Representation learning for automated spectroscopic redshift estimation

J. Frontera-Pons^{1,2}, F. Sureau¹, B. Moraes^{3,4}, J. Bobin¹, and F. B. Abdalla^{3,5}

¹ Laboratoire AIM, CEA, CNRS, Université Paris-Saclay, Université Paris Diderot, Sorbonne Paris Cité, 91191 Gif-sur-Yvette, France

e-mail: joana.frontera-pons@cea.fr

² DR2I, Institut Polytechnique des Sciences Avancées, 94200 Ivry-sur-Seine, France

³ Department of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, UK

⁴ Instituto de Física, Universidade Federal do Rio de Janeiro, 21941-972 Rio de Janeiro, Brazil

⁵ Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown 6140, South Africa

Received 21 September 2018 / Accepted 7 March 2019

ABSTRACT

Context. Determining the radial positions of galaxies up to a high accuracy depends on the correct identification of salient features in their spectra. Classical techniques for spectroscopic redshift estimation make use of template matching with cross-correlation. These templates are usually constructed from empirical spectra or simulations based on the modeling of local galaxies.

Aims. We propose two new spectroscopic redshift estimation schemes based on new learning techniques for galaxy spectra representation, using either a dictionary learning technique for sparse representation or denoising autoencoders. We investigate how these representations impact redshift estimation.

Methods. We first explored dictionary learning to obtain a sparse representation of the rest-frame galaxy spectra modeling both the continuum and line emissions. As an alternative, denoising autoencoders were considered to learn non-linear representations from rest-frame emission lines extracted from the data. In both cases, the redshift was then determined by redshifting the learnt representation and selecting the redshift that gave the lowest approximation error among the tested values.

Results. These methods have been tested on realistic simulated galaxy spectra, with photometry modeled after the Large Synoptic Survey Telescope (LSST) and spectroscopy reproducing properties of the Sloan Digital Sky Survey (SDSS). They were compared to Darth Fader, a robust technique extracting line features and estimating redshift through eigentemplates cross-correlations. We show that both dictionary learning and denoising autoencoders provide improved accuracy and reliability across all signal-to-noise (S/N) regimes and galaxy types. Furthermore, the former is more robust at high noise levels; the latter is more accurate on high S/N regimes. Combining both estimators improves results at low S/N.

Conclusions. The representation learning framework for spectroscopic redshift analysis introduced in this work offers high performance in feature extraction and redshift estimation, improving on a classical eigentemplates approach. This is a necessity for next-generation galaxy surveys, and we demonstrate a successful application in realistic simulated survey data.

Key words. methods: data analysis – techniques: spectroscopic – galaxies: distances and redshifts

1. Introduction

Galaxy redshift surveys are among the main observational tools to probe cosmological models. The leading methods measure the distance scale imprinted in the large-scale distribution of galaxies by oscillations in the primordial baryon-photon plasma (Kazin et al. 2014; Alam et al. 2017; Bautista et al. 2018). This baryonic acoustic oscillation (BAO) sound horizon can be used as a standard ruler to characterize the expansion rate of the Universe at different times, thereby providing constraints on cosmological parameters such as the total matter and dark energy densities (Blake & Glazebrook 2003; Seo & Eisenstein 2003). A precise measurement of the redshifts of galaxies is fundamental to extract this cosmological information from large galaxy surveys, and it is also key to the supplementary goals of constraining models of galaxy formation and evolution. To achieve these aims, most current and upcoming surveys such as the extended Baryon Oscillation Spectroscopic Survey (eBOSS) and the Dark Energy Spectroscopic Instrument (DESI) choose to observe the spectroscopic energy distribution (SED) of galaxies in the

optical or near-infrared wavelength range with multiplexed fiber spectrographs (Dawson et al. 2016; DESI Collaboration 2016).

Spectroscopic redshift estimation methods are typically based on the identification and fitting of spectral features – such as emission and absorption lines from electronic transitions in different elements – or of distinctive continuum features – such as the 4000 Å break due to the absorption of high-energy photons from metals in stellar atmospheres and the reduced number of hot blue stars in old galaxies (e.g., Baldry et al. 2004; Hutchinson et al. 2016). Especially for bluer, higher redshift galaxies without a prominent continuum break, one of the main hurdles for redshift estimation is the identification of relevant spectral features; high noise levels may introduce features that could be interpreted as physical lines if the analysis is too sensitive to noise, or true features might not be identified if attempts are made to mitigate false positive detections (Machado et al. 2013).

Another significant difficulty is ensuring that the spectral templates upon which many fitting methods depend are physically consistent and sufficiently representative of the observed

galaxy population in a particular survey (Bautista et al. 2018). Redshifts are generally determined by cross-correlation or χ^2 fitting between observed spectra and a reference set of spectroscopic templates (Glazebrook et al. 1998). These template-matching methods strongly rely on a catalog of galaxy spectra at zero redshift to which the unknown redshift galaxies will be compared, namely the template set. A template set can be prohibitively large if we wish to ensure correct retrieval of most of the significant features of an observed spectrum; additionally, there might be many degeneracies in the information spread throughout the full template set.

Former approaches exploit principal component analysis (PCA) to reduce the dimensionality of the problem and summarize the most relevant signal features in a set of principal components (most famously Glazebrook et al. 1998). One can then choose to retain a certain number of derived eigentemplates based on an “energy” metric that summarizes the amount of information captured. However, the resulting representation would only be efficient if the whole catalog shared common features (continuum, emission lines) that could be encoded efficiently with a few orthogonal templates. In other words, each eigentemplate used to represent a galaxy spectrum will probably contain a combination of features not specific to the galaxy we need to represent.

Moreover, modern large-scale surveys observe increasingly large data sets of galaxy spectra. In this context, although visual identification of the key features in the spectrum constitutes the most common method for validating galaxy redshift estimation (Hinton et al. 2016), the huge amount of data impels the development of robust and fully automated data-processing schemes to analyze the data and extract useful information such as the redshift.

We explore in this article unsupervised feature extraction from galaxy spectra through modern learning techniques. Notably, we investigate dictionary learning (DL) for sparse decomposition and denoising autoencoders (DAE) for spectra representation. Compared to PCA, dictionary learning with sparse representation is much more efficient to capture features that are not shared among the training data (such as combination of emission lines for instance) or different structures in the data (e.g., lines and breaks). It is therefore a good candidate for robust representation of structures specific to the tested spectra, leading to robust redshift estimation. Denoising autoencoders were selected for their ability to capture complex non-linear features present in the data, as already illustrated in Frontera-Pons et al. (2017). Ultimately, we exploit these new representations for spectroscopic redshift estimation. We assess the relative performance of the two resulting algorithms by comparing them with the redshift estimation code DARTH FADER (DF; Machado et al. 2013), based on cross-correlating estimated line features with eigentemplates learnt from a training set. We also investigate whether combining the results of both proposed estimators improves the performance in redshift estimation.

This paper is organized as follows. In Sect. 2, the DARTH FADER algorithm that is used for comparison is briefly recapped. The Sect. 3 is devoted to presenting dictionary learning for spectra representation and its application for redshift estimation. In Sect. 4, we detail the denoising autoencoder architecture and its corresponding redshift estimation scheme. Section 5 describes the simulated galaxy spectroscopic data used in the analysis. Section 6 describes the different code configurations and analyzes the results of the runs, comparing the performances of the different methods. Section 7 summarizes the results of the paper.

2. Redshift estimation with DARTH FADER

DARTH FADER is a robust redshift estimation code (Machado et al. 2013). Line features are firstly extracted from the galaxy SED, and then cross-correlated with eigentemplates as in traditional redshift estimation methods (Glazebrook et al. 1998). The robustness of the method comes from the robustness of the line feature extraction step, as well as a criterion to estimate redshift only when a sufficient number of line features are detected. This significantly improves redshift measurement performance and is one of the main advantages of DARTH FADER over the alternatives (Machado et al. 2013). Lines are estimated from the spectra using wavelet filtering and sparsity to remove continuum emission and represent line features in a galaxy SED (Machado et al. 2013). Wavelets are particularly suited for these tasks, given that measured SEDs are composed of a slowly varying continuum with mostly uncorrelated high-frequency noise and a few very sharp emission and absorption features. In the following we describe the main steps and features used in this approach to estimate the redshift.

2.1. Spectra modeling

To perform feature extraction and denoising, DARTH FADER assumes that it can model spectra as a combination of lines, noise, and continuum,

$$S = L + N + C, \quad (1)$$

where after continuum subtraction lines can further be broken down between emission and absorption lines, $L_e > 0$ and $L_a < 0$. It further assumes that line features are only important on small and intermediate wavelength scales and that the continuum possesses solely large-scale information. These assumptions are not rigorously true; in particular, strong lines can spread to larger scales and contribute to the continuum. Conversely, weak lines can be confused with noise in a low signal-to-noise (S/N) regime. In the next paragraphs, we describe how DARTH FADER deals with these issues.

2.1.1. Continuum subtraction

To subtract the continuum, DARTH FADER first identifies strong emission and absorption lines, and extracts them from the original spectrum using a pyramidal median transform (Starck et al. 1996). The reason for this choice of transform is that strong lines will be flux outliers compared to the continuum, hence the advantage of a median transform. Furthermore, it is a multiscale transform, which means that it filters features of varying widths. Applying the transform, outliers are identified and removed. The remaining spectrum is a good representation of continuum and noise. DARTH FADER then applies a starlet transform (Starck et al. 2015) to this representation to identify and remove the continuum. The starlet transform is a particular form of a wavelet transform – an undecimated isotropic wavelet transform – which decomposes the signal as follows:

$$S_\lambda = c_J(\lambda) + \sum_{j=1}^J w_j(\lambda), \quad (2)$$

where the w_j are the details at scale 2^{-j} . The largest scale coefficient, c_J , will be the best representation of the continuum, which can then be subtracted from the original SED.

2.1.2. Line feature denoising

Once the continuum is subtracted, the key step is to separate lines from noise. DARTH FADER employs sparsity constraints in a wavelet representation to achieve this goal in low S/N regimes. The shape of the spectral line features suggests that a particular choice of basis functions – in this case, a family of wavelets – will ensure that the decomposed signal contains at most a few significantly non-zero coefficients. Imposing this sparsity condition by minimizing a ℓ_1 norm on the wavelet-transformed line signal, and enforcing additional constraints on line and input wavelet coefficients, we can reconstruct the solution $\hat{L}$ that best matches the input data. More rigorously, we consider the following minimization problem to consecutively estimate emission and absorption lines:

$$\min_L \|\hat{W}L\|_1 \quad \text{s.t.} \quad S \in C, \quad (3)$$

where $\hat{W}$ is the wavelet transform operator, $\|\cdot\|_1$ is the ℓ_1 norm promoting sparsity – so that overall we look for a sparse solution in the wavelet domain, – and C is a convex set of constraints. We denote C the intersection of a set of positive (resp. negative) constraints on emission lines L_e (resp. absorption lines L_a) and a data fidelity constraint built as

$$\left| w_j^{[S]}(\lambda) - w_j^{[L]}(\lambda) \right| \leq \varepsilon_j, \quad \forall (j, \lambda) \in \mathcal{M}, \quad (4)$$

where $\mathcal{M}$ is the multiresolution support. This multiresolution support is first built by detecting significant wavelet coefficients at scale j and wavelength λ in the continuum-free spectrum using a prescribed threshold based on false discovery rate (FDR) for each wavelet scale (Starck & Murtagh 2006). This method ensures that, on average, false positives generated by noise will be kept at a chosen level, which has been shown to be a more efficient method for detection of features in low S/N data regimes than alternatives such as $K\sigma$ clipping.

2.2. Construction of eigentemplates and redshift estimation

To estimate redshift using these extracted line features, eigentemplates are first constructed from a training set of rest-frame line features by removing the continuum of a set of noise-free galaxy spectra as previously described. The tested spectra processed for line feature extraction as described above are then cross-correlated with these eigentemplates to derive a redshift estimate as in traditional redshift estimation methods (Glazebrook et al. 1998). For the purposes of providing benchmark results for the new methods developed in this paper, we will not preselect the galaxies by counting the number of line features in the spectra as was done in Machado et al. (2013), to illustrate the raw performance of the methods.

3. Redshift estimation with dictionary learning

The first learning technique we propose for redshift estimation relies on learning a representation for the full galaxy spectrum (i.e., continuum and emission lines) with a dictionary learning approach, assuming that spectra can be sparsely decomposed in such a dictionary (i.e., only a few atoms are used or similarly a few coefficients are non-zero for each decomposition). Estimating the redshift is then performed by finding the redshift where the sparse decomposition in this dictionary leads to the lowest approximation error. In the following, we first present

our motivation for using dictionary learning to represent galaxy spectra, then describe how we perform dictionary learning, and explain how such an adaptive dictionary can be used for redshift estimation.

3.1. Motivation

Dictionary learning techniques were proposed in the early 2000s (Olshausen & Field 1996; Engan et al. 1999; Aharon et al. 2006) and have since been applied to many restoration problems (e.g., Elad & Aharon 2006; Mairal et al. 2008, 2009; Zhang & Li 2010). Contrary to methods relying on PCA (e.g., Glazebrook et al. 1998; Machado et al. 2013) where template information is compressed in several orthogonal eigentemplates learnt from data or simulations, these techniques rather learn correlated templates assuming the observed spectra can be sparsely represented in a dictionary obtained from the data. Such techniques are therefore adapted to learn features (such as combination of emission lines for instance) or different structures in the data (e.g., lines and breaks) that are not necessarily common to all data but representative of a subset of it and are potentially correlated, whereas PCA rather extracts orthogonal features common to all data. This makes dictionary learning a good candidate for galaxy spectrum representation and redshift estimation, considering that it is unlikely that we can obtain a close sparse approximation of a tested spectrum if we redshift this learnt representation using an incorrect redshift value.

3.2. Dictionary learning for galaxy spectrum representation

To fix our notations, a spectrum $\mathbf{x} \in \mathbb{R}^{w_s}$ is approximated by a sparse decomposition $\mathbf{D}\boldsymbol{\alpha}$ in a dictionary $\mathbf{D} \in \mathbb{R}^{w_s \times n_a}$ with n_a atoms and with only a few coefficients of $\boldsymbol{\alpha}$ different from zero. The dictionary $\mathbf{D}$ is derived from a training set of n_t examples $\mathbf{X} \in \mathbb{R}^{w_s \times n_t}$ by solving the bilinear minimization problem

$$\underset{\mathbf{D} \in \mathcal{D}, \mathbf{A}}{\operatorname{argmin}} \|\mathbf{X} - \mathbf{D}\mathbf{A}\|_F^2 \quad \text{s.t.} \quad \forall i, \|\boldsymbol{\alpha}_i\|_0 \leq \tau, \quad (5)$$

where $\mathbf{A} \in \mathbb{R}^{n_a \times n_t}$ is the matrix containing the coefficients $\{\boldsymbol{\alpha}_i\}_{i=1 \dots n_t}$ as columns for each training example, $\|\cdot\|_F$ denotes the Frobenius norm, $\|\cdot\|_0$ counts the number of non-zero entries of a vector, τ enforces a targeted sparsity degree, and $\mathcal{D}$ designates the set of dictionaries with atoms in the unit ℓ_2 ball.

The training set used for learning can either be derived from real or simulated data. In practice, the critical point lies in the choice of a representative training set in terms of spectral variety and an observed wavelength range large enough to encompass the band of wavelength for testing the entire probed redshift range. The training set is then constructed by blueshifting these spectra to obtain rest-frame data that are used to learn a dictionary.

In the training phase, the joint non-convex problem described in Eq. (5) is typically handled by using an alternate minimization strategy, alternating sparse coding steps with dictionary updating steps as illustrated in Algorithm 1. The former is performed by minimizing Eq. (5) with respect to $\mathbf{A}$ for $\mathbf{D}$ fixed to its previous estimate. The latter corresponds to minimizing Eq. (5) with respect to $\mathbf{D}$ for $\mathbf{A}$ fixed to the previously estimated codes. Both steps can be achieved using standard algorithms. In this work, we will use the classical dictionary learning method of optimal directions (MOD) as detailed in Engan et al. (1999) with orthogonal matching pursuit (OMP; Mallat & Zhang 1993; Pati & Krishnaprasad 1993) as a sparse coder.

Because the problem described in Eq. (5) is non-convex, initialization of the dictionary is important so as not to obtain a non-meaningful local minimum. Furthermore, we would like to learn both continuum and line features on spectra to preserve as much information as possible (to be robust to noise and to reduce line confusion). In order to capture such a variety of features, we propose to use the procedure described in Algorithm 2 to initialize the learning algorithm. We first separate lines from continuum in our rest-frame training set by masking known line emission bands and extrapolating the resulting data in the region of the mask (step 1). A dictionary is then learnt for the line features and a second one for the extracted continuum (step 2). Finally, we concatenate the two dictionaries to initialize the learning procedure for the dictionary to represent both continuum and lines (step 3). The global dictionary is learnt, with a targeted sparsity degree τ given by the sum of the targeted sparsity degrees selected to derive the two sub-dictionaries.

Algorithm 1 Dictionary learning with MOD (Engan et al. 1999)

- 1: Initialization: Choose the number of atoms n_a , the targeted sparsity degree τ , initialize the dictionary. Choose the number of iterations N_{it} .
 - 2: **for** $n = 0$ to N_{it} **do** ▷ Main Learning Loop
 - 3: **for** $i = 1..n_t$ **do** ▷ Sparse Coding
 - 4: Compute the sparse code α_i using OMP with stopping criterion $\|\alpha_i\|_0 < \tau$
 - 5: **end for**
 - 6: Update $\mathbf{D}$ using MOD ▷ Dictionary Update
 - 7: **end for**
 - 8: **return** $\mathbf{D}$
-

Algorithm 2 Dictionary learning for galaxy spectra representation

Initialization step:

- 1: **Line/Continuum separation:** From the original training set $\mathbf{X}$, obtain two training sets: $\mathbf{X}_L$ for line features and $\mathbf{X}_C$ for continuum extraction.
 - 2: **Sub-dictionary learning:** Choose the number of atoms M_L (resp. M_C), a targeted sparsity degree τ_L (resp. τ_C), and a number of iterations N_L (resp. N_C). Use Algorithm 1 to learn a dictionary for lines $\mathbf{D}_L$ based on $\mathbf{X}_L$ and a dictionary for continuum $\mathbf{D}_C$ based on $\mathbf{X}_C$, with a dictionary initialized by randomly picking training examples.
 - 3: **Concatenation:** concatenate $\mathbf{D}_L$ and $\mathbf{D}_C$ to obtain a dictionary $\mathbf{D}_T$ with $M_L + M_C$ atoms.
- Dictionary Learning:**
- 4: Use Algorithm 1 to learn a dictionary $\mathbf{D}$ from the original training set $\mathbf{X}$, setting the number of atoms $M_L + M_C$ and the targeted sparsity degree $\tau_L + \tau_C$, with N_T iterations, with the initial dictionary $\mathbf{D}_T$.
- return** $\mathbf{D}$
-

3.3. Redshift estimation

Once the dictionary has been built, the redshift for the tested spectra can be estimated using a cross-matching procedure. For a tested redshift value, the atoms of the dictionary are redshifted and the best sparse decomposition is computed using the same targeted sparsity degree as in the training phase. Finally, for each spectrum, the redshift is chosen as the one providing the best sparse approximation among all the evaluated redshift values.

More precisely, for an observed spectrum $\mathbf{x}_z$, at a certain redshift z , our estimate is

$$\hat{z} = \underset{t \in \mathcal{T}, \alpha^{(t)}}{\operatorname{argmin}} \frac{\|\mathbf{x}_z - \mathbf{D}^{(t)} \alpha^{(t)}\|_2^2}{\|\mathbf{x}_z\|_2^2} \quad \text{s.t.} \quad \|\alpha^{(t)}\|_0 \leq \tau, \quad (6)$$

where $\mathbf{D}^{(t)}$ is the dictionary computed previously whose atoms have been redshifted by t accounting for the observed wavelength range, $\alpha^{(t)}$ are the corresponding sparse coefficients estimated using OMP with the same targeted sparsity degree τ as in the training phase, and $\mathcal{T}$ is a grid of tested redshift values that should be sufficiently finely sampled to typically avoid line confusion.

This approach therefore assumes that whenever the tested redshift is incorrect, a sparse decomposition in the redshifted dictionary cannot adequately approximate the signal, because all features captured in the dictionary do not match the observed spectrum. We note that the extreme case of a sparsity degree of one in the testing phase (which we will not use) would lead to a selection of the best matching atom in our dictionary.

4. Redshift estimation with denoising autoencoders

In this section, we investigate another type of representation for spectroscopic data built with a deep learning architecture, namely denoising autoencoders. Autoencoder architectures define a direct encoding function that transforms the input into a more suitable representation and a decoding function that reconstructs the corresponding input signal (Bourlard & Kamp 1988). More suitable representations preserve a significant amount of information, which enable reconstruction of the original signal. We detail in this section the application of denoising autoencoders for unsupervised line feature extraction in spectroscopic data. Ultimately, the learnt features will be used for redshift estimation.

4.1. Motivation

Recent advances in machine learning and deep learning techniques have shown their capability in solving supervised tasks. They provide state of the art results in classification for computer vision (e.g., Krizhevsky et al. 2012), speech recognition (e.g., Hinton et al. 2012; Dahl et al. 2012), natural language processing (e.g., Collobert et al. 2011), galaxy surface analysis (e.g., Tuccillo et al. 2017), and other applications. Moreover, representation learning methods have been praised as a powerful tool to derive unsupervised data-driven representations (Bengio et al. 2013). These methods allow us to design features that efficiently unfold complex underlying structures contained in the data. Unsupervised feature-extraction techniques such as denoising autoencoders have been successfully exploited and compared to PCA for galaxy SED representation in Frontera-Pons et al. (2017). In this work, the denoising autoencoders' ability to capture useful information, such as the redshift, has been highlighted, which motivates the study of this model for galaxy spectrum representation and redshift estimation.

4.2. Denoising autoencoders for template representation

In the classical autoencoder framework, the *encoder* f_θ provides the representation vector from the input galaxy spectrum $\mathbf{h}_i = f_\theta(\mathbf{x}_i)$, where $\mathbf{x} = [x_1, \dots, x_m]^T \in \mathbb{R}^m$ corresponds to the spectrum with only extracted line features for each galaxy in the considered population $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, $\mathbf{h}_i \in \mathbb{R}^{\text{nhid}}$ is the feature vector or

code, and $nhid$ is the number of hidden units or the dimension of the representation vector. Analogously, the *decoder* g_θ projects from the code space back into the input space, yielding a reconstruction of the original spectrum, $\hat{\mathbf{x}}_i = g_\theta(\mathbf{h}_i)$. More specifically, these functions are usually written as affine transformations, typically followed by a non-linearity, $f_\theta(\mathbf{x}) = s_f(\mathbf{b}_f + \mathbf{W}_f \mathbf{x})$ and $g_\theta(\mathbf{h}) = s_g(\mathbf{b}_g + \mathbf{W}_g \mathbf{h})$, where s_f and s_g are the encoder and decoder activation functions, and θ is the set of parameters that characterize the encoder and the decoder. Common options for the activation functions include the element-wise sigmoid, the hyperbolic tangent non-linearity, or the identity function, if staying linear, among others. The parameters $\mathbf{b}_f$ and $\mathbf{b}_g$ are the bias vectors of the encoder and decoder respectively, and $\mathbf{W}_f$ and $\mathbf{W}_g$ are the encoder and decoder weight matrices. Different weight matrices in the encoder and decoder are permitted in the architecture. However, weight-tying, in which one defines $\mathbf{W}_g = \mathbf{W}_f^T$, is most often adopted and so it will be assumed hereafter. Moreover, the bias vectors $\mathbf{b}_f$ and $\mathbf{b}_g$ have not been considered in this work to construct the representation.

The optimization of the parameters is performed to minimize the reconstruction error for the galaxy spectra, $L(\mathbf{x}, \hat{\mathbf{x}})$ over all the samples in the training population. Therefore, the cost function can be written according to

$$J_{AE}(\theta) = \sum_{i=1}^N L(\mathbf{x}_i, g_\theta(f_\theta(\mathbf{x}_i))). \quad (7)$$

This minimization is generally carried out by stochastic gradient descent. The choice of the reconstruction error measure $L(\cdot)$ depends on the input data domain range and nature. In other words, it is selected so that $L(\cdot)$ returns a negative log-likelihood for the observed value of $\mathbf{x}$. The mean squared error loss has been used for feature extraction, $L(\mathbf{x}, \hat{\mathbf{x}}) = \|\mathbf{x} - \hat{\mathbf{x}}\|^2$.

It is worth mentioning that with this configuration the basic autoencoders could learn the identity function to perfectly reconstruct the input. In order to avoid this trivial solution, some regularization constraints should be included during the training stage. Some studies, like Rifai et al. (2011) and Alain & Bengio (2014), underline the improvement brought by regularized autoencoders compared to the basic autoencoder framework. The purpose of this regularization is to render the representation invariant to local variations in the input.

In this article we focus on denoising autoencoders. In this case, the regularization by denoising makes the whole transformation robust and insensitive to small random perturbations in the input. Other variations could be explored such as under-complete representations that allow for a compression of the input data, or over-complete representations imposing sparsity on the code (Coates et al. 2011).

Denoising autoencoders were originally introduced by Vincent et al. (2008). In this approach, the training objective in Eq. (7) is modified to recover a clean input spectrum from an artificially corrupted version of it. Specifically, the cost function to be minimized becomes

$$J_{DAE}(\theta) = \sum_{i=1}^N \mathbb{E}_{q(\tilde{\mathbf{x}}|\mathbf{x}_i)} [L(\mathbf{x}_i, g_\theta(f_\theta(\tilde{\mathbf{x}})))], \quad (8)$$

where $\mathbb{E}_{q(\tilde{\mathbf{x}}|\mathbf{x}_i)}[\cdot]$ denotes the expectation over all the corrupted samples in the training population and J_{DAE} is optimized by stochastic gradient descent. Therefore, the recovered signal does not seek a perfect reconstruction of the original galaxy spectrum $\mathbf{x}$, but to retrieve the mean of the distribution that might generate $\mathbf{x}$. The different corruption processes discussed in

Vincent et al. (2008, 2010) involve additive Gaussian noise, salt and pepper noise, or masking noise. If some prior knowledge about the kind of perturbation the data might encounter is available, it can be incorporated into the corruption stage to make the model robust against this perturbation. Otherwise, the above-mentioned corruptors are useful in most scenarios. Moreover, the underlying structure and the information contained in the galaxy population have to be retained by the scheme in order to undo the effect of the corruption process, that is, to perform denoising. Good generalization of the model translates to a low reconstruction error for galaxies with similar characteristics to those in the training population, while yielding high reconstruction error for most other configurations.

4.3. Redshift estimation

The denoising autoencoders presented above are used to learn representations from rest-frame spectroscopic data. Then, the redshift is estimated for each spectra as the value providing the smallest reconstruction error from the model that has been redshifted in order to match the observed test spectra (Fig. 1). Specifically, the model is built from a catalog of galaxy spectra at zero redshift, named the training set as for the dictionary learning framework. The denoising autoencoder architecture is optimized in order to minimize the reconstruction error for the samples in the training set. These samples have to be representative of the expected galaxy population in the analysis.

After training, for every redshift value evaluated the model parameters are redshifted. Thereupon, the tested spectra are projected to the representation space through the encoder function defined by the denoising autoencoder and projected back to the input space with the decoder, leading to an approximation of the input signal. The redshift is estimated as the value minimizing this approximation error. In other words, we hope that, when the discriminating features of the test spectra are aligned with their rest-frame counterparts used for the training stage, the model will be able to reconstruct the input signal with a small error while yielding a large reconstruction error in any other case. The error is computed using an Euclidean metric, and for each test sample the redshift is obtained according to

$$\hat{z} = \underset{t \in \mathcal{T}}{\operatorname{argmin}} \frac{\|\mathbf{x}_z - g_{\theta^{(t)}}(f_{\theta^{(t)}}(\mathbf{x}_z))\|_2^2}{\|\mathbf{x}_z\|_2^2}, \quad (9)$$

where $\theta^{(t)}$ denotes the denoising autoencoder model redshifted at t and $\mathcal{T}$ is the grid of tested redshift values. In other words, the columns of the weighting matrix W are redshifted and treated similarly as the atoms in a dictionary described in Sect. 3.

5. Data

In order to assess how these two new methods perform for spectroscopic redshift estimation in a realistic setting, we wish to use a simulated data set consisting of a combination of photometric and spectroscopic data mimicking modern galaxy surveys. We require that this data obey the following constraints:

- (i) There is a realistic distribution of galaxy types, photometric properties, and redshifts corresponding to an idealized selection function for a state-of-the-art photometric galaxy survey.
- (ii) Each galaxy is consistently matched to a corresponding SED from a template library containing realistic continuum, emission, and absorption features. The matching must ensure that the integrated flux through broadband filters corresponding

Fig. 1. Example of test set spectra with various S/N (measured as in Darth Fader paper) in the white Gaussian noise scenario.

to the original photometric sample is consistent with the original observations.

(iii) The SEDs are resampled, integrated, and noise is added to simulate realistic spectra from existing or planned galaxy spectroscopic surveys.

To generate a realistic distribution of galaxies with simulated redshift values, SEDs, and broadband photometry, we employ the COSMOSSNAP simulation code (Jouvel et al. 2009). COSMOSSNAP uses real data from the 30-band COSMOS photometric redshift catalog as a basis (Ilbert et al. 2008), thereby ensuring that realistic relationships between galaxy type, color, size, and redshift are taken into account. The catalog was originally generated from a combination of observations from astronomical surveys covering the spectral range from the UV (Galaxy Evolution Explorer), through the optical (Subaru) and to near- and far-infrared bands (CFHT, UKIRT, *Spitzer*). This data set is matched to *Hubble* Advanced Camera for Surveys imaging data, thus including realistic size-magnitude distributions from high-quality shape measurements originally made for weak lensing applications (Leauthaud et al. 2007).

From this seed catalog, and assuming that the measured photometric redshifts are “true” redshifts, COSMOSSNAP can create simulated catalogs for any broadband photometric survey. The procedure is as follows: based on each galaxy’s properties, COSMOSSNAP chooses a spectral template from a predefined library, such that the integrated fluxes through the 30 broadband filters provide the best fit to the observations. It uses the Coleman Extended library, which includes four spectral types – Elliptical, Sbc, Scd, and Irregular (Coleman et al. 1980). It extends the spectral range into the UV and IR using synthetic spectra from the Galaxy Isochrone Synthesis Spectral Evolution Library (Bruzual & Charlot 1993) and adds an extra fifth type to represent starburst galaxies. To add realistic spectral features atop the original templates, galaxy emission line fluxes are calculated based on continuum properties of each galaxy. From the UV rest-frame luminosity of a given galaxy, a star formation rate is inferred using a calibration from Kennicutt (1998). This is then translated into an [OII] line flux, a relation which is valid for different galaxy types. Additional emission line fluxes are calculated relative to the [OII] flux, based on observations

(Moustakas et al. 2006). The final SED of each galaxy is then corrected by host extinction (i.e., dimming due to dust within the galaxy itself, estimated from the photometric properties) and redshifted following the best-fit photometric redshift value.

At the end of this process, each galaxy has a “true” redshift and its associated SED model. Given an arbitrary choice of broadband filter and a model of the full filter throughput – including atmospheric transmission, telescope optical effects, and more – the SED can be integrated to calculate simulated noiseless magnitudes. A realistic two-component model of magnitude errors with tunable observational properties is applied for each galaxy in the catalog. The resulting magnitude and error distributions can reproduce closely those of current and future large-scale galaxy surveys. For this analysis, we decide to create a photometric catalog modeled after the expected throughput of the six Large Synoptic Survey Telescope¹ (LSST) broadband filters commonly referred to as “ugrizY”, and with the expected depth properties of the science-ready “Gold” sample (Abell et al. 2009). Hence we exclude from the catalog galaxies fainter than 25.3 AB magnitude and with $S/N < 10$ in the *i*-band. We obtain 218966 galaxies in an effective area of 1.24 deg² with realistic photometric properties, together with best-fit spectral templates with realistic continuum and emission line properties.

COSMOSSNAP produces SEDs with a chosen wavelength resolution for the continuum and absorption lines. The emission lines are added at higher resolution, to ensure that their shapes and amplitudes are fully characterized. To work with realistic spectra, we need to resample, integrate, and add noise to the best-fit SEDs. On a real fiber-fed spectrograph such as the ones designed for the BOSS (Smee et al. 2013) and DESI (DESI Collaboration 2016) surveys, the resolution is a variable property that depends on the characteristics of the instrument, in particular on the interplay between the 1D point spread function full width at half maximum of the spectrograph and the pixel size of the CCD. Noise on the 1D spectra is mainly due to Poisson sampling of photons from the source and CCD readout noise, among other effects. For our purposes of evaluating the performance of redshift algorithms, we assume a constant resolution

¹ <http://www.lsst.org/>

$R \equiv \lambda/\Delta\lambda$ – which implies a wavelength binning constant in logarithmic scale – and uncorrelated Gaussian noise with constant variance σ on all wavelength bins. The SEDs are log-binned with a constant bin size of $2.17 \times 10^{-4} \log 10 \text{ \AA}$ – corresponding to a resolution of $R \sim 850$ – and integrated within those bins. We add noise of different S/N levels, where the level is defined according to the spectral energy flux integrated in the r -band², following Bolton et al. (2012) and Machado et al. (2013):

$$S/N_r = \text{median} \left[\frac{\text{flux}}{\sigma} \right]_{5600 \text{ \AA}}^{6760 \text{ \AA}}. \quad (10)$$

We work with $S/N \in \{2, 5, 20\}$, where $S/N = 20$ is our “clean” case.

For the remainder of this paper, we work with a training data set that includes 2000 clean resampled spectra randomly sampled from the original COSMOSSNAP data from a redshift range $[0, 1.7]$. These training SEDs are blueshifted to the rest frame and cropped to the wavelength range $[1250 \text{ \AA}, 10\,500 \text{ \AA}]$. For testing, we randomly sample galaxies of different types. COSMOSSNAP classifies its spectra in 36 classes, organized in four groups: E11S0, SaSc, SdSm, and SB. We build a test set by randomly sampling 5000 galaxies from each of the groups to investigate potential systematic effects in our methods depending on the galaxy type. This galaxy type information was not included in the training phase. All test galaxy spectra are cropped to the wavelength range $[3000 \text{ \AA}, 10\,500 \text{ \AA}]$.

6. Experimental results

Let us now present the results for redshift estimation obtained with the proposed dictionary learning and denoising autoencoder representation learning frameworks. These two approaches have been compared to the DARTH Fader algorithm (Machado et al. 2013), based on robustly extracting line features and cross-correlating them with eigentemplates to infer a redshift. We start by describing how we set this algorithm up for this comparison and then describe the parameters for the two proposed methods.

6.1. DARTH Fader

We run DARTH Fader on the 4×5000 test galaxy spectra for all S/N levels. Our configuration choices mostly follow the standard setup. We briefly describe them now, and refer the reader to Sect. 2 and to Machado et al. (2013) for more details. We derive eigentemplates from the clean training data set described in Sect. 5. We set the threshold for the principal components so that the eigentemplates retained contain 99.93% of total eigenvalue weight, as in Machado et al. (2013). We keep 26 eigentemplates as a result. If the preserved percentage of the variance and the number of retained eigentemplates are not sufficient for reconstruction, the performances of the redshift estimation scheme can degrade. On the other hand, if all the energy was intended to be retained by the representation, a larger number of eigentemplates may yield a drop in performance for noisy scenario. According to our experiments, these factors worsen the redshift estimation but not significantly. For denoising the test spectra, we restrict the multiscale transform to six scales and keep the regularisation at 0.01. For redshift estimation, we cross-correlate the eigentemplates with the denoised version of the test

spectra to avoid the misclassification of noise features as spectral lines, even though denoising may result in removing physical features in low S/N regimes. Contrary to what is recommended for optimal use of DARTH Fader, we do not clean the resulting catalog with FDR thresholding. Although this results in sub-optimal performance metrics, we wish to compare the performance of the algorithms on the full galaxy set, and not only on those relatively few galaxies where DARTH Fader successfully retrieves many features.

6.2. Dictionary learning

As illustrated in Algorithm 2, the first step in our analysis involves constructing a meaningful initial dictionary for the subsequent learning phase. This implies constructing two sub-dictionaries for line features and continuum, and therefore the separation between line features and continuum emissions in the rest-frame training set. The line features have been removed with a mask centered on known wavelength value for each emission line, and extended to a width of $80 \text{ \AA}$ in order to ensure that no emission line energy is leaking into the continuum. Then, a multiscale iterative inpainting technique is used to extrapolate the continuum inside this region: we iteratively keep only low-frequency scale coefficients in the mask while enforcing values outside the masked region not to be affected by the procedure. Some statistics on the separation are summarized in Fig. 2, which illustrates its good performance.

A “continuum” template dictionary of 40 atoms with a targeted sparsity degree of three and 100 iterations was learnt using the training data set with continuum features. Similarly for the set with line features, a “line” template dictionary of 20 atoms with a targeted sparsity degree of three and 100 iterations was learnt. The number of atoms and sparsity degree were heuristically fixed from the overall complexity observed in the training data (using more atoms leads to high correlation preventing the learning) and as a trade-off between estimation error, efficient learning, and robust subsequent redshift estimation. Indeed, even though increasing the sparsity degree would decrease the overall approximation error, it would also increase the risk of line confusion by potentially selecting as atoms more isolated features in the spectra.

The atoms learnt for the two sub-dictionaries are represented in the upper panels of Fig. 3. Several things can be understood from this figure: first, in both cases, the atoms are correlated contrary to what would have been obtained with PCA; second, in particular because of the small targeted sparsity degree, the atoms learnt for lines all contained a combination of line features, which would help to avoid line confusion; third, the $4000 \text{ \AA}$ break feature is captured by most but not all atoms in the continuum sub-dictionary. Learning was finally performed on the original training data set by initializing the dictionary with the 60 atom dictionary obtained by concatenating the two previous ones, using ten iterations and a sparsity degree of six. The two components of the final dictionary are displayed in the lower panels of Fig. 3, showing that some leakage between the two components has been introduced during the final learning phase, which was actually beneficial to reduce approximation errors. This also leads the continuum part to contain more high-frequency features, which could result in better redshift estimation by combining several correlated line and continuum features in the atoms. For spectroscopic redshift estimation, the redshift grid $\mathcal{T}$ tested was built by uniformly sampling the range from 0 to 1.7 with a density of 0.001 and sparse coding was performed for all galaxy spectra at the given sampled redshifts.

² The LSST r -band filter is not identical to the SDSS one. Therefore, our definition of S/N is not identical to that in Machado et al. (2013). This does not have any impact on our analysis and conclusions as they are self-consistent.

Fig. 2. Statistics on extracted line (*left*) and continuum (*right*) features. The input median signal in the original training set is displayed, along with the median and standard deviation computed over the line (*left*) and continuum (*right*) training data sets. In the first case, the emission lines are well extracted, with signal variability only in regions of emission or absorption lines and without offset due to continuum leakage into line estimation. In the second case, the continuum does not display either discontinuities or acute signal variability in regions of strong emission or absorption lines.

Fig. 3. *Top left:* dictionary learnt from training data set with line features. Atoms typically contain a mixture of line emissions, with different sign and resolution. *Top right:* dictionary learnt from training data set with continuum features. Atoms are correlated, with breaks such as the 4000 Å break. *Bottom left:* component of the final dictionary corresponding to the original line dictionary. Some fluctuations from the continuum are added. *Bottom right:* component of the final dictionary corresponding to the original continuum dictionary. Extra line emissions are added to the continuum.

6.3. Denoising autoencoders

Firstly, the model was constructed using TensorFlow (Abadi et al. 2016). TensorFlow is an open source software library for numerical computation using data flow graphs. It was developed by Google and tailored for machine learning. Extensive documentation describing all the functionalities can be found on the website below³.

³ <https://www.tensorflow.org/>

The denoising autoencoder was trained with the training set described above composed of $N = 2000$ samples spanning from 1250 Å to 10 500 Å resulting in a vector of dimension $m = 4258$. The training samples followed the same preprocessing as in the DARTH Fader scheme for continuum line subtraction, and the continuum of the spectra has been removed through wavelet filtering as explained in Sect. 2.1.1. The number of visible units is fixed to $m = 4258$ in agreement with the size of the input data and the number of hidden units is a free parameter.

Fig. 4. Weight filters learned by the DAE for $nhid = 100$ and $S/N = 5$.

The representation size has to be large enough in order to retain sufficient information for reconstruction. As the representation size increases, the approximation error will decrease. This improvement will be significant only for the galaxy true redshift and will not impact on its estimation. As the computation time for the redshift estimation algorithm is proportional to $nhid$, we chose a value providing a sufficiently reliable reconstruction of the original signal for the redshift estimation, but greater values could be used instead. We have investigated $nhid = 20, 100, 200, 1000$ and chosen an architecture with $nhid = 100$. Smaller values did not provide a useful reconstruction of the spectra and greater values did not improve the redshift estimation step. Thus, this choice is a trade-off between approximation error and computation time.

The weights are randomly initialized from a uniform distribution and are tied across all the experiments. Moreover, the activation function for both the encoder and the decoder is set to be a hyperbolic tangent. The input is artificially contaminated with a Gaussian noise such that the S/N will be constant and equal to five through all the training set. Furthermore, the optimization of the parameters is performed through stochastic gradient descent; and the reconstruction cost criteria to be minimized is the squared reconstruction error. The learning rate has been set to 10^{-4} and the batch size to 100. The training stops after 500 epochs. The choice of the learning rate, the batch size, and the number of epochs is done in order to ensure a convergence on the training stage. Figure 4 displays the encoder weights learned by the denoising autoencoder. This illustrates what kind of features the model is sensitive to and gives an idea about how the input data are coded. The earmarks highlighted by the filters are in agreement with the position of the absorption and emission lines dominating the training samples. Moreover, Fig. 5 displays one sample belonging to the training set (a), its representation (b), and its corresponding reconstruction (c). From the way the information is coded in Fig. 5, it is hard to give a straightforward interpretation. Due to the mixing performed by the encoder, the information is distributed over all the code and nothing can be said about hidden units individually.

Let us now illustrate the results for redshift estimation with the denoising autoencoders. Figure 6 illustrates some approximation error profiles over all the investigated redshift values. The redshift has been uniformly sampled from 0 to 1.7 with 0.001 steps. From this figure, it is clear that the proposed method correctly finds the redshift in the studied scenario. Moreover, the technique described in this work is highly sensitive to the true redshift value and its precision depends on the redshift sampling grid. However, some strong oscillations due to the matching of the different features present in the spectra show

that some line confusion could occur in the subsequent redshift estimation.

6.4. Comparison of results

In this section, we compare the results of all three methods, paying special attention to dictionary learning and denoising autoencoders. These methods improve performance in all S/N regimes and for all galaxy types when compared to Darth Fader, but each has its own advantages and drawbacks.

Figure 7 compares the estimated to true redshift values for all three methods in three different S/N regimes. For $S/N = 20$, all methods perform well, with small dispersion around the truth and very few catastrophic outliers. Nonetheless, there are qualitative differences in the distribution of these outliers. Dictionary learning encounters small difficulties with most types of galaxies, in particular at intermediate redshifts. Darth Fader and DAE, on the other hand, show a common pattern of secondary linear features, which is due to feature confusion: when an insufficient number of features is present, the redshift estimation process confuses between features in a predictable matter. For example, an $H\alpha$ hydrogen transition line can be misidentified as an OII oxygen line. The phenomenon is stronger with Darth Fader, which in addition has another cluster of outliers at low-estimated/high-true redshift.

At $S/N = 5$, Darth Fader performance is significantly degraded. DAE and DL outlier rate increases slightly, but the patterns of errors remain the same for them. However, DAE starts showing a similar cluster of outliers that were already present in Darth Fader's $S/N = 20$ results. This trend is even stronger when DAE reaches $S/N = 2$. There are striking similarities between its results in this regime and Darth Fader's at $S/N = 5$, seeming to indicate that both algorithms are reacting to the same underlying patterns in the data. The situation is markedly different for dictionary learning. In the $S/N = 2$ regime, the number of outliers increases but they still follow a pattern of clustered groups of outliers relatively close to the true redshift values.

Figure 8 investigates the distribution of catastrophic outliers per redshift range and galaxy type for the same three S/N cases. Failure rates are clearly dependent on these variables. In particular, galaxy type strongly influences the performance of the algorithm. This is to be expected: the SED of starburst galaxies, for example, typically includes very strong emission lines due to intense star-formation activity, whereas the SED of elliptical galaxies includes a strong break on the 4000 Å rest-frame continuum as its most prominent feature. At $S/N = 20$, most DF failures take place at the highest redshifts, for most types. DAE demonstrates almost perfect performance, while DL has a relatively low rate of failure, mostly concentrated on redshifts larger than one. Ellipticals and lenticulars (E/LS0) at intermediate redshifts are a curious case: Darth Fader and – in particular – DAE maintain a low rate of catastrophic outliers for this subclass. However, DL shows bad performance for this subclass, even at high $S/N = 20$. What we observe for this type of galaxy and this redshift bin is confusion of features in the continuum in the representation learnt via dictionary learning. Indeed, these galaxies represent only 5% of the training set, dominated by starbursts (SB) galaxies (about 86%), so learning is mainly driven by approximation of SB galaxies rather than elliptical galaxies. Errors in modeling the continuum of E/LS0 are expected therefore to be larger, which could translate to incorrect redshift estimation if this affects discriminative features for redshift estimation. Furthermore, for this redshift bin, the 4000 Å break is close to the end of the observable range, and confusion of

Fig. 5. Panel a: training sample. Panel b: representation. Panel c: reconstructed spectra.

Fig. 6. Approximation error profiles as a function of the redshift for four galaxies of the test set. The true redshift is displayed as a dotted vertical line for each galaxy.

features arise with other discontinuities in the continuum located before or after this break.

For $S/N = 5$ and $S/N = 2$, both DAE and DL markedly outperform Darth Fader in all galaxy classes. Comparison between DAE and DL methods is more ambiguous. DAE achieves excellent results for all galaxy classes in higher S/N regimes, as seen in the top and middle panels of the middle column in Fig. 8; catastrophic outlier rates are nearly zero everywhere. However, performance rapidly degrades when S/N decreases, as seen in the lower panel. Moreover, comparison to the middle lower panel of Fig. 7 shows that those failures are random within a band of estimated redshift, which signals a pathological behaviour. DL, on the other hand, keeps a consistent pattern of failures, maintaining lower than 20% catastrophic outlier rates for most galaxy types and redshift bins. In that sense, both methods are complementary. In terms of galaxy types, all methods perform best with starburst galaxies; even in the low $S/N = 2$ regime, Darth Fader still succeeds in keeping the catastrophic outlier rate of this class below 20% for redshifts smaller than $z = 1$.

Table 1 presents a quantitative overview of the catastrophic outlier rates. In addition to the total success rates for each algorithm in the different S/N regimes, we also investigate the partial rates when different combinations of algorithms succeed in measuring precise redshifts. We see, for example, that it is exceedingly rare for Darth Fader to succeed when both dictionary learning and denoising autoencoders fail. Even if only one of them succeeds, it is still quite rare for Darth Fader to succeed also. In other words, once we have DL and DAE results, DF results are superfluous; they measure accurate redshifts mostly when the other two methods also do. Focusing on the $S/N = 2$ case, we see the potential advantage of combining the methods. The total success rate when both DL and DAE succeed is 71.4% (i.e., $(5203 + 9082)/20\,000$), which is already significant

for a low S/N regime. However, if we identify a way of selecting the best method when one or the other succeeds, we can reach $18\,739/20\,000 = 93.7\%$ success rate. In the next subsection, we will develop an algorithm to select the best redshift possible in each case and study its accuracy.

We now turn to investigating the dispersion of redshift estimates around the true values after excluding catastrophic outliers from consideration. Figure 9 compares the dispersion distribution of DAE and DL for all galaxy types in the $S/N = 5$ and $S/N = 2$ cases. Yet again, the presence of sharp features clearly influences the performance of both methods, with dispersion of estimated redshifts decreasing from elliptical to starburst galaxies. Additionally, DAE is more precise (i.e., smaller σ) than DL (except in the particular case of ellipticals at $S/N = 5$, where performance is comparable). As could be expected, overall dispersion of the estimated values is higher for $S/N = 2$ than $S/N = 5$, but this is the only quantitative difference between the two noise levels. In summary, inclusion of continuum for DL allows us to obtain a more consistent redshift estimation when noise increases compared to DAE (lower confusion), but adding this essentially low-frequency information also degrades the precision in estimating the redshift, since modeling errors on continuum (giving low precision in redshift estimation) may dominate over modeling errors of line features (giving high precision in redshift estimation).

The results described in this section suggest a clear strategy for leveraging the strengths of the different methods: for cases with high S/N in the continuum, DAE redshifts are more precise and contain less catastrophic outliers. With lower S/N , DL remains more robust and should be preferred. These results may depend on the noise characteristics, among other SED properties, and should be re-evaluated for each separate application.

6.5. Defining a “best” redshift from a DAE/DL combination

As we discussed in the past subsections, the DAE and DL algorithms show complementary performance, indicating that a method for combining their results based on observational properties can increase the accuracy of redshift estimation. In this subsection, we devise an algorithm to take advantage of their strengths, and assess the performance of the resulting redshifts. The two main observational properties that we consider are the estimated redshifts from each method and the galaxy types. The latter are not an observational property per se. Nevertheless, a broad division in four types, such as the one we are using, can be approximated by color cuts in broadband magnitudes from the optical galaxy targeting surveys that serve as a base for spectroscopic surveys. We will postpone a more realistic analysis of this particular aspect to future work.

Fig. 7. True versus estimated redshift values for DARTH Fader (*top row*), dictionary learning (*middle row*), and denoising autoencoders (*bottom row*) in the $S/N = 20, 5$, and 2 cases (*left, middle, and right columns*, respectively). Galaxies are divided by type: ellipticals and lenticulars (ELISO, blue circles), regular spirals (SaSc, orange squares), broken/irregular spirals (SdSm, green triangles), and starbursts (SB, red crosses).

Algorithm 3 Method to choose a best redshift estimate z_{best} .

```

for all galaxies do
  if DAE/DL redshifts agree then
    Take DAE  $z_{\text{est}}$ .
  else
    for DAE/DL methods do
      Associate galaxy to a type/ $z_{\text{est}}$  bin.
      Get catastrophic outlier rate  $f_c$  in that bin.
    end for
    Take  $z_{\text{est}}$  of method with lower  $f_c$ .
  end if
end for

```

The method we propose is described in Algorithm 3. For each galaxy, we assess whether DAE and DL estimated redshifts agree within a precision threshold. We define it as $\Delta z = 0.003$, which is more strict than the catastrophic outlier rate threshold due to the absence of a true-redshift dependance. Whenever those values agree, we choose the DAE redshift, which has been shown to have smaller dispersion around the true redshifts. If the redshift values do not agree, we resort to using true redshifts to define catastrophic outlier rates in estimated redshift and type bins. For each galaxy, we locate it in a redshift-type bin, and then choose the method for which this bin has a lower catastrophic outlier rate. The catastrophic

Fig. 8. Catastrophic failure rate (measured as $|\Delta z/(1+z)| > 0.003$) by galaxy type and redshift bins, Darth Fader (*top row*), dictionary learning (*middle row*), and denoising autoencoders (*bottom row*) in the $S/N = 20, 5$, and 2 cases (*left, middle, and right columns*, respectively). Performance degrades as noise levels increase for most galaxy types and redshifts, especially in Darth Fader’s case. Denoising autoencoders perform particularly well in the high S/N regimes, while dictionary learning is more stable across S/N regimes.

outlier rates are defined with estimated, not true, redshifts in those cases, meaning that there will be two different binning schemes.

Figure 10 shows the results of applying this algorithm to the three different galaxy samples. Compared to the results of the individual methods shown in Fig. 7, the improvement is manifest. In the $S/N = 20$ and $S/N = 5$ cases, where DAE redshifts are very accurate, only a handful of galaxies is selected with DL, which brings marginal improvements to the global catastrophic outlier rates. In the $S/N = 2$ case, however, the improve-

ment is non-negligible. Of the galaxies, 71.4% have redshifts in agreement with each other. After combination, we obtain a global success rate of 93.15%. If we had only adopted DL redshifts, the global success rate would have been 89.1% – a few percentage points lower – with the larger DL scatter. If instead we wanted to use DAE redshifts to retain a lower scatter, we would be restricted to only 77.7% of the sample. Figure 11 shows the catastrophic outlier rates for each galaxy type and redshift bin for each S/N . While $S/N = 20$ and $S/N = 5$ are mostly equivalent to DAE results, the $S/N = 2$ figure shows

Table 1. Success rate of the redshift estimation algorithms in the different S/N cases, where failure is defined as $z_{\text{est}} - z_{\text{true}} > 0.003(1 + z_{\text{true}})$.

$S/N = 20$				$S/N = 5$				$S/N = 2$			
Total DF success: 19479 (97.4%)				Total DF success: 13125 (65.6%)				Total DF success: 5586 (27.9%)			
Total DL success: 18668 (93.3%)				Total DL success: 18694 (93.5%)				Total DL success: 17822 (89.1%)			
Total DAE success: 19954 (99.8%)				Total DAE success: 19880 (99.4%)				Total DAE success: 15545 (77.7%)			
Total “best” success: 19945 (99.7%)				Total “best” success: 19895 (99.5%)				Total “best” success: 18630 (93.1%)			
DF	DL	DAE		DF	DL	DAE		DF	DL	DAE	
✓	✓	✓	18185	✓	✓	✓	12238	✓	✓	✓	5203
✓	✓	×	14	✓	✓	×	18	✓	✓	×	87
✓	×	✓	1264	✓	×	✓	856	✓	×	✓	256
×	✓	✓	455	×	✓	✓	6357	×	✓	✓	9082
×	×	×	16	✓	×	×	13	✓	×	×	40
×	✓	×	14	×	✓	×	81	×	✓	×	3450
×	×	✓	50	×	×	✓	429	×	×	✓	1004
×	×	×	2	×	×	×	8	×	×	×	878

Notes. The “best” algorithm combines DL/DAE results as described in Sect. 6.5 and shown in Fig. 10. The lower half of the table shows the partial statistics for different combinations of algorithm success; for example, the first line counts the number of galaxies for which the three algorithms succeeded in retrieving the true redshift, the second line counts the number of galaxies where DF and DL succeeded while DAE failed, and so forth.

Fig. 9. Dispersion between true and estimated redshift values for all galaxy types in the $S/N = 5$ (top row) and $S/N = 2$ (bottom row) cases. Denoising autoencoders (blue histograms) have generally smaller dispersion than dictionary learning (orange histograms).

more specifically how the combination of results helps to reduce the number of catastrophic outliers. In particular, the intermediate redshift elliptical galaxy failures by DL are replaced by DAE better redshifts; conversely, in the lower and higher redshift regimes for the other galaxy types, where DAE fails much more frequently, DL redshifts are retained, bringing catastrophic outlier rates significantly down in those bins when compared to DAE redshift performance. These results clearly demonstrate the advantage of combining the two methods. The exact values of the improvement will depend on the specifics of each simulation or data, but the complementarity is related to the different algorithms and should be qualitatively similar in other data settings.

7. Conclusion

In this paper, we introduced two new methods of spectroscopic redshift estimation, and benchmarked them on simulated data against a reference method based on line feature estimation and cross-correlation with eigentemplates. Both new methods rely on deriving an efficient representation for the data that is then used for redshift estimation. The first one uses the MOD dictionary learning technique to obtain a sparse representation for the full galaxy spectra (continuum and lines), which is then used to estimate redshifts from noisy spectra by searching for the lowest sparse approximation error among all tested redshift values. The second method applies denoising autoencoders for

Fig. 10. Best estimated redshift versus true redshift for three S/N cases. Blue and orange dots indicate that the redshift was chosen from the DAE and DL method, respectively. In the higher S/N cases, most estimated redshifts come from the DAE method, due to its almost perfect accuracy. At lower S/N values, where the method often starts failing, our algorithm increases the proportion of DL values, which are more robust but with higher variance.

Fig. 11. Catastrophic failure rate (measured as $|\Delta z|/(1+z) > 0.003$) by galaxy type and redshift bins for the best redshift estimation algorithm, in the $S/N = 20, 5$, and 2 cases (left, middle, and right columns, respectively). Degradation of performance at all noise levels has mostly been remedied due to the complementarity of DL and DAE strengths.

non-linear unsupervised feature extraction, learning the features from rest-frame spectra, and deriving the best-fit redshift value by passing the test spectra through the autoencoder and minimizing the reconstruction error on the input signal.

Both methods show significant improvement over the original Darth Fader pipeline, being able to recover redshift values with high accuracy and precision at high S/N regimes, with markedly less line confusion. Moreover, the more pronounced the line features on SEDs – as characterized by galaxy type –, the more precise results are. As S/N is reduced, measurement dispersion and catastrophic outlier rates increase as expected. In all sub-cases investigated, denoising autoencoders achieve smaller dispersion around the true redshift value. However, the catastrophic outlier rate increases rapidly as S/N is lowered. On the other hand, the catastrophic outlier rate from sparse dictionary learning is more resilient to the effects of noise, outperforming denoising autoencoders for $S/N = 2$ and below.

Given those complementary strengths, we designed an algorithm to combine DAE and DL results in an optimal way. If they both measure the same correct redshift value, we favor DAE values due to their smaller intrinsic scatter. If they do not agree, we use catastrophic outlier rates – calibrated with true values – to decide which value to pick. This strategy yields much-improved

results: in the $S/N = 2$ regime, we ensure that 71.4% of the galaxies can be identified whose redshifts agree, and this robust sample has a 0.5% catastrophic outlier rate. If completeness is a priority, we obtain a global galaxy sample with 6.85% catastrophic outlier rate, which is an improvement of $\sim 5\%$ over the DL method alone.

These results are encouraging. Fully automated spectroscopic redshift estimation methods that perform in a robust manner would be of great benefit to upcoming large-scale spectroscopic galaxy surveys such as DESI and Euclid, especially if they work in low S/N regimes. A next step would be to investigate the performance of the algorithms in simulations that fully reproduce the expected data quality from those surveys. Additionally, a combination of methods for different regimes, reliant solely on observed properties, can potentially produce extremely clean and robust redshift catalogs, although this is dependent on the specific properties of the data and noise, and will need to be investigated further.

Acknowledgements. This work is funded by the DEDALE project (contract no. 665044) and LENA (ERC StG no. 678282) within the H2020 Framework Program of the European Commission. BM thanks Julien Zoubian, Stéphanie Jouvel, Adrienne Leonard, and Paul Nathan for useful discussions.

References

- Abadi, M., Barham, P., Chen, J., et al. 2016, [OSDI](#), **16**, 265
- Abell, P. A., Allison, J., Anderson, S. F., et al. 2009, ArXiv e-prints [arXiv:[0912.0201](#)]
- Aharon, M., Elad, M., & Bruckstein, A. 2006, [IEEE Trans. Signal Process.](#), **54**, 4311
- Alain, G., & Bengio, Y. 2014, [J. Mach. Learn. Res.](#), **15**, 3563
- Alam, S., Ata, M., Bailey, S., et al. 2017, [MNRAS](#), **470**, 2617
- Baldry, I. K., Glazebrook, K., Brinkmann, J., et al. 2004, [ApJ](#), **600**, 681
- Bautista, J. E., Vargas-Magaña, M., Dawson, K. S., et al. 2018, [ApJ](#), **863**, 110
- Bengio, Y., Courville, A., & Vincent, P. 2013, [Pattern Analysis and Machine Intelligence, IEEE Transactions on](#), **35**, 1798
- Blake, C., & Glazebrook, K. 2003, [ApJ](#), **594**, 665
- Bolton, A. S., Schlegel, D. J., Aubourg, É., et al. 2012, [AJ](#), **144**, 144
- Bourlard, H., & Kamp, Y. 1988, [Biol. Cybern.](#), **59**, 291
- Bruzual, A. G., & Charlot, S. 1993, [ApJ](#), **405**, 538
- Coates, A., Ng, A. Y., & Lee, H. 2011, [International Conference on Artificial Intelligence and Statistics](#), 215
- Coleman, G. D., Wu, C.-C., & Weedman, D. W. 1980, [ApJS](#), **43**, 393
- Collobert, R., Weston, J., Bottou, L., et al. 2011, [J. Mach. Learn. Res.](#), **12**, 2493
- Dahl, G. E., Yu, D., Deng, L., & Acero, A. 2012, [Audio, Speech, and Language Processing, IEEE Transactions on](#), **20**, 30
- Dawson, K. S., Kneib, J.-P., Percival, W. J., et al. 2016, [AJ](#), **151**, 44
- DESI Collaboration (Aghamousa, A., et al.) 2016, ArXiv e-prints [arXiv:[1611.00037](#)]
- Elad, M., & Aharon, M. 2006, [IEEE Trans. Image Process.](#), **15**, 3736
- Engan, K., Aase, S. O., & Husoy, J. H. 1999, [Acoustics, Speech, and Signal Processing, 1999. Proceedings, 1999 IEEE International Conference on \(IEEE\)](#), **5**, 2443
- Frontera-Pons, J., Sureau, F., Bobin, J., & Le Floch, E. 2017, [A&A](#), **603**, A60
- Glazebrook, K., Offer, A. R., & Deeley, K. 1998, [ApJ](#), **492**, 98
- Hinton, G., Deng, L., Yu, D., et al. 2012, [IEEE Signal Process Mag.](#), **29**, 82
- Hinton, S. R., Davis, T. M., Lidman, C., Glazebrook, K., & Lewis, G. F. 2016, [Astron. Comput.](#), **15**, 61
- Hutchinson, T. A., Bolton, A. S., Dawson, K. S., et al. 2016, [AJ](#), **152**, 205
- Ilbert, O., Capak, P., Salvato, M., et al. 2008, [ApJ](#), **690**, 1236
- Jouvel, S., Kneib, J.-P., Ilbert, O., et al. 2009, [A&A](#), **504**, 359
- Kazin, E. A., Koda, J., Blake, C., et al. 2014, [MNRAS](#), **441**, 3524
- Kennicutt, Jr., R. C. 1998, [ARA&A](#), **36**, 189
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, [Advances in Neural Information Processing Systems](#), 1097
- Leauthaud, A., Massey, R., Kneib, J.-P., et al. 2007, [ApJS](#), **172**, 219
- Machado, D. P., Leonard, A., Starck, J. L., Abdalla, F. B., & Jovel, S. 2013, [A&A](#), **560**, A83
- Mairal, J., Elad, M., & Sapiro, G. 2008, [IEEE Trans. Image Process.](#), **17**, 53
- Mairal, J., Ponce, J., Sapiro, G., Zisserman, A., & Bach, F. R. 2009, [Advances in Neural Information Processing Systems](#), 1033
- Mallat, S., & Zhang, Z. 1993, [IEEE Trans. Signal Process.](#), **41**, 3397
- Moustakas, J., Kennicutt, Jr., R. C., & Tremonti, C. A. 2006, [ApJ](#), **642**, 775
- Olshausen, B., & Field, D. 1996, [Vision Res.](#), **37**, 3311
- Pati, Y. C., & Krishnaprasad, P. S. 1993, [IEEE Trans. Neural Networks](#), **4**, 73
- Rifai, S., Vincent, P., Muller, X., Glorot, X., & Bengio, Y. 2011, [Proceedings of the 28th International Conference on Machine Learning \(ICML-11\)](#), 833
- Seo, H.-J., & Eisenstein, D. J. 2003, [ApJ](#), **598**, 720
- Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, [AJ](#), **146**, 32
- Starck, J. L., & Murtagh, F. 2006, [Astronomical Image and Data Analysis \(Berlin, Heidelberg: Springer-Verlag\)](#)
- Starck, J.-L., Murtagh, F., Pirenne, B., & Albrecht, M. 1996, [PASP](#), **108**, 446
- Starck, J. L., Murtagh, F., & Fadili, J. 2015, [Sparse Image and Signal Processing: Wavelets and Related Geometric Multiscale Analysis, 2nd edn. \(Cambridge: Cambridge University Press\)](#)
- Tuccillo, D., Huertas-Company, M., Decencière, E., et al. 2017, [MNRAS](#), **475**, 894
- Vincent, P., Larochelle, H., Bengio, Y., & Manzagol, P. A. 2008, [Proceedings of the 25th International Conference on Machine Learning \(ACM\)](#), 1096
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., & Manzagol, P.-A. 2010, [J. Mach. Learn. Res.](#), **11**, 3371
- Zhang, Q., & Li, B. 2010, [Computer Vision and Pattern Recognition \(CVPR\), 2010 IEEE Conference on \(IEEE\)](#), 2691