

Finding Young Stellar Populations in Elliptical Galaxies from Independent Components of Optical Spectra

Ata Kabán*

Louisa A. Nolan†

Somak Raychaudhury†

Abstract

Elliptical galaxies are believed to consist of a single population of old stars formed together at an early epoch in the Universe, yet recent analyses of galaxy spectra seem to indicate the presence of significant younger populations of stars in them. The detailed physical modelling of such populations is computationally expensive, inhibiting the detailed analysis of the several million galaxy spectra becoming available over the next few years. Here we present a data mining application aimed at decomposing the spectra of galaxies into several coeval stellar populations, without the use of detailed physical models. This is achieved by performing a linear independent basis transformation that essentially decouples the initial problem of joint processing of a set of correlated spectral measurements into that of the independent processing of a small set of prototypical spectra. Two methods are investigated: (1) A fast projection approach is derived by exploiting the correlation structure of neighboring wavelength bins within the spectral data. (2) A factorisation method that takes advantage of the positivity of the spectra is also investigated. The preliminary results show that typical features observed in stellar population spectra of different evolutionary histories can be convincingly disentangled by these methods, despite the absence of input physics. The success of this basis transformation analysis in recovering physically interpretable representations indicates that this technique is a potentially powerful tool for astronomical data mining.

1 Introduction

The optical spectrum of a galaxy is a linear superposition of the spectra of its billions of constituent stars. Yet, since large populations of stars form in galaxies at definite periods of its lifetime, and the atomic and nuclear physics of the evolution of stellar populations,

though complex, are well understood, the detailed modelling of composite spectra of stellar populations can be used to yield a wealth of information about the history of a galaxy from its spectrum.

The spectrum of a star can be modelled as a function of three parameters— its mass, its age and its composition (since it is made mostly of hydrogen and helium, the last parameter is characterised by the relative abundance of other elements, and is known as the “chemical abundance”). Elliptical galaxies, which account for about 20% for all galaxies in the Universe, are believed to consist predominantly of a single coeval stellar population (e.g. [3, 12, 22]), all formed at an early epoch in the Universe. This implies that an elliptical galaxy can be modelled as a system of stars, all of the same age and chemical abundance, evolving together, if validated assumptions can be made about the distribution of stellar masses. However, as a result of detailed spectral studies conducted in the last decade (e.g. [10]), it now transpires that elliptical galaxies are more complex objects, at least some of which have undergone more recent bursts of substantial star formation, and consequently are likely to contain more than one stellar population component.

The determination of the star formation history of a galaxy has important implications for the still-controversial issue of the formation and evolution of galaxies. Until recently, the analysis of a large statistical sample of stellar populations of galaxies would not have been possible since only small ensembles of galaxy spectra were available. However, the development of data mining tools for automating parts of the analysis is becoming more and more essential in the light of the rapid increase in the availability of data that is approaching.

Recent and ongoing galaxy spectral surveys (2dFGRS, www.mso.anu.edu.au/2dFGRS/ (completed in 2003) and SDSS, www.sdss.org/) will produce more than two million galaxy spectra in the next few years, which is to be integrated into a more ambitious database of publicly-available astronomical data, incorporating Grid technology (the *Virtual Observatory*, www.ivoa.net). Since the detailed physical modelling of stel-

*School of Computer Science, The University of Birmingham, Birmingham B15 2TT, UK. Email: A.Kaban@cs.bham.ac.uk.

†School of Physics and Astronomy, The University of Birmingham, Birmingham B15 2TT, UK. Email: {lan, somak}@star.sr.bham.ac.uk

lar populations is numerically expensive, even a simple question like “what fraction of elliptical galaxies contain a significantly younger stellar population?” will take years to address by conventional modelling techniques using stellar population synthesis. The timely extraction of useful knowledge, such as the characteristics of the star formation history (ages, chemical abundances and masses of the component stellar populations) of galaxies, from these data will largely depend on developing appropriate data analysis tools that are able to complement more specialised astrophysical analyses.

The astrophysical questions motivating this study are:

1. Can we disentangle major stellar population components of elliptical galaxies without the use of detailed physical models?
2. How do the results from a data-driven analysis of observed galaxy spectra correlate with the parameters of star formation history determined via a completely independent model fitting technique used in astrophysics?

These questions have not been addressed before in a data driven manner — that is, based on the data characteristics only, without any specialised physical input.

1.1 Roadmap In this paper, we discuss and investigate statistical methods that attempt to solve the described inverse modelling problem by relating multivariate observations to lower-dimensional vectors of statistically independent unobserved variables through the use of a linear model. The required statistical assumptions will be derived from general characteristics of the data, in order to employ these methods in an appropriate manner.

The preliminary results presented in the next sections are based on the data described in Section 2. A projection approach that exploits the correlation structure of the spectra is presented in Section 3. In this approach, the required assumption for solving the inverse modelling is derived from exploiting the correlation structure between neighboring wavelength bins, which comes naturally with spectral data. The independent spectral components obtained turn out to be also physically interpretable and exhibit typical features of spectra of the young and mature stellar populations. We then compare the results with a positivity-based single stage approach, presented in Section 4, that has been often employed for analysing spectral data in different domains [9, 18]. We provide a simple probabilistic reformulation of this method that highlights its implicit assumptions, links it to the methods developed in [13] and

also allows us to potentially incorporate measurement errors (if known from domain knowledge) into the algorithm. In Section 5, the results are presented, discussed and comparatively assessed, first in a data-driven manner and then, more importantly, from the astrophysical perspective. Finally, our conclusions are summarised in the last section.

2 The data and model setting

The data we use here represent the *observed* optical spectra of 21 nearby elliptical galaxies, compiled by blending together 5855 measurements over the range 2000-8000 Å from various observatories on ground and in space. These represent the following galaxies: NGC 0205, NGC 0224, NGC 1052, NGC 1400, NGC 1407, IC 1459, NGC 1553, NGC 3115, NGC 3379, NGC 3557, NGC 3605, NGC 3904, NGC 3923, NGC 4374, NGC 4472, NGC 4621, NGC 4697, NGC 5018, NGC 5102, NGC 7144 and NGC 7252. The spectra have been corrected for redshift (i.e. converted from their observed wavelengths to their emitted wavelengths), and the fluxes are normalised to unity in the region 5020-5500 Å.

Since these spectra are compiled from sources with varying spectral coverage, the resulting data matrix has 1453 missing values, which are first imputed using a KNN imputation [21] from synthetic data. We preferred this non-parametric procedure here, as the missing data mechanism may not be random — an assumption made by most of other imputation schemes. The validity of the ‘missing at random’ assumption in the case of the analysed data set will need further study, simply because in some wavelength regions it is consistently hard or impossible to take a measurement.

In addition, for each measurement, an error value is also provided from known instrumental characteristics and uncertainty in calibration.

2.1 How many stellar populations? According to existing domain knowledge, it is likely that there are (at least) two components of interest [10]. However, the first eigenvector explains more than 95% of the data. Therefore, prior to deciding that a 2D representation space is justified (i.e. that at the given noise levels there is enough useful information in the data and we are not attempting to model the noise in a second component), we perform some simple, data-driven rank tests. We use the error matrix to derive thresholds for these tests. As shown on Figure 1, both a 2-norm test and an F-norm test [19] suggest that at the given error levels, the ‘clean’ matrix of spectra has rank 2. Although it is known that these perturbation bounds often tend to underestimate the rank [19], there are no records of over-estimation, so




Figure 1: Rank-tests for the matrix of stellar population spectra. The threshold values represent the 2-norm and the F-norm respectively of the matrix of known measurement errors.

we proceed to searching for a suitable two-dimensional latent representation space.

2.2 The model Each stellar population spectrum is essentially a vector over a binned wavelength range, that represents the flux (in arbitrary units) per unit wavelength in the bin centered on the wavelength. The data matrix, having T spectra in rows will be denoted by $\mathbf{Y} \in \mathcal{R}^{T \times N}$. Single elements of this matrix will be referred to as y_{tn} , the t -th row is denoted by \mathbf{y}_t or more explicitly $\mathbf{y}_{t,1:N}$ and likewise, the n -th column is denoted by \mathbf{y}_n or $\mathbf{y}_{1:T,n}$. The whole matrix \mathbf{Y} will also be referred to as $\mathbf{y}_{1:N}$. Similar notational convention will also apply to other variables in the model.

To account for the generation mechanism described in the previous section, namely that the observed optical spectrum of a galaxy is a linear superposition of the stars in the galaxy, a linear factor model will be assumed in this study. That is, we hypothesise that the T observations can be explained as a superposition of $K < T$ latent underlying component spectra $\mathbf{s}_k \in \mathcal{R}^K$ (sometimes termed also as factors or hidden causes [17, 13, 6]) that are not observable directly but only through an unknown linear mapping $\mathbf{A} \in \mathcal{R}^{T \times N}$. Formally, this can be written as the following.

$$(2.1) \quad \mathbf{y}_n = \mathbf{A}\mathbf{s}_n + \boldsymbol{\epsilon}$$

In eq. (2.1), the first term of the r.h.s. is the so called systematic component and $\boldsymbol{\epsilon}$ is the noise term or the stochastic component of this model. The noise term $\boldsymbol{\epsilon}$ is assumed to be zero-mean i.i.d. Gaussian, where the

diagonal structure of the covariance accounts for the notion that all dependencies that exist in \mathbf{Y} should be explained by the underlying hidden components.

The K components will be assumed statistically independent, this being a standard assumption of independent factor models [6, 13]. In the present application, this assumption is also cosmologically plausible, as there is little (no) interaction between stellar populations at different ages in a galaxy. The linearity of the mixture is physically justified as the fluxes of the hypothesised different subpopulations mix in an additive way.

3 Independent projections of stellar population spectra

A projection based approach is presented in this section. This will be accomplished in stages. A linear dimensionality reduction will first be performed. We then proceed at identifying independent directions in the low-dimensional projection space. This multi-stage projection approach is well-suited as a first attempt. It allows us to formulate sub-tasks in statistical terms, and given the 2D nature of the problem, it also allows us to benefit from a visual control over the data representation obtained at various stages.

3.1 Dimensionality reduction using SVD Dimensionality reduction is a useful preprocessing stage for both computational convenience and de-noising. It is well known from linear algebra [4, 19] that the best rank- K approximation of a matrix under any unitarily invariant norm is its rank- K SVD (Singular Value Decomposition) approximation. This is given by $\mathbf{Y} \approx \mathbf{U}\mathbf{D}\mathbf{V}^T$, where \mathbf{U} is the $T \times K$ matrix of left singular vectors, \mathbf{D} is the $K \times K$ diagonal matrix of singular values, \mathbf{V} is the $N \times K$ matrix of right singular vectors, and $\mathbf{U}^T\mathbf{U} = \mathbf{V}^T\mathbf{V} = \mathbf{I}_K$, where \mathbf{I}_K is the K -dimensional identity matrix. The projection is then simply obtained as $\mathbf{X} = \mathbf{U}^T\mathbf{Y}$.

It should be pointed out that the scope of an SVD-based projection is to identify an optimal (in the sense of any unitarily invariant norm) subspace of the data space. However, generally, individual singular vectors or eigenvectors are not interpretable separately, as they are not independent from each other. The same is true for PCA [17], for much the same reasons.

3.2 Finding non-orthogonal informative directions using contextual ICA We now turn to the key part of our analysis, where the directions of independent projection need to be found. Approaches with this aim are known under the name of Independent Component Analysis (ICA) [7, 6, 14]. A vast number of ICA

algorithms have been developed over the last decade, each having different built-in assumptions. In general terms, we can write the data likelihood of the desired basis transformation as the following

$$(3.2) \quad p(\mathbf{x}_{1:N}|\mathbf{B}) = \int d\mathbf{s}_{1:N} p(\mathbf{x}_{1:N}|\mathbf{s}_{1:N}, \mathbf{B}) p(\mathbf{s}_{1:N}).$$

where \mathbf{B} is the $K \times K$ unknown linear mapping (squared mixing matrix) that transforms the latent components \mathbf{S} into \mathbf{X} . That is, as standard in ICA, instead of inferring \mathbf{S} from \mathbf{Y} , it is easier to infer them from \mathbf{X} .

Assuming that the SVD projection performed as described in the previous section has removed the noise, then the noise term is a delta function

$$(3.3) \quad p(\mathbf{x}_{1:N}|\mathbf{s}_{1:N}, \mathbf{B}) = \delta(\mathbf{x}_{1:N} - \mathbf{B}\mathbf{s}_{1:N})$$

where \mathbf{B} is a squared $K \times K$ parameter matrix (mixing matrix) that contains the desired new bases in its columns and $\mathbf{s}_{1:N}$ are the independent representations in the new basis — both having to be estimated from the data. Thus, (3.2) reduces to the simple form below

$$(3.4) \quad p(\mathbf{x}_{1:N}) = |\det \mathbf{B}|^{-N} \prod_{k=1}^K p(s_{k,1:N})$$

$$(3.5) \quad = |\det \mathbf{B}^{-1}|^N \prod_{k=1}^K p((\mathbf{B}^{-1})_k \mathbf{x}_{1:N}).$$

Standard in squared ICA problems, it is easier to optimise for the inverse of \mathbf{B} . That is, instead of the ‘top-down’, or ‘generative’ transform \mathbf{B} , we estimate the ‘bottom-up’ or ‘projection’ transform \mathbf{B}^{-1} . However, without knowing $p(\mathbf{s}_{1:N})$, this is still an ill-posed problem. Clearly, a mechanical application of any ICA algorithm, out of the hundreds of existing ones, would produce different results, although any of these would be somewhat arbitrary. What we need is a well motivated prior distribution $p(\mathbf{s}_{1:N})$. However, as in most data mining applications of ICA, there is no such information explicitly available.

3.2.1 Exploiting correlations within the spectral data Let us observe, however, that in spectral data, there is a natural correlation structure between flux values in neighbouring wavelength bins. This is what we exploit here, by capturing it in a form of a

contextual (predictive) model.

$$\begin{aligned} p(s_{k,1:N}) &= \prod_{n=1}^N p(s_{kn}|s_{k,1:n-1}) \\ &= \prod_{n=1}^N p(s_{kn} - E[s_{kn}|s_{k,1:n-1}]) \\ &= \prod_{n=1}^N p((\mathbf{B}^{-1})_k(\mathbf{x}_n - E[\mathbf{x}_n|\mathbf{x}_{1:n-1}])) \end{aligned}$$

$\forall k = 1 : K$. The advantage of doing so is that now we only need to specify the form of density of the residual projections. Assuming a good enough predictor, then the residual is likely to have a heavy tailed (termed also super-Gaussian [7] or kurtotic) form of density. Indeed, using just the simplest first order predictor, which is an identity function

$$(3.6) \quad E[s_{kn}|s_{k,1:n-1}] \equiv s_{k,n-1}, \forall k = 1 : K$$

the difference process $\mathbf{x}_n - \mathbf{x}_{n-1}$ of the data already becomes highly kurtotic, as shown on Figure 2.

A similar approach has been previously taken and successfully demonstrated in the context of face image separation [5], where neighbouring pixel values of an image do also exhibit significant correlations.




Figure 2: Histograms of the difference process. Kurtosis values are 11.0503 and 33.9364 respectively.

Let us denote $\mathbf{r}_n = \mathbf{x}_n - E[\mathbf{x}_n|\mathbf{x}_{1:n-1}]$ and $u_{kn} = s_{kn} - E[s_{kn}|s_{k,1:n-1}]$. The data likelihood is then the following.

$$(3.7) \quad p(\mathbf{x}_{1:N}) = |\det \mathbf{B}^{-1}|^N \prod_n \prod_k p((\mathbf{B}^{-1})_k \mathbf{r}_n)$$

where now we know that $p(u_{kn})$ is a super-Gaussian density. Maximisation of this likelihood can now be

accomplished by employing any standard ICA algorithm — over $\mathbf{r}_{1:N}$ rather than $\mathbf{x}_{1:N}$. As the predictor may not be very accurate (we just used an identity predictor in our experiments), it is preferable to chose a robust approximation of the generalised exponential density, that grows relatively slowly in $|u_{kn}|$. Following the arguments in [7], in our experiments we have used the following:

$$(3.8) \quad \log p(u_k) \propto \exp(-u_k^2/2),$$

and the optimisation has been performed using the Newton method implemented in the FastICA routines [6], employing the faster deflationary approach. This has a cubic convergence [7]. Indeed, highly kurtotic independent projections have been found (kurtosis: 33.8796 and 12.7484 respectively) on the data investigated, in about ten iterations only.

A geometric illustration of the procedure just described is shown on Fig. 3. The SVD-compressed data are shown as dots and indeed, informative directions would be difficult to determine directly from the data. The scatter-plot of the difference process is shown as crosses. A star-like structure is apparent, with two main, non-orthogonal linear directions of high data density. These are the new bases (columns of \mathbf{B}) that are determined by the ICA procedure. Indeed, the two directions defined by the new bases found by the algorithm are highlighted on the plot as dark lines. The PCA axes of the data are also shown on the same plot for comparison. Interestingly, one of the axes is almost identical to one of the independent directions. The second is, however, just orthogonal to the first, while the ICA axes are not orthogonal to each other but do follow the two main directions of high density in the data.

To obtain the component spectra from the ICA procedure described, we simply compute the projections of the individual flux values of all galaxies at all wavelength bins onto the new bases (which is the composition of the two linear transforms performed during the analysis process described):

$$(3.9) \quad \mathbf{s}_{1:N} = \mathbf{B}\mathbf{x}_{1:N} = \mathbf{A}\mathbf{y}_{1:N}$$

where \mathbf{B} and \mathbf{A} now denote the recovered mixing parameters. Specifically, after \mathbf{B} is found, \mathbf{A} is computed as the matrix product \mathbf{UB} .

The component-wise reconstruction of the 21 individual stellar population spectra from their independent components are shown on Figure 4. The physical interpretability of these components will be assessed in Section 5, however, as a data-driven observation, it is interesting to note that the recovered spectral components turned out to be positive valued — although positivity

has not been artificially imposed at this stage during the analysis process. We note that indeed negative values of the flux would be difficult to interpret, therefore in the next section we discuss a different approach, where the required latent density $p(\mathbf{s}_{1:N})$ is derived from a positivity constraint.

4 A Positivity-based approach

The use of positive factorisation of positive matrices to replace PCA for analysing positive data, such as spectral data dates back to work reported in [9]. Positive (more exactly non-negative) factor models have been further developed in [18]. However, in the absence of either a density-based or a geometric interpretation, the implicit assumptions are not clear and therefore the interpretation of the results may not be straightforward. Somewhat related, in [13], a fully Bayesian formulation of a positive factorisation model is given and variants with sparse positive priors are applied to synthetic stellar population spectra in a different context of investigations than ours. Retaining the probabilistic framework, that allows us to make all assumptions explicit, we present a simpler version of their algorithm, based on maximum a posteriori (MAP) / maximum likelihood (ML) estimation, which will highlight the link with positive factorisation algorithms [18]. A MAP or ML estimation is sufficient for our purposes as we are concerned with a data explanation task for a fixed data set only. Once the most appropriate model is found, the full Bayesian machinery remains available to derive fully generative models that are able to better generalise on new data.

Retaining the positivity of the representation, and the fact that the overall transform in the approach adopted in the last section has been linear, the following linear model can be formulated.

$$(4.10) \quad p(\mathbf{y}_n|\mathbf{A}, \mathbf{E}_n) = \int d\mathbf{s} p(\mathbf{y}_n|\mathbf{A}\mathbf{s}, \mathbf{E}_n)p(\mathbf{s})$$

A Gaussian measurement noise will be assumed. Furthermore, according to prior knowledge about the levels of imprecision of the physical instruments, which vary independently for each stellar population and each wavelength bin, we will have individual diagonal variances \mathbf{E}_n^2 at each wavelength bin n , $p(\mathbf{y}_n|\mathbf{s}, \mathbf{A}, \mathbf{E}_n) \sim \mathcal{N}(\mathbf{y}_n|\mathbf{A}\mathbf{s}_n, \mathbf{E}_n^2)$. In rest, we use the same notation as before, \mathbf{y}_n refers to the relative flux values observed at the n -th wavelength bin, \mathbf{A} is the unknown mixing matrix parameter and $p(\mathbf{s})$ is the distribution of the latent components.

Now the latent prior needs to be specified. Apart from its positive support we don't have much information in this respect. Therefore we formulate a vague




Figure 3: Geometric illustration of the described ICA procedure. The number of points shown as dots equals the number of different wavelength bins in the set of measurements, each point being the 2D compressed representation of the fluxes at one of these bins. The differences between these 2D vectors at consecutive wavelength bins are marked with 'x'. The PCA basis of the data is superimposed (in light color) and also translated to the origin for comparison with the ICA basis found on the difference process. The ICA basis is shown in dark color.

exponential prior, that is

$$(4.11) \quad p(\mathbf{s}) \propto \prod_k \exp(-\alpha_{kn} |s_k|)$$

where $\alpha_{kn} = \alpha, \forall k, n$ is a small positive constant. If $\alpha \rightarrow 0$, then the prior becomes non-informative but also improper and in this case the MAP estimation procedure given below becomes ML.

To obtain a MAP estimate, the posterior needs to be maximised $p(\mathbf{s}_n | \mathbf{y}_n, \mathbf{A}, \mathbf{E}_n) \propto p(\mathbf{y}_n | \mathbf{s}_n, \mathbf{A}, \mathbf{E}_n) p(\mathbf{s}_n)$ which is proportional to the complete data likelihood.

Positivity of the elements of both \mathbf{A} and \mathbf{S} are then imposed by adding Lagrangian terms [20] to the complete data log likelihood.

$$\begin{aligned} \mathcal{L} &= \sum_n \{ \log p(\mathbf{y}_n | \mathbf{s}_n, \mathbf{A}, \mathbf{E}_n) + \log p(\mathbf{s}_n) \} \\ &+ Tr \mathbf{L}_1^T \mathbf{A} + Tr \mathbf{L}_2^T \mathbf{S} \end{aligned}$$

where \mathbf{L}_1 and \mathbf{L}_2 are a set of non-negative Lagrange multipliers and Tr denotes the trace of a matrix.

From the stationary equations w.r.t. \mathbf{A} and \mathbf{S} the Lagrange multipliers are obtained.

$$\begin{aligned} \mathbf{L}_1 &= \sum_n \mathbf{E}_n^{-2} \mathbf{A} \mathbf{s}_n \mathbf{s}_n^T - \sum_n \mathbf{E}_n^{-2} \mathbf{y}_n \mathbf{s}_n^T \\ \mathbf{L}_2 &= \sum_n \mathbf{A}^T \mathbf{E}_n^{-2} \mathbf{A} \mathbf{s}_n - \sum_n \mathbf{A}^T \mathbf{E}_n^{-2} \mathbf{y}_n + \alpha \end{aligned}$$

Now from the Karush-Kuhn-Tucker conditions [20] $L_{tk} A_{tk} = 0$ and $L_{kn} s_{kn} = 0, \forall t, k, n$, we have two fixed point equations which provide the convergent alternating iterative algorithm of the multiplicative form below.

$$\begin{aligned} \mathbf{A} &= \mathbf{A} \odot (\mathbf{E}^{-2} \odot \mathbf{Y}) \mathbf{S}^T \oslash [\mathbf{E}^{-2} \odot \mathbf{A} \mathbf{S}] \mathbf{S}^T \\ \mathbf{S} &= \mathbf{S} \odot \mathbf{A}^T (\mathbf{E}^{-2} \odot \mathbf{Y}) \oslash \{ \mathbf{A}^T [\mathbf{E}^{-2} \odot \mathbf{A} \mathbf{S}] + \alpha \} \end{aligned}$$

where \odot denotes element-wise multiplication and \oslash denotes element-wise division. If $\alpha = 0$, the iterative algorithm above is identical to the least-squares based non-negative factorisation algorithm proposed in [18] from a non-probabilistic starting point, with the only difference that now we also have the \mathbf{E} terms to account for known measurement errors.

5 Evaluation

5.1 Data driven evaluation Here we assess the effectiveness of the methods presented above according to two indicators: (1) data reconstruction and (2) the mutual information (MI) between the components of the representation created.

Method	Min	Median	Max
SVD, cICA	0	3.36×10^{-4}	0.294
NMF	10^{-10}	3.64×10^{-4}	0.317
NMFe	0	4.08×10^{-4}	0.351

Table 1: Data reconstruction errors under the L2 norm. cICA = contextual ICA, NMF = Non-negative matrix factorisation with $\alpha=0$, NMFe = NMF with exponential prior having $\alpha = 0.1$

Table 1 shows the data reconstruction results across all $N \times T$ measurements for the various methods. These

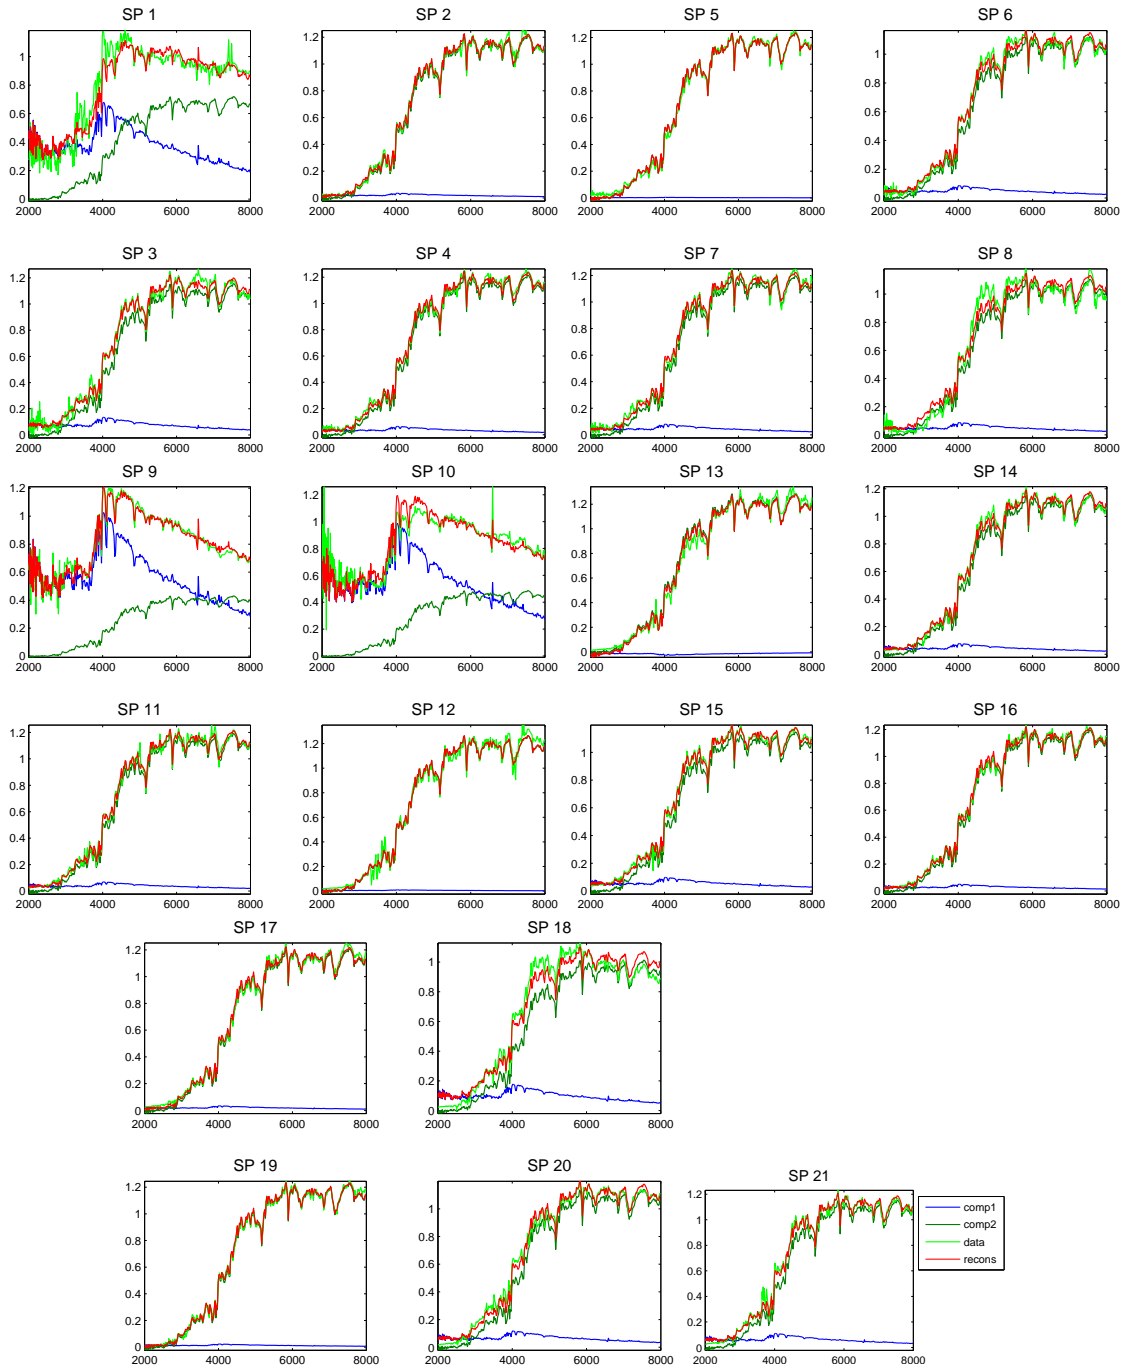


Figure 4: The reconstruction of stellar population spectra using the projection based contextual ICA. Here, our algorithm has decomposed the observed spectrum of each galaxy into two inferred populations (blue and dark green — the two darkest colors, in black-and-white printing), whose sum is given in red, superimposed with the data in light green (intermediate grey and light grey respectively, in black-and-white printing). In many cases (e.g. #1,9,10,18), a significant younger stellar population is found, which is confirmed by our detailed physical modelling. The numbering of the spectra corresponds to the enumeration order of the corresponding galaxies as given in Section 2.

are measured as the squared distances between the data and the reconstruction, which is in accordance with the Gaussian noise assumption. The reconstruction error of cICA are identical to that of the SVD, since the ICA transform does not reduce the dimensionality further. For NMF and NMFe, 15 randomly initialised runs were performed and the one with the highest likelihood value was selected in this evaluation, in order to avoid the possibility of getting trapped in a local optimum. Indeed, note that the positivity-based single stage approach involves a non-convex optimisation whereas the SVD-based preprocessing is a convex problem. The median of the error appears to be the smallest in the case of SVD+cICA, however, a pairwise application of the non-parametric Wilcoxon rank sum test to the whole sample distribution of the individual $T \times N$ reconstruction errors returned that the difference of the medians for cICA and NMF is not statistically significant at the 0.05% level. In the case of all other pairs, the differences between medians were found statistically significant — as expected, the additional term that enforces the latent prior distribution in NMFe causes a slight increase in data reconstruction error. However, this small difference on its own would not be a practically concerning issue here, as we deliberately formulated constraints in our models in order to obtain representations that are statistically independent to a higher degree — in the hope that these may then be interpretable and capable of being independently further processed. More interesting is therefore to evaluate to what extent do these methods achieve statistically independent representations.

The information theoretic quantity that measures the degree of statistical dependence is the Mutual Information (MI) [1]. This is a non-negative value that vanishes if the component densities are perfectly independent (the smaller the MI the better). It can be shown [14] that maximising the ICA log likelihood in the noise-free case is equivalent to minimising MI between the representation components. Here we compute the sample-based MI of the components, as estimated according to the procedure described in [2]. The comparative values obtained by various methods for this data are shown in Table 2. For the contextual ICA method, that works on predictive residuals, two values are given: the first value is computed from the projections of the residuals whereas the second value is computed from the projections of the data, taken as it was iid. (just for obtaining a value to compare with those obtained with the rest of the methods). It is apparent from the table that in both cases the contextual ICA model achieves lower MI (greater independence of the components). The positivity-based method with non-

Method	MI
NMF	0.5923
NMFe	0.6019
cICA	(u) 0.00010394 (s) 0.5583
PCA	1.0931

Table 2: Sample-based mutual information estimates of the two components obtained with the various methods for the set of stellar population spectra investigated. Smaller values signify higher independence achieved after the estimated basis transform.

informative improper prior is the next best performing method, whereas employing sparse priors (greater α values) does not lead to components that are more independent, for this data. PCA is used as a baseline, as we know that it doesn’t produce an independent representation. Visually, the positivity based components do not look much different from the ones obtained from cICA (Fig. 4), although they are slightly more noisy. The PCA-based decomposition, however, in general, are theoretically not guaranteed to be interpretable, as already discussed in Section 3.1 (see [17] for details). We now turn to evaluate the interpretability of the obtained results from the astrophysical perspective.

5.2 Astrophysical evaluation Here, we compare the results from the linear independent basis transformation analysis with an entirely independent determination of the star formation history, based on detailed astrophysical models of the evolution of stellar populations.

The 21 observed spectra have been analysed by matching them with synthetic stellar population spectra. For each of the observed spectra, a two-stellar population component model [8] was fitted. The age, chemical abundance and relative mass fraction of each component were allowed to vary freely. The best fit in each case was determined by a minimum χ^2 test [15, 11].

The principle for creating synthetic stellar population spectra is simple, although the input physics is complex. The spectral energy distribution of a star evolves according to its initial mass and chemical abundance. If the initial mass distribution and the chemical abundance of a stellar population is known, and the spectral evolution of each individual star in this initial population may be modelled, the stellar spectra may be summed over the mass distribution at any point in time to give the integrated spectrum of the population at that age. The ingredients for a stellar population model are therefore: stellar evolutionary tracks; a library of stellar




Figure 5: Comparison of the derived components with physical models of the stellar population spectra. *Top two plots:* Synthetic stellar population spectra according to the physical models of [8]. *Right:* Spectra of a population of age 10 Gyr, where chemical abundance, from bottom to top, 0.2, 1.0 and 2.5 times solar; *Left:* age = 0.7 Gyr, same chemical abundances. The dotted lines mark some of the absorption features in the spectrum which are typically strong in young stellar populations, and the dashed lines mark some of the absorption features which are typically strong in old stellar populations. From left to right, the absorption line species are: MgII (2799 Å), H ϵ (3970 Å), H δ (4102 Å), H γ (4340 Å), H β (4861 Å), Mgb (5175 Å), NaD (5893 Å), H α (6563 Å), TiO (7126 Å). *Bottom two plots:* the 2 components found from the various different linear independent basis transformation analyses, from bottom to top: cICA, NMF, NMFc, PCA. (The spectra are shifted along the vertical axis for the sake of clarity.) The recovered spectra are convincingly disentangled into one component with young stellar population features (MgII, H ϵ , H δ , H γ , H β , H α : dotted lines) and shape, and a second with the features (Mgb, NaD, TiO: dashed lines) and shape of an old, high chemical abundance stellar population.

spectra; a method of calibrating the theoretical luminosity and effective temperature, determined from the evolutionary sequence, so that the appropriate atmosphere may be assigned to each star at each time-step in its evolution. In contrast, the linear independent basis transformation method employs no knowledge of the underlying physics in the observed spectra.

Model spectra [8] for a young (0.7 Gyr) and an old (10 Gyr) stellar population, with three different chemical abundances (0.2, 1.0 and 2.5 times solar abundance) are shown in Fig. 5, together with the spectra recov-

ered from the linear independent basis transformation analyses. The similarity is most apparent. The data-driven linear analyses, whilst not reproducing any of the physical model spectra precisely (which would not be expected anyway), extract many of the important identifying characteristics of these two categories of model spectra, which are indeed quite different from each other both in their overall shape and details. Some of the most important features, the so-called absorption-lines, are marked with dashed (Mgb, NaD and TiO, typically strong features in old, high chemical abundance stel-

lar population spectra) and dotted (MgII, H ϵ , H δ , H γ , H β and H α , typically strong in the spectra of young (\lesssim 1 Gyr) stellar populations) vertical lines on the figure (Fig. 5), which demonstrate the physical interpretability of the representations created by the basis transformation analyses.

We correlate the star formation history parameters derived from fitting the two-component model spectra to the observed spectra (i.e. a physical analysis approach) with the weight of the contributions from the linear basis transformation analyses, by defining these weights as $c_k = a_{tk} / \sum_{k'} a_{tk'}$ for any given spectrum t . Here, a_{tk} is the (t, k) -th element of the matrix \mathbf{A} of the new basis and $k = 1 : 2$. Fig. 6 shows the results of the correlations, and Fig. 7 graphically shows some of these correlations.

From Figs. 6 and 7 we can conclude that, for the ICA and NMF analyses, c_1 (and hence also c_2 , as $c_1 + c_2 = 1$) correlates with the proportion of young (\lesssim 1 Gyr) stellar population component present in the observed spectrum, regardless of their chemical abundance.

6 Conclusions

We have presented a scientific data mining application that searches for linear independent basis transformations of galaxy spectra to find the spectra of individual stellar populations characterised by age and chemical abundance. We have shown that characteristic stellar population components of elliptical galaxies can be disentangled from the observed spectra of these galaxies, without the use of detailed physical models. The components returned by the linear basis transformation analyses are clearly physically interpretable, with one component displaying the shape and many of the absorption-line features typical of a young stellar population, and the second component having the over-all shape and typical absorption features of an old, high chemical abundance stellar population. The weights of the contributions from the linear basis transformation analyses correlate well with both the ages of the younger stellar populations and the mass fractions of the smaller stellar populations determined from the (completely independent) detailed physical modelling of the observed galaxy spectra.

The computational demand of the projection approach presented is essentially that of the SVD computation, so it is expected that the method is easily applicable to large sets of measurements as they become available. The positivity based approach, per iteration, has a comparable scaling, however, in all our experiments the number of iterations to convergence was of an order of magnitude larger for the positivity based single

stage approach. Further study is necessary to investigate models with other types of positively supported priors as well as refining the best performing models and algorithms to be able to deal with previously unseen data.

The use of the data analysis presented in this paper, integrated with the more complex process of astrophysical analysis will be detailed elsewhere [16]. We intend to investigate the effectiveness of these data-driven methods on larger sets of UV-optical spectra as they become available, where more comprehensive statistical evaluation will be possible. From the analysis of large archives of galaxy spectra using this technique, we hope to address some of the fundamental questions in astrophysics, that of when and how galaxies form and evolve.

Acknowledgement

This research was partly supported by a Paul & Yuanbi Ramsay research award from the School of Computer Science of the University of Birmingham.

References

- [1] T.Cover and J. Thomas, Elements of Information Theory. John Wiley and Sons, Inc.,1991.
- [2] G.A. Darbellay and I. Vajda, Estimation of the information by an adaptive partitioning of the observation space. IEEE Trans. Information Theory, Vol. 45, no. 4, pp. 1315-1321, May 1999.
- [3] O.J. Eggen, D. Lynden-Bell, A.R. Sandage, Evidence from the motions of old stars that the Galaxy collapsed, *Astrophys. J.*, 136, 748, 1962
- [4] G.H. Golub and C.F. Van Loan, Matrix Computations. Johns Hopkins University Press, 1989.
- [5] A. Hyvarinen, Independent Component Analysis for Time-dependent Stochastic Processes, Proc. Int. Conf. on Artificial Neural Networks (ICANN'98) 1998, pp. 541-546.
- [6] A. Hyvarinen, Fast and Robust Fixed Point Algorithms for Independent Component Analysis. IEEE Transactions on Neural Networks 10(3): 626-634, 1999.
- [7] A. Hyvarinen, One-unit Contrast Functions for Independent Component Analysis: A Statistical Analysis. Proc IEEE Neural Networks for Signal Processing VII, 1997, pp. 388-397.
- [8] R. Jimenez, P. Padoan, F. Matteucci, A. Heavens, Galaxy formation and evolution: low-surface-brightness galaxies *Mon. Not. R. Astron. Soc.* 299, 123, 1998
- [9] M. Juvela, K. Lehtinen, P. Paatero, The Use of Positive Matrix Factorisation in the analysis of Molecular Line Spectra. *Mon. Not. R. Astron. Soc.*, 280 pp.616-626, 1996.

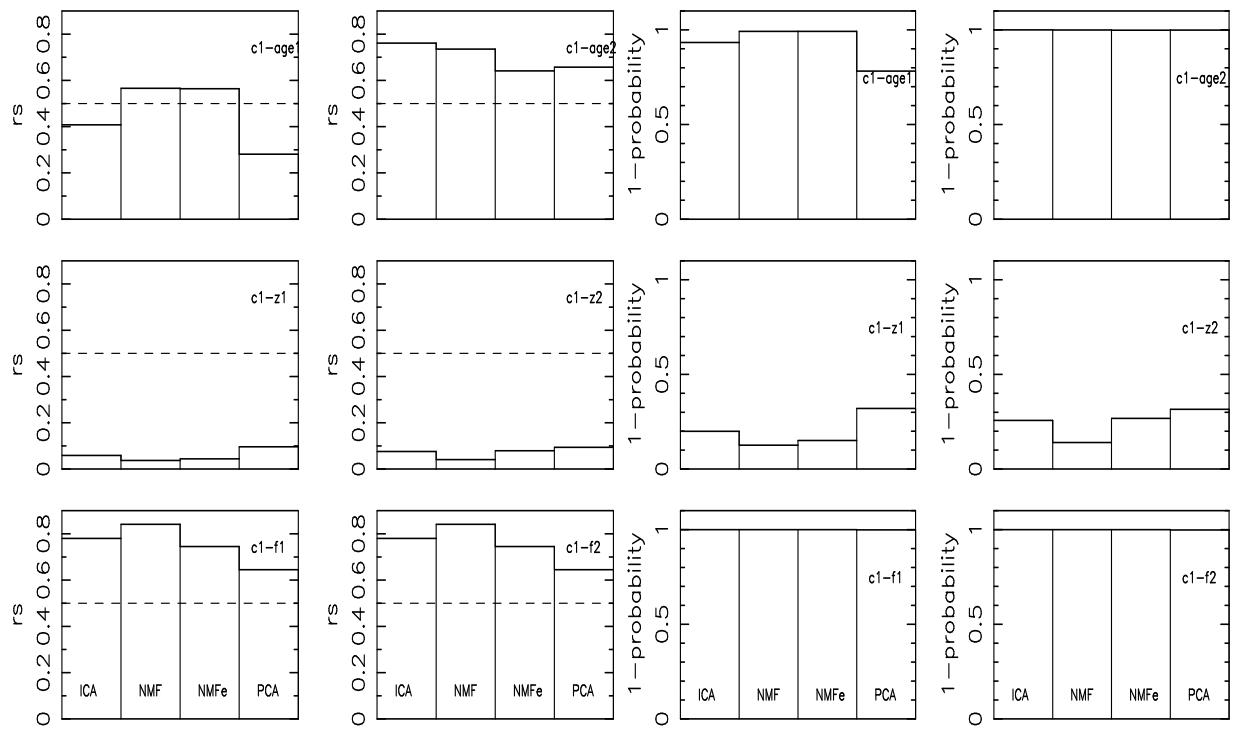


Figure 6: *Left*: The modulus of Spearman’s rank order correlation coefficient rs for the weight of the first component of the various linear basis transformation analyses, $c1$, correlated with the parameters recovered from the physical analysis i.e. a two-component spectral model fitting to the observed galaxy spectra. The parameters investigated in this correlation analysis are the age, chemical abundance z and relative mass fraction f for each component. Values greater than 0.5 indicate a strong correlation. *Right*: the significance of rs . High values of (1-probability) indicate a high level of significance.

- [10] G. Kauffmann, S. White, B. Guideroni, The Formation and Evolution of Galaxies Within Merging Dark Matter Halos, *Mon. Not. R. Astron. Soc.* 264, 201, 1993
- [11] R P Kraft, L A Nolan, T J Ponman, C Jones, S Raychaudhury, A Chandra observation of the nearby lenticular galaxy NGC 5102: where are the x-ray binaries? *Astrophys. J.*, submitted, 2004
- [12] R. B. Larson, Models for the formation of elliptical galaxies, *Mon. Not. R. Astron. Soc.* 173, 671, 1975
- [13] J. Miskin, Ensemble Learning for Independent Component Analysis. PhD Thesis. University of Cambridge, 2000.
- [14] T.-W. Lee, M. Girolami, A.J. Bell and T.J. Sejnowski, A Unifying Framework for Independent Component Analysis. *Computers and Math. with Applications*, vol. 39, no. 11, pp.1–21, 2000.
- [15] L.A. Nolan et al., The star-formation history of galaxies in the GEMS groups, in preparation.
- [16] L.A. Nolan, M.O. Harva, A. Kabán, S. Raychaudhury, Finding Young Stellar Population in Elliptical Galaxies from Independent Components of their UV-optical spectra, in preparation for submission to *Mon. Not. R. Astron. Soc.*
- [17] M. Tipping and C. Bishop, Probabilistic Principal Component Analysis, *Journal of the Royal Statistical Society, Series B*, 61, Part 3, pp. 611–622, 1999.
- [18] D. Lee and S. Seung, Algorithms for Non-Negative Matrix Factorisation, *Advances in Neural Information Processing Systems* 13, 556–562, MIT Press, 2001.
- [19] G. W. Stewart, Perturbation Theory for the Singular Value Decomposition. Appeared in *SVD and Signal Processing, II*, R. J. Vacarro ed., Elsevier, 1991.
- [20] H.A. Taha, *Operations Research – An Introduction*. 1997, Prentice-Hall.
- [21] O. Troyanskaya, M. Cantor, G. Sherlock, P. Brown, T. Hastie, R. Tibshirani, D. Botstein and R.B. Altman, Missing Value Estimation Methods for DNA Microarrays. *Bioinformatics* Vol. 17 no 6 2001.
- [22] S. D. M. White & M. J. Rees, Core condensation in heavy halos - A two-stage theory for galaxy formation and clustering, *Mon. Not. R. Astron. Soc.* 183, 341, 1978

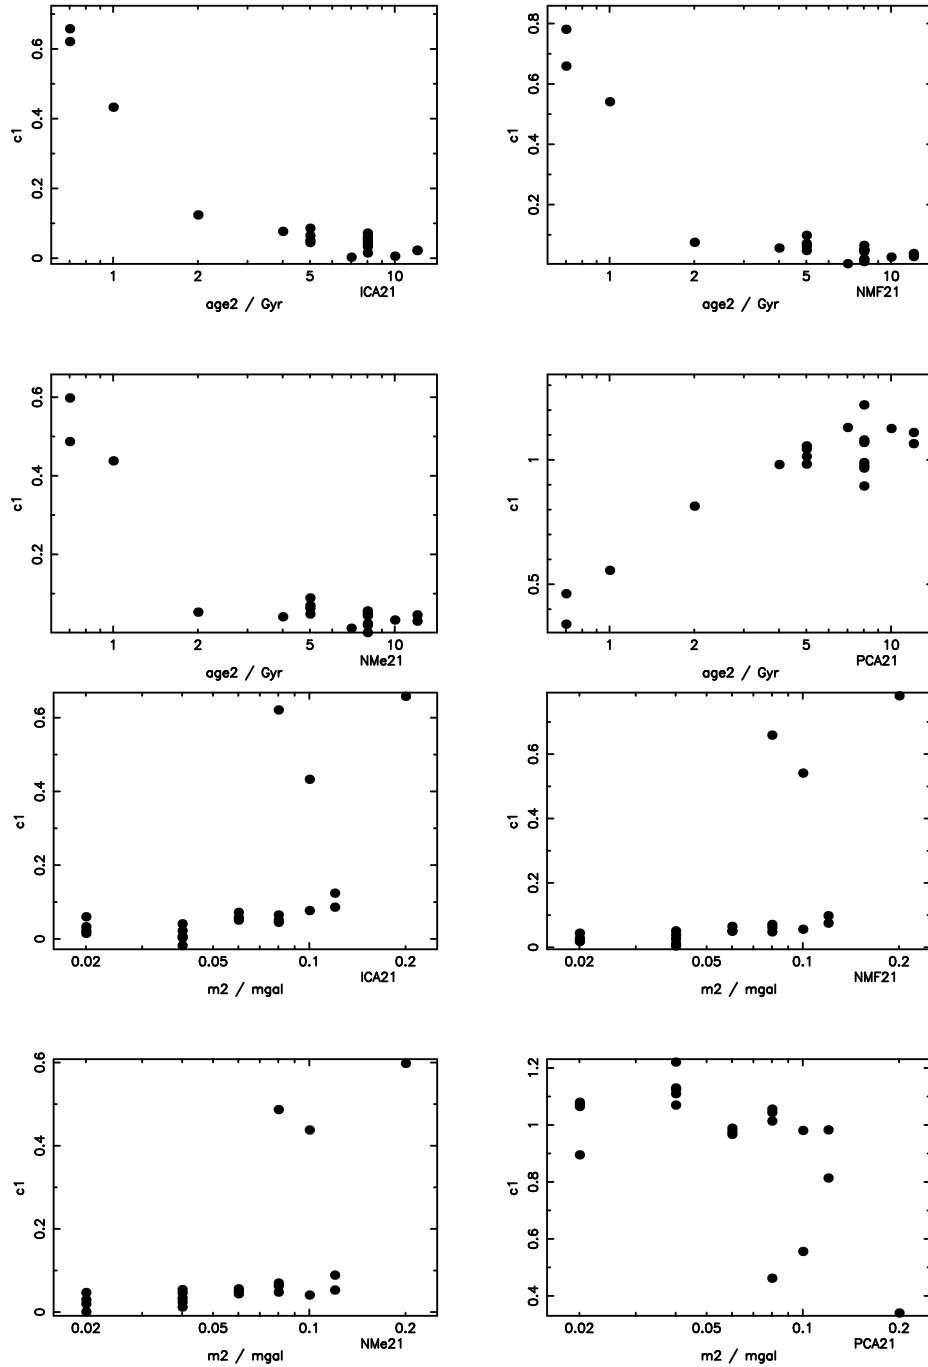


Figure 7: Scatter-plots showing the correlation of (*top*) the age of the younger stellar population and (*bottom*) the mass fraction of the smaller stellar population determined from the model spectra fitting with the weight of the first component of the various linear basis transformation analyses (c_1). A high value (low for PCA) of c_1 clearly corresponds to a substantial young ($\lesssim 1$ Gyr) stellar population.