

Rebutting fake news on full spectral fitting

Roberto Cid Fernandes¹★

¹*Departamento de Física-CFM, Universidade Federal de Santa Catarina, C.P. 476, 88040-900, Florianópolis, SC, Brazil*

1 June 2018

ABSTRACT

A recent paper by Ge et al. performs a series of experiments with two full spectral fitting codes, pPXF and STARLIGHT, finding that the two yield consistent results when the input spectrum is not heavily reddened. For $E(B - V) \gtrsim 0.2$, however, they claim STARLIGHT leads to severe biases in the derived properties. Counterintuitively, and at odds with previous simulations, they find that this behaviour worsens significantly as the signal-to-noise ratio of the input spectrum increases. This communication shows that this is entirely due to an $A_V < 1$ mag condition imposed while initializing the Markov chains in the code. This choice is normally irrelevant in real-life galaxy work but can become critical in artificial experiments. Alleviating this usually harmless initialization constraint changes the Ge et al. results completely, as was explained to the authors before their publication. We replicate their spectral fitting experiments, finding much smaller biases. Furthermore both bias and scatter in the derived properties all converge as S/N increases, as one would expect. We also show how the very output of the code provides ways of diagnosing anomalies in the fits. The code behaviour has been documented in careful and extensive experiments in the literature, but the biased analysis of Ge et al. is just not representative of STARLIGHT at all.

Key words: methods: data analysis – galaxies: stellar content

1 INTRODUCTION

Full spectral synthesis techniques have blossomed after Bruzual & Charlot (2003) released a suite of evolutionary population synthesis models for the spectra of simple stellar populations (SSP) as a function of age and metallicity. That paper updated the theory to a spectral resolution which observational work on galaxies had achieved long before, fostering the development of methods to match observed galaxy spectra with combinations of the SSP models. These methods became generally known as “full spectral fitting”, highlighting the λ -by- λ nature of how data and models are compared.

Reviews on this subject can be found in Walcher et al. (2011) and Conroy (2013), while Cid Fernandes (2006, 2007) gives outdated but useful reviews reflecting the early days of the field. The introduction section in the recent paper by Ge et al. (2018, hereafter GYCMLL) is also a good source of references. The authors then proceed to compare two publicly available full spectral fitting codes: pPXF (Cappellari & Emsellem 2004; Cappellari 2017), and STARLIGHT (Cid Fernandes et al. 2005).

Comparing different methods is a tedious yet useful exercise if carried out with due care. Unfortunately, GYCMLL

completely misrepresents the performance of STARLIGHT. Their results convey the false idea that the code leads to significant biases in derived properties such as $E(B - V)$, mean stellar age, and metallicity. Particularly disastrous results are obtained for young and dusty systems, specially when assigned exquisite signal-to-noise ratios.

The sole purpose of this communication is to set the record straight. We present a revised version of the same spectral fitting experiments performed by GYCMLL. The results, however, could hardly be more different. The nature of the difference lies on an irrelevant and trivially fixable technicality which GYCMLL were fully aware of. As a result they end up painting a distorted and biased picture of STARLIGHT’s performance. This calls for clarification, and this is what is presented here.

2 SPECTRAL FITTING EXPERIMENTS

The experiments performed by GYCMLL follow the usual Monte Carlo logic of: (1) take an input spectrum of known properties, (2) perturb it with gaussian noise, (3) process the perturbed spectrum through the spectral fitting code, and (4) compare the output properties with the input ones.

Their main set of experiments uses solar metallicity SSP spectra with five different ages ($\log t/\text{yr} = 8.0, 8.5, 9.0, 9.5,$

★ e-mail:cid@astro.ufsc.br

10.0) and attenuated by a foreground screen of dust producing $E(B - V)$ values of 0.0, 0.1, 0.2, 0.3, 0.4 and 0.5 (the Calzetti et al. 2000 reddening curve is adopted). The stellar population library is that of Vazdekis et al. (2010) for a Salpeter IMF and “Padova 2000” evolutionary tracks (Girardi et al. 2000). Each of these 30 input spectra (5 ages \times 6 extinctions) was then perturbed with gaussian noise with a λ -independent amplitude adjusted to produce signal-to-noise ratios of $S/N = 10, 18, 32, 56, 100, 178,$ and 316 in the 5490–5510 Å interval. In order to estimate the biases and uncertainties in derived properties 50 incarnations of each input spectrum were built, each with a different realization of the noise.

We have replicated these simulations, generating our own set of mock spectra exactly as described above. We then fit them with STARLIGHT using a spectral base of 150 SSPs spanning 25 logarithmically-spaced ages between $t = 63$ Myr and 15.8 Gyr and 6 metallicities ($[M/H] = -1.7, -1.3, -0.7, -0.4, 0.0,$ and $+0.2$) also from Vazdekis et al. (2010), and assuming a Calzetti extinction curve, as done by GYCMLL.

2.1 Results

Fig. 1 shows our results. The plot has the same layout and axis scales as figure 2 of GYCMLL, to which it should be compared. Results for the five input ages are shown in different columns (as labeled in the top row), while the input $E(B - V)$ values are coded by different colors, as indicated in the bottom left panel.

Panels in the top row show the difference between the output and input values of $E(B - V)$. The curves connect the mean values of the bias $\Delta E(B - V) = E(B - V)_{\text{out}} - E(B - V)_{\text{in}}$ obtained from the 50 perturbed versions of each input spectrum. Error bars are only shown for the runs with $E(B - V)_{\text{in}} = 0.5$ (in red), but the scatter is about the same for other values of $E(B - V)_{\text{in}}$.

The contrast with the results of GYCMLL is striking. Take the red curve for the 100 Myr model, for instance. While they obtain a bias of $\Delta E(B - V) = -0.20 \pm 0.02$ at $S/N = 316$ we obtain an insignificant $\Delta E(B - V) = -0.011 \pm 0.005$. Moreover, as expected on the basis of pure common sense, but contrary to what they find, the bias decreases as the S/N improves.

Other panels in Fig. 1 show the statistics of the output minus input values of the mass to light ratio, the luminosity weighted mean age and metallicity. The results again thoroughly contradict those reported by GYCMLL. Focusing again on the runs for $t_{\text{in}} = 100$ Myr and $E(B - V)_{\text{in}} = 0.5$, they find that $\Delta \log(t_L)$ (see equation 3 of their paper) starts from ~ 0.3 dex at $S/N = 10$ and grows to ~ 0.5 dex at $S/N = 316$. Similarly, the bias in metallicity starts off badly and only worsens as the S/N improves. The real situation is shown by the red curves in the first column of Fig. 1, which show modest biases of about 0.1 dex at low S/N that decrease to negligible values as the S/N increases.

Before explaining the reason for these dramatic differences let us also replicate their figure 5, where they show that STARLIGHT fails to fit the spectrum and recover the parameters of an input 100 Myr, solar metallicity SSP reddened by $E(B - V)_{\text{in}} = 0.5$. Fig. 2 shows our version of that plot. The top panel shows the input spectrum (in black) and the STARLIGHT fit in red. The difference between the two is only

visible in the residual spectrum (shown in the panel below), and even then it is hardly noticeable. In the units of the plot, where the flux at 5635 Å is $\equiv 1$, the rms of the residual is a tiny 0.0009, and even the peak value is just 0.0128 (over the CaII K line). The fit shown by GYCMLL is much worse, with residuals an order of magnitude larger.

The bottom panels in Fig. 2 show the corresponding STARLIGHT solution in terms of light and mass fractions. The light distribution (at the 5635 Å normalization wavelength, left panel) peaks at the correct age, and 88% of it is contained within ± 0.1 dex of $\log(t/\text{yr}) = 8.0$. The output $E(B - V)$ is 0.493, and the metallicity is mostly solar (73% of the light in $[M/H] = 0.0$ and most of the rest in $[M/H] = +0.2$). Not identical to the input parameters, but pretty close by any real-world standard. This fit is obtained with the same configuration file used by GYCMLL¹. There are several ways to play with the configuration parameters to further improve the fit (see Fig. 4), although one should wonder how necessary or relevant it is to better a fit which already produces such tiny residuals. Also, as discussed in Section 2.3, the code itself points to a better solution.

The bottom right panel translates light to mass fractions. Again, the distribution peaks in the input population. Inevitably, the scatter is larger, since the input model is nearly the youngest in the base, with a $M/L_{\lambda 5635}$ ratio dozens of times smaller than the oldest ones. This explains why the insignificant 1% of light attributed to the 1 Gyr, $[M/H] = +0.2$ population gets inflated to 5% of the mass, and similarly for the other components. The same effect results in a bias of about +0.1 dex in the mass-to-light ratio even at high S/N , as seen in the first column of the $\Delta \log M/L_r$ panels in Fig. 1. As should be evident from its very name, STARLIGHT fits light, not mass fractions. It should be equally evident even to non STARLIGHT users that components accounting for such tiny light fractions should be treated as noise, not signal.²

The comparison of our Fig. 2 with figure 5 in GYCMLL reveals drastic differences. In their version the STARLIGHT solution is scattered all over the age-metallicity plane. They only show results in terms of mass fractions, which boosts the scatter due to the highly non-linear $M-L$ relation of stars. In terms of light fractions their results would not look as scattered, though still much worse than those seen in the bottom left panel of Fig. 2.

How come the results reported here are so strikingly different from those of GYCMLL?

2.2 The futile reason for the discrepancy

The reason for these completely different results lies on a hitherto deemed unimportant line of the code which stipulates a maximum A_V of 1 mag while initializing the Markov

¹ GYCMLL incorrectly associate this configuration file with that used in the “state-of-the-art analysis of the CALIFA dataset by de Amorim et al. (2017)”. That would not even be possible, as the version of STARLIGHT used in that paper differs from the public one in several details, including the configuration file.

² The analysis of which components are significant is usually done a posteriori, while analyzing the output of a fit, but the code also offers ways of setting a priori significance thresholds by adjusting the corresponding parameters in the configuration file.




Figure 1. A revised version of figure 2 in GYCMLL. The different rows show the output minus input biases in $E(B - V)$, mass-to-light ratio, mean stellar age, and metallicity as a function of S/N . Each column shows results for a different input age, indicated in the top row. Colours code the input value of $E(B - V)$, as labeled in the bottom-left panel. For clarity, $\pm 1\sigma$ error bars are only shown for the extreme case of $E(B - V)_{\text{in}} = 0.5$ (red).

chains, instead of allowing it to reach the upper limit established in the configuration file. The logic behind this single if/then line of the code was that real galaxies seldom suffer so much extinction, so starting the parameter chains from A_V as large as 4 (as in the example configuration file distributed with the code and used by GYCMLL) would be a waste of time and an unnecessary hindrance to the progress of the chains.

In any case, if a spectrum does require more dust to be fitted the chains should travel to high A_V regions of the parameter space even if initialized at low values. This indeed happens in the $E(B - V)_{\text{in}} \geq 0.3$ (equivalent to $A_{V,\text{in}} \geq 1.2$) STARLIGHT fits of GYCMLL, which do cross the $A_V < 1$ limit of the initial chains. In some of their simulations, however, the chains do not run for long enough to reach the correct values. Since $A_V/E(B - V) = 4.05$ for the reddening law used in these experiments, $A_V < 1$ implies $E(B - V) < 0.25$, so that their conclusion that STARLIGHT works well when $E(B - V) \leq 0.2$ is a mere reflection of this initialization strategy.

The difference between our figures and their versions in GYCMLL boils down to removing this single line of the

code. Clearly this otherwise harmless and innocent technical condition has nefarious effects for this particular experiment.

The reason why their results are particularly disastrous when young and highly reddened SSPs are used as input to the code is rather technical, though not difficult to grasp. Old SSPs are already naturally red, and can only get substantially redder by increasing A_V , so the code recovers their properties even if the chains start far from the right parameters. A highly reddened young population, however, has a spectrum whose overall shape can be well matched by combinations of older stellar populations with less A_V . Indeed, even the bad fit example shown by GYCMLL does match the continuum shape, failing only in the absorption lines. These regions of the parameter space are visited first by the Markov chains, specially if they start from low A_V . In order to reach the extreme corner of the parameter space containing the true solution, the chains must be allowed to run until they realize there are solutions where not only the continuum shape can be fitted, but also the absorption lines. The setup parameters in the example configuration files provided in the public distribution of STARLIGHT do not let the chains run for long enough to reach A_V values as high as needed in

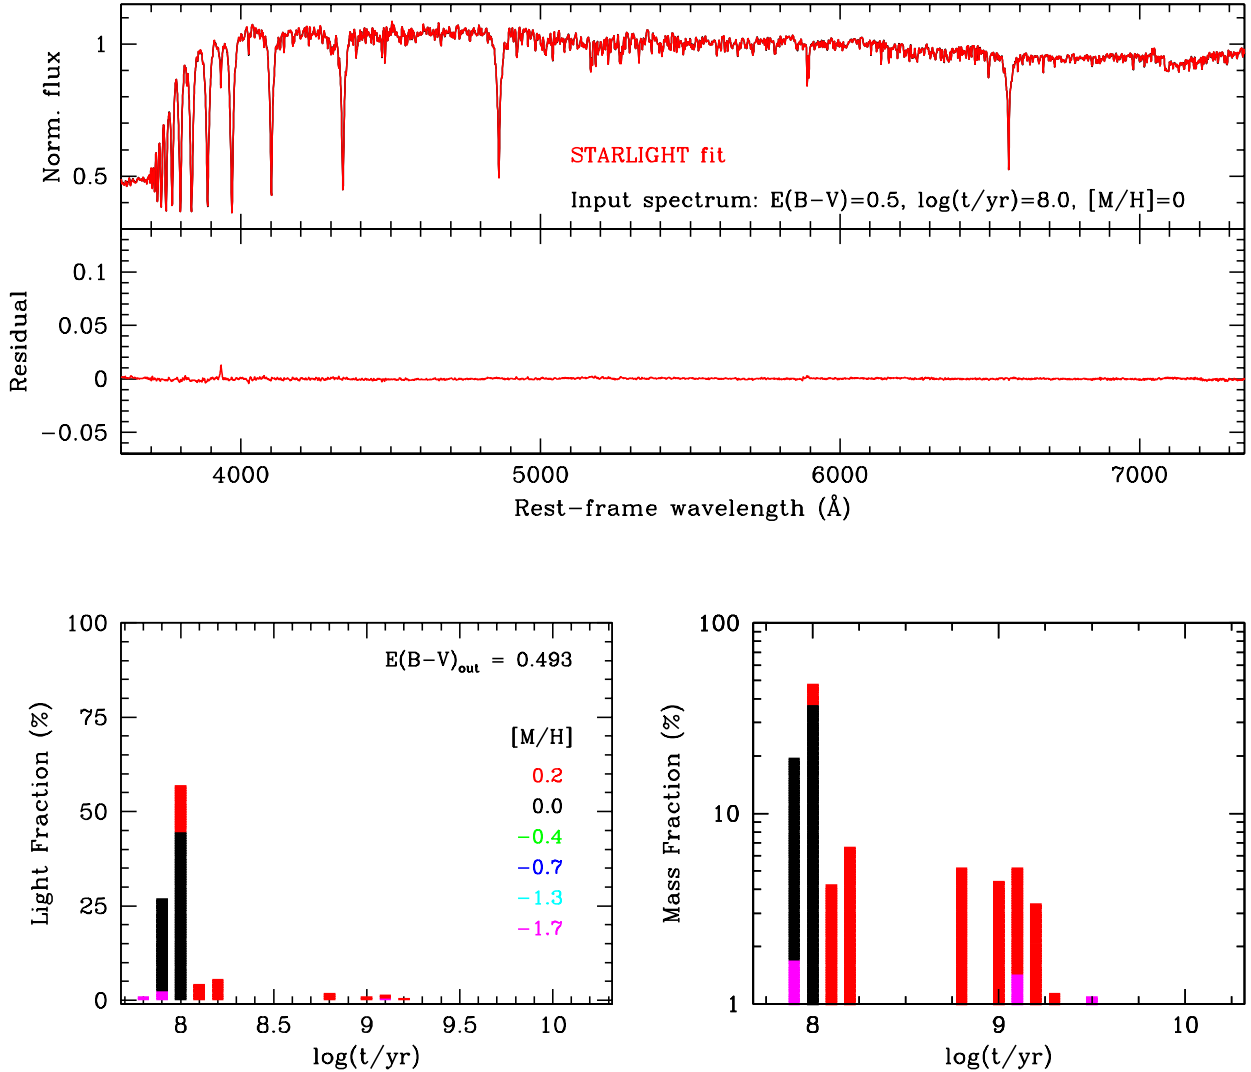


Figure 2. A revised version of figure 5 in GYCMLL. The top panels show the input spectrum (in black) of a solar metallicity 100 Myr old SSP extinguished by $E(B-V) = 0.5$, the ~ indistinguishable STARLIGHT fit (in red), and the corresponding output minus input residual spectrum on a zoomed scale. Bottom panels show the mixture of SSP spectra behind the model fit in terms of light (left) and mass (right) fractions.

this case. The higher the S/N the longer it takes to reach this region, which explains why GYCMLL find the counterintuitive and illogical result that biases worsen as S/N increases, as well as the large computing times.³

Though it is possible to overcome the biases reported by GYCMLL by playing with the several technical knobs and weighting options provided by STARLIGHT (an alternative explored in their so called “slow mode” fits, also explored here in appendix A), simply removing the $A_V < 1$ condition

³ Inspection of the (usually ignored) run-time screen output of the code reveals signs of anomalous behaviour of the chains, like very small step sizes, and in some cases warnings that the “rapid χ^2 tricks” explained in the user manual are struggling.

already leads to perfectly satisfactory results. Indeed, Figs. 1 and 2 use the same configuration used by GYCMLL.

All of this was made abundantly clear to Ge et al. through private communications with the third author in December/2017. Their paper does in fact echo some of the explanations given then and expanded here. Instead of using the revised code, however, they chose to focus their whole study on experiments which exploit what is in practice an irrelevant and trivially fixable technical vulnerability of the code. The result is a biased mischaracterization of STARLIGHT’s performance which serves no purpose but to cloud the field with meaningless simulations.⁴

The revised version of the code is available at

⁴ The running terminology for this is “fake news”.




Figure 3. Biases in color excess and age (as derived from simulations like those of GYCMLL and with the unrevised code) against the “problem detector” δ_{SCF} index. Positive values of δ_{SCF} indicate that a single component provides a better spectral fit than the one found for a mixed stellar population. All points are for an input spectrum composed of an SSP of 100 Myr. The S/N is 316 and 18 in the top and bottom panels, respectively. Circles are coloured according to $E(B-V)_{\text{in}}$ following the same palette shown in the bottom left panel of Fig. 1. Gray crosses mark the best single component fit. Mixed population solutions producing $\delta_{\text{SCF}} > 0$ should be rejected.

www.starlight.ufsc.br for users who suspect their analysis may have been affected by the same issue. As discussed in section 2.4, where we test to which extent previous fits of SDSS and CALIFA spectra were affected by this, the likelihood of this happening in real galaxy work is tiny, so most users should not worry about this issue. If the object under study is as extreme as the worst cases explored by GYCMLL, where all the light is concentrated in a highly reddened single population, then biases like those they report may indeed happen. However, if this is your case, then you should be using STARLIGHT in a different way. In fact, as explained next, the output of the original code already contains a much better solution than the one used by GYCMLL.

2.3 Single component fits

As documented in its user manual, besides the light fractions associated with the smallest χ^2 model found by the Markov chains (the “official” best solution), STARLIGHT also fits the input spectrum with each of the base components individually. In principle this feature should only interest users working with star clusters, for whom the question is not what are the mixture parameters, but what is the single component which best fits the data and what is the implied extinction. These single component fits (SCF) can however

be used to perform a very basic test of the reasonability of the best fit mixed model.

The test goes like this: If any of the χ_{SCF}^2 ’s is smaller than the χ^2 of the best-fit mixed population, then either you are fitting a star cluster or something is not right (e.g., the chains have not run for long enough, problems with the input data, etc.). In either case, when $\chi_{\text{SCF}}^2 < \chi^2$ the mixed model should either be ruled out or replaced by the best SCF.

Had GYCMLL applied this sanity check to their simulations they would have rejected most of their STARLIGHT fits, particularly those at high S/N . To demonstrate this we have run the same simulations as before, but now with the original code with the problematic $A_V < 1$ initialization.

The results are shown in Fig. 3, where we plot the bias in $E(B-V)$ and $\log(t_L)$ against $\delta_{\text{SCF}} \equiv (\chi^2 - \chi_{\text{SCF}}^2)/N$, where N is the number of wavelengths. This index measures how much worse the mixed solution is with respect to the SCF. For clarity only results for $t_{\text{in}} = 100$ Myr (the worst case) are shown. Top panels are for runs with $S/N = 316$ (the largest in their simulations) and bottom ones for $S/N = 18$. Circles are coloured according to the input value of $E(B-V)$ as in Fig. 1. As reported by GYCMLL, biases are larger the larger $E(B-V)_{\text{in}}$ is. Gray crosses show the difference between the best SCF and input properties. As expected, the SCF recover the input parameters exactly.⁵

At high S/N all points have $\delta_{\text{SCF}} > 0$, signaling that the single component, and not the mixed population fit should be considered the best solution. More importantly, the correspondence between δ_{SCF} and the bias is evident, confirming that this index can be used to track problematic fits. At lower S/N δ_{SCF} also tracks the bias, but with more scatter. The values of δ_{SCF} are also much smaller because the χ^2 ’s scale as $(S/N)^2$. Unlike at high S/N , negative values of δ_{SCF} appear in the bottom panels. This happens because noise may make a mixed population fit slightly better than the one obtained with the correct model. In any case, a $\delta_{\text{SCF}} > 0$ criterion does identify the most biased solutions.

The point to take from this experiment is that δ_{SCF} is a useful problem detector for STARLIGHT runs.

2.4 Testing previous runs

As explained in Section 2.2, the biases reported by GYCMLL derive from an $A_V < 1$ initialization condition which should not affect real galaxy analysis with STARLIGHT, but leads to very slow progress of the Markov chains in extreme situations like when A_V is high and all the light is concentrated in a single population. The likelihood of real galaxy spectra triggering this anomalous behaviour of the code is small, but one would still like to have a way of testing this, and the δ_{SCF} index suits this need.

To illustrate its use, we have applied this test to our ~ 10 yr old STARLIGHT fits of SDSS DR7 spectra, used in several publications by our group and others. The information on

⁵ Technically speaking, STARLIGHT did give GYCMLL the correct solution. It just so happens that because of the particularities of their simulations this solution was not in the best mixed model column of the output, but in the one reporting the single component fits.

the SCF was gathered from the STARLIGHT’s output. Out of 614667 galaxies in the SDSS main galaxy sample with $S/N \geq 10$ just 798 ($\sim 0.1\%$) have $\delta_{\text{SCF}} > 0$. Inspection of these cases shows that some are just bad input, but most are early-type galaxies where the mixed fit already puts most of the light in the best SCF population anyway. Similar results are obtained for the 363253 CALIFA spectra in the PYCASSO database (de Amorim et al. 2017), 0.3% of which have $\delta_{\text{SCF}} > 0$. STARLIGHT users can easily replicate this test with their own runs, though judging from these results this is hardly necessary.

These basic tests reinforce our conclusion that the initialization which leads to the catastrophic results of GYCMLL has negligible effect in real life. Though the code is far from being perfect, it is evident that biases as large as they report are in no measure typical of STARLIGHT.

Finally, at the risk of stating the obvious, one should always check the spectral fits. Not one-by-one, as this is not viable with the huge datasets currently available, but at least filtering those with suspicious output, like bad quality fit, low S/N , $\delta_{\text{SCF}} > 0$, outliers in relations between derived properties, etc. For instance, a one second look at figure 5 of GYCMLL tells a minimally experienced eye that there is something wrong with the fit. Indeed, that very figure prompted an inquire from the authors to the developer of STARLIGHT, and all explanations given here were then offered in detail, including the revised code itself and a version of Fig. 2. Yet, for unbeknownst reasons the authors chose to feature that misleading plot as an example of STARLIGHT.

2.5 Other experiments

All results reported above were obtained with the same “standard” STARLIGHT setup used in the main part of GYCMLL’s paper, but they also present experiments with a “slow mode” configuration for the Markov chains in an appendix. We have replicated those experiments using the revised version of the code. Results are shown in appendix A. While they obtain that these fits yields much better results than those with the “default” configuration, we find no significant differences.

GYCMLL further explore cases where the input spectrum is not a single SSP, but a linear combination of two SSPs of different ages. This more interesting test was also replicated, and results are shown in appendix B. While they find that biases in the derived properties increase to unacceptable levels as $E(B - V)_{\text{in}}$ increases, we obtain much smaller biases and no dependence on $E(B - V)_{\text{in}}$.

In both cases the differences between what is reported here and there stem from the otherwise unimportant $A_V < 1$ initialization limit of which GYCMLL were informed.

2.6 A note on extinction

In GYCMLL one reads that one of their motivations to study what they call the “algorithm bias” is that Cid Fernandes et al. (2005) “found that the dust extinction had a clear difference with that of the MPA/JHU group” (sic).

That is a shocking misunderstanding of section 4.1 and figure 8 in Cid Fernandes et al. (2005), where the A_z values derived with STARLIGHT were compared to those in the value

added catalogues by the MPA/JHU group (Kauffmann et al. 2003; Brinchmann et al. 2004) for SDSS galaxies. What is demonstrated and explicitly said there is exactly the opposite: “We thus conclude that there are no substantial differences between the MPA/JHU and our estimates of the stellar extinction other than those implied by differences in the reddening laws adopted in the two studies.” Moreover, both Cid Fernandes et al. (2005) and Asari et al. (2007) report convincing empirical evidence that the stellar extinction estimated by STARLIGHT is reasonable.

Still regarding extinction, and given that the biases reported by GYCMLL occur for their simulations with large $E(B - V)_{\text{in}}$, it is fit to recall that most full spectral fitting codes (STARLIGHT included) treat the effects of dust in a blatantly over-simplistic way. Users dealing with dusty sources should seriously ponder how this might affect their analysis. The attenuation model assumed might have a far more relevant role than the spectral fitting algorithm.

3 DISCUSSION

Despite the rebuttal of the GYCMLL results presented in this communication, it is fair to recognize that their experiments lead to the identification of a previously unnoticed aspect of the code’s behaviour. As shown above, this should have no practical effect in the analysis of real galaxy data. A “quick-and-dirty” method to test this by comparing the best mixed model fit with the best SCF (given in the output, though seldom used) was presented and shown to corroborate this conclusion.

The code has been previously subjected to simulations qualitatively like those of GYCMLL, where the output is compared to a known input (e.g. Cid Fernandes et al. 2005, 2014), but no serious biases were found. Recently Magris et al. (2015) compared the performance of four inverse population synthesis codes, finding that they all lead to similar accuracies in derived properties. Regarding extinction, for instance, their tests with STARLIGHT yield a bias in A_V of just -0.03 mag and a precision of ± 0.06 . As in our own previous simulations, Magris et al. (2015) use as input spectra computed with star formation histories and dust content designed to mimic real galaxies.

The GYCMLL simulations, on the other hand, are not based on galaxies, but reddened SSPs. This in itself already makes things difficult for a Markov chain based algorithm, as it places the right solution in an extreme corner of a vast parameter space, which may take long to reach, particularly at high S/N . What really triggers the anomalous behaviour of the chains, however, are the large values of extinction used in some of their tests. It is this combination of unfavorable and unrealistic conditions that sometimes turns the otherwise irrelevant $A_V < 1$ initialization limit into a critical bottleneck, hindering the performance of the algorithm to the point that in order to obtain adequate results the setup chain parameters require adjustments, brutally slowing the code. The previous tests of the code were made for far more realistic conditions, which explains why the potential effects of the apparently (and in most cases truly) harmless $A_V < 1$ initialization constraint were not realized before.

The only difference between the results shown in Figs. 1 and 2 with respect to their twins in GYCMLL is that

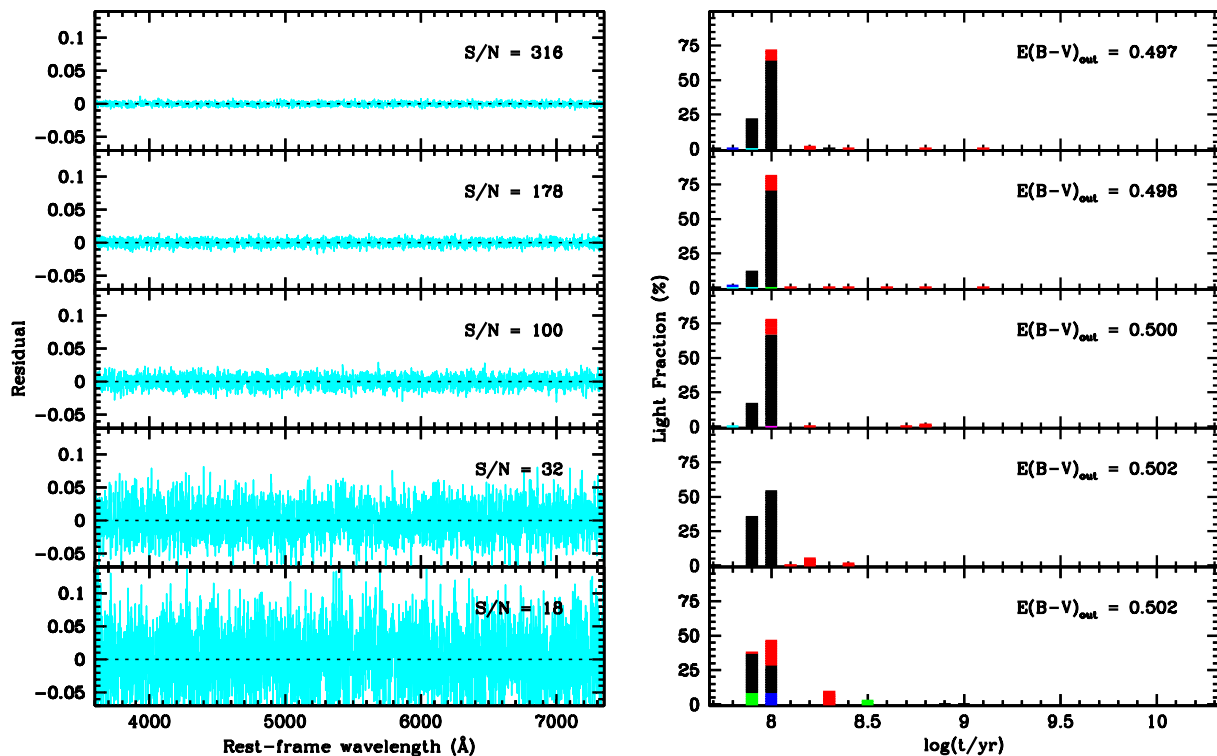


Figure 4. Output minus input residual spectra (left) and the corresponding light fractions age-distribution (right) for the same input model as in Fig. 2 (a 100 Myr solar SSP with $E(B - V) = 0.5$) for “slow mode” fits (see text). Each row is for a different level of noise.

they were run without the $A_V < 1$ condition. This banal modification heals all the problematic behaviour incorrectly used to characterize STARLIGHT’s performance.

Though that seems hardly necessary given the already excellent results shown in these plots, even better fits can be obtained by adjusting the configuration parameters. Fig. 4 exemplifies this by showing a series of “slow mode” fits obtained for the poster child case explored by GYCMLL, namely an input SSP of 100 Myr, $[M/H] = 0$ and $E(B - V) = 0.5$ observed at different S/N . This plot looks the same whether using the original or the revised version of the code, confirming that with appropriate parameters the Markov chains can beat the $A_V < 1$ initialization even in this extreme case. The only difference is in the computing times, which are an order of magnitude faster with the revised code.

4 CONCLUDING REMARKS

A common nightmare among those who offer their codes to public use is that they will be misused or unfairly characterized. It is, after all, the user’s responsibility to understand the workings and the limitations of the code and to use it reasonably. Unfortunately none of these unspoken rules-of-the-trade were followed by GYCMLL. Despite being explicitly informed on the practical irrelevance of the unexpected effect they came across, they proceeded to publish a paper which seeds unfounded doubts not only on STARLIGHT users

but also on the reliability of all results obtained with the code. This is what prompted this (otherwise unnecessary) rebuttal article.

None of what was presented here implies that STARLIGHT is better or worse than any other full spectra fitting code. STARLIGHT does have known limitations, as consistently pointed out in our own publications. Its built-in tendency to break the solution among many components, for instance, needs to be dealt with carefully to avoid over-interpretations, as is its option to work on an observationally-oriented light fractions space instead of the more theory-guided alternative to fit for mass fractions. Smoothing the output age and metallicity arrays or censoring components with small light fractions are examples of post-processing steps to handle this issue. Other codes tackle this problem not a posteriori, but imposing smooth solutions to begin with, as in the pPXF-based study by [McDermid et al. \(2015\)](#), where regularization in the mass fractions space is imposed. MOPED ([Heavens et al. 2000](#); [Panter et al. 2003](#)), STECKMAP ([Ocvirk et al. 2006](#)), sedfit ([Walcher et al. 2006](#)), VESPA ([Tojeiro et al. 2007](#)), ULySS ([Koleva et al. 2009](#)), TGASPEX and DynBaS3D ([Magris et al. 2015](#)), BEAGLE ([Chevallard & Charlot 2016](#)), and FIREFLY ([Wilkinson et al. 2017](#)) are examples of full spectral fitting codes which offer alternative takes on this and numerous other issues involved in the non-trivial task of

turning spectroscopic data on galaxies onto information on their stellar populations and dust content.

The inherent astrophysical and mathematical complexity of the subject, coupled to the numerous technicalities and philosophical issues involved in any one full spectral fitting software make the comparison of different codes a challenging task. We concur with GYCMML in that such comparisons should be stimulated, but, as demonstrated here, this requires a level of care that is just not met by their own study.

We close with a sobering thought. The whole area of full spectral fitting is nowhere near the standard of code-comparison achieved after decades of work by the photoionization community (Ferland et al. 2016), for instance. In fact, it is not at all clear that such a standard can even be met given the built-in differences in extant approaches to the problem. On the other hand, the field does have perspectives other than algorithm testing. A particularly promising line of work is to extend the wavelength range. To mention a recent example, López Fernández et al. (2016) perform a combined analysis of spectra and photometry, simultaneously fitting optical spectra from the CALIFA survey and UV photometry from GALEX. The extended baseline allows one to harness valuable information from the λ -dependent behaviour of different stellar populations, ultimately leading to better constrained parameters. Studies seeking to improve upon the over-simplistic way dust attenuation is treated in full spectral fitting codes should also be encouraged.

ACKNOWLEDGEMENTS

The author is in debt with all those who helped with STARLIGHT over the years, from the students who patiently tested it to all collaborators whose input (specially when critical) contributed to keeping it up and running for so long. Support from CNPq is duly acknowledged.

REFERENCES

- Asari N. V., Cid Fernandes R., Stasińska G., Torres-Papaqui J. P., Mateus A., Sodr e L., Schoenell W., Gomes J. M., 2007, *MNRAS*, **381**, 263
- Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, **351**, 1151
- Bruzual G., Charlot S., 2003, *MNRAS*, **344**, 1000
- Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, *ApJ*, **533**, 682
- Cappellari M., 2017, *MNRAS*, **466**, 798
- Cappellari M., Emsellem E., 2004, *PASP*, **116**, 138
- Chevallard J., Charlot S., 2016, *MNRAS*, **462**, 1415
- Cid Fernandes R., 2006, Bolet n de la Asociaci n Argentina de Astronom a La Plata Argentina, **49**, 228
- Cid Fernandes R., 2007, in Vazdekis A., Peletier R., eds, IAU Symposium Vol. 241, Stellar Populations as Building Blocks of Galaxies. pp 461–469 ([arXiv:astro-ph/0701899](https://arxiv.org/abs/astro-ph/0701899)), [doi:10.1017/S1743921307008794](https://doi.org/10.1017/S1743921307008794)
- Cid Fernandes R., Mateus A., Sodr e L., Stasińska G., Gomes J. M., 2005, *MNRAS*, **358**, 363
- Cid Fernandes R., et al., 2014, *A&A*, **561**, A130
- Conroy C., 2013, *ARA&A*, **51**, 393
- Ferland G., et al., 2016, preprint, ([arXiv:1603.08902](https://arxiv.org/abs/1603.08902))
- Ge J., Yan R., Cappellari M., Mao S., Li H., Lu Y., 2018, *MNRAS*,
- Girardi L., Bressan A., Bertelli G., Chiosi C., 2000, *A&AS*, **141**, 371
- Heavens A. F., Jimenez R., Lahav O., 2000, *MNRAS*, **317**, 965
- Kauffmann G., et al., 2003, *MNRAS*, **341**, 33
- Koleva M., Prugniel P., Bouchard A., Wu Y., 2009, *A&A*, **501**, 1269
- L pez Fern andez R., et al., 2016, *MNRAS*, **458**, 184
- Magris C. G., Mateu P. J., Mateu C., Bruzual G., Cabrera-Ziri I., Mej a-Narv ez A., 2015, *PASP*, **127**, 16
- McDermid R. M., et al., 2015, *MNRAS*, **448**, 3484
- Ocvirk P., Pichon C., Lan on A., Thi baut E., 2006, *MNRAS*, **365**, 46
- Panther B., Heavens A. F., Jimenez R., 2003, *MNRAS*, **343**, 1145
- Tojeiro R., Heavens A. F., Jimenez R., Panther B., 2007, *MNRAS*, **381**, 1252
- Vazdekis A., S nchez-Bl zquez P., Falc n-Barroso J., Cenarro A. J., Beasley M. A., Cardiel N., Gorgas J., Peletier R. F., 2010, *MNRAS*, **404**, 1639
- Walcher C. J., B ker T., Charlot S., Ho L. C., Rix H.-W., Rossa J., Shields J. C., van der Marel R. P., 2006, *ApJ*, **649**, 692
- Walcher J., Groves B., Budav ri T., Dale D., 2011, *Ap&SS*, **331**, 1
- Wilkinson D. M., Maraston C., Goddard D., Thomas D., Parikh T., 2017, *MNRAS*, **472**, 4297
- de Amorim A. L., et al., 2017, *MNRAS*, **471**, 3727

APPENDIX A: SLOW MODE RUNS

GYCMML also explore “slow mode” STARLIGHT fits, meaning fits where the technical parameters controlling the behaviour of the Markov chains are set to allow for more chains and also to let them run for longer. They find that with this more thorough exploration of the parameter space “the spectral fitting results show significant improvements”, particularly for runs with $E(B - V)_{in} > 0.2$.

Whilst that is in fact true with the original version of the code, by simply removing the $A_V < 1$ imposed upon initialization one finds that these improvements go from significant to negligible. This is shown in Fig. A1 (to be compared with figure A1 in GYCMML), which replicates Fig. 1, except that now the the slow mode configuration described by GYCMML was used in the fits.

We note in passing that, despite what is said in GYCMML, clipping⁶ has nothing to do with any of this.

APPENDIX B: FITS FOR COMBINATIONS OF TWO DIFFERENT SSPS

GYCMML also fitted mock spectra generated by adding two SSPs scaled to have the same flux in the 5490–5510   window, i.e., a fifty-fifty percent mixture in terms of light fraction. Despite its simplicity, seen in comparison with the single SSP mock spectra discussed in the main text this test goes one step further in the sense of emulating realistic galaxy spectra. The 13 SSPs used in this experiment

⁶ Clipping is a STARLIGHT feature designed to ignore pixels which “refuse” to be well matched by the model, like when an emission line is not properly masked, or some bad pixel is not appropriately flagged as such




Figure A1. A revised version of figure A1 in GYCMLL. The plot is identical to Fig. 1 in the main text, except that a slow mode configuration for the Markov chains was used for the STARLIGHT fits. Differences with respect to the results in Fig. 1 are however negligible.

have solar metallicity and ages from 63 Myr to 15.8 Gyr in logarithmic steps of 0.2 dex. Spectra for each of the resulting 78 unique combinations were perturbed by noise and processed through STARLIGHT.

Fig. B1 shows our results for this same experiment. The plot replicates figure 11 in their paper, and shows biases in the derived properties as a function of the input $E(B-V)$. Each panel shows 78 lines connecting the average values of the bias obtained for 50 realizations of each of the two-SSP combinations. The color reflects the mean log age of the two SSPs, weighted either by light (left panels) or mass (right).

Once again, and for the same reasons, our results contrast with those of GYCMLL, who identify large biases in derived properties and a strong tendency for them to increase with increasing $E(B-V)_{\text{in}}$. We find small biases and no tendency with input reddening. Only $\Delta E(B-V)$ (top panels) slightly increases with $E(B-V)_{\text{in}}$ when the mean age of the pair of SSPs leans towards the smaller values allowed (lines in the magenta–blue range), but even in this case the asymptotic bias is a tiny 0.02 mag, much smaller than obtained by GYCMLL. The average and standard deviation of $\Delta E(B-V)$ over all 78×50 combinations of input spectrum and noise are -0.004 ± 0.009 , as show by the open circle and error bar

towards the right end of the top panels in Fig. B1. For the other properties in Fig. B1 we find $\Delta \log M/L_r = 0.03 \pm 0.05$, $\Delta \log(t_L) = 0.02 \pm 0.04$, and $\Delta[M/H]_L = -0.04 \pm 0.07$. For the mass weighted mean ages and metallicities (shown in the right panels) we obtain $\Delta \log(t_M) = 0.08 \pm 0.11$, and $\Delta[M/H]_M = -0.002 \pm 0.056$.

Fig. B2 repeats the analysis of Fig. B1, but for slow mode runs, as in figure A4 of GYCMLL. The differences with respect to Fig. B1 are marginal, as can be appreciated visually comparing the two plots. In numerical terms, these longer runs yield essentially the same bias and scatter statistics: $\Delta E(B-V) = -0.005 \pm 0.008$, $\Delta \log M/L_r = 0.02 \pm 0.04$, $\Delta \log(t_L) = 0.02 \pm 0.03$, $\Delta[M/H]_L = -0.03 \pm 0.05$, $\Delta \log(t_M) = 0.06 \pm 0.09$, and $\Delta[M/H]_M = -0.003 \pm 0.051$.




Figure B1. A revised version of figure 11 in GYCMLL, showing biases resulting from fits to a combination of two SSPs of different ages. Each line connects the average $\Delta X = X_{\text{out}} - X_{\text{in}}$ values as a function of $E(B - V)_{\text{in}}$, where X is a given derived property. The statistics is done with 50 realizations of the noise (at $S/N = 60$). Colours code for the mean $\log(t/\text{yr})$ of the two populations, running chromatically from 7.9 (magenta) to 10.2 (red). In the left panels the colors are coded according to the luminosity weighted mean log age, while on the right the colors map the mass weighted log age. The open circle and error bar at the right end of the plots mark the mean and ± 1 sigma range of ΔX over all 78 combinations and 50 realizations.




Figure B2. As Fig. B1, but for slow mode runs (analogous to figure A4 in GYCMLL). Improvements with respect to the “standard mode” runs in Fig. B1 are only marginal.