must have when interpreting SFHs in spatially resolved data. For instance, while the nucleus stands out in absolute SFR units at all times, most of its mass was formed long ago, such that it is now forming stars at a lower pace than

average (ie, $b < 1$), while the outer regions are growing at a pace comparable to our higher than their lifetime average.

Another feature seen in Fig. 11 is that the integrated light and mass growth curves are remarkably similar, ie, the analysis of the total yields essential the same results of the analysis of the sum of the parts. This is good news for studies based on integrated light, but, since we are talking about a single galaxy, it remains to be seen how general this results

Figure 12. Rxt : L, M & SFR... a la EPerez. NO Z(t) panel!!

Figure 13. Rxt : x, mu FRACTIONS ... a la EPerez+Cid

is. The reason for this homogeneity is that, as seen in Fig ??, the A_V is not a strongly varying function of position. In the limit of $A_V = \text{constant}$, one indeed expects the total spectrum to have the same properties as the sum of its parts ..

5.9 Space \times time diagrams: SFHs in 2D

As is evident at this point of the paper, one of the main challenges involved in the analysis of the multi dimensional data build from the combination of the spatial dimensions with the t and Z dimensions opened by population synthesis is how to visualize the results. All examples shown so far project two or more of these axes.

Fig. 12 shows an attempt to visualize galaxy evolution as a function of both space and time. The trick is to compress xy into R , and collapse the Z axis, producing a radially averaged SFH(R, t). The right panel show the luminosity density \mathcal{L} at each radial position R and for each lookbacktime t . As with the other plots in this section, the original \mathcal{L}_{tZyx} array from which this map was derived was smoothed in time using a $FWHM(\log t) = ??$ gaussian, marginalized over Z and collapsed into one spatial dimension. Unlike in Fig. 10, we now use the first type of radial averaging discussed in Section 5.7, but the results do not depend strongly on this choice. The light distribution in R and t show ...

The middle panel in Fig. 12 shows the mass surface density as a function of R and t . It shows BLABLABLA The right panels show Despite its different appearance, this map results from a simple mathematical transformation of the middle panel...

DO FIG AND THEN WRITE! CAN HAVE 6 PANELS: L, M, SFR (top) and x, μ and sSFR (bottom)...

5.10 Space vs. time data cubes

There is so much one can plot in a single image. To go beyond the t - R diagrams Fig. 12 and 13 one needs more dimensions than a sheet of paper can accommodate. An alternative is shown in Fig. 14 shows t -slices through the cumulative \mathcal{M}_{xyt} cube....

Age sequences like those in Fig. 14 can be constructed for several, but not all properties inferred through STARLIGHT or any other fossil method. One can, for instance, replace mass, by intrinsic luminosity, SFR or stellar metallicity, in cumulative or differential form. However, it is obviously impossible to reconstruct maps of A_V, σ_*, v_* as a function of t . The panels in Fig. 14 are *not* snapshots of a movie, ie.,

Figure 14. snapshots ... Replace upper panel by M($R;t$) profiles?? (like in my Fig06.* plots??)

Figure 15. Residual cube in a 3D rendering. The main emission lines are indicated. EPEREZ DO IT!!

they are not pictures of the galaxy as it appeared at different look-back times. Instead, these are maps of where stars of a given age t are located nowadays, a sort of motion-less movie.

Simulators should take notice of these simple facts. IFS data plus fossil methods provide an extremely rich, yet inherently limited form of time-travel. Illustrative and beautiful as they are, movies of stars and gas particles moving through time and space will never be directly compared to anything observational. In other words, simulations should be "convolved" through this "reality" filter. The observationally relevant predictions are the distribution of stellar ages and metallicities as a function of xy .

5.11 Emission line maps

Independent of the stellar population applications described throughout this paper, the STARLIGHT fits provide a residual spectrum R_λ which is (inasmuch as possible) free of stellar light. This is of course of great aid in the measurement of emission lines, and hence in the derivation of a series of diagnostics of the warm ISM properties, such as its kinematics, nebular extinction and chemical abundance. The completeness of nebular and stellar analysis has been amply explored in integrated spectra (??REFS??), and the addition of 2 spatial dimensions to this analysis holds great promise.

Fig. 15 shows a $R_{\lambda,xy}$ residuals cube, with some of the main emission lines marked. The quality of the residuals cannot be well appreciated in this 3D plot, so the reader is referred to GD et al for clearer plots.

+ BLABLABLA? About care with H β and etc??

6 DISCUSSION AND SUMMARY

We have presented

NGC ??? (CALIFA ???) was used as a show case. It was not the goal of this article to analyse this galaxy, so we only briefly mention the results: It has this and that...

simulations should generate things as they can be observed.

Raidal probs are ok, but would be more interesting partitioniioning things like the bar, spiral arms ...

ACKNOWLEDGMENTS

This work has been funded with support from the Spanish Ministerio de Ciencia e Innovation for the support to this project through the grants AYA2010-15081, and AYA2010-22111-C03-03 that are also co-financed with FEDER funds.

CAPES ...

RGB to bibtex things!!!

REFERENCES

Bruzual G., Charlot S., 2003, MNRAS, 344, 1000

Cappellari, M., & Copin, Y. 2003, MNRAS, 342, 345
 González Delgado , R. M.; Cerviño, M., Martins, L. P.,
 Leitherer, C.; Hauschildt, P. H., 2005, MNRAS, 357, 945
 Sánchez, S.F., Kennicutt, R.C., Gil de Paz, A., et al., 2012,
 A&A, 538, 8
 Vazdekis, A., Snchez-Blzquez, P.; Falcn-Barroso, J.; Ce-
 narro, A. J.; Beasley, M. A.; Cardiel, N.; Gorgas, J.;
 Peletier, R. F., 2010, MNRAS, 404, 1639
 !⊙j⊙! Cappellari, M., Copin, Y., 2003, MNRAS, 342, 345
 !⊙j⊙! Sánchez, S.F., et al., 2012, å, 538, 8