

# Fermilab

Identifying Anomalous DESI Galaxy Spectra with a Variational Autoencoder

FERMILAB-PUB-25-0117-PPD

arXiv:2506.17376

This manuscript has been authored by Fermi Forward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

# Identifying Anomalous DESI Galaxy Spectra with a Variational Autoencoder

C. Nicolaou,<sup>1</sup> R.P. Nathan,<sup>1</sup> O. Lahav,<sup>1</sup> A. Palmese,<sup>2</sup> A. Saintonge,<sup>1</sup> J. Aguilar,<sup>3</sup> S. Ahlen,<sup>4</sup> C. Allende Prieto,<sup>5,6</sup> S. Bailey,<sup>3</sup> S. BenZvi,<sup>7</sup> D. Bianchi,<sup>8,9</sup> A. Brodzeller,<sup>3</sup> D. Brooks,<sup>1</sup> T. Claybaugh,<sup>3</sup> A. de la Macorra,<sup>10</sup> J. Della Costa,<sup>11,12</sup> Arjun Dey,<sup>12</sup> P. Doel,<sup>1</sup> J. E. Forero-Romero,<sup>13,14</sup> E. Gaztañaga,<sup>15,16,17</sup> S. Gontcho A Gontcho,<sup>3</sup> G. Gutierrez,<sup>18</sup> K. Honscheid,<sup>19,20,21</sup> C. Howlett,<sup>22</sup> M. Ishak,<sup>23</sup> R. Kehoe,<sup>24</sup> D. Kirkby,<sup>25</sup> T. Kisner,<sup>3</sup> A. Kremin,<sup>3</sup> A. Lambert,<sup>3</sup> M. Landriau,<sup>3</sup> L. Le Guillou,<sup>26</sup> A. Meisner,<sup>12</sup> R. Miquel,<sup>27,28</sup> J. Moustakas,<sup>29</sup> S. Nadathur,<sup>16</sup> F. Prada,<sup>30</sup> I. Pérez-Ràfols,<sup>31</sup> G. Rossi,<sup>32</sup> E. Sanchez,<sup>33</sup> M. Schubnell,<sup>34,35</sup> M. Siudek,<sup>6,17</sup> D. Sprayberry,<sup>12</sup> G. Tarlé,<sup>35</sup> B. A. Weaver,<sup>12</sup> and H. Zou,<sup>36</sup>

<sup>1</sup>Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, UK

<sup>2</sup>Department of Physics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

<sup>3</sup>Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA

<sup>4</sup>Physics Dept., Boston University, 590 Commonwealth Avenue, Boston, MA 02215, USA

<sup>5</sup>Departamento de Astrofísica, Universidad de La Laguna (ULL), E-38206, La Laguna, Tenerife, Spain

<sup>6</sup>Instituto de Astrofísica de Canarias, C/Vía Láctea, s/n, E-38205 La Laguna, Tenerife, Spain

<sup>7</sup>Department of Physics & Astronomy, University of Rochester, 206 Bausch and Lomb Hall, P.O. Box 270171, Rochester, NY 14627-0171, USA

<sup>8</sup>Dipartimento di Fisica “Aldo Pontremoli”, Università degli Studi di Milano, Via Celoria 16, I-20133 Milano, Italy

<sup>9</sup>INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy

<sup>10</sup>Instituto de Física, Universidad Nacional Autónoma de México, Circuito de la Investigación Científica, Ciudad Universitaria, Cd. de México C. P. 04510, México

<sup>11</sup>Department of Astronomy, San Diego State University, 5500 Campanile Drive, San Diego, CA 92182, USA

<sup>12</sup>NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA

<sup>13</sup>Departamento de Física, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia

<sup>14</sup>Observatorio Astronómico, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia

<sup>15</sup>Institut d’Estudis Espacials de Catalunya (IEEC), c/ Esteve Terradas 1, Edifici RDIT, Campus PMT-UPC, 08860 Castelldefels, Spain

<sup>16</sup>Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth, PO1 3FX, UK

<sup>17</sup>Institute of Space Sciences, ICE-CSIC, Campus UAB, Carrer de Can Magrans s/n, 08913 Bellaterra, Barcelona, Spain

<sup>18</sup>Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, USA

<sup>19</sup>Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

<sup>20</sup>Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

<sup>21</sup>The Ohio State University, Columbus, 43210 OH, USA

<sup>22</sup>School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia

<sup>23</sup>Department of Physics, The University of Texas at Dallas, 800 W. Campbell Rd., Richardson, TX 75080, USA

<sup>24</sup>Department of Physics, Southern Methodist University, 3215 Daniel Avenue, Dallas, TX 75275, USA

<sup>25</sup>Department of Physics and Astronomy, University of California, Irvine, 92697, USA

<sup>26</sup>Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), FR-75005 Paris, France

<sup>27</sup>Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain

<sup>28</sup>Institut de Física d’Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Edifici Cn, Campus UAB, 08193, Bellaterra (Barcelona), Spain

<sup>29</sup>Department of Physics and Astronomy, Siena College, 515 Loudon Road, Loudonville, NY 12211, USA

<sup>30</sup>Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain

<sup>31</sup>Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 10, 08930 Barcelona, Spain

<sup>32</sup>Department of Physics and Astronomy, Sejong University, 209 Neungdong-ro, Gwangjin-gu, Seoul 05006, Republic of Korea

<sup>33</sup>CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain

<sup>34</sup>Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, USA

<sup>35</sup>University of Michigan, 500 S. State Street, Ann Arbor, MI 48109, USA

<sup>36</sup>National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Rd., Chaoyang District, Beijing, 100012, P.R. China

Accepted XXX. Received YYY; in original form ZZZ

## ABSTRACT

The tens of millions of spectra being captured by the Dark Energy Spectroscopic Instrument (DESI) provide tremendous discovery potential. In this work we show how Machine Learning, in particular Variational Autoencoders (VAE), can detect anomalies in a sample of approximately 200,000 DESI spectra comprising galaxies, quasars and stars. We demonstrate that the VAE can compress the dimensionality of a spectrum by a factor of 100, while still retaining enough information to accurately reconstruct spectral features. We then detect anomalous spectra as those with high reconstruction error and those which are isolated in the VAE latent representation. The anomalies identified fall into two categories: spectra with artefacts and spectra with unique physical features. Awareness of the former can help to improve the DESI spectroscopic pipeline; whilst the latter can lead to the identification of new and unusual objects. To further curate the list of outliers, we use the *Astronomaly* package which employs Active Learning to provide personalised outlier recommendations for visual inspection. In this work we also explore the VAE latent space, finding that different object classes and subclasses are separated despite being unlabelled. We demonstrate the interpretability of this latent space by identifying tracks within it that correspond to various spectral characteristics. For example, we find tracks that correspond to increasing star formation and increase in broad emission lines along the Balmer series. In upcoming work we hope to apply the methods presented here to search for both systematics and astrophysically interesting objects in much larger datasets of DESI spectra.

**Key words:** techniques: spectroscopic – methods: statistical – methods: data analysis – galaxies: peculiar

## 1 INTRODUCTION

Astronomical spectra are highly informative, as they provide a way to infer the composition, temperature and physical processes that occur in the observed object. Galaxy spectra are crucial for understanding the properties and evolution of galaxies. Spectroscopic surveys are currently gathering such data at an unprecedented rate providing astronomers with a wealth of information. The Dark Energy Spectroscopic Instrument (DESI, [DESI Collaboration et al. 2016a](#)) has already gathered  $\sim 40$  million spectra (including repeat observations) surpassing the number of spectra previously available, aiming at an eventual total of around 35 million unique spectra. The increased volume of incoming data have accelerated the shift from traditional data analysis and visualisation tools (e.g. [Lahav et al. 1995](#)) to big-data techniques (e.g. [Baron 2019](#); [Huertas-Company & Lanusse 2022](#)).

Hidden within these huge new datasets are likely to be large numbers of anomalies. An anomaly is defined as an object or observation that is rare and differs significantly from the majority of the other data accompanying it, suggesting that it has been generated by a different underlying mechanism ([Hawkins 1980](#)). Anomalies can be rarely occurring events, extreme objects, errors due to instrument malfunction, or objects that are completely new and have the potential to unveil new physics. The last of these, also sometimes referred to as "unknown unknowns", are arguably the most interesting but also most challenging to find in large datasets. Historically, astronomers relied on visual inspection of data to make such scientific discoveries, but as volumes of data increase, visual inspection alone becomes infeasible. Astronomers must therefore turn to automated tools that can provide a smaller, curated subset of the data containing those observations most likely to be outliers.

In particular, machine learning has become increasingly popular as a powerful tool for automating the detection of atypical or anomalous observations. Two types of anomaly detection that are important to differentiate are outlier detection and novelty detection. Outlier detection refers to the case where the training data contains both typical data (inliers), which constitute the majority of the dataset, and anom-

alous data (outliers). The goal is to identify the outliers present in the dataset. In novelty detection, the training dataset contains only typical data, and the goal is to identify whether a new observation is an outlier. In this context an outlier is also called a novelty. Novelty detection is semi-supervised whereas outlier detection is unsupervised ([Ruff et al. 2020](#)).

Three broad categories of anomaly detection methods are deviation based, proximity based and statistical. Deviation-based anomaly detection uses the reconstruction error from a model trained on the data as the anomaly score. The model learns a lower-dimensional *latent* representation of the data and then attempts to use this to reconstruct the original. Anomalies will tend to have a high reconstruction error as they have features that are atypical and which the model is not able to reproduce. Anomaly-detection methods based on proximity assume that anomalous data are isolated in the latent representation from the majority of the data. Approaches to define how isolated a sample is include clustering, distance-based measurements and density estimation. Statistical anomaly detection makes use of parametric or non-parametric models to define a probability distribution assumed to describe the data. A sample is anomalous if it has a low probability of being generated by the model ([An & Cho 2015](#); [Reis et al. 2019](#)). In this work, we make use of both deviation-based and proximity-based methods.

Deviation-based methods for spectra rely on the intuition common to many fields of study, that though the data are typically captured in high dimension, they can be represented with good fidelity using a much lower-dimensional manifold (the "manifold hypothesis"). For example, DESI spectra have 7,800 bins (spectral channels) and therefore effectively 7,800 dimensions. Due to the astrophysical processes underlying the spectra, spectral lines and features like the continuum are highly correlated. Deviation-based methods consequently tend to rely on dimensionality-reduction techniques. Because of their simplified representation of the data being studied, such techniques often have the added benefit of providing an unsupervised encoding of significant physical characteristics.

Principal Component Analysis (PCA) has been one of the most widely and successfully used algorithms for dimensionality reduction and for the unsupervised extraction of useful features from spectro-

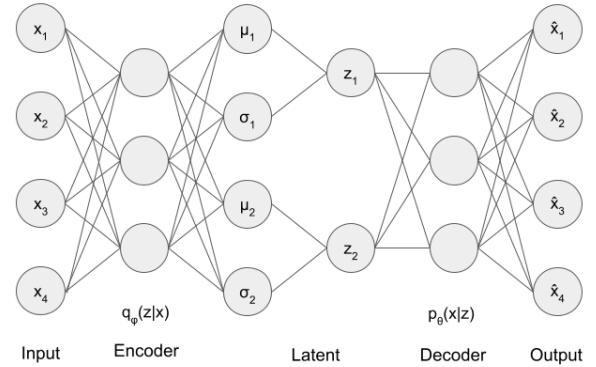
\* E-mail: constantina.nicolaou.17@ucl.ac.uk

scopic data. [Boroson & Lauer \(2010\)](#) used the PCA reconstruction error to detect anomalies in a sample of Sloan Digital Sky Survey (SDSS) quasar spectra. [Folkes et al. \(1999\)](#) and [Madgwick et al. \(2003a\)](#) used PCA to effectively compress galaxy spectra to a few components, finding correlations between the PCA components and physical attributes such as star formation rate and morphological type. Spectral classification derived from PCA was also used for galaxy clustering by [Madgwick et al. \(2003b\)](#). [Yip et al. \(2004a\)](#) used PCA to classify a sample of SDSS galaxy spectra while [Rogers et al. \(2007\)](#) analysed a sample of SDSS early-type galaxies using PCA in order to infer differences in the average age of stellar population and star formation history. Following [Slonim et al. \(2001\)](#) on the information bottleneck method, [Ferreras et al. \(2022\)](#) extracted the stellar population content of galaxy spectra by measuring the entropy of spectra in order to quantify the amount of information encoded.

Being a linear method, however, PCA can be limited when non-linear features are present. [Yip et al. \(2004b\)](#) showed that 50 PCA components are necessary to acceptably reconstruct a typical SDSS quasar spectrum exhibiting broad emission lines. Autoencoders (AEs) and Variational Autoencoders (VAEs) are non-linear deep-learning methods which have been successfully employed to address these limitations. In particular, [Portillo et al. \(2020\)](#) used a VAE on a sample of SDSS galaxy spectra. They showed that a VAE with six latent dimensions could satisfactorily reconstruct spectra whose original dimension was 1000 and outperform PCA with the same number of components. The VAE latent space showed separation of classes and they demonstrated the interpretability of the latent space. Finally, they identified outliers by estimating the local density of samples in the latent space using the local outlier factor (LOF) algorithm and ranking spectra according to how isolated they were in the latent space. The top 10 outliers identified were stars, spectra with the majority of the fluxes masked, low signal-to-noise ratio (S/N) spectra, bad instrument calibration and contamination by stellar light. [Ichinohe & Yamada \(2019\)](#) used a VAE to rank objects based on the reconstruction error. They approached this in the framework of novelty detection for a sample of simulated high-resolution X-ray spectra. Most recently, [Liang et al. \(2023\)](#) used an AE to compress galaxy spectra from the DESI Early Data Release (EDR, [DESI Collaboration et al. 2023](#)) into a redshift-invariant latent space, and a normalizing flow to detect outliers. Additionally, [Scourfield et al. \(2023\)](#) used a VAE model to demonstrate their ability to denoise spectra.

Proximity-based approaches to anomaly detection in spectroscopic data include [Meusinger et al. \(2012\)](#) who used self-organising maps (SOM, [Kohonen 1982](#)) – a clustering based algorithm – on a sample of SDSS quasar spectra. Unusual spectra were defined as those points to be found in the low density regions of representation created via the SOM. Outliers mainly included broad absorption lines, unusual red continua, weak emission lines and conspicuously strong iron emission lines. [Fustes et al. \(2013\)](#) also used SOM for outlier detection on a simulated sample of spectra based on SDSS observations. [Baron & Poznanski \(2017\)](#) used unsupervised random forests to identify outliers in SDSS spectra. Some interesting outliers identified include galaxies which host supernovae, high ionisation spectral lines, galaxies with unusual gas kinematics and galaxy-galaxy gravitational lenses. [Suárez-Pérez & Forero-Romero \(2022\)](#) used Uniform Manifold Approximation and Projection (UMAP) as a technique to assess the data quality of DESI. They used UMAP to project DESI nightly data into a 2-dimensional space where they were able to identify instrumentation outliers.

In this study, we follow a broadly similar approach to [Portillo et al. \(2020\)](#) and perform outlier detection using a VAE, but we



**Figure 1.** Example architecture of a Variational Autoencoder (VAE). The encoder is tasked with mapping the input,  $x$  onto the lower dimensional latent space,  $z$ , as a normal distribution described by two vectors containing the latent mean  $\mu$  and latent standard deviation  $\sigma$ . The decoder is tasked with decompressing samples from the latent representation and producing reconstructions,  $\hat{x}$ , of the original data.

apply it to astronomical spectra obtained from the DESI survey and utilise much larger datasets. Although [Liang et al. \(2023\)](#) probed a similar group of spectra for outliers, our VAE-driven approach is more straightforward as it does not seek to be redshift-invariant and instead makes direct use of the spectroscopic redshifts assigned by the DESI pipeline. In our work, we explore two approaches for identifying outliers. The first method uses the reconstruction error of the VAE as an anomaly score. The second method is proximity based, and is based on the position of spectra in the VAE latent space as anomalous spectra are expected to be isolated from the majority of the data.

We find that outlier spectra fall into a variety of categories, with the main distinction being between those which are anomalous because of physical features and those due to instrumental artefacts. Depending on the science question, certain types of outliers will be relevant while others will present a nuisance. We make use of Astronomaly ([Lochner & Bassett 2021](#)), a software package which combines the input of a human expert with the processing power of machine learning to provide a curated list of outliers which are most relevant to the user.

This paper is organised as follows. An introduction to VAEs is provided in Section 2. In Section 3 we outline the specific VAE architecture we used, with a description of the DESI data and the preprocessing we applied in Section 4. The results are then split into four sections. In Section 5.1 we examine the quality of the VAE reconstructions. In Section 5.2 we present the outliers obtained from the reconstruction error and latent space approaches. In Section 5.3 we present the application of Astronomaly. In Section 5.4 we explore the VAE latent space and the interpretability of the latent representations. Finally, in Section 7 we discuss our results, along with limitations and improvements for future work.

## 2 VARIATIONAL AUTOENCODERS

VAEs ([Kingma & Welling 2013](#)) are probabilistic, generative models based on the encoder-decoder architecture. VAEs are tasked with learning efficient low-dimensional representations of the data in an unsupervised manner. Figure 1 depicts an example architecture of a VAE.

In a standard autoencoder (AE) ([Kramer 1991](#)), the encoder is typ-

ically a feedforward neural network and is tasked with mapping the input data,  $x$ , onto a low-dimensional hidden representation,  $z$ , called the latent space. As the dimension of  $z$  is smaller than  $x$ , the encoder performs dimensionality reduction. This presents a bottleneck to the model and the encoder must learn an efficient compression that preserves as much of the information as possible by capturing the most meaningful factors of variation in the data. The decoder is also a feedforward neural network, typically mirroring the architecture of the encoder. Using just the latent space representations, it is tasked with reconstructing the original data. The reconstructed output,  $\hat{x}$  is compared to the original data through the loss function which is typically given by the  $L^2$  function,  $\|x - \hat{x}\|$ . This measures how far the reconstructed values are from the original ones in Euclidean terms. During training, this error is backpropagated, informing the weights update so as to achieve a lower loss value. Hence, the AE attempts to find the best encoder-decoder pair so as to maximise the information encoded in the latent space and learn the underlying feature representations of the data.

While AEs are able to capture non-linearities in the data and learn powerful representations in low dimensions, the latent space they create is unregularised. While the encoder is able to usefully map similar inputs together in latent space forming clusters, this subspace may be discontinuous, with gaps between clusters that do not correspond to any data points. Using the decoder to try and generate from these gaps, may therefore result in unrealistic outputs. This limits the generative usefulness of the decoder. Additionally, the lack of a well-defined structure in the latent space may result in less reliable representations of input data, especially for outliers, leading to poor separation between normal and anomalous data.

The latent space of VAEs, on the other hand, is by design continuous and therefore allows for the decoder to randomly sample from any point in the space and generate meaningful outputs. Instead of encoding each input in a vector, the encoder produces a probability distribution over  $z$ ,  $q_\phi(z|x)$ , where  $\phi$  represents the parameters of the encoder network. The encoded distribution is given by a normal distribution and hence can be described by two vectors containing the latent mean,  $\mu$ , and latent variance,  $\sigma^2$ . A point from the latent space is then sampled from the encoded distribution according to

$$z \sim q_\phi(z|x) = \mathcal{N}(\mu_{z|x}, \sigma_{z|x}^2). \quad (1)$$

The sampled  $z$  are then passed through the decoder. As the sampling procedure is stochastic, each time the same input is passed through the encoder during training, a slightly different sample in latent space will be propagated through the decoder. Hence the decoder will not learn just a single point in latent space that characterises the input, but it will learn that all points that fall within a small area in latent space defined by the variance of  $q_\phi$  are representations of the input.

The decoder also outputs a probability distribution, inferring  $x$  from  $z$ . The probability distribution is given by  $p_\theta(x|z)$ , where  $\theta$  are the parameters of the decoder network. The decoder distribution is chosen to reflect the type of input data. For example, in the case of continuous input data a normal distribution is used resulting in the outputs being given by

$$\hat{x} \sim p_\theta(x|z) = \mathcal{N}(\mu_{x|z}, \sigma_{x|z}^2). \quad (2)$$

The VAE objective function (or loss function) is given by the evidence lower bound (ELBO) which is made up of two terms defined as follows:

$$\begin{aligned} \log p(x) &\geq \mathbb{E}_{z \sim q_\phi(z|x)} [\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x) || p(z)) \\ &\geq \mathcal{L}(x) \end{aligned} \quad (3)$$

where  $\mathcal{L}(x)$  is known as ELBO as it represents the lower bound on  $\log p(x)$ .

The ELBO is tractable and thereby the VAE is able to apply back-propagation and optimise a lower bound on the likelihood of the data by finding the parameters of the VAE network ( $\phi$  and  $\theta$ ) that maximise the lower bound (or equivalently minimise the negative ELBO) (Rezende et al. 2014).

The first term of the ELBO is the expectation of  $p_\theta(x|z)$ , which is the decoder output, over the latent variables,  $z$ , sampled from the output of the encoder. By maximising this we are essentially maximising the likelihood of the observations given  $z$  and hence this term represents the reconstruction likelihood which encourages the VAE to produce accurate reconstructions.

The second term is a Kullback-Leibler divergence (KL divergence or  $D_{KL}$ , Kullback & Leibler 1951) between  $q_\phi(z|x)$  and  $p(z)$ . The KL divergence is a non-symmetric measure of similarity between two probability distributions. By maximising the ELBO, the second term is minimised which imposes that the approximate posterior distribution  $q_\phi(z|x)$  is close to the prior distribution  $p(z)$ . The prior distribution is chosen to be a standard Gaussian  $p(z) = \mathcal{N}(0, I)$  where  $I$  is the identity matrix. As the prior on  $z$  is diagonal, it encourages the latent variables to be independent and thereby encode different interpretable factors of variation. It is evident that this term acts as a regularisation term, enforcing the output of the encoder, and hence the latent variables  $z$ , to follow a standard Gaussian. This not only enables the generative process but also results in a regularised latent space that clearly separates normal from anomalous data points.

### 3 VAE IMPLEMENTATION

We implemented a VAE using `tensorflow-probability` (Abadi et al. 2015). The encoder consists of the input layer which is made up of 1000 nodes, followed by 4 hidden layers with dimensions 800, 600, 500, 300. The architecture of the decoder mirrors that of the encoder. To illustrate this approach, the latent dimension is set to 10. This is chosen based on effective reconstruction results demonstrated in the literature (e.g. Portillo et al. (2020)) but more rigorous testing is required to find the optimal number of latent dimensions (by carrying out a hyperparameter optimisation procedure for example) and to understand the effects of different dimension sizes (Scourfield et al. 2023). Dropout is used in both the encoder and decoder to regularise the VAE and avoid overfitting. The dropout rate is set to 0.2. The rectified linear unit (ReLU) activation function given by

$$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases} \quad (4)$$

is adopted in all layers except for the output layers of the encoder and decoder where a linear activation function is used. We use a Gaussian distribution for the output of the decoder and thus the reconstruction loss term is given by the Gaussian log-likelihood. The first term in the loss function of the VAE (see eq. 3), is given by

$$\mathbb{E}_z [\log p(x|z)] = \mathbb{E}_z \left[ \sum_i \frac{1}{2} - \log 2\pi\sigma_{x|z,i}^2 - \frac{(x_i - \mu_{x|z,i})^2}{2\sigma_{x|z,i}^2} \right] \quad (5)$$

where  $i$  indicates the  $i^{th}$  element of the vectors (Yu 2020), and  $\log$  represents the natural logarithm. The inverse variance associated with each flux measurement of each spectrum is used in the reconstruction loss, acting as a weighting term. This allows us to explicitly incorporate heteroscedastic uncertainties and element-wise masking. Bad pixels have their weight in the loss set to zero which is equivalent to infinite uncertainty.

We train the VAE using Adam optimiser with the starting learning rate set to  $10^{-3}$ . The learning rate is set to be reduced by a factor of 10 if the loss function on the validation set does not improve for 5 epochs. We use a batch size of 512 and trained the model for 50 epochs. For the purposes of this proof-of-concept study, it was felt sufficient to select VAE hyperparameters through a limited random search of the hyperparameter space. Optimising the VAE in a more rigorous manner, either by using a more extensive random search or a grid search (Bergstra & Bengio 2012), should reveal a truly optimised architecture which might be employed in future work.

## 4 DESI DATA

The Dark Energy Spectroscopic Instrument (DESI) is a robotic spectroscopic surveyor that operates on the Mayall 4-meter telescope at Kitt Peak National Observatory (DESI Collaboration et al. 2022). DESI, which can obtain simultaneous spectra of almost 5000 objects over an approximately 3-degree field (DESI Collaboration et al. 2016b; Silber et al. 2023; Miller et al. 2023), is currently conducting a five-year survey of about a third of the sky. This campaign will obtain spectra for approximately 40 million galaxies and quasars (DESI Collaboration et al. 2016a).

The primary goal of the Dark Energy Spectroscopic Instrument (DESI) is to determine the nature of dark energy through the most precise measurement of the expansion history of the universe ever obtained (Levi et al. 2013). DESI was designed to meet the definition of a Stage IV dark energy survey with only a 5-year observing campaign. Forecasts for DESI (DESI Collaboration et al. 2016b) predict a factor of approximately five to ten improvement on the size of the error ellipse of the dark energy equation of state parameters  $w_0$  and  $w_a$  relative to previous Stage-III experiments.

The sheer scale of the DESI experiment necessitates multiple supporting software pipelines and products, which include significant imaging from the DESI Legacy Imaging Surveys (Dey et al. 2019), an extensive spectroscopic reduction pipeline (Guy et al. 2023), a template-fitting pipeline to derive classifications and redshifts for each targeted source (Redrock; Bailey et al. 2024 (in preparation), and for the special case of QSOs, Brodzeller et al. 2023), and multiple pipelines with specific functions, e.g. to assign fibers to targets and to tile the survey and plan and optimize observations as the campaign progresses (Schlafly et al. 2023).

The dataset consists of DESI spectra from the Early Data Release (EDR) of the Bright Galaxy Survey (BGS, Hahn et al. 2022; Myers et al. 2022; Ruiz-Macias et al. 2021; Juneau et al. 2024). During the 5-year DESI observing run, the BGS will observe 10 million galaxies over the redshift range  $0 < z < 0.6$ . The BGS targets galaxies with a limit in the r-band magnitude of  $r < 19.5$  (BGS Bright), a secondary lower priority sample defined by the magnitude range  $19.5 < r < 20.175$  (BGS Faint) and a small low-z quasar sample (BGS AGN).

The redshift and spectral classifications are obtained from the

DESI redshift pipeline, named **redrock**<sup>1</sup> which is a template-fitting code, tasked with finding an optimal fit to a spectrum via linear combinations (PCA) of the template components derived from galaxy models (Guy et al. 2023). The fitting is performed over a specified redshift range for three template classes: “star”, “galaxy”, and “quasar”.

Using the redrock information, we select objects with redshifts in the range  $0 \leq z \leq 0.3$ , and convert the observed wavelength vector of each object to the rest-frame. The analysis is performed over the rest-wavelength range of 3600 Å to 7556 Å, which is fully sampled by all the objects in the chosen wavelength range. The analysis could be repeated over higher redshift intervals. We resample the spectra to 1000 wavelength bins ensuring that the total integrated flux remains consistent before and after resampling. While this leads to some small-scale information being lost as the DESI spectral resolution is much higher, the resampling is done to reduce the computational cost. This resampling corresponds to a resolution of 4 Å which is sufficient to avoid blending between important spectral lines such as H $\alpha$   $\lambda 6565$  and [N II]  $\lambda 6550, \lambda 6585$ . Note that we have opted for a simple resampling procedure which did not account for covariances between resampled bins and therefore this may lead to a slight underestimation of the uncertainties. Additionally, before the resampling procedure, pixels with missing values, zero flux or zero inverse variance are removed and a simple linear interpolation is used to fill in the values.

Each resampled spectrum has an indicator mask vector that denotes bad pixels e.g. pixels that are affected by CCD (charged-coupled device in the spectrograph cameras) defects or cosmic rays. Bad pixels are removed and iterative PCA (Yip et al. 2004a) is used to infill the missing values. The dataset does not suffer from heavily masked spectra as only 0.01% of the spectra have more than 10% of the pixels masked. The spectra are normalised individually to have unit norm. The median pixel signal-to-noise ratio (S/N) is calculated for each spectrum. We only keep the spectra that have a median pixel S/N above 5.

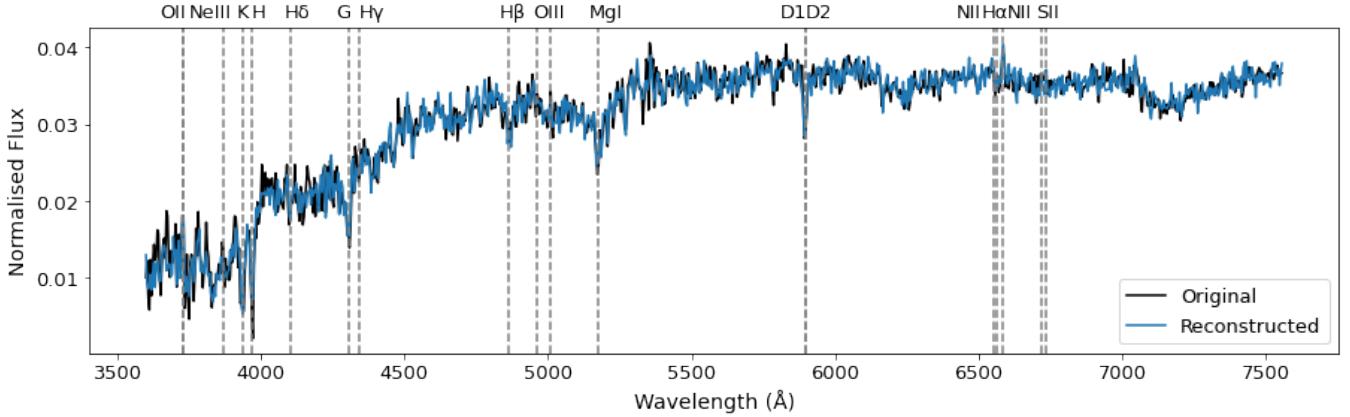
The resulting dataset consists of  $\sim 208,000$  spectra, of which  $\sim 156,000$  form the training set and the rest form the validation set. In future work, the study presented here can be scaled to a larger dataset by making use of more spectra, which will also have a greater potential to lead to a novel object being identified. Using the spectral classifications provided by **redrock**, the dataset is composed of around 97.8% galaxy spectra, 0.7% quasar spectra and 1.5% star spectra. Note however that these are not true labels and while **redrock** performs well on galaxies and stars it misclassified  $\sim 10 - 15\%$  of true quasars, in particular low-redshift ones (Alexander et al. 2023).

## 5 RESULTS

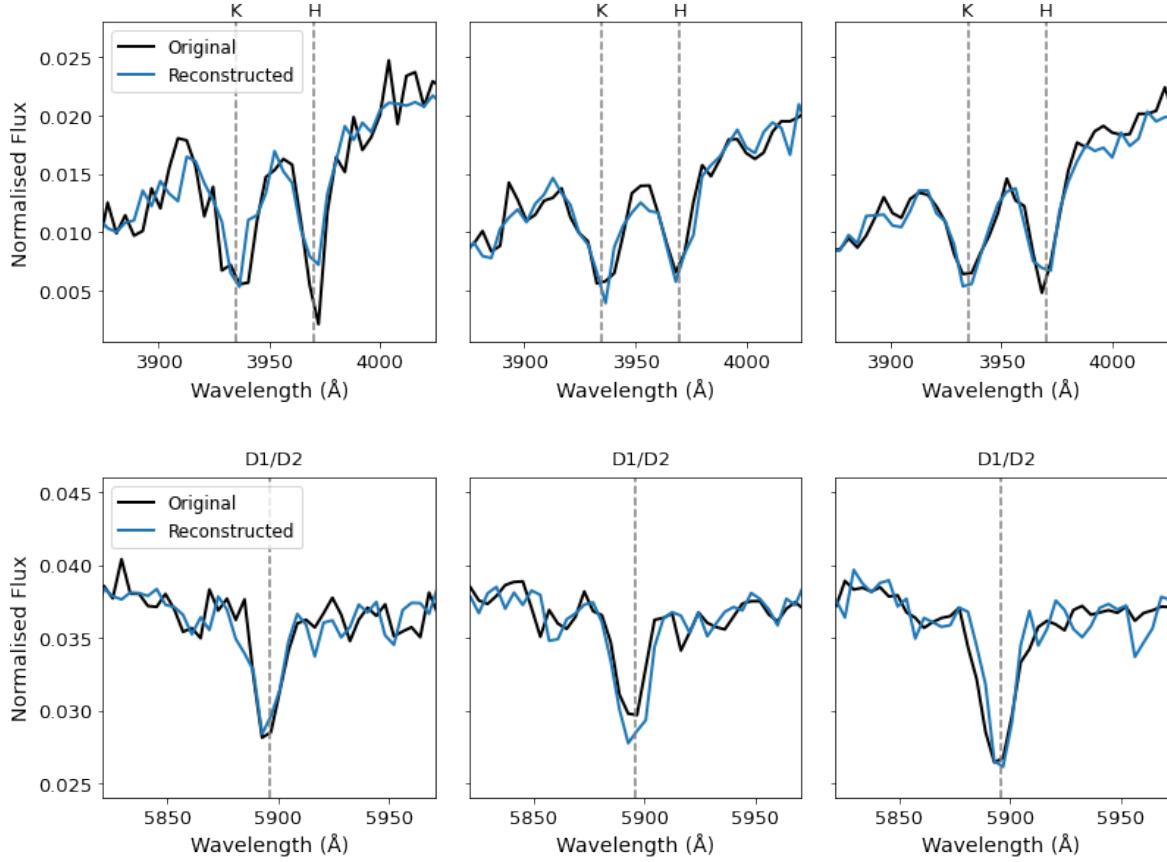
### 5.1 Reconstruction Accuracy

First let us consider how well the VAE is able to reconstruct the spectra of various objects. Even though the VAE was not optimised to achieve the lowest possible reconstruction error, it is still important to assess whether it is able to reconstruct the spectra to a reasonable degree. To quantify the performance of the VAE on the task of reconstruction, we calculate the weighted mean squared error (MSE) between the original and reconstructed spectrum. The weighted MSE for the  $i^{th}$  spectrum is defined as

<sup>1</sup> <https://github.com/desihub/redrock>



**Figure 2.** Original (black) and reconstructed (blue) spectrum based on 10 latent variables (in this and and subsequent figures) of an **elliptical galaxy** from the validation set. The median pixel S/N of the galaxy is 33. The spectrum is in the rest frame. The dashed grey lines indicate typical spectral lines labelled at the top of the plot. Target ID: 39627836461419062.

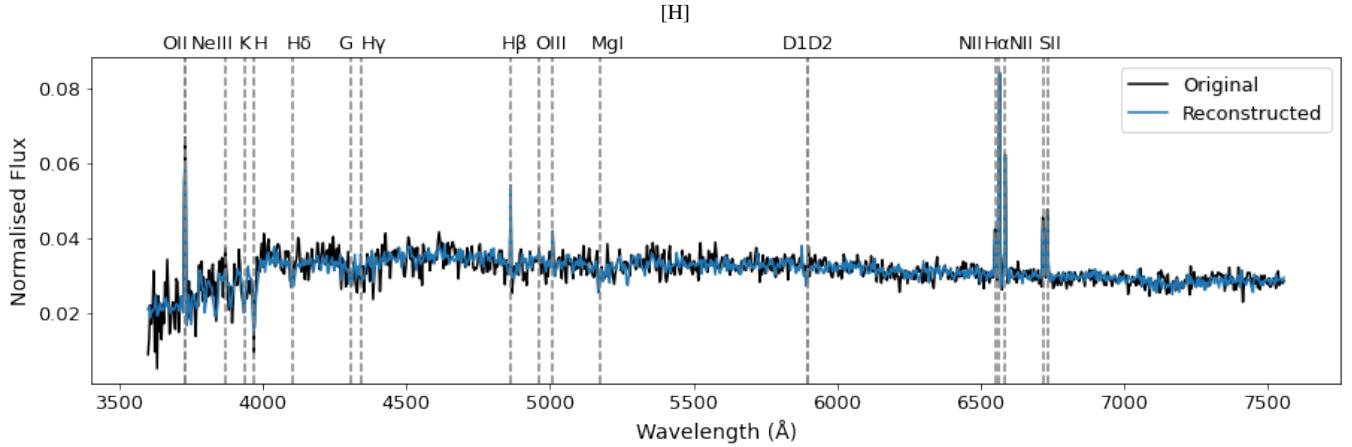


**Figure 3.** Original (black) and reconstructed (blue) spectra of **three elliptical galaxies**. The top row shows the VAE reconstruction of the Ca II K and H  $\lambda 3935, \lambda 3970$  lines and the bottom row the Na I D-lines  $\lambda 5896, \lambda 5898$ . Target IDs left to right: 39627836461419062, 39633321180793652 and 39628435806491872.

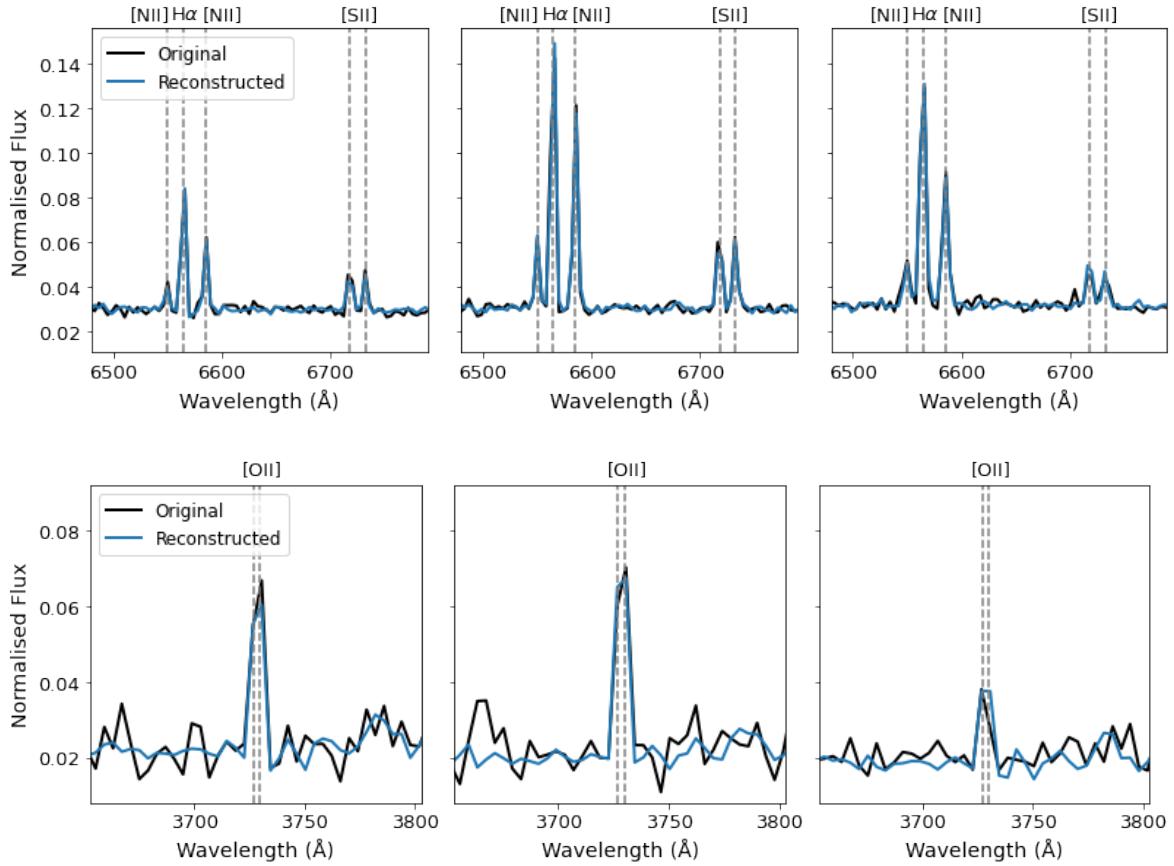
$$\text{weighted MSE}_i = \frac{1}{M} \sum_{m=1}^M \mathbf{w}_i \cdot (\mathbf{x}_i - \hat{\mathbf{x}}_i)^2 \quad (6)$$

where  $\mathbf{x}_i$  is the original spectrum,  $\hat{\mathbf{x}}_i$  is the the reconstructed spectrum and  $\mathbf{w}_i = \frac{1}{\sigma_i^2}$  where  $\sigma_i$  is the noise estimate from the data. All three vectors have dimensions  $M = 1000$  corresponding to the 1000 flux

measurements. The mean weighted MSE on the validation set, as defined in Eq.6, is 1.12. Broken down to the three classes, galaxies have the lowest average MSE at 1.09, stars have a MSE of 1.87 and quasars 2.57. Since our dataset is mostly composed of galaxies the VAE was able to see many examples of galaxy spectra and hence learn to reconstruct them well. The VAE however, struggles to reconstruct stars and quasars in comparison to galaxies, as these are rare in



**Figure 4.** Original (black) and reconstructed (blue) spectrum of an **emission-line galaxy** from the validation set. The median pixel S/N of the galaxy is 14. The spectrum is in the rest frame. The dashed grey lines indicate typical spectral lines labelled at the top of the plot. Target ID: 39627974831507323.

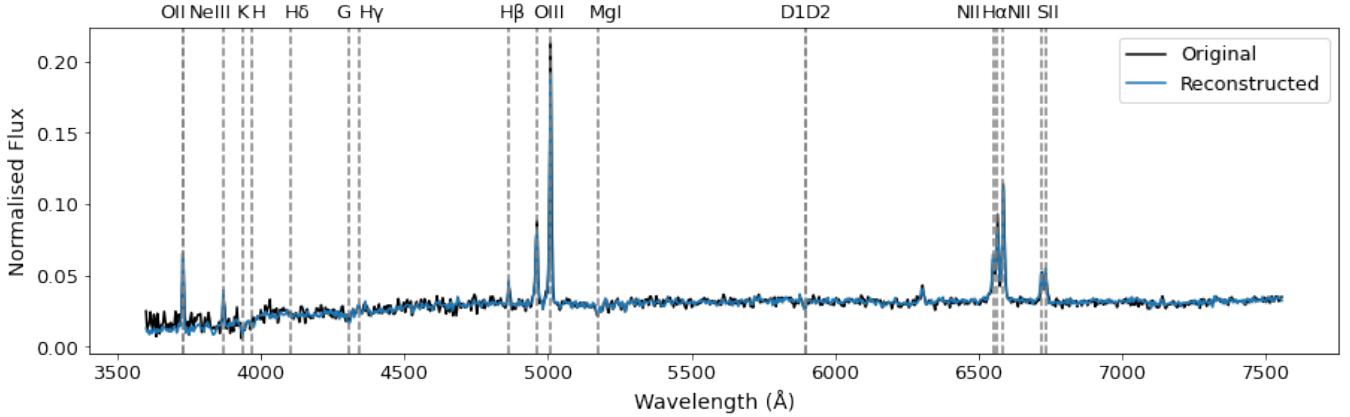


**Figure 5.** Original (black) and reconstructed (blue) spectra of **three emission-line galaxies**. The top row shows the VAE reconstruction of the  $H\alpha \lambda 6565$ ,  $[N\text{ II}] \lambda 6550, \lambda 6585$  and  $[S\text{ II}] \lambda 6718, \lambda 6733$  lines and the bottom row the  $[O\text{ II}] \lambda 3727, \lambda 3730$  lines. Target IDs left to right: 39627974831507323, 39633342945036446 and 39628521886189863.

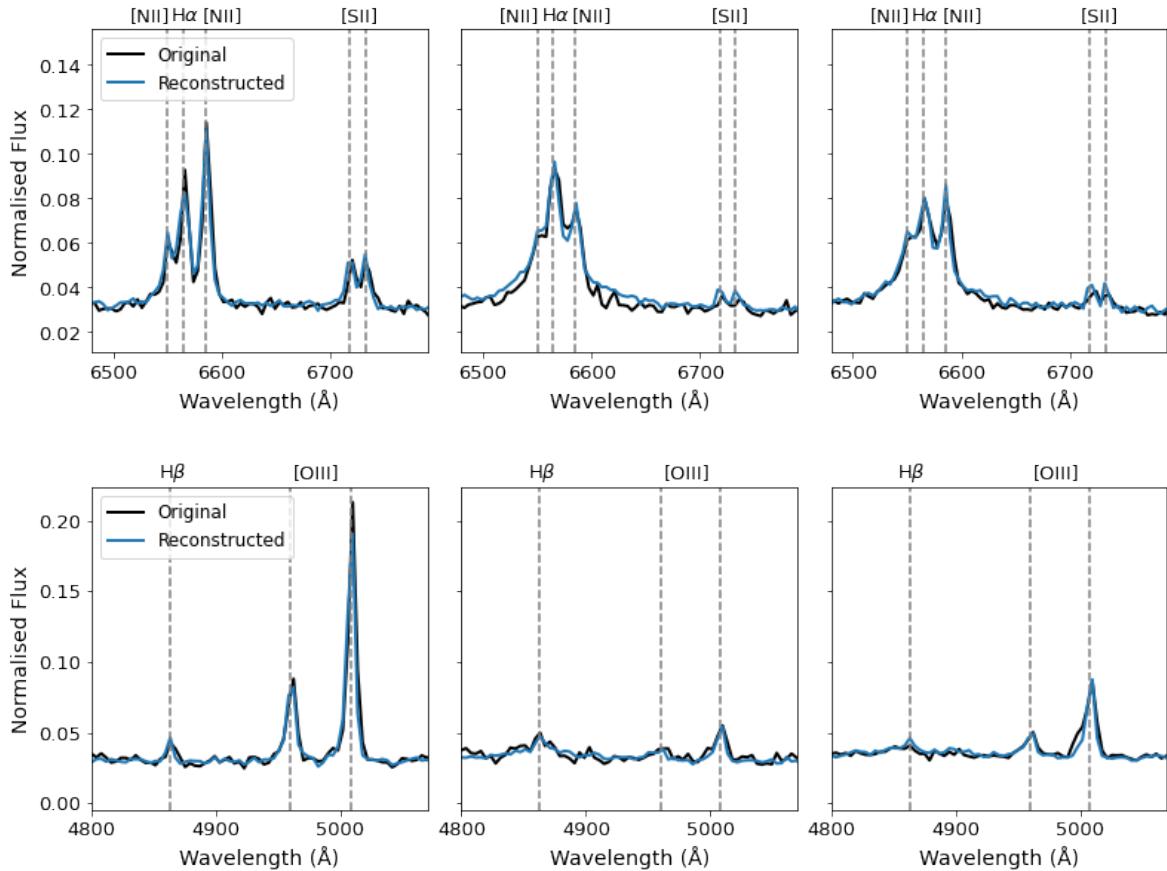
the dataset. The VAE hasn't seen enough star and quasar spectra examples in order to learn the underlying correlations and the most efficient representation in the latent space. This demonstrates that the reconstruction error is an effective estimate of anomaly score.

### 5.1.1 Galaxy Spectra Reconstructions

Figure 2 shows the original and reconstructed spectrum of a quiescent galaxy obtained from the validation set. The median pixel S/N of the spectrum is 33 and the weighted MSE is 1.19. The VAE is able to match the continuum and absorption lines of the spectrum well. The 4000 Å break is present and captured well by the VAE. In Figure 3 we zoom in on regions of the spectrum that contain spectral



**Figure 6.** Original (black) and reconstructed (blue) spectrum of an AGN from the validation set. The median pixel S/N of the galaxy is 17. The spectrum is in the rest frame. The dashed grey lines indicate typical spectral lines labelled at the top of the plot. Target ID: 39633118658822194.

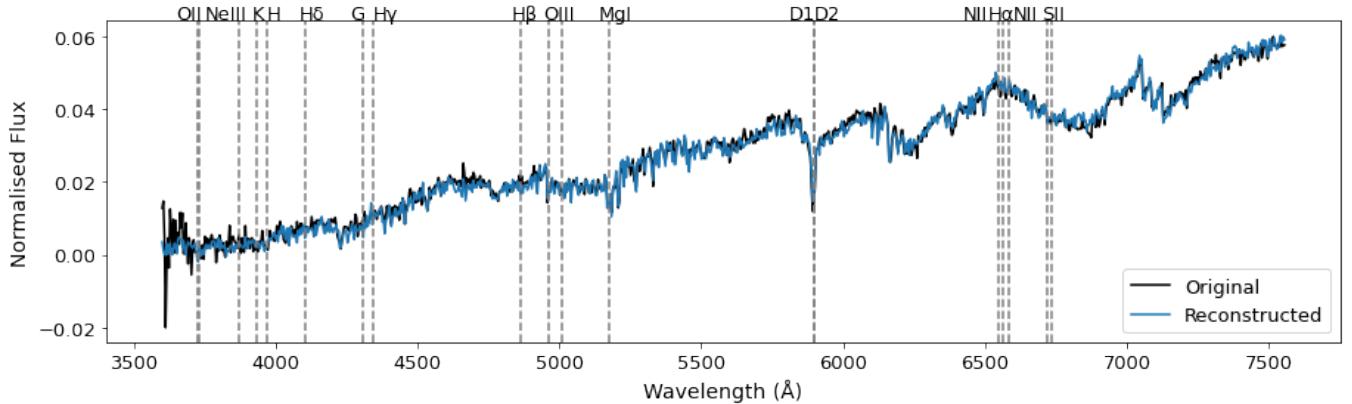


**Figure 7.** Original (black) and reconstructed (blue) spectra of three AGNs. The top row shows the VAE reconstruction of the  $H\alpha \lambda 6565$ ,  $[N\text{ II}] \lambda 6550, \lambda 6585$  and  $[S\text{ II}] \lambda 6718, \lambda 6733$  lines and the bottom row the  $H\beta \lambda 4863$  and  $[O\text{ III}] \lambda 4960, \lambda 5008$  lines. Target IDs left to right: 39633118658822194, 39628133594305704 and 39627794468045508.

lines. We show the  $\text{Ca II K}$  and  $H \lambda 3935, \lambda 3970$  which are prominent absorption lines in the spectra of solar-type and cooler stars due to singly ionised calcium and the  $\text{Na I D}$ -lines  $\lambda 5896, \lambda 5898$  which occur in the yellow region of the spectrum due to sodium. The VAE is able to reconstruct well the absorption lines, capturing well both the width and amplitude but in some cases the VAE reconstruction undershoots or overshoots slightly, failing to exactly meet the amplitude

of the spectral line. Getting the best possible reconstruction accuracy is beyond the scope of this work but to achieve this we would suggest using higher resolution spectra and a more curated sample as well as optimising the VAE hyperparameters.

Figure 4 shows the original and reconstructed spectrum of an emission-line galaxy with a median pixel S/N of 14. The VAE is able to reconstruct it well with a weighted MSE of 0.84. The spectrum



**Figure 8.** Original (black) and reconstructed (blue) spectrum of a **star** (M type) from the validation set. The median pixel S/N of the spectrum is 27. The spectrum is in the rest frame. The dashed grey lines indicate typical spectral lines labelled at the top of the plot. Target ID: 39627582181739826.

shows strong emission lines in  $\text{H}\alpha$ ,  $[\text{N II}]$ ,  $[\text{S II}]$ ,  $\text{H}\beta$  and  $[\text{O II}]$  indicating that star formation is occurring. In Figure 5 we zoom in on two regions of interest to observe more closely the  $\text{H}\alpha$ ,  $[\text{N II}]$ ,  $[\text{S II}]$   $\lambda 6718, \lambda 6733$ ,  $\text{H}\beta$   $\lambda 4863$  and  $[\text{O II}]$   $\lambda 3727, \lambda 3730$  emission lines. The VAE is able to reconstruct these very well, accurately mapping the width and height of the peaks and does not show any blending of the lines that are closely located.

### 5.1.2 AGN Spectra Reconstructions

Figure 6 shows the reconstructed spectrum of an active galactic nucleus (AGN), selected according to the **redrock** SPECTYPE of QSO, with median pixel S/N of 17. The VAE is able to reconstruct the AGN spectrum with a weighted MSE of 1.68. In Figure 7 we zoom in on regions of interest. We observe the significant broadening of emission lines especially around the  $\text{H}\alpha$  region. Due to its non-linear nature, the VAE is able to capture well the width of the emission lines. Portillo et al. (2020) have also shown that VAEs successfully reconstruct broad emission lines overcoming the limitations of PCA which requires more components, compared to a VAE, to adequately reconstruct the width of broad lines.

### 5.1.3 Star Spectrum Reconstruction

The reconstructed spectrum of a star is shown in Figure 8. The spectrum has a median pixel S/N of 27 and the VAE is able to reconstruct it with a weighted MSE of 1.8. Even though there were comparatively few stars within the training set, the VAE still manages to reconstruct the stellar spectra with only slightly higher mean error than the galaxies and quasars.

## 5.2 Outlier Identification

VAEs can be used to identify anomalies in a dataset. Outliers are expected to be badly reconstructed leading to high MSE or be placed in an isolated region in the latent space, i.e. the lower-dimensional representation produced by the VAE encoder. An anomalous spectrum may have both a high MSE and be isolated in latent space, but this is not necessarily true for all outliers as the latent space position does not only depend on the reconstruction error but depends on the entire loss function of the VAE. For example, an outlier, might be reconstructed badly because the VAE placed it together with a group

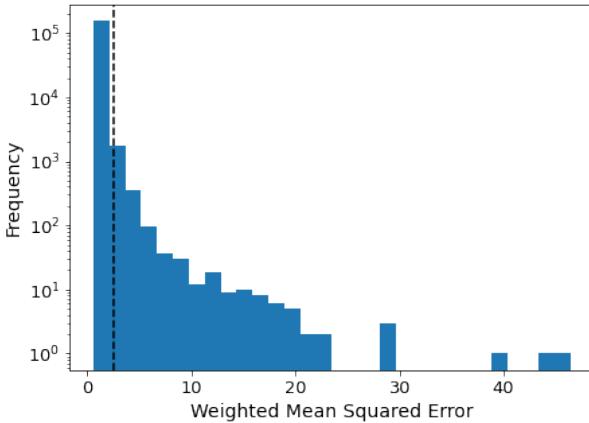
of typical spectra in the latent space. In this case, the VAE will attempt to reconstruct the anomalous spectrum as a normal one. This will lead to a bad reconstruction and a high MSE, but the spectrum will not be located in an isolated region. The spectrum will have a similar local density to its neighbours (as it is in a populated area) but it will have a much higher MSE than its neighbours. A case where a spectrum is located in an isolated region of the latent space but does not have a high MSE might be a low S/N spectrum where the pixel errors are high. Since the MSE is weighted by the inverse variance this will lead to a lower MSE. In this study we explore both approaches for identifying outliers and compare the results.

### 5.2.1 Reconstruction Error Approach

First, we explore the reconstruction error approach for detecting outliers. For each spectrum in the training set we calculate the weighted MSE between the original and reconstructed spectrum. For the purposes of this work we have chosen to identify the most extreme 1% of points as outliers for investigation. Figure 9 shows a histogram of the weighted MSE from all spectra. As expected, most spectra have a low MSE. The 1% threshold is indicated by the vertical dashed line and corresponds to a MSE of 2.55. This threshold suggests that spectra with an MSE higher than 2.55 are outliers while the rest are inliers, resulting in 1,560 outliers and 154,392 inliers. It is interesting to note that the average MSE of quasars reported earlier in Section 5.1 is 2.57 which falls within the 1% threshold. This is because quasars are rare in the training sample and as such they are identified as outliers by the VAE.

In Figure 10 we plot the top four outliers ranked according to highest MSE from top to bottom, together with images retrieved from the Legacy Survey sky viewer (data release 9).<sup>2</sup> The spectrum with the highest reconstruction error is shown in Figure 10 (a) and has a MSE of 46.47. It has an  $\text{H}\alpha$  emission line of extreme strength. The flux ratio of  $\text{H}\alpha$  to any other emission line present is extremely large, which is uncommon in the dataset. The spectrum exhibits emission lines in the Balmer series,  $[\text{O II}]$  and  $[\text{N II}]$  primarily, with weaker emission lines in  $[\text{O III}]$   $\lambda 4960, \lambda 5008$ ,  $\text{He I}$   $\lambda 3889$  and  $[\text{S II}]$  (shown in the magnified inset). This suggests intense star formation in a low-metallicity starbursting region as part of a galaxy. This results in

<sup>2</sup> <https://datalab.noirlab.edu/ls/ls.php> – images maybe retrieved for objects with a given Target ID, RA and dec.



**Figure 9.** Histogram of the weighted MSE for all the spectra in the training set. The vertical dashed line represents the 1% threshold and corresponds to an MSE value of 2.55. Spectra with a MSE higher than this threshold constitute the top 1% outliers in the dataset. Note the y scale is logarithmic.

lines which are very strong and also narrow. Such extreme emission line ratios are atypical in the dataset and thereby this spectrum is flagged as an outlier. From the image next to the spectrum we can observe that is not a galaxy but rather a fiber that has been placed on a particularly bright region inside of a very extended nearby galaxies. It is promising that the VAE is able to identify this as an outlier as it can be used as a tool for filtering bad spectra and providing a clean sample of galaxies.

Figure 10 (b) shows the spectrum of a quasar. `redrock` classifies this spectrum as a galaxy with an estimated redshift of 0.257. However the spectrum corresponds to a quasar found at a much higher redshift, around 2. Applying the correct redshift would set the rest-frame spectrum in the ultraviolet part of the EM spectrum. The correct rest-frame spectrum corresponding to this object is not plotted here as it is well outside the wavelength range of this study, but we annotate three main features ( $\text{Mg II}$ ,  $\text{C III}$ ] and  $\text{C IV}$ ) on the plot to demonstrate the correct labelling of features and aid with visualisation. The VAE is unfamiliar with these out-of-place features and hence does not reconstruct them well resulting in a MSE of 44.72. This example shows, that using the VAE we are able to identify spectra that have been assigned a wrong class and redshift by the DESI spectroscopic pipeline. There has been much work in the AGN community to amend these types of errors using traditional methods, so this is an example of how VAE-based methods might be able to complement these efforts.

Figure 10 (c) shows the spectrum of an AGN with both broad and narrow features, suggesting a Seyfert 1 galaxy. As discussed earlier, AGN spectra constitute less than 1% of the dataset with Seyfert 1 being a fraction of that. As these are rare in the dataset, they are picked up as anomalous.

Conversely to the case of the outlier spectrum in Figure 10 (b) where `redrock` underestimated the redshift of the source, in (d) the redshift is overestimated. The spectrum exhibits the shape of a black body with multiple absorption lines suggesting the source as a star. However the spectrum was labelled as a galaxy by `redrock` with a redshift of 0.0692. The true redshift of the star is smaller ( $z \sim 0$ ). The correct rest-frame spectrum is shown in purple where the two strong absorption lines at the blue end of the spectrum now match the  $\text{Ca II K}$  and  $\text{H}$  lines and the 4000 Å break is matched to the correct wavelength.

Inspecting by eye the 100 highest MSE spectra, the outliers fall

into one or more of the following broad categories: extreme emission line(s), stars, AGNs, discontinuities arising from miscalibration between neighbouring cameras, wrong assigned redshift, cataclysmic variables, and very broad emission lines (with the full width at half maximum spanning several hundred Angstroms). The top 100 outliers have a median MSE of 13 and median S/N of 42.

### 5.2.2 Local Outlier Factor Approach

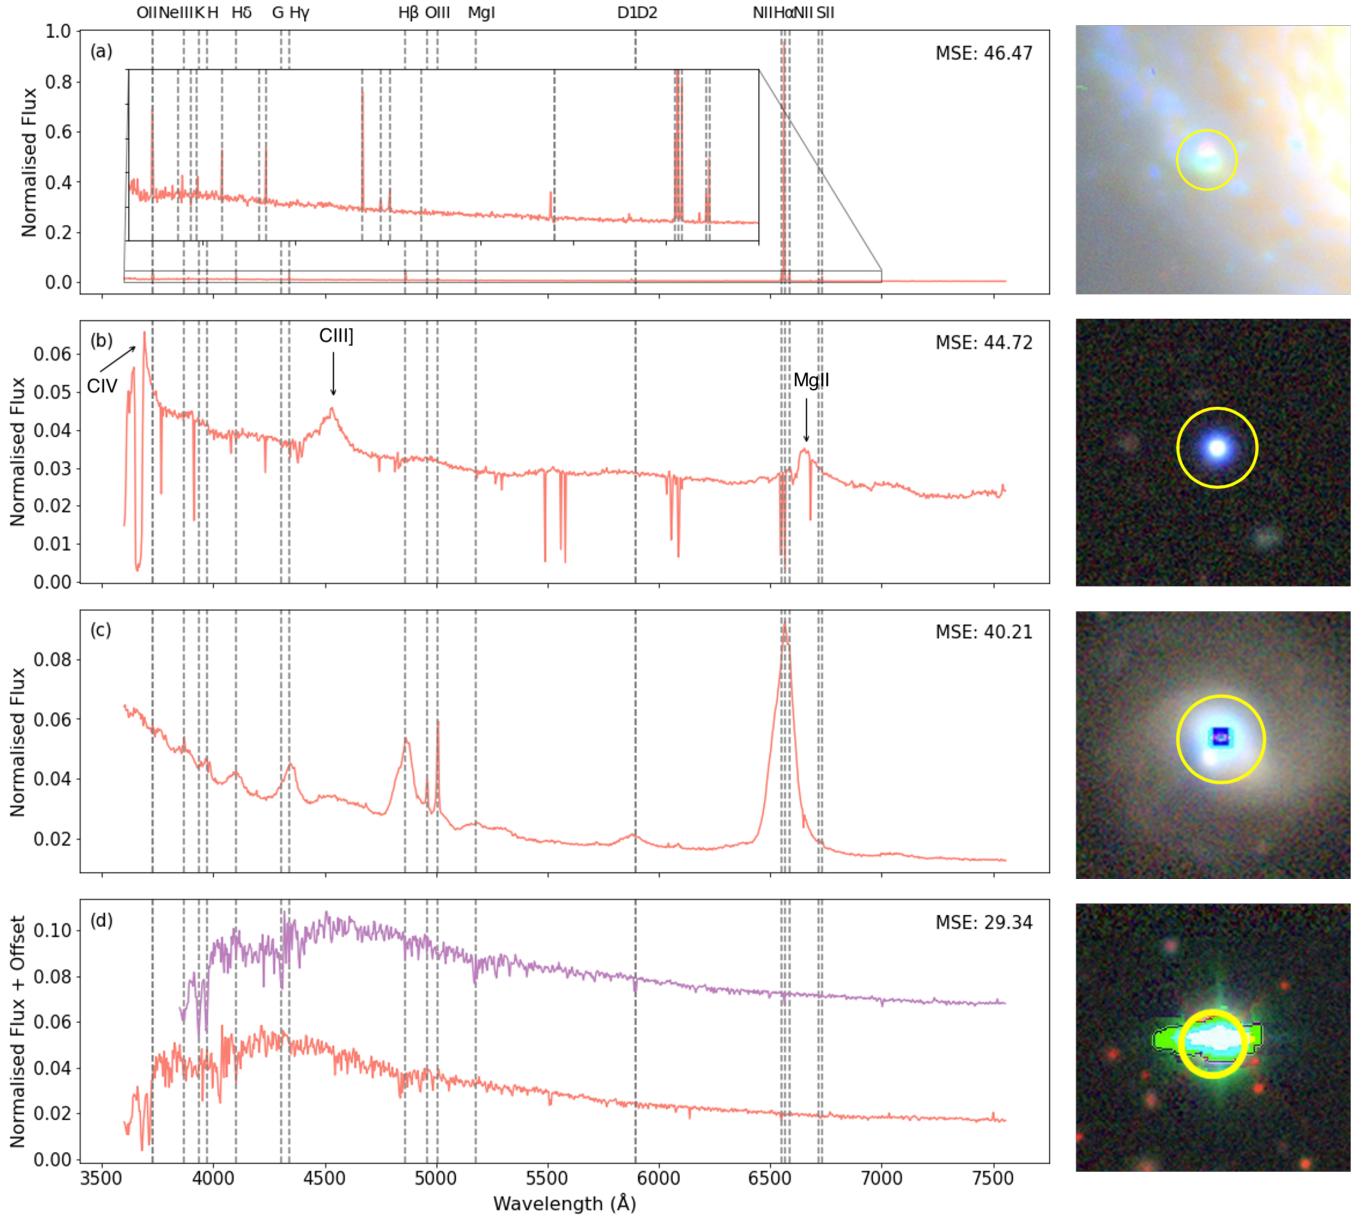
In Section 5.4 we shall explore how the VAE latent space is organised and how it arranges data points of different object type. For now, we use the VAE latent space to find outliers, specifically by identifying those points which are in relatively isolated areas of the latent space. To do this, we use the local outlier factor (LOF) algorithm (Breunig et al. 2000). This unsupervised algorithm estimates the local density of a given spectrum in the latent space and compares it to the local density of its  $k$  nearest neighbours. A spectrum is considered an outlier if its local density in the VAE latent space is substantially lower than that of its neighbours. Here we use the LOF algorithm with  $k = 20$  nearest neighbours. The LOF algorithm computes the negative outlier factor (NOF) for each spectrum. Inliers tend to have a NOF close to  $-1$  while outliers tend to have more negative NOF scores. Figure 11 shows a histogram of the NOF of all training samples. Most spectra have a NOF close to  $-1$  indicating that these are inliers. The 1% threshold, shown by the vertical dashed line, corresponds to a NOF of  $-1.44$  suggesting that spectra with a NOF lower than  $-1.44$  are considered as outliers.

In Figure 12 we plot four outliers of different type according to lowest NOF. Figure 12 (a) and (b) have the lowest and second lowest NOF, (c) has the fifth and (d) the tenth lowest score. In order to present a greater variety of outliers, similar spectra have been skipped in this selection. The spectrum in Figure 12 (a) has a NOF of  $-4.13$ . This galaxy spectrum has a large dip in the flux with negative flux values. This is probably attributed to a bad sky subtraction or bad pixels which were not caught by the DESI spectroscopic pipeline. We also suspect that the assigned redshift is underestimated. The arrow points to what we suspect is the misplaced 4000 Å break.

The spectrum in Figure 12 (b) is of an M star with particularly high flux in the red part of the spectrum. The spectrum in Figure 12 (d) is also of a star.

The spectrum in Figure 12 (c) shows extreme emission in  $\text{H}\alpha$  which is similar to the top outlier from the MSE approach. The inset provides a magnification of the weaker emission lines. This spectrum has a MSE of 8.46 corresponding to the 103rd highest MSE. In fact, the spectrum with the highest MSE (Figure 10 (a)) has a NOF of  $-3.38$  which corresponds to the seventh lowest NOF score. All 4 outliers presented here obtained using the LOF approach are also in the top 1% of outliers with the highest MSE but their relative ranking is lower.

Inspecting by eye the top 100 outliers resulting from the LOF approach, we observe that these fall into one or more of the following broad categories: bad spectral regions which were not flagged by the mask vector and hence were not treated in the preprocessing, low S/N, extreme emission line(s), stars, AGNs, discontinuities due to bad calibration and wrong redshift. The top 100 outliers have a median NOF of  $-2.25$  and a median S/N of 7. While the top 100 outliers of the two methods have some common categories, it is important to note the apparent difference in the median S/N. This suggests that increased noise affects the spectrum's position in the latent space. Lower S/N spectra have smaller weights and hence the deviation of the reconstructed spectrum from the original one is damped, keeping the calculation of the weighted MSE low. However, the noise



**Figure 10.** Top four outliers according to the highest reconstruction error of the VAE. The MSE of each spectrum is given in the upper right corner of each spectrum panel. The spectrum in panel (a) exhibits extreme H $\alpha$  emission strength. The inset provides a magnification of difficult to see emission lines of the spectrum. The spectrum in (b) is of a high redshift quasar whose redshift was underestimated by redrock. We demonstrate the correct feature labels in the annotations. The spectrum in (c) is of an AGN exhibiting broad emission lines. The spectrum in (d) plotted in coral is of a star whose redshift was overestimated. The correct rest-frame spectrum is overplotted in purple. The object corresponding to each spectrum is shown in the images on the right hand side. The yellow circle is a marker indicating the objects of interest of various angular sizes. The images were obtained from the Legacy Survey sky viewer (data release 9). Target IDs from top to bottom: 39627763589581458, 39627817943567887, 39627818572710520 and 39633152305530204.

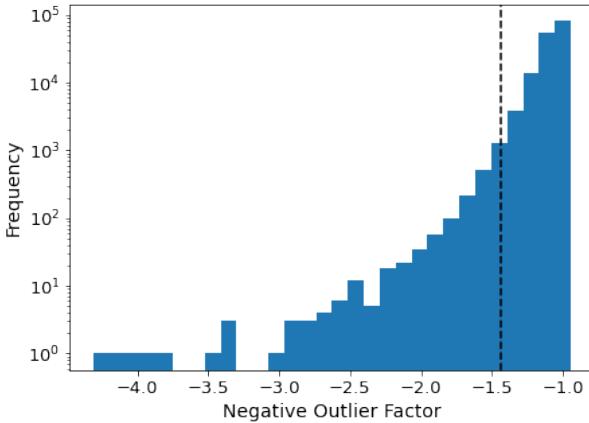
in the spectrum is still propagated through the encoder introducing some uncertainty to the spectrum's position in the latent space. As a result that the LOF approach is more sensitive to the S/N.

### 5.3 Astronomy

While the VAE has drastically reduced the number of spectra an astronomer would need to go through to identify anomalies, a common issue remains. As we have seen, the anomalies fall into many categories. An astronomer who is interested in spectra with transient event signatures may have little interest in outliers that were

flagged due to artefacts or wrongly assigned redshifts, whereas an astronomer in charge of data quality will be interested in the latter. The relevance of the outliers is inherently subjective and is based on the scientific goal. The solution we adopted in this study is Astronomaly (Lochner & Bassett 2021). Astronomaly offers a solution to the outlier relevance problem by using Active Learning to offer personalised curation of outliers to the user.

There are three main steps to the Astronomaly framework: feature extraction, anomaly detection, and Active Learning. Feature extraction refers to the dimensionality reduction of high dimensional data to a few components that contain the most critical information. In



**Figure 11.** Histogram of the negative outlier factor (NOF) for all the spectra in the training set calculated using the local outlier factor (LOF) algorithm. Spectra with scores significantly lower than  $-1$  are considered outliers. The vertical dashed line represents the 1% threshold and corresponds to a NOF of  $-1.44$ . Spectra with a NOF score lower than this threshold constitute the top 1% outliers in the dataset. Note the y scale is logarithmic.

our case, this is the latent space provided by our VAE and hence the feature extraction in Astronomaly's framework is skipped. Next, Astronomaly applies the LOF algorithm to the latent space provided, similarly to the methodology we demonstrated earlier in this study. The spectra are ranked according to lowest NOF and presented in this sorted order to the user in Astronomaly's interface. We adapted the interface to tailor it to our use case, for optimal plotting of spectra. Astronomaly also allows the user to observe where spectra reside in the latent space by using t-distributed Stochastic Neighbor Embedding (t-SNE, [van der Maaten & Hinton \(2008\)](#)) to project the 10D latent space to 2D. This second step of Astronomaly which involves calculating the anomaly score using the LOF algorithm could potentially be replaced with the reconstruction error, or more appropriately reconstruction probability method. Alternatively, the two approaches could also be combined taking into account both the position on the latent space and reconstruction probability.

In the third step, the expert can visually inspect the sorted spectra (from highest to lowest anomaly score) and provide a label (0-5) that reflects how relevant, or not, the specific outlier spectrum is to the user. The novelty Astronomaly introduces to our analysis is through this last step. Active Learning is used to adjust the original anomaly scores provided by LOF according to the relevance score provided by the user labels, learning in this way which anomalies the user is interested in. To predict the relevance score for the remainder of the dataset, Astronomaly uses a random forest regression algorithm ([Breiman 2001](#)). The original anomaly score of a data point is updated based on the relevance score predictions and its distance in latent space to the nearest human-labelled neighbour. For regions with few to no human labels, the score reverts to the original anomaly score provided by LOF which means that the user gets the opportunity to see and label the object, so new anomalies are less likely to be missed. In regions where confidence is high, the predicted relevance score dominates. When the Active Learning is trained on the human labels, the outliers are reordered according to the new relevance anomaly scores prioritising in this way outliers that are more relevant to the user. The process can be iterated as more human labels are added to provide improved outlier recommendations.

We applied Astronomaly to the VAE encoded DESI spectra. To qualitatively examine the effect of the Active Learning process, we

assigned a high relevance score label to spectra with a discontinuity due to miscalibration of the blue and red channels and to extreme emission lines (similarly to Figure 10 (a) and Figure 12 (c)). Providing labels to only the first 50 spectra shown and training the Active Learning step on these, we observed that the updated sorted list of outliers is now mostly composed of these two classes of spectra.

#### 5.4 Interpretation of the Latent Space

In addition to anomaly detection, one of the great advantages of VAEs is that they enable us to learn more about the bulk of the distribution through examination of the latent space. In Figure 13 we plot the latent representation of the training spectra across all 10 VAE latent dimensions using a corner plot. While these pairwise 2D projections of the 10D latent space do not capture the distribution of spectra in its totality, the scatter plots are still highly informative and useful when trying to interpret the latent space. The data are colour-coded according to their spectral classification. We observe that the VAE successfully separates the three classes without ever having seen the labels. The labels are strictly used for colour-coding plots only. The separation is particularly good for stars, evident in the fifth latent dimension for example. The scatter plot of VAE 5 and VAE 10 shows two distinct groups, the galaxies in blue and the stars in the green extending arm. Quasars are also separated, shown more clearly in the seventh and eighth dimensions. The scatter plot of VAE 5 and VAE 8 clearly shows the quasars represented in pink separating from the galaxies and stars.

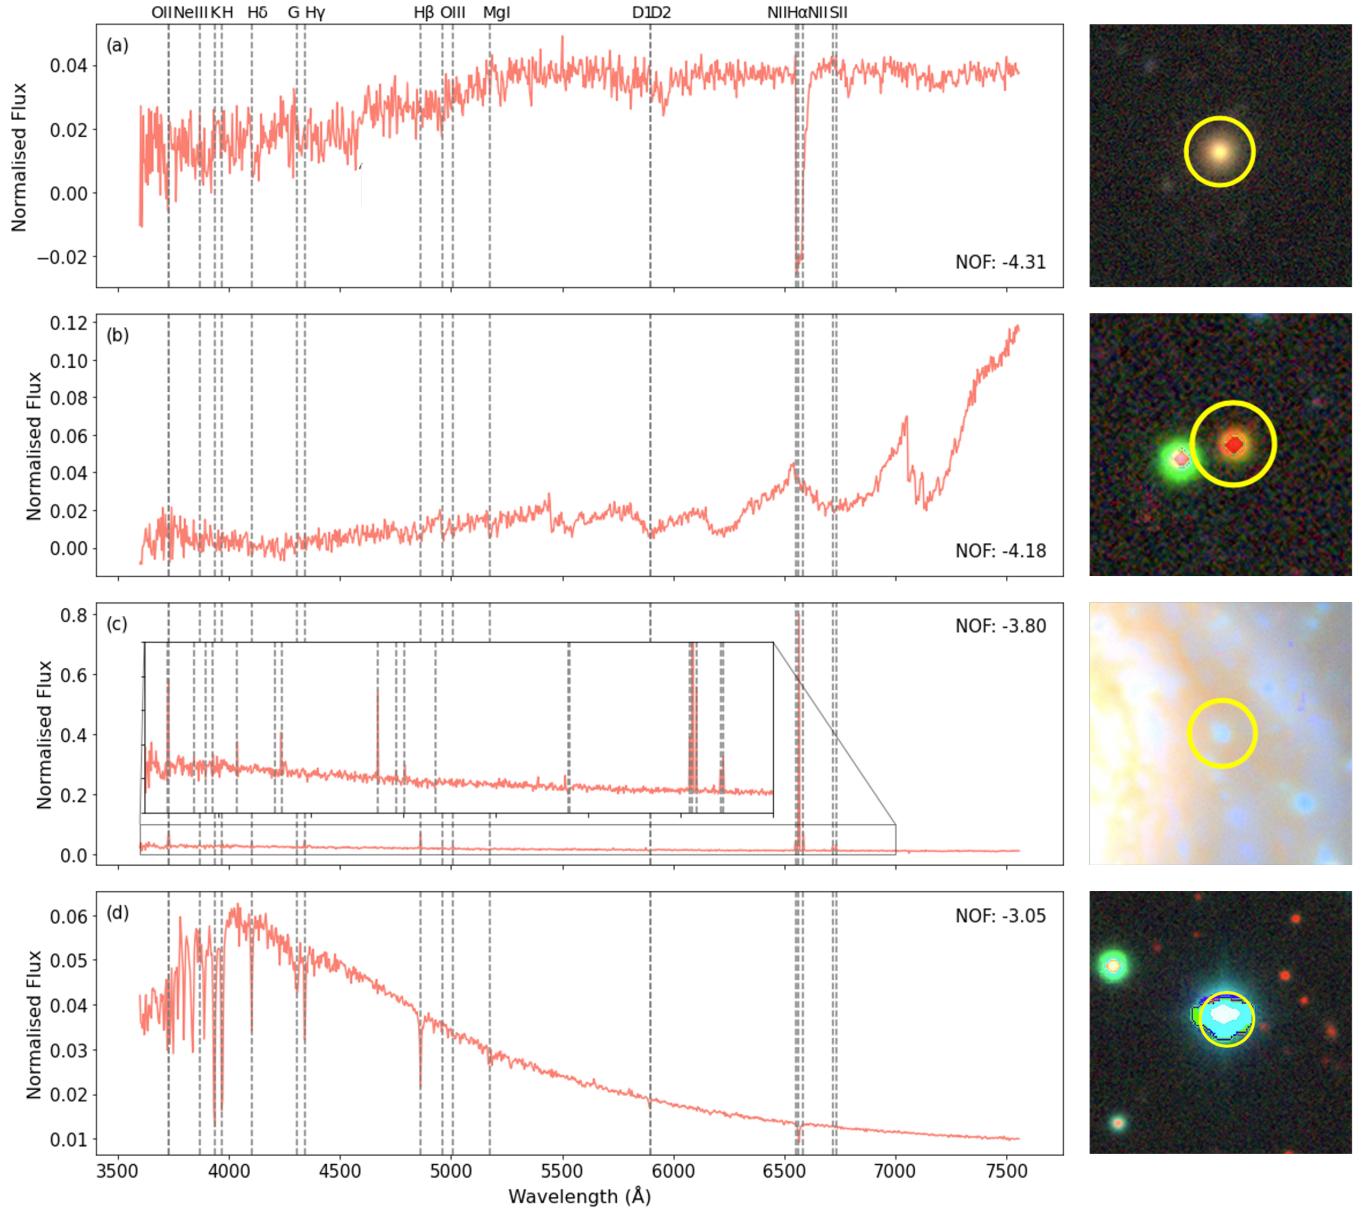
Beyond examining the class separation in latent space, it is insightful to also understand how sub-classes of spectral classes are distributed in the latent space. As we have not included any information on sub-classes, the way that we approach this is by examining the pairwise 2D projections of the 10D latent space focusing on a single class each time and extracting spectra, along a specified track in the latent space.

##### 5.4.1 Galaxy Tracks

Figure 14 shows the 2D scatter plot of VAE 5 and VAE 7 for the galaxy spectra only. From this plot we can observe a main tail extending to the lower right side of the plot (indicated by following the "B" labels) and a weaker extension to the upper right side of the plot (indicated by following the "R" labels). The tracks have been labelled R and B because the constituent spectra exhibit redder and bluer continua correspondingly. We extract an indicative galaxy spectrum from each point along the paths B1-B5 and R1-R5 and plot these in Figure 15 (top and bottom correspondingly).

We make sure that the galaxy spectrum corresponding to a point in latent space is representative of the majority of the spectra surrounding that point by inspecting several spectra that lie close to it. Another approach to ensure this, would have been to plot the stacked spectra of the galaxy spectra that fall within a small radius from the specified point in the latent space. Stacking spectra helps to better study the average properties of the galaxies involved and reveal fainter spectral signatures.

The R1 spectrum has absorption lines at Ca II H and K and Na I (D-lines) and a red continuum with a 4000 Å break indicating an old metal-rich stellar population. All spectra in the R-track show a decrease in flux in the blue part of the continuum. R2 shows emission lines in [O II], H $\alpha$  and [N II]. As we follow the R track these emission lines are strengthened and joined by [S II], [O III] and the rest of the Balmer series indicating star formation. The H $\alpha$

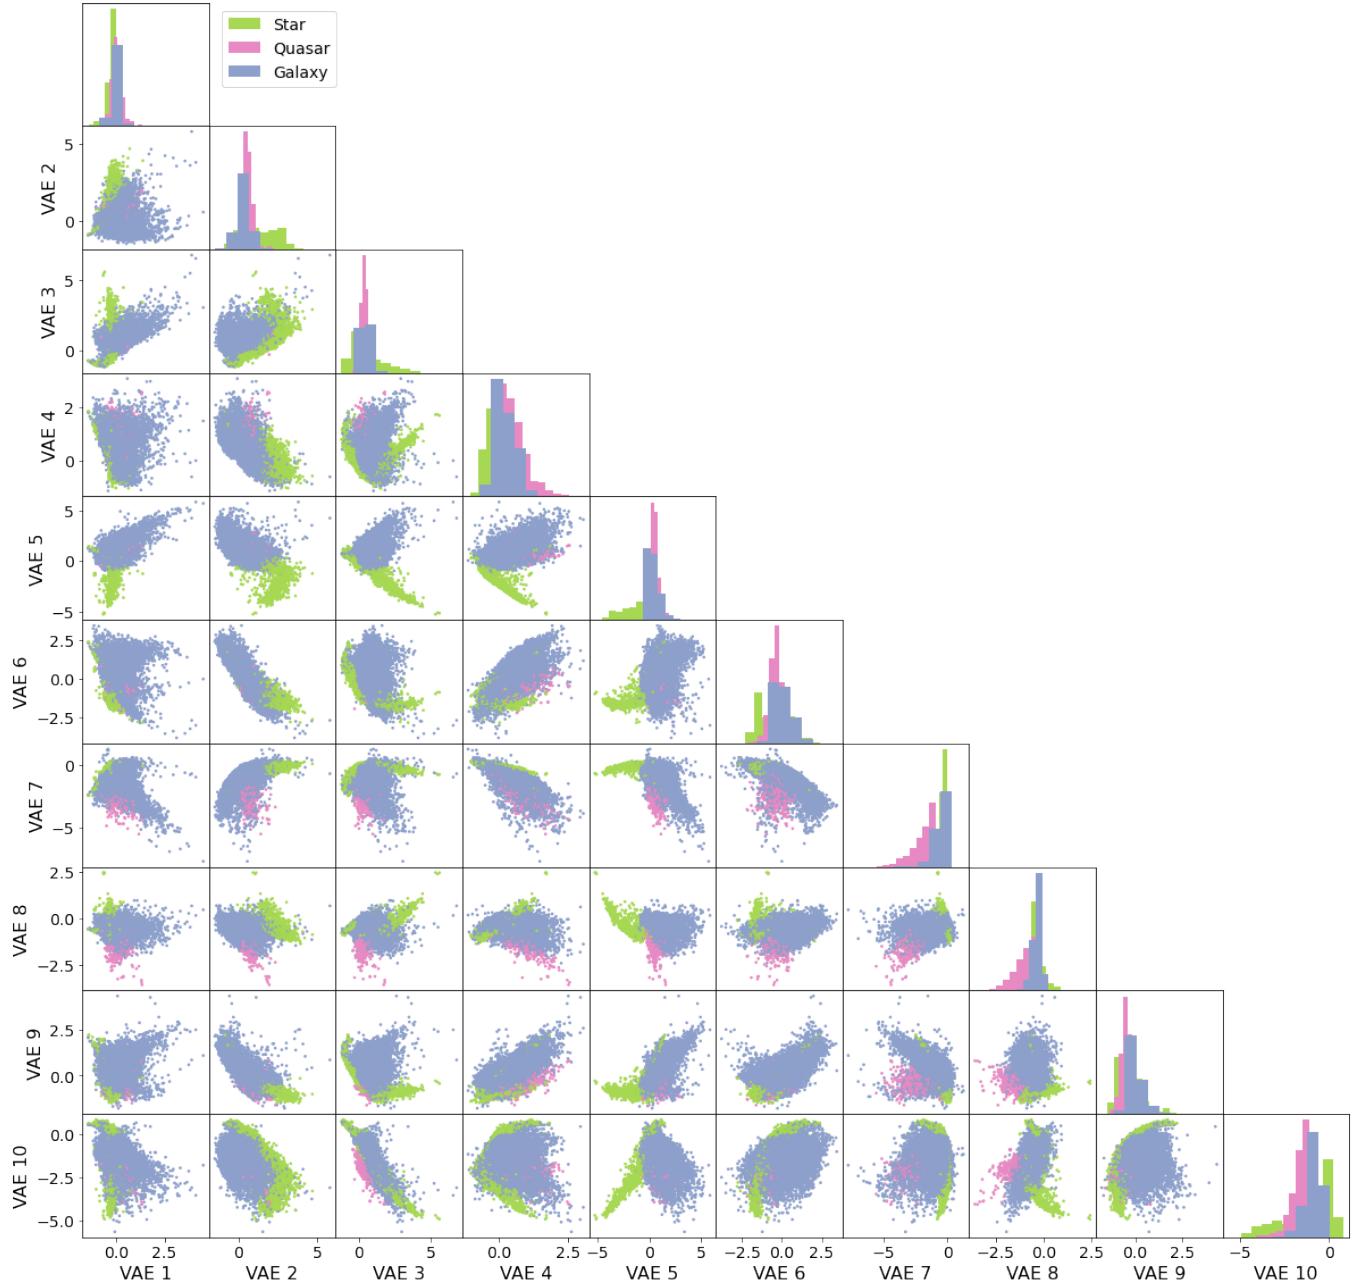


**Figure 12.** Four unique outliers according to the lowest negative outlier factor (NOF). The NOF of each spectrum is found at the upper or lower right corner of each panel. (a) has the lowest NOF, (b) the second lowest, (c) the fifth and (d) the tenth. Similar spectra have been skipped in order to avoid repeatability and present a greater variety of outliers. The galaxy spectrum in panel (a) exhibits a large dip in flux likely due to a bad sky subtraction. The arrow in (a) points to what we suspect is the misplaced 4000 Å break, suggesting that the redshift of the galaxy is underestimated. The spectrum in (b) is of a M star with particularly high flux in the red part of the spectrum. The spectrum in (c) exhibits an extreme emission strength in H $\alpha$ . The inset in (c) provides a magnification of the weaker emission lines of the spectrum. (d) shows the spectrum of a star. The object corresponding to each spectrum is shown in the images on the right hand side. The yellow circle is a marker indicating the objects of interest of various angular sizes. The images were obtained from the Legacy Survey sky viewer (data release 9).

emission line increases in strength as we follow the R track. The R5 spectrum shows a strong discontinuity around 5000 Å which is in the wavelength band where the data from the blue and red cameras of the DESI spectrograph are coadded. After inspecting the individual camera data we observed that the continuum flux around 5000 Å recorded by blue camera was shifted down relative to that of the red camera and therefore this discontinuity is likely an artefact arising from a miscalibration between the two channels. Interestingly, the R5 spectrum is isolated in the latent space indicating that the VAE has picked up on this atypical feature. Inspecting the spectrum closest

to R5 in the latent space we observed that it also exhibits a, slightly weaker, discontinuity at 5000 Å.

The B1 galaxy spectrum similarly to the R1 spectrum, does not have any strong emission features. In contrast to the R1 spectrum, its continuum is bluer and it does not exhibit the 4000 Å break indicating a younger stellar population. B2 is also a blue spectrum and has emission lines in [O II], H $\beta$ , H $\alpha$ , [N II] and [S II]. Following the B track, these lines increase in strength and are joined by the rest of the Balmer emission lines, [O III] and a weak He I emission line. The blue part of the spectra shows several emission lines. B3-5 exhibit strong H $\alpha$  emission with H $\alpha$   $\gg$  [N II] suggesting low-metal



**Figure 13.** Corner plot of the training spectra embedded in the 10 dimensional VAE latent space. The spectra are colour-coded according to their spectral classification.

starburst galaxies (Almeida et al. 2012). Again we can see that the  $H\alpha$  emission increases in strength as we follow the B track.

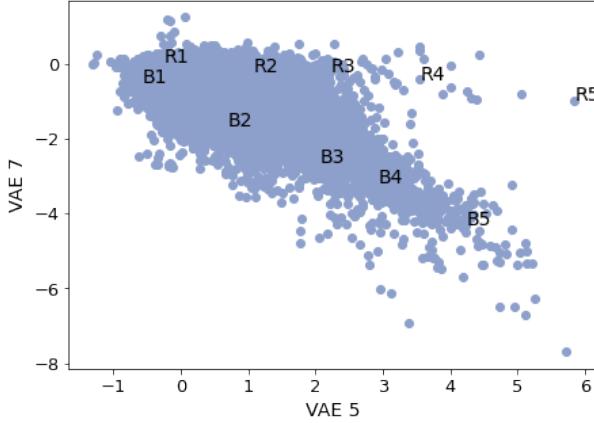
#### 5.4.2 Quasar Tracks

We repeat the analysis for quasar spectra. In Figure 16 we plot the 2D scatter of VAE 6 and VAE 8 for the quasar spectra only. We draw out two tracks, labelled Br (Broad) and N (Narrow) corresponding to two extensions from the main collection of scatter points. The spectra extracted along the two tracks are plotted in Figure 17, where the top plot shows spectra from the Br track and the bottom plot shows spectra from the N track.

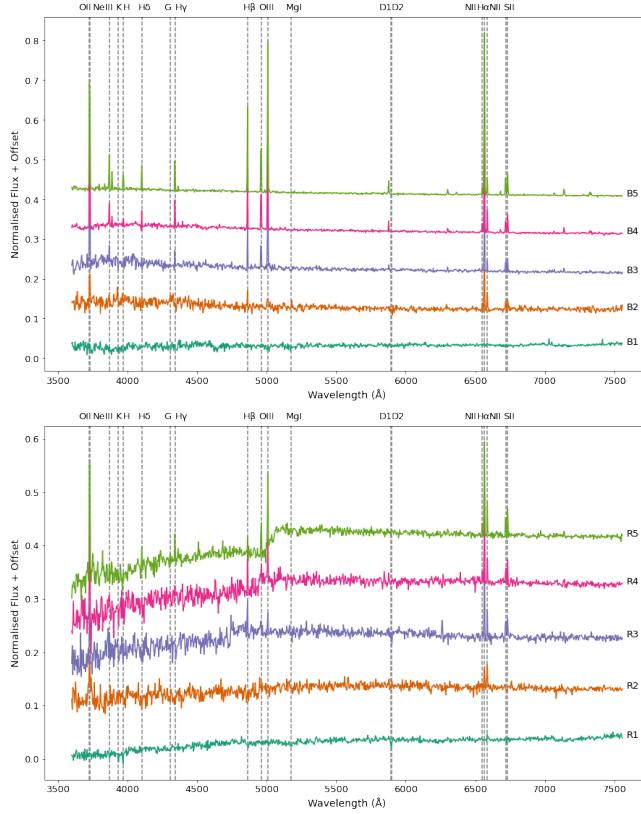
The Br1 spectrum shows emission lines in [O II], [O III],  $H\alpha$

and [N II]. We can observe that there is broadening in the spectrum around the  $H\alpha$  emission line. As we follow the Br track, the  $H\alpha$  and [O III] emission lines, in particular, increase in strength and decrease in broadening. Broad emission lines for the rest of the Balmer series are also visible, especially in Br4 and 5. We also observe that the spectra exhibit a blue continuum. As the spectra in the Br track have both broad and narrow spectral lines the most likely source is Seyfert 1 galaxies.

The N track on the other hand, has only narrow emission lines. The N1 spectrum has emission lines in [O III],  $H\alpha$  and [N II]. Subsequent spectra in the N track show increasing strength in these lines as well as emission lines in [O II], [Ne III]  $\lambda 3870$ ,  $H\beta$  and [S II]. The [O III] emission is particularly strong with  $H\beta$  being much weaker.



**Figure 14.** Scatter plot of the fifth and seventh VAE dimensions of the **galaxy** spectra only. The labels draw out two tracks, B and R, which correspond to bluer and redder continua in spectra.

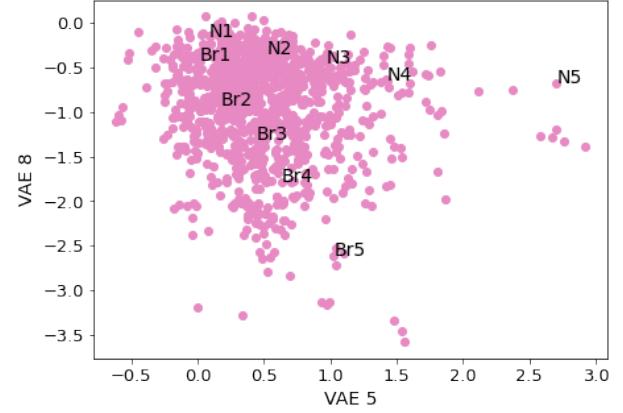


**Figure 15.** Galaxy spectra along the B (top) and R (bottom) tracks drawn in the VAE latent space shown in Figure 14. The normalised spectra are plotted with an arbitrary offset for visibility. The R track shows redder spectra whereas the B track shows bluer spectra.

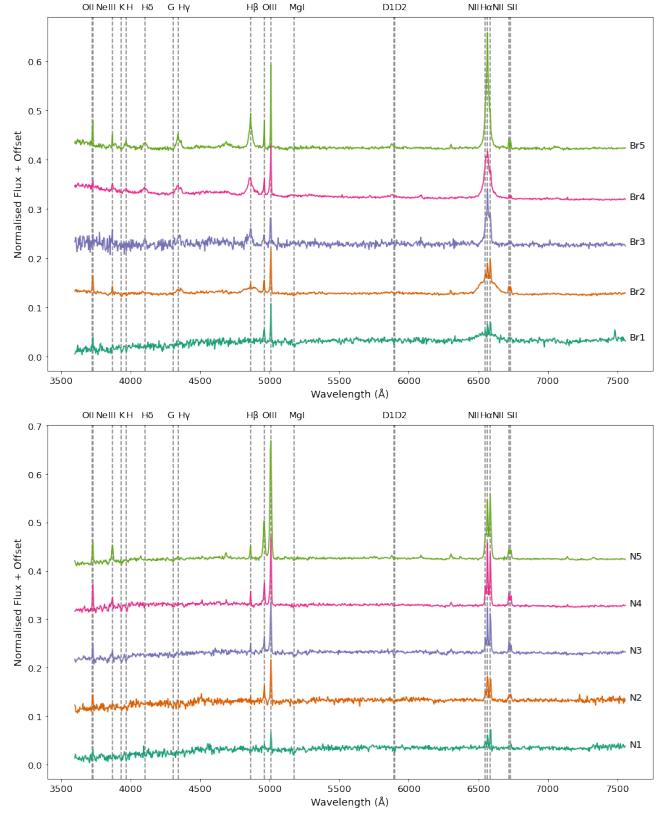
We also notice that for these galaxies  $\text{H}\alpha$  is of similar strength to  $[\text{N II}]$ . These characteristics suggest that these spectra correspond to Seyfert 2 galaxies.

#### 5.4.3 Star Tracks

Finally we examine the 2D scatter plot of VAE 5 and VAE 10 for the star spectra only, shown in Figure 18. The star spectra projected in



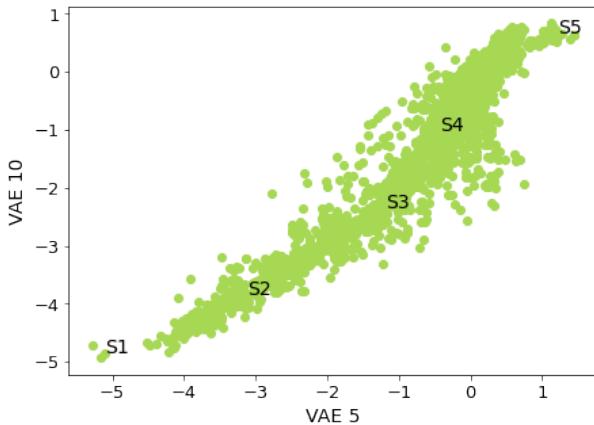
**Figure 16.** Scatter plot of the fifth and eighth VAE dimension of the **quasar** spectra only. The labels draw out two tracks, Br and N, which correspond to broad and narrow spectra.



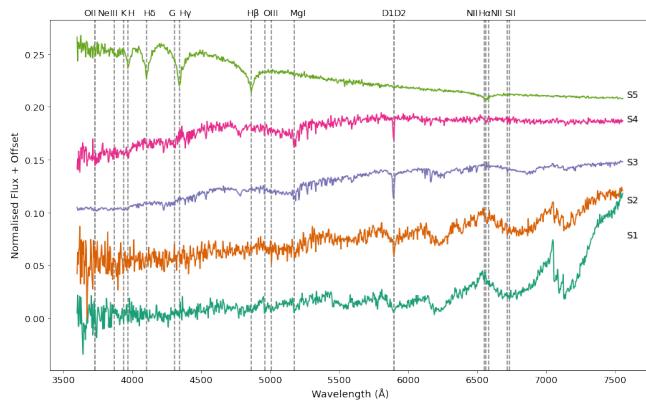
**Figure 17.** AGN spectra along the Br (top) and N (bottom) tracks drawn in the VAE latent space shown in Figure 16. The normalised spectra are plotted with an arbitrary offset for visibility. The Br track shows broad spectra whereas the N track shows narrower spectra.

this 2D latent space are arranged in an elongated structure. Here we examine a single track, labelled S, following the distribution of star spectra from the bottom left to the top right of the plot.

The spectra are plotted in Figure 19. The S1 and S2 spectra are from low temperature stars with a rightward rising profile that peaks in the red/infrared. We can also observe absorption in the  $\text{Na I D}$ -lines and the  $\text{TiO}$  (Titanium oxide) bands at  $\lambda 4955, 7150$ . These features are typical of M stars. Note how similar the spectrum in Figure 19



**Figure 18.** Scatter plot of the fifth and tenth VAE dimension of the star spectra only. The labels draw out a single track labelled S.



**Figure 19.** Star spectra along the S track drawn in the VAE latent space shown in Figure 18. The normalised spectra are plotted with an arbitrary offset for visibility.

in dark green, which corresponds to the S1 point in Figure 18, is to the M star with high flux in the red part of the spectrum identified in Figure 12 (b). S3 and S4 show stronger absorption in the Na I D-lines and other neutral metals such as Mg I. The continuum again shows a rise in the red part of the spectrum but the gradient is less steep than that of M stars. The spectra at S3 and S4 are typical of K stars. The latent space distribution of the stellar spectra suggests that the VAE has successfully separated K and M stars, with M star spectra found at the bottom half of the elongated structure in the latent space and K stars at the upper part. Finally we inspect the region found at the very top of the plot that is slightly separated from the main body, represented by S5. The spectrum at S5 is significantly bluer and exhibits strong absorption lines mainly from the hydrogen atom (Balmer series). The absorption lines are significantly broadened, indicating that this is a white dwarf. While inspecting the surrounding area of S5 in the latent space, we also observed stellar spectra similar to the one corresponding to a white dwarf but without broadening of the spectral lines. These are spectra originating from hot, blue stars such as B and A stars. These have been placed by the VAE close to the white dwarf spectra in latent space, as they all exhibit very a blue spectral profile with strong absorption lines from the hydrogen atom.

#### 5.4.4 VAE Generated Spectra

In order to better understand the information captured by individual VAE components, we generate spectra, using the trained VAE by sampling from the latent space along individual VAE latent dimensions. The coordinates of the generated spectra in the latent space are fixed to the mean value of each latent dimension, with only one dimension being varied at each time. We repeat this for all VAE latent dimensions. The results are plotted in Figure 20. The middle spectrum of each subplot in Figure 20 plotted in purple is common to all and represents a spectrum generated from the centroid of all spectra's latent means.

To assist in the interpretation of Figure 20, in Table 1 we summarise the changes to the main spectral features that occur, as we move along each latent dimension (in the increasing direction). The continuum is observed to remain unchanged in shape and amplitude in most latent dimensions. The red region of the spectrum is observed to decrease in VAE 5, 6 and 10 with 6 also exhibiting a weakening in the 4000 Å break along with an increase in the blue region of the spectrum. In contrast, VAE 7 shows a decrease in the blue part of the spectrum which we observed earlier in Figure 14 where redder spectra had higher values in this dimension. These dimensions capture information about the age of the dominant stellar population. VAE 6 and 7 are also able to describe broadened emission lines as we can observe larger widths in the Balmer series lines. VAE 2 and 10 both show decrease in all emission lines leading to a spectrum with mostly absorption lines. VAE 4 also shows decrease in all emission lines, except H $\alpha$  and [N II] which increase in strength. Conversely, VAE 5 shows an increase in all emission lines. Both VAE 5 and 10 show a strong TiO band at  $\lambda$ 7150 which is expected as low values of both VAE 5 and 10 are used to describe M stars. We tentatively observe also that decreasing VAE 5, seems to decrease star formation and increase stellar age.

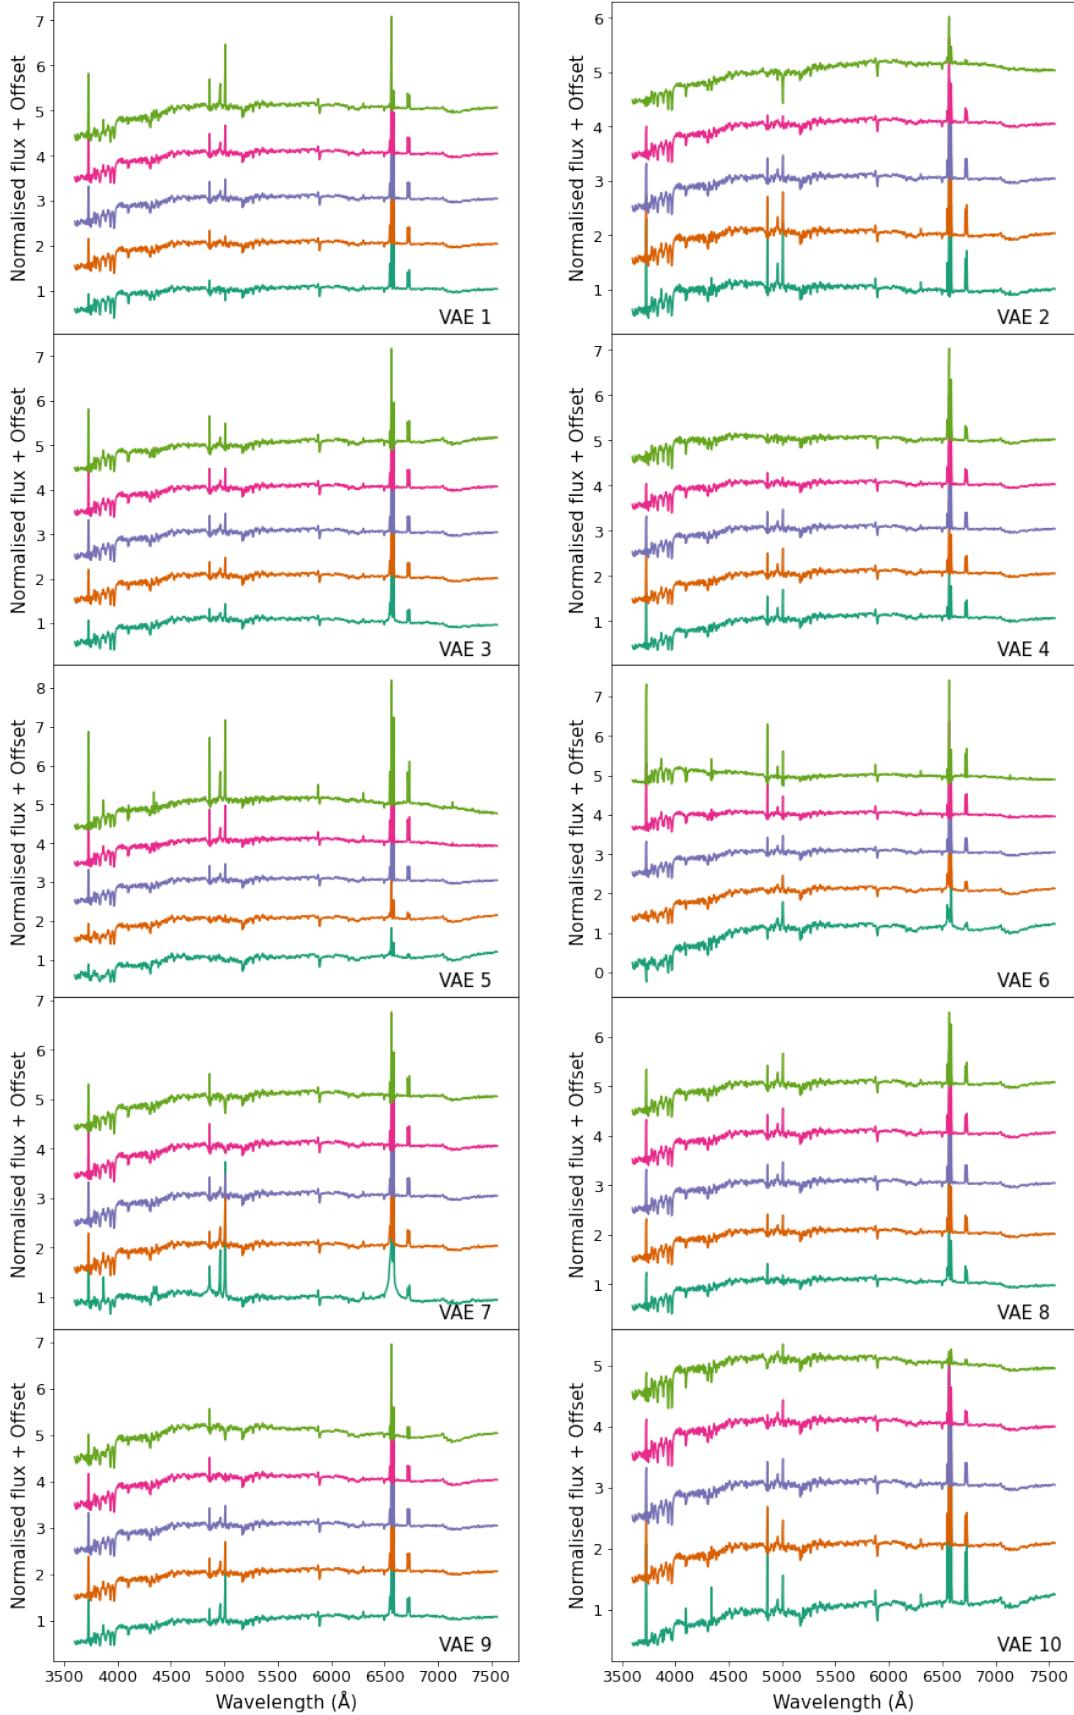
Studying the VAE latent dimensions individually allows us to interpret these and relate them to the physical attributes of the spectra. However, we should be aware that these interpretations change when multiple dimensions are changed from their mean at the same time. Collectively, the VAE dimensions, give the VAE the flexibility to reconstruct complex spectra by making use of non-linear combinations. In addition, any physical trends suggested by this work are intended to be indicative and would need to be confirmed by traditional techniques such as stellar population synthesis fitting.

## 6 DISCUSSION AND CONCLUSIONS

In this paper we have:

- used a VAE to detect anomalous DESI spectra, finding them as either outliers in the VAE latent space or those which are reconstructed poorly compared with their original representation;
- used Astronomaly to curate these sets of outliers further making visual inspection more effective; and
- explored the VAE latent space, showing it can be used for unsupervised classification and that it can encode physical characteristics of the dataset.

The increasing complexity and volume of astronomical spectra, requires the use of data intensive techniques. In this work, we have shown that a VAE is capable of effectively compressing the information found in spectra by a factor of 100, whilst retaining enough information to accurately reconstruct the different features present in spectra as shown in Section 5.1. PCA has also been shown to effectively compress spectra (Ronen et al. 1999; Madgwick et al. 2003a;



**Figure 20.** Synthetic spectra generated using the VAE, along each VAE latent dimension (in the increasing direction) while keeping the rest of the VAE dimensions fixed to their means. The middle spectrum in each subplot is common in all and corresponds to the centroid of the VAE latent means. The spectra are plotted with offset for visibility.

**Table 1.** Summary of how the main spectral features change as we move along the specified VAE latent dimension (L. D.), in the increasing direction, as depicted in Figure 20. The dash indicates that the feature remains approximately constant for that VAE dimension, the  $\uparrow$  indicates that the feature increases in strength and the  $\downarrow$  indicates that it decreases.

| VAE L. D. | Continuum   | H $\alpha$   | [N II]       | H $\beta$    | [O III]      | [O II]       | Broadening |
|-----------|---|--------------|--------------|--------------|--------------|--------------|------------|
| 1         | -   | $\uparrow$   | $\downarrow$ | $\uparrow$   | $\uparrow$   | $\uparrow$   | No         |
| 2         | -   | $\downarrow$ | $\downarrow$ | $\downarrow$ | $\downarrow$ | $\downarrow$ | No         |
| 3         | -   | $\uparrow$   | -            | $\uparrow$   | -            | $\uparrow$   | No         |
| 4         | -   | $\uparrow$   | $\uparrow$   | $\downarrow$ | $\downarrow$ | $\downarrow$ | No         |
| 5         | red $\downarrow$                                    | $\uparrow$   | $\uparrow$   | $\uparrow$   | $\uparrow$   | $\uparrow$   | No         |
| 6         | blue $\uparrow$ , red $\downarrow$ , 4000 Å weakens | $\uparrow$   | $\downarrow$ | $\uparrow$   | -            | $\uparrow$   | Yes        |
| 7         | blue $\downarrow$                                   | -            | -            | -            | $\downarrow$ | -            | Yes        |
| 8         | -   | $\downarrow$ | -            | -            | $\uparrow$   | -            | No         |
| 9         | -   | $\uparrow$   | $\downarrow$ | -            | $\downarrow$ | $\downarrow$ | No         |
| 10        | red $\downarrow$                                    | $\downarrow$ | $\downarrow$ | $\downarrow$ | $\downarrow$ | $\downarrow$ | No         |

Portillo et al. 2020), although, Yip et al. (2004b) demonstrated its limitation when non-linear features, such as broad emission lines, are present. Optimising the VAE hyperparameters used in this study by way of a random or grid search should lead to better performance and more accurate reconstructions.

The VAE successfully managed to group together spectra belonging to the same class and separate the different classes in latent space, without having ever seen the corresponding labels. The VAE encoder produces a latent space that is interpretable. By inspecting the latent space we found the VAE was able to distinguish between bluer and redder continua and broad and narrow emission lines. Traversing different tracks, we observed that the spectra are distributed in such a way that there is a gradual evolution in the strength of emission and absorption lines. The changes between adjacent points in the latent space are not abrupt, providing evidence that the VAE latent space created is indeed regularized. We found tracks that correspond to increasing star formation and increase in broad emission lines along the Balmer series. The VAE also successfully grouped different sub-groups of stellar spectra. Using star sub-classes and line-ratio classifications such as BPT (Baldwin et al. 1981) or its updates to obtain galaxy sub-classes will further aid in the interpretation of the latent space. Unlike, line-ratio diagnostics, the VAE also captures the information present in the continuum.

The generative properties of the VAE allow us to probe different points in the latent space and interpret the information encoded in each dimension. When generating from the latent space care should be taken to primarily sample from places in the latent space that are occupied by training examples. Sampling far away from the encoded training spectra, forces the VAE to extrapolate often resulting in non-meaningful outputs.

A non-linear dimensionality reduction algorithm such as t-SNE or UMAP (Uniform Manifold Approximation and Projection, McInnes et al. 2018) which can preserve the global structure of the 10D latent space, could be used to further compress the encoded information to 2D. Such methods facilitate visualisation of high dimensional data and allow for a more efficient exploration of the latent space than inspecting a corner plot of all latent dimensions. Pat et al. (2022) applied a Probabilistic Autoencoder (PAE, Boehm & Seljak 2022; Böhm et al. 2023) to reduce the dimensionality of SDSS spectra from

1000 to 10 dimensions. UMAP was then applied to the compressed data to further reduce the dimensionality to 2D which allowed for an efficient interpretation of the projected latent space and identification of patterns that translate to the physical properties of spectra. Performing this two stage compression yielded improved class separation compared to using only UMAP.

We have shown that VAEs are capable of detecting anomalous spectra in the dataset using the weighted MSE between the original and reconstructed spectrum. An improvement to the weighted MSE approach would be to use the weighted reconstruction probability instead. As the VAE is a stochastic generative model that outputs probabilities, using the reconstruction probability as a proxy for anomaly detection is more principled and objective. The reconstruction probability is the Monte Carlo estimate of the reconstruction loss in the VAE’s loss function (An & Cho 2015). It takes into account the variability of the distributions from the sampling procedure.

Spectra were also scored according to their position on the latent space using the LOF algorithm. While some spectra were identified by both methods as outliers, LOF identified on average lower S/N outlier spectra. This suggests that a spectrum’s position in the latent space is affected by the noise present. Similarly, bad pixels in spectra will influence the spectrum’s position in the latent space, even if their weights are set to zero which prevents them from affecting the reconstruction probability calculation in the loss function. This highlights the importance of treating known bad pixels during the preprocessing by removing them and infilling the missing values using imputation methods, as done in this study. Failing to do so would result in these spectra being identified as outliers due to artefacts we are already aware of. A more comprehensive comparison of the two methods for identifying outliers (MSE and LOF) would be insightful for understanding any trends of outliers identified by each method and their sensitivity to S/N.

In addition to identifying outlier spectra related to rare objects or events, the VAE is also capable of identifying spectra with artefacts or with a wrongly assigned redshift. Both outlier methods employed in this study identified such spectra. Lan et al. (2022) show that redrock has a high redshift recovery rate from a sample of robust visually inspected redshifts from main survey spectra, with failures being  $< 2\%$ . This shows that redshift failures are rare and VAEs can

facilitate the visual inspection process which is crucial for identifying unexpected problematic features affecting the quality of the data. A spectrum that is flagged as an outlier due to a redrock redshift failure or CCD issues which were not identified and properly masked, can help inform of the shortcomings of the DESI spectroscopic pipeline. Based on this, improvements can be made so as to catch any such cases in the future, leading to cleaner data with the majority of outliers being due to physical spectral features.

The two methods identified atypical spectra in the dataset, including spectra with a wrong assigned redshift, bad pixels, stars, AGNs, miscalibration of the red and blue channels and others. Whilst we have identified the broad categories corresponding to the top 100 identified outliers from both methods, domain expertise is required to visually inspect these and determine the significance of the outliers, particularly whether these contain novel events, objects or phenomena that are unknown to scientists. To further curate the list of outliers that an expert would need to visually inspect, we made use of Astronomaly which employs Active Learning to provide personalised outlier recommendations based on the user's feedback on which outliers are relevant for a given science question.

Finally, we would observe that VAEs are just one of a family of dimensionality-reduction techniques which could be employed to identify anomalous spectra. In upcoming work, we hope to compare and contrast some of these techniques, in the process proposing enhancements to the standard approaches; and ultimately applying them to find anomalies at scale within the DESI dataset. And we note that similar techniques could be applied to other spectroscopic surveys (e.g. Euclid, Roman, Prime Focus Spectrograph, Wide Field Spectroscopic Telescope) and indeed to domains other than astronomy.

## 7 DATA AVAILABILITY

The data underlying the analyses in this paper were drawn from the Early Data Release of the Dark Spectroscopic Instrument (DESI EDR, DESI Collaboration et al. 2023). The full dataset is available at <https://data.desi.lbl.gov/doc/releases/edr>.

## ACKNOWLEDGEMENTS

CN and RPN have been supported by the STFC UCL Centre for Doctoral Training in Data Intensive Science. OL acknowledges STFC Consolidated Grant ST/R000476/1 and visits to All Souls College and the Physics Department, University of Oxford.

The material in this paper is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science and Technology of Mexico (CONAHCYT); the Ministry of Science, Innovation and Universities of Spain (MICIU/AEI/10.13039/501100011033), and by the DESI Member Institutions: <https://www.desi.lbl.gov/>

**collaborating-institutions.** Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U. S. National Science Foundation, the U. S. Department of Energy, or any of the listed funding agencies.

The authors are honored to be permitted to conduct scientific research on Iolkam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation.

## REFERENCES

Abadi M., et al., 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, <http://tensorflow.org/>

Alexander D. M., et al., 2023, *The Astronomical Journal*, 165, 124

Almeida J. S., Terlevich R., Terlevich E., Fernandes R. C., Morales-Luis A. B., 2012, *The Astrophysical Journal*, 756, 163

An J., Cho S., 2015, Special Lecture on IE, 2

Baldwin J. A., Phillips M. M., Terlevich R., 1981, *Publications of the Astronomical Society of the Pacific*, 93, 5

Baron D., 2019, arXiv e-prints, [p. arXiv:1904.07248](https://arxiv.org/abs/1904.07248)

Baron D., Poznanski D., 2017, *XXXXXX*, 465, 4530

Bergstra J., Bengio Y., 2012, *J. Mach. Learn. Res.*, 13, 281–305

Boehm V. M., Seljak U., 2022, *Transactions of Machine Learning Research*

Boroson T. A., Lauer T. R., 2010, *The Astronomical Journal*, 140, 390

Breiman L., 2001, *Machine Learning*, 45, 5

Breunig M. M., Kriegel H.-P., Ng R. T., Sander J., 2000, *SIGMOD Rec.*, 29, 93–104

Brodzeller A., et al., 2023, *AJ*, 166, 66

Böhm V., Kim A. G., Juneau S., 2023, Fast and efficient identification of anomalous galaxy spectra with neural density estimation ([arXiv:2308.00752](https://arxiv.org/abs/2308.00752)), <https://arxiv.org/abs/2308.00752>

DESI Collaboration et al., 2016a, arXiv e-prints, [p. arXiv:1611.00036](https://arxiv.org/abs/1611.00036)

DESI Collaboration et al., 2016b, arXiv e-prints, [p. arXiv:1611.00037](https://arxiv.org/abs/1611.00037)

DESI Collaboration et al., 2022, *AJ*, 164, 207

DESI Collaboration et al. 2023, The Early Data Release of the Dark Energy Spectroscopic Instrument, doi:10.5281/ZENODO.7964161, <https://zenodo.org/record/7964161>

Dey A., et al., 2019, *AJ*, 157, 168

Ferreras I., Lahav O., Somerville R. S., Silk J., 2022, arXiv e-prints, [p. arXiv:2208.05489](https://arxiv.org/abs/2208.05489)

Folkes S., et al., 1999, *Monthly Notices of the Royal Astronomical Society*, 308, 459

Fustes D., Manteiga M., Dafonte C., Arcay B., Ulla A., Smith K., Borrachero R., Sordo R., 2013, *A&A*, 559, A7

Guy J., et al., 2023, *AJ*, 165, 144

Hahn C., et al., 2022, arXiv e-prints, [p. arXiv:2208.08512](https://arxiv.org/abs/2208.08512)

Hawkins D. M., 1980, Introduction. Springer Netherlands, Dordrecht, pp 1–12, doi:10.1007/978-94-015-3994-4\_1, [https://doi.org/10.1007/978-94-015-3994-4\\_1](https://doi.org/10.1007/978-94-015-3994-4_1)

Huertas-Company M., Lanusse F., 2022, arXiv e-prints, [p. arXiv:2210.01813](https://arxiv.org/abs/2210.01813)

Ichinohe Y., Yamada S., 2019, *Monthly Notices of the Royal Astronomical Society*, 487, 2874

Juneau S., et al., 2024, arXiv e-prints, [p. arXiv:2404.03621](https://arxiv.org/abs/2404.03621)

Kingma D. P., Welling M., 2013, arXiv e-prints, [p. arXiv:1312.6114](https://arxiv.org/abs/1312.6114)

Kohonen T., 1982, *Biological Cybernetics*, 43, 59

Kramer M. A., 1991, *AIChE Journal*, 37, 233

Kullback S., Leibler R. A., 1951, *Ann. Math. Statistics*, 22, 79

Lahav O., et al., 1995, *Science*, 267, 859–862

Lan T.-W., et al., 2022, arXiv e-prints, [p. arXiv:2208.08516](https://arxiv.org/abs/2208.08516)

Levi M., et al., 2013, arXiv e-prints, [p. arXiv:1308.0847](https://arxiv.org/abs/1308.0847)

Liang Y., Melchior P., Hahn C., Shen J., Goulding A., Ward C., 2023, Outlier Detection in the DESI Bright Galaxy Survey ([arXiv:2307.07664](https://arxiv.org/abs/2307.07664))

Lochner M., Bassett B. A., 2021, *Astronomy and Computing*, 36, 100481

Madgwick D. S., Somerville R., Lahav O., Ellis R., 2003a, *Monthly Notices of the Royal Astronomical Society*, 343, 871

Madgwick D. S., et al., 2003b, *Monthly Notices of the Royal Astronomical Society*, 344, 847

McInnes L., Healy J., Melville J., 2018, arXiv e-prints, p. [arXiv:1802.03426](https://arxiv.org/abs/1802.03426)

Meusinger H., Schalldach P., Scholz R.-D., in der Au A., Newholm M., de Hoon A., Kaminsky B., 2012, *A&A*, 541, A77

Miller T. N., et al., 2023, arXiv e-prints, p. [arXiv:2306.06310](https://arxiv.org/abs/2306.06310)

Myers A. D., et al., 2022, arXiv e-prints, p. [arXiv:2208.08518](https://arxiv.org/abs/2208.08518)

Pat F., et al., 2022, arXiv e-prints, p. [arXiv:2211.11783](https://arxiv.org/abs/2211.11783)

Portillo S. K. N., Parejko J. K., Vergara J. R., Connolly A. J., 2020, *AJ*, 160, 45

Reis I., Rotman M., Poznanski D., Prochaska J. X., Wolf L., 2019, arXiv e-prints, p. [arXiv:1911.06823](https://arxiv.org/abs/1911.06823)

Rezende J. D., Mohamed S., Wierstra D., 2014, arXiv e-prints, p. [arXiv:1401.4082](https://arxiv.org/abs/1401.4082)

Rogers B., Ferreras I., Lahav O., Bernardi M., Kaviraj S., Yi S. K., 2007, *Monthly Notices of the Royal Astronomical Society*, 382, 750

Ronen S., Aragon-Salamanca A., Lahav O., 1999, *XXXXXX*, 303, 284

Ruff L., Kauffmann J. R., Vandermeulen R. A., Montavon G., Samek W., Kloft M., Dietterich T. G., Müller K.-R., 2020, arXiv e-prints, p. [arXiv:2009.11732](https://arxiv.org/abs/2009.11732)

Ruiz-Macias O., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 502, 4328

Schlafly E. F., et al., 2023, *AJ*, 166, 259

Scourfield M., Saintonge A., de Mijolla D., Viti S., 2023, *Monthly Notices of the Royal Astronomical Society*, 526, 3037–3050

Silber J. H., et al., 2023, *AJ*, 165, 9

Slonim N., Somerville R., Tishby N., Lahav O., 2001, *Monthly Notices of the Royal Astronomical Society*, 323, 270

Suárez-Pérez J. F., Forero-Romero J. E., 2022, in SciOps 2022: Artificial Intelligence for Science and Operations in Astronomy (SCIOPS). p. 7, doi:10.5281/zenodo.6554524

Yip C. W., et al., 2004a, *The Astronomical Journal*, 128, 585

Yip C. W., et al., 2004b, *The Astronomical Journal*, 128, 2603

Yu R., 2020, arXiv e-prints, p. [arXiv:2006.10273](https://arxiv.org/abs/2006.10273)

van der Maaten L., Hinton G., 2008, *Journal of Machine Learning Research*, 9, 2579

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