

# UC Irvine

## UC Irvine Previously Published Works

### Title

Archetype-based Redshift Estimation for the Dark Energy Spectroscopic Instrument Survey

### Permalink

<https://escholarship.org/uc/item/730113t3>

### Journal

The Astronomical Journal, 168(3)

### ISSN

0004-6256

### Authors

Anand, Abhijeet

Guy, Julien

Bailey, Stephen

et al.

### Publication Date

2024-09-01

### DOI

10.3847/1538-3881/ad60c2

































### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



# Archetype-based Redshift Estimation for the Dark Energy Spectroscopic Instrument Survey

Abhijeet Anand<sup>1</sup> , Julien Guy<sup>1</sup> , Stephen Bailey<sup>1</sup> , John Moustakas<sup>2</sup> , J. Aguilar<sup>1</sup>, S. Ahlen<sup>3</sup> , A. S. Bolton<sup>4</sup> , A. Brodzeller<sup>1,5</sup> , D. Brooks<sup>6</sup>, T. Claybaugh<sup>1</sup>, S. Cole<sup>7</sup> , A. de la Macorra<sup>8</sup> , Biprateep Dey<sup>9</sup> , K. Fanning<sup>10,11</sup> , J. E. Forero-Romero<sup>12,13</sup> , E. Gaztañaga<sup>14,15,16</sup>, S. Gontcho A Gontcho<sup>1</sup> , G. Gutierrez<sup>17</sup>, K. Honscheid<sup>18,19,20</sup>, C. Howlett<sup>21</sup> , S. Juneau<sup>4</sup>, D. Kirkby<sup>22</sup> , T. Kisner<sup>1</sup> , A. Kremin<sup>1</sup> , A. Lambert<sup>1</sup>, M. Landriau<sup>1</sup> , L. Le Guillou<sup>23</sup> , M. Manera<sup>24,25</sup> , A. Meisner<sup>4</sup> , R. Miquel<sup>25,26</sup>, E. Mueller<sup>27</sup>, G. Niz<sup>28,29</sup> , N. Palanque-Delabrouille<sup>1,30</sup> , W. J. Percival<sup>31,32,33</sup> , C. Poppett<sup>1,34,35</sup>, F. Prada<sup>36</sup> , A. Raichoor<sup>1</sup> , M. Rezaie<sup>37</sup> , G. Rossi<sup>38</sup>, E. Sanchez<sup>39</sup> , E. F. Schlafly<sup>40</sup> , D. Schlegel<sup>1</sup>, M. Schubnell<sup>41,42</sup>, D. Sprayberry<sup>4</sup>, G. Tarlé<sup>42</sup> , C. Warner<sup>43</sup>, B. A. Weaver<sup>4</sup>, R. Zhou<sup>1</sup> , and H. Zou<sup>44</sup> 

<sup>1</sup> Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA; [AbhijeetAnand@lbl.gov](mailto:AbhijeetAnand@lbl.gov)

<sup>2</sup> Department of Physics and Astronomy, Siena College, 515 Loudon Road, Loudonville, NY 12211, USA

<sup>3</sup> Physics Dept., Boston University, 590 Commonwealth Avenue, Boston, MA 02215, USA

<sup>4</sup> NSF NOIRLab, 950 N. Cherry Avenue, Tucson, AZ 85719, USA

<sup>5</sup> Department of Physics and Astronomy, The University of Utah, 115 South 1400 East, Salt Lake City, UT 84112, USA

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

<sup>7</sup> Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK

<sup>8</sup> Instituto de Física, Universidad Nacional Autónoma de México, Cd. de México C.P. 04510, México

<sup>9</sup> Department of Physics & Astronomy and Pittsburgh Particle Physics, Astrophysics, and Cosmology Center (PITT PACC), University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, PA 15260, USA

<sup>10</sup> Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Menlo Park, CA 94305, USA

<sup>11</sup> SLAC National Accelerator Laboratory, Menlo Park, CA 94305, USA

<sup>12</sup> Departamento de Física, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia

<sup>13</sup> Observatorio Astronómico, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia

<sup>14</sup> Institut d'Estudis Espacials de Catalunya (IEEC), 08034 Barcelona, Spain

<sup>15</sup> Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth, PO1 3FX, UK

<sup>16</sup> Institute of Space Sciences, ICE-CSIC, Campus UAB, Carrer de Can Magrans s/n, 08913 Bellaterra, Barcelona, Spain

<sup>17</sup> Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, USA

<sup>18</sup> Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

<sup>19</sup> Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

<sup>20</sup> The Ohio State University, Columbus, 43210 OH, USA

<sup>21</sup> School of Mathematics and Physics, University of Queensland, 4072, Australia

<sup>22</sup> Department of Physics and Astronomy, University of California, Irvine, 92697, USA

<sup>23</sup> Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), FR-75005 Paris, France

<sup>24</sup> Departament de Física, Serra Hünter, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain

<sup>25</sup> Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra Barcelona, Spain

<sup>26</sup> Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain

<sup>27</sup> Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, UK

<sup>28</sup> Departamento de Física, Universidad de Guanajuato—DCI, C.P. 37150, Leon, Guanajuato, México

<sup>29</sup> Instituto Avanzado de Cosmología A.C., San Marcos 11—Atenas 202, Magdalena Contreras, 10720. Ciudad de México, México

<sup>30</sup> IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France

<sup>31</sup> Department of Physics and Astronomy, University of Waterloo, 200 University Avenue W, Waterloo, ON N2L 3G1, Canada

<sup>32</sup> Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON N2L 2Y5, Canada

<sup>33</sup> Waterloo Centre for Astrophysics, University of Waterloo, 200 University Avenue W, Waterloo, ON N2L 3G1, Canada

<sup>34</sup> Space Sciences Laboratory, University of California, Berkeley, 7 Gauss Way, Berkeley, CA 94720, USA

<sup>35</sup> University of California, Berkeley, 110 Sproul Hall #5800 Berkeley, CA 94720, USA

<sup>36</sup> Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain

<sup>37</sup> Department of Physics, Kansas State University, 116 Cardwell Hall, Manhattan, KS 66506, USA

<sup>38</sup> Department of Physics and Astronomy, Sejong University, Seoul, 143-747, Republic of Korea

<sup>39</sup> CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain

<sup>40</sup> Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

<sup>41</sup> Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

<sup>42</sup> University of Michigan, Ann Arbor, MI 48109, USA

<sup>43</sup> Department of Astronomy, University of Florida, 211 Bryant Space Science Center, Gainesville, FL 32611, USA

<sup>44</sup> National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Road, Chaoyang District, Beijing, 100012, People's Republic of China

Received 2024 May 29; revised 2024 June 27; accepted 2024 July 7; published 2024 August 21

## Abstract

We present a computationally efficient galaxy archetype-based redshift estimation and spectral classification method for the Dark Energy Survey Instrument (DESI) survey. The DESI survey currently relies on a redshift fitter and spectral classifier using a linear combination of principal component analysis-derived templates, which is very efficient in processing large volumes of DESI spectra within a short time frame. However, this method occasionally

yields unphysical model fits for galaxies and fails to adequately absorb calibration errors that may still be occasionally visible in the reduced spectra. Our proposed approach improves upon this existing method by refitting the spectra with carefully generated physical galaxy archetypes combined with additional terms designed to absorb data reduction defects and provide more physical models to the DESI spectra. We test our method on an extensive data set derived from the survey validation (SV) and Year 1 (Y1) data of DESI. Our findings indicate that the new method delivers marginally better redshift success for SV tiles while reducing catastrophic redshift failure by 10%–30%. At the same time, results from millions of targets from the main survey show that our model has relatively higher redshift success and purity rates (0.5%–0.8% higher) for galaxy targets while having similar success for QSOs. These improvements also demonstrate that the main DESI redshift pipeline is generally robust. Additionally, it reduces the false-positive redshift estimation by 5%–40% for sky fibers. We also discuss the generic nature of our method and how it can be extended to other large spectroscopic surveys, along with possible future improvements.

*Unified Astronomy Thesaurus concepts:* [Galaxy spectroscopy \(2171\)](#); [Astronomical methods \(1043\)](#); [Redshift surveys \(1378\)](#); [Astronomy software \(1855\)](#); [Astronomy data analysis \(1858\)](#)

## 1. Introduction

The Dark Energy Survey Instrument (DESI) is a Stage-IV large spectroscopic survey (Levi et al. 2013; DESI Collaboration et al. 2016) that will observe more than  $\sim 40$  million spectra of galaxies, stars, and quasars in the wavelength range 3600–9800 Å with an average spectral resolution ( $R = \lambda/\Delta\lambda$ ) of 2000 at the shortest wavelengths to 5500 at the longest wavelengths (DESI Collaboration et al. 2022; Miller et al. 2024; Silber et al. 2023). The main survey started in 2021 May and will continue for 5 yr, resulting in the largest 3D map of the Universe ever. The early spectroscopic data of more than  $\sim 1$  million galaxies, quasars, and stars was released in early 2023 June (DESI Collaboration et al. 2024a). The DESI early data release<sup>45</sup> (EDR) consists of raw images, reduced spectra, spectral classifications, and redshifts of each target observed with DESI during the survey validation (SV) phase. This requires a large number of software developments to process such large data sets efficiently and has motivated the development of classical machine-learning- and neural-network-based approaches to analyze the data.

One of the principal aims of these surveys is to construct an accurate 3D map of the large-scale structure of the Universe with these tracer galaxies and quasars. These maps allow us to measure baryon acoustic oscillations and the redshift-space distortions that help us understand the expansion history and growth of structure(s) in the Universe. This requires precise redshift measurements for galaxies and quasars, as understanding the systematics in redshift efficiency is crucial for accurately modeling the number density of tracers that contribute to these signals. These systematics also impact the total error budget of the best-fit cosmological parameters (see Krolewski et al. 2024; Yu et al. 2024, for DESI systematics).

At the same time, precise redshift measurements are also important for noncosmological scientific analyzes. For example, stacking observed spectra to detect the radio HI line from radio surveys (Anand et al. 2019) or other weak emission lines (Maddox et al. 2013) that are typically undetectable in individual spectra relies on accurate redshift information. Similarly, robust redshift estimates help reduce the misidentification of galaxy groups, which usually have similar redshifts that differ only by their peculiar velocities (Wang et al. 2020). Furthermore, very precise redshifts are necessary to model and understand the gas kinematics in the halos of galaxies using quasar absorption lines (Tumlinson et al. 2017). Owing to the importance of precise redshifts in astronomy, it is imperative to develop fast, optimal,

and robust redshift measurement pipelines for the large ongoing and upcoming spectroscopic surveys.

Given its straightforward and simple mathematics, principal component analysis (PCA) has demonstrated remarkable success in astronomy. The method’s popularity stems from its deterministic and quickly convergent nature when solving for linear coefficients. Recent progress has allowed us to improve PCA-based spectral fitting algorithms, incorporating astronomical data uncertainties while constructing PCA templates (Bailey 2012; Tsalmantza & Hogg 2012). These templates have been used to model astronomical spectra of galaxies, stars, and quasars and measure their redshifts. An example of such a method is the Sloan Digital Sky Survey (SDSS) redshift pipeline (Bolton et al. 2012), employed for modeling SDSS BOSS/eBOSS spectra and measuring their redshifts. It has worked extremely well on millions of spectra, which helped the SDSS collaboration construct one of the largest 3D sky maps at the time (Ahumada et al. 2020). At the same time, PCA templates have been combined with stellar population synthesis models to measure galactic properties from their spectra (see Chen et al. 2012, for details).

However, one issue with PCA-based templates is their inability to present the physical properties of galaxy spectra in their eigenvectors alone. One commonly encountered problem in PCA fits involves fitting noise as negative and unphysical fluxes, particularly in low signal-to-noise ratio (SNR) spectra that yield inaccuracies in measuring their redshifts and associated physical properties. This drawback is important to tackle in light of the ongoing large spectroscopic surveys that are poised to accumulate millions of spectra in the coming years.

To tackle these issues, several other modeling techniques have been used in recent years. One such method is nonnegative matrix factorization (NMF), which is also a dimension reduction technique akin to PCA. It factorizes a large input matrix (typically the fluxes as a function of wavelength) into two smaller nonnegative matrices,<sup>46</sup> namely eigenspectra, and coefficients, which can reproduce the input matrix fairly well. Furthermore, these eigenspectra exhibit a greater resemblance to the physical characteristics, like emission or absorption features in astronomical spectra. Lee & Seung (1999) presented a simple iterative update rule to derive these smaller matrices while ensuring a non-increasing nature for the associated cost function. While NMF presents

<sup>45</sup> EDR is publicly available at <https://data.desi.lbl.gov/doc/releases/edr/>.

<sup>46</sup>  $V \approx WH$ , where  $W \geq 0$  and  $H \geq 0$  are found by minimizing the Frobenius norm  $\|V - WH\|_F$ .

inherent complexity as there is no optimal algorithm for finding the global minimum for the cost function associated with these two matrices, it is still extremely useful for astronomical spectra where the fluxes are always positive. Consequently, NMF has emerged as an alternative tool for modeling galaxy and quasar continua (see Zhu & Ménard 2013; Anand et al. 2021; Napolitano et al. 2023), especially to search for metal absorbers in quasar spectra in the context of circumgalactic/intergalactic media or intracluster medium studies (Anand et al. 2021, 2022).

However, very few NMF algorithms exist that may work well with low-SNR data. Recently, Green & Bailey (2023) have extended the NMF algorithm to construct eigenspectra and coefficients for astronomical data with negative fluxes. This is an important result for low-SNR spectra coming from low-resolution large spectroscopic surveys, where even carefully calibrated sky subtraction may result in negative fluxes in reduced spectra.

An alternative method that has been previously employed involves utilizing correlation peaks in the cross-correlation function of the high-pass-filtered input spectrum with a series of templates. This technique was initially developed for the Galaxy And Mass Assembly (GAMA) survey galaxies, as described in Baldry et al. (2014). However, it is readily adaptable to forthcoming surveys such as DEVILS (Davies et al. 2021) on Anglo-Australian Telescope and WAVES<sup>47</sup> on the Vista telescope with 4MOST (Driver et al. 2016). This approach offers two advantages over traditional PCA-based redshift determination methods. First, it provides highly reliable redshifts with uncertainties as low as  $<50 \text{ km s}^{-1}$ . The DESI’s PCA-based redshift fitter also achieves an uncertainty of  $<40 \text{ km s}^{-1}$  in redshift measurements, similar to the correlation peak approach. Second, it facilitates the identification of multiple peaks in the cross-correlation function, which helps identify overlapping galaxies (such as those in clusters or groups, e.g., Holwerda et al. 2015, 2022) or strongly lensed systems (e.g., the SLACS program initiated by Bolton et al. 2008).

The recent progress in neural networks and machine learning methods have also played a crucial role in analyzing astronomical data sets, encompassing tasks from image and spectral classification to the prediction of cosmological parameters (see Baron 2019, for a review). An example of such an approach is QuasarNET, a convolutional neural network (CNN)–based method developed for the classification and redshift measurement of quasars (Busca & Balland 2018) observed with surveys such as eBOSS and DESI. Consequently, the astronomical community relies on these approaches to analyze large numbers of spectra.

To address the problem of unphysical line fitting in PCA approach, some studies have explored redshift fitting (using a minimum  $\chi^2$  approach) of galaxy spectra using physical galaxy models (also known as the *archetypes*) and polynomials (Cool et al. 2013). An example of such automated software is redmonster (Hutchinson et al. 2016), which was developed for the eBOSS program of the SDSS. The software yielded a high-redshift success rate ( $\sim 90.5\%$ ) for luminous red galaxies (LRGs). However, this model had the freedom to take negative coefficients in order to achieve the optimal solution, which made it susceptible to unphysical modeling of negative features

as real absorption features. Though it worked fine for early SDSS data sets, it was not computationally efficient and, therefore, not suited for a survey as large as DESI, which aims to observe 10 times more spectra than SDSS.

In this paper, our objective is to improve the existing method for classifying and measuring redshifts of the DESI spectra. We aim to integrate the existing DESI redshift fitter with a suite of physically motivated archetypes, combined with a set of polynomials to construct spectral fits that are more physical. This also allows us to maintain the efficiency intrinsic to PCA. We introduce several new features in our model, such as a new state-of-the-art set of archetypes, Legendre polynomials, Gaussian priors on polynomial coefficients, and enforcing positive coefficients to the archetypes to avoid any unphysical fittings of the spectra.

The structure of our paper is as follows: we describe the spectral data and DESI-specific methods and techniques in Section 2. In Section 3, we state the problem and illustrate our methodology and implementation. Section 4 presents extensive tests and the performance evaluation of our method using existing DESI data. Finally, in Section 5, we discuss the implications of our method and potential future avenues of exploration.

## 2. DESI Data and Classification Procedure

### 2.1. DESI Spectra

The DESI focal plane is segmented into 10 wedges or petals, each connected to a dedicated spectrograph. These spectrographs consist of three cameras denoted as *b* (3600–5800 Å), *r* (5760–7620 Å), and *z* (7520–9824 Å), each exhibiting a resolving power ranging from 2000–5000. Each petal is equipped with 500 fibers, each directed toward unique sky positions. Consequently, DESI can simultaneously observe spectra from 5000 targets. The fiber assignment<sup>48</sup> algorithm is run throughout the night on the fly, and specific targets (along with their R.A. and decl.) are assigned to each of the 5000 fibers, constituting a single observing “tile.” Furthermore, a subset of the 5000 fibers (usually  $\sim 20$ ) is specifically allocated for observing blank sky regions to measure the sky emission model for spectral reduction. These designated fibers are referred to as “sky fibers.” These are pre-selected to target specific regions of the sky that are empty. The sky fibers are the same for each tile, though they can point in different directions for each tile. These tiles are designed for specific observation programs, the two main ones being the dark and bright. As their name indicates, they are optimized for different observation conditions, and the observations switch from one to the other depending on a combination of moon brightness, *r*-band sky magnitude, sky transparency, and seeing (see Myers et al. 2023; Schlafly et al. 2023, for more details). Moreover, certain targets are observed multiple times, and their spectra are combined across exposures to construct very high-SNR spectra.

The instrumental design of DESI allows it to observe many galaxies and quasar targets with broad physical properties, such as stellar populations, emission lines, and dust content over a wide range of redshifts. DESI targets three broad classes of

<sup>48</sup> The fiber assignment algorithm basically selects which fibers are to be assigned to the targets, blank sky location, and calibration stars for each tile. The priority of targets means they will be assigned first or last while assigning the fibers. E.g., QSO targets will always be assigned first on dark tiles followed by luminous red galaxy (LRG) and emission-line galaxy (ELG) targets (DESI Collaboration et al. 2022).

<sup>47</sup> Wide Area Vista Extragalactic Survey.



**Table 1**  
DESI Target Selection, where  $grz$  Are Magnitudes from the Legacy Survey (LS)

Target Class	Magnitudes	Redshift	References
BGS_BRIGHT	$r < 19.5$	$< 0.4$	(a)
BGS_FAINT	$19.5 < r < 20.175$	$< 0.4$	(a)
ELG_LOP	$g_{\text{fiber}} < 24.1$	$\in (1.1, 1.6)$	(b)
ELG_VLO	$g_{\text{fiber}} < 24.1$	$\in (0.6, 1.1)$	(b)
LRG	$z_{\text{fiber}} < 21.6$	$\in (0.4, 1)$	(c)
QSO	$r < 23$	$> 1.6$	(d)

**Note.**  $g_{\text{fiber}}$  and  $z_{\text{fiber}}$  are the  $g$ - and  $z$ -band fiber magnitudes, i.e., the magnitude corresponding to the expected flux within a DESI fiber.

**References.** (a) Hahn et al. (2023), (b) Raichoor et al. (2023), (c) Zhou et al. (2023), (d) Chaussidon et al. (2023).

galaxies: bright galaxy sample (BGS; Hahn et al. 2023), luminous red galaxies (LRGs; Zhou et al. 2023), and emission-line galaxies (ELGs; Raichoor et al. 2023). Galaxies in the DESI BGS sample are low redshift, and they span from massive and quenched galaxies with evolved stellar populations to low-mass and star-forming galaxies with young stellar populations. They are further classified as BGS\_BRIGHT and BGS\_FAINT depending on their brightness and magnitudes (Myers et al. 2023). Galaxies in the DESI BGS sample are observed in bright conditions on bright tiles. Next, the LRGs are targeted based on their  $grzW1$  photometry and 4000 Å break (see Zhou et al. 2023, for detailed selection cuts) and are dominated by a metal-rich stellar population and low star formation rates. They are also more massive than BGSs and targeted in the redshift range of  $0.4 < z < 1$ .

Next, the ELG sample includes primarily star-forming galaxies, and their spectra show several high equivalent width emission lines. The principal emission line that is aimed to resolve for ELGs in the DESI survey is the [O II]  $\lambda\lambda 3727, 29$  doublet. They are targeted in the redshift range of  $0.6 < z < 1.6$  with three main selections: ELG\_LOP, ELG\_HIP, and ELG\_VLO. The main difference between ELG\_LOP and ELG\_VLO is based on  $g-r$  and  $r-z$  cuts such that these two sets are disjoint and target different redshift bins. ELG\_LOP has higher priority in fiber assignment and covers galaxies in  $1.1 < z < 1.6$ , where other DESI tracers are less dense, while ELG\_VLO has lower priority and targets in  $0.6 < z < 1.1$  (see Raichoor et al. 2023, for more details). ELG\_HIP is a 10% random subsample of ELG\_LOP and ELG\_VLO and has the same fiber assignment priority as LRGs. DESI has also been designed to observe millions of quasars in the redshift range  $z > 1.6$  to understand the large-scale structure and constrain the cosmological parameters. The QSOs are targeted with optical and optical+IR color cuts in  $g-z$  versus  $grz-W1$  color space (see Chaussidon et al. 2023, for more details). QSOs have the highest fiber assignment priority followed by LRGs, ELG\_LOP, and ELG\_VLO. These targets are observed in dark times on dark tiles only.

The apparent magnitudes and redshift range of the main DESI targets are compiled in Table 1. We also present the typical DESI spectra in their observed frame corresponding to each target class in Figure 1. The top panel shows the spectrum of a representative low-redshift BGS galaxy featuring multiple emission lines. The second panel displays the spectrum of an LRG, showing the clear Ca H and K absorption features and 4000 Å break (redshifted at  $\sim 6000$  Å). An example spectrum of ELG is presented in the third panel, highlighting the clearly

visible [O II] doublet line redshifted at  $\sim 8600$  Å. Finally, in the last panel, we present a typical DESI quasar spectrum, where broad emission lines (hydrogen and metal lines) are distinctly visible on the top of the quasar power-law continuum.

DESI observations are divided into two categories: survey validation (SV) phase and main survey (MS). The SV phase ran between 2020 December and 2021 May, during which several tiles were observed, details of which are fully described in Myers et al. (2023). The SV tiles were designed to observe distinct target classes (e.g., LRG, BGS, ELG, QSO) on several nights encompassing different exposure times (hence varying SNR), weather, and observation conditions. This comprehensive data set provides a unique opportunity to thoroughly assess and quantify the performance of the DESI instruments and the spectral reduction and redshift estimation pipeline. The SV phase full data sets were publicly released as part of DESI EDR. Additional details regarding these observations can be found in DESI Collaboration et al. (2024b). The main survey started in 2021 May and will continue for 5 yr until 2026 May. The data set will be publicly released by the DESI collaboration. The Year 1 (Y1) data of the main survey that encompasses observations spanning approximately  $\sim 1$  yr, from 2021 May to 2022 June, will be released first.

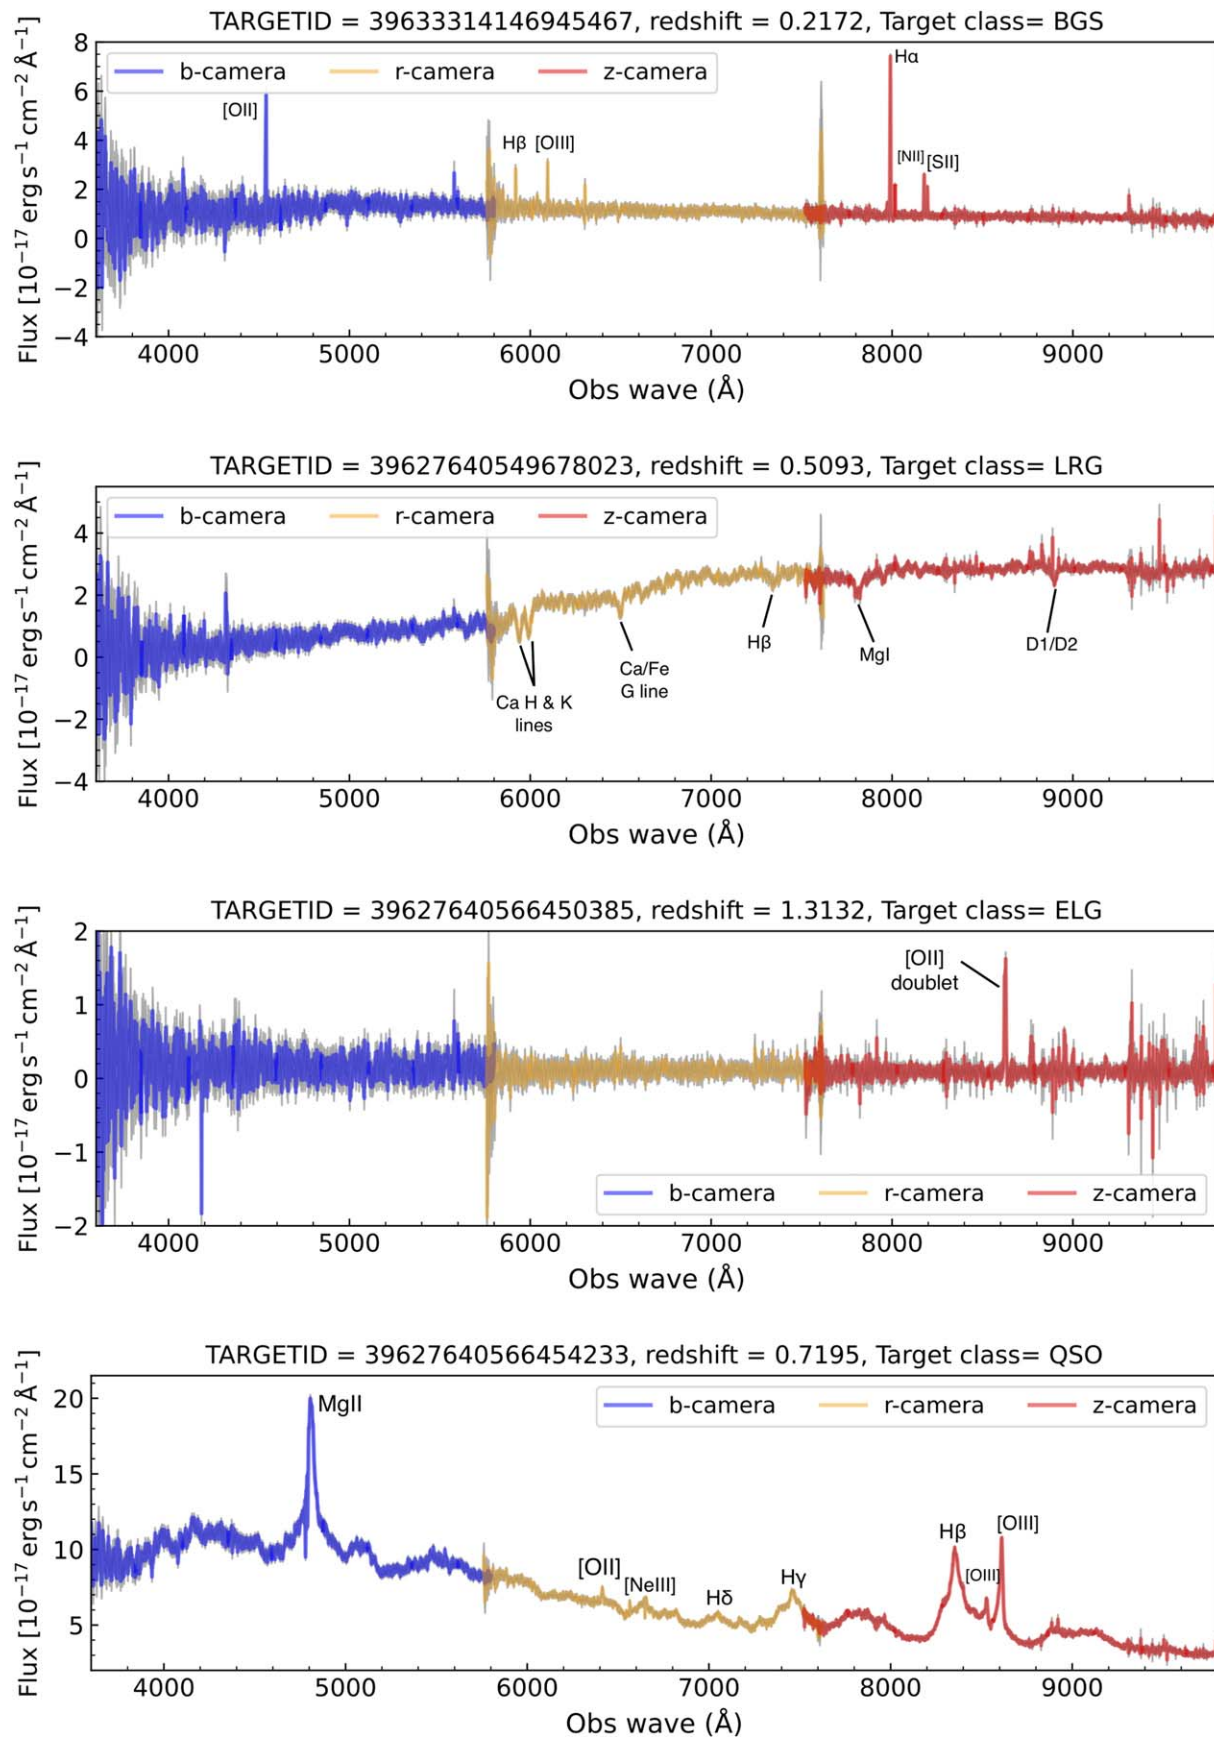
In this manuscript, we use targets observed during the SV phase and Y1 to assess the efficacy of our archetype-based method for redshift fitting specifically for galaxy spectra. To achieve this, we carefully constructed a set of “galaxy archetypes” (described in Section 3.4) and applied our approach (described in Section 3) to DESI targets. We quantitatively evaluate the performance of our algorithm and provide and shed light on its effectiveness and potential impact, as detailed in the subsequent Sections.

## 2.2. Current DESI Redshift Fitter

Redrock<sup>49</sup> (S. J. Bailey et al. 2024, in preparation) is the principal spectral fitter and redshift estimator software for the DESI spectra. It is very computationally efficient and can be run on both CPUs and GPUs. Though Redrock is run specifically on DESI spectra, the methods and implementation are quite generic. The underlying algorithm selects the least  $\chi^2$  from PCA-based templates fit over a range of redshifts for three main spectral types: galaxy, stars, and QSOs. The QSO and stellar templates are further divided into different subtypes to cover their diversity. The redshift and spectral class solution corresponding to the lowest  $\chi^2$  is the final redshift and spectral type of the input spectrum. The templates were generated from a combination of real and synthetic spectra of astronomical targets using an iterative principal component generator, *empca* (Bailey 2012), which also takes uncertainties of the data into account.

S. J. Bailey et al. (2024, in preparation) provides the details of each template and its performance. The principal components for galaxies were generated using a set of 20,000 synthetic spectra (described in Section 3.3). In contrast, the QSO templates were constructed using quasars observed with the eBOSS program of SDSS (Brodzeller et al. 2023). On the other hand, for stars, the PCA templates were generated using synthetic stellar spectra in six different  $T_{\text{eff}}$  (corresponding to stellar subtypes: *B*, *A*, *F*, *G*, *K*, and *M*) bins to account for the broad diversity among stars (Allende Prieto et al. 2018; Cooper

<sup>49</sup> <https://github.com/desihub/redrock/>



**Figure 1.** DESI example spectra of BGS, LRG, ELG, and QSO targets. The colored lines show the three DESI cameras. The uncertainties in fluxes are shown in gray in each panel. We also label the expected location of absorption and emission lines in each spectrum.

et al. 2023). In addition, separate PCA templates were generated for cataclysmic variables (CVs) and white dwarfs (WDs) as their physical properties differ significantly from main-sequence stars. The CV and WD templates were generated using the same archetype technique described in Bolton et al. (2012).

Redrock is very efficient and has performed well on SV and Y1 data sets. However, it occasionally yields negative flux models for some low-SNR spectra, where it fits noise or poorly subtracted sky signals as negative features (see Section 4.1). Currently, if the method identifies negative flux at the location of a forbidden line, such as [O II], it implements an ad hoc method to correct these unphysical fits. It adds a prior to the  $\chi^2$  based on the model flux<sup>50</sup> around the [O II] line. This correction forces the final  $\chi^2$  to be higher and modifies the ranking of other models. However, this lacks a solid mathematical basis and does not work for other emission lines in the current version (0.19.0).<sup>51</sup>

Redrock also lacks algorithmic features capable of accurately absorbing the calibration errors that may occasionally show up in the reduced spectra. It is difficult to quantify the occurrence of such spectral artifacts (described below) as Redrock does not raise any warning bits for such issues. However, they still need attention, as even a relatively small failure rate can be large in absolute numbers, given the survey size.

There are two primary calibration issues (see Guy et al. 2023, for more detailed discussion) considered in this work. First, the slightly different bias or zero level in the CCD quadrants causes discontinuity at the boundaries, from where the electrons are read by the amplifiers present at all four corners (see Figures 2 and 55, and Appendix E of Guy et al. 2023). Similar discontinuities can arise from fluctuations of dark current for some CCD columns. This causes the entire spectrum to shift up or down in one or more DESI cameras (more precisely, at half of the camera), corresponding to one of the amplifiers for some of the fibers). These discontinuities can still show up even after the flux calibration, as these effects are additive, and the flux calibration is multiplicative. The flux calibration is obtained with bright standard stars, so the relative effect of an offset is barely visible.

Second, the flux calibration corrects for the difference in throughput from one camera to another. However, as described in Guy et al. (2023), there remain wavelength and focal plane position-dependent calibration errors caused by the chromatic distortions of the corrector (see also Figure 39 of Guy et al. 2023). Fibers are positioned to get the best throughput in the  $r$  band at the cost of higher losses at other wavelengths. For this reason, we consider additional corrections beyond a pure offset for redshift fitting. In general, this correction is less significant than the first one.

To address these issues in the current DESI redshift fitting pipeline, we propose a different approach. Our proposed method combines PCA-based redshift scanning with physically motivated galaxy spectra (*archetypes*; more details in Section 3.4) to model the DESI spectra and measure the redshifts. Additionally, we introduce polynomial modeling in each camera to accurately model these calibration errors and absorb them in the best-fit redshift model. We provide the

mathematical description of our model in the next section, followed by an extensive demonstration of how it naturally resolves these issues and performs on the DESI survey validation phase and Y1 data sets.

### 3. Method

#### 3.1. Modeling Galaxy Spectra

As elaborated in Section 2.2, Redrock suffers from certain challenges. There are three main issues we are trying to resolve with our new approach, as detailed in Section 2.2. (1) Unphysical fitting of negative fluxes in DESI spectra, (2) discontinuities in the spectrum due to CCD bias and zero issues, and (3) gradient-like throughput offsets caused by chromatic distortions of the corrector. In the current formulation, we use Legendre polynomials to model these two spectral defects and archetypes to yield more physical models for the spectrum. Given that our archetypes span the parent galaxy sample fairly well (more discussion in Section 3.4), we model the input spectrum with a single archetype. This is mathematically less complex and also computationally easier to implement. Finally, our model can be written as follows:

$$S_G(\lambda) = \alpha_k(RT_{G,k})(\lambda) + \sum_j \sum_i a_{i,j}(RP_i)(\lambda');$$

$j = \text{cameras, e.g., } b, r, z \text{ for DESI}$   
 $\alpha_k > 0 \text{ and } a_{i,j} \in \mathbb{R}; \lambda \equiv \lambda_{\text{obs}},$  (1)

where  $\lambda$  is observed wavelength,  $T_{G,k}(\lambda)$  are  $k$ th galaxy archetype,  $P_i(x)$  are Legendre polynomials with degree  $i$ ,  $\lambda'$  is reduced wavelength,<sup>52</sup> and  $\alpha_k, a_{i,j}$  are linear coefficients for archetypes and Legendre polynomials, respectively, and  $R$  is the resolution matrix,<sup>53</sup> which is used to decorrelate the data and accounts for the resolution of the spectrographs (see Guy et al. 2023, for more details). Applying the  $R$  matrix on templates is a relatively slower step in Redrock as it involves a large number of matrix operations. Note that  $\alpha_k > 0$  is fit over entire spectra, so for each spectrum, there is just one archetype coefficient. At the same time, there are  $3(d+1)$  coefficients (“3” is for three cameras in DESI) for polynomial terms where  $d$  is the highest degree of polynomial used in modeling. This fitting approach is applied to all of the cameras simultaneously, ensuring accurate modeling of spectra in each of the wavelength bands. In total, our model has  $1 + 3(d+1)$  coefficients.

Moreover, we also want to leverage the efficiency and computational speed of PCA and fit the spectra; therefore, after scanning for redshifts using PCA templates, we select the  $N$ -best redshifts (or  $N_{\text{zbest}}$ , corresponding to the  $N$ -least  $\chi^2$  (see Figure 2 for the method schematic) regardless of redrock spectral classifications and refine the model fit for each redshift with a linear combination of physical spectra (to model the shape and features of spectra) and Legendre polynomials (to model pipeline errors), which can result in different ranking in terms of best-fit  $\chi^2$ . Here, it is important to understand that our model is not independent of current redrock as we run our method on a list of redshifts ( $N_{\text{best}}$ ) obtained with the initial PCA-based method. Therefore, if the true redshift does not fall

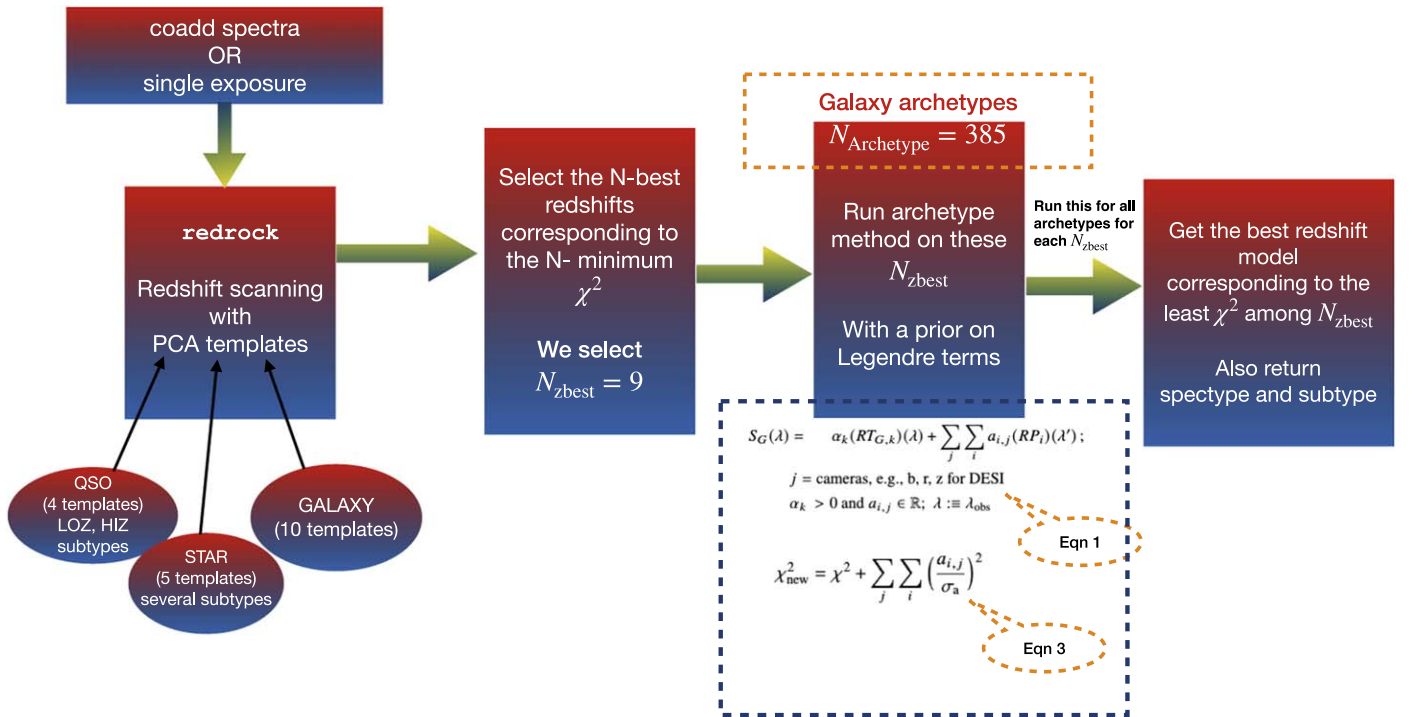
<sup>50</sup>  $[\text{O II}]_{\text{model flux}} = |\sum_k a_k T_k(\lambda_i)|$ ,  $3724 < \lambda_i < 3733$ , where  $\lambda_i$  is rest-frame wavelength (around [O II] line) in Å and  $T_k$  and  $a_k$  are PCA templates and coefficients.

<sup>51</sup> <https://github.com/desihub/redrock/releases/tag/0.19.0>

<sup>52</sup>  $\lambda' = 2 \cdot \frac{\lambda - \lambda_{\text{min}}}{(\lambda_{\text{max}} - \lambda_{\text{min}})} - 1 \in [-1, 1]$ , where  $\lambda_{\text{min}}$  and  $\lambda_{\text{max}}$  are minimum and maximum observed wavelengths in each band.

<sup>53</sup> It is used to convert the correlated fluxes into uncorrelated ones:  $F' = RF$ ,  $F'$  is final flux that is used in calculating  $\chi^2$ .





**Figure 2.** A schematic of our archetype based per-camera polynomial fitting. The coadded or single-exposure spectra are input to the redrock that uses previously estimated PCA-based templates of several spectral types to perform the redshift scan. We retain the  $N_{zbest}$  redshifts corresponding to the least  $N \chi^2$  values after the preliminary scan. Subsequently, we model the input spectra using Equation (1) at each of these  $N_{zbest}$  redshift (in a loop), employing each archetype (in another loop) to determine the final best  $z$  associated with the least  $\chi^2$ . The final output is the best-fit redshift, along with the spectral type and index of the best-fit archetype.

within this list, our method will also fail. Therefore, the choice of  $N_{zbest}$  is also a free parameter in our method, and we explored this space (see Section 4.5) to achieve an optimal value. We also emphasize that the statistical precision of the redshifts in the archetype method will be the same as that of the Redrock. Our focus in the paper is more on improving catastrophic failures than redshift accuracy. We aim to perform an independent test of the archetype procedure as a redshift classifier in the future.

Next, to model the CCD discontinuities and throughput offsets in the spectra, we use only the first two Legendre polynomials, i.e., a constant and slope term ( $a_{j,0}$ ,  $a_{j,1}$ ). However, the code can accommodate several Legendre terms, and we could explore this possibility in the future. These per-camera coefficients are also fit simultaneously. Finally, we have seven coefficients (a principal archetype and one pair of Legendre coefficients for each of the three cameras). We then solve for the coefficients with `scipy`'s `optimize.lsq_linear`<sup>54</sup> module. To summarize, this algorithm solves for the coefficients subject to the specified bounds (described in Equation (1)) through an iterative least-square solving method known as bounded-value least-squares (BVLS), initially proposed by Stark & Parker (1995). BVLS has been shown to converge to a solution in a similar way the nonnegative least-squares (NNLS) converge, as it is modeled on the NNLS iterative approach as developed by Lawson & Hanson (1995). We use this method described above and choose the final redshift that corresponds to the least  $\chi^2$  (see Figure 2 for detailed schematic). For a spectrum given its flux,  $F(\lambda_i)$  and variance  $\sigma^2(\lambda_i)$ , the  $\chi^2$  from its best-fit model ( $S_G$ ) is calculated

as

$$\chi^2 = \sum_i \left( \frac{F(\lambda_i) - S_G(\lambda_i)}{\sigma(\lambda_i)} \right)^2; \quad i = \text{wavelength pixels}. \quad (2)$$

Finally, if a spectrum is typical of galaxies and there are pipeline defects in any of the cameras, the new model should have a lower  $\chi^2$  than the best PCA fit (i.e., without archetypes and Legendre polynomial), as we allow Legendre polynomials to absorb these discontinuities. Hence, the final redshift should correspond to the archetype model rather than the PCA-based fit. In contrast, if the spectrum is vastly different from a typical galaxy spectrum, our method should yield a larger  $\chi^2$  than the PCA-only (without archetypes) model.

The model  $\chi^2$  is calculated for each archetype at each redshift (from best  $N_{zbest}$  redshifts) in a loop, so there are  $N_{archetype} \times N_{zbest}$  loops in total. Then, we select the final redshift of the input spectra corresponding to the least  $\chi^2$  and also store the corresponding archetype subtype (ELG, BGS, or LRG) as the best-fit model for the input spectrum, in case the archetype  $\chi^2$  is smaller than PCA-only (without archetypes) model. The final output redshift file also includes a  $\Delta\chi^2$ , the difference between the best-fit  $\chi^2$  and second best-fit  $\chi^2$ . This difference in  $\chi^2$  is a key indicator of a good redshift, as it indicates if the second best-fit model is significantly far from the best solution.

Using this  $\Delta\chi^2$  parameter, good redshift criteria ( $Q_o$ ) is derived from an extensive analysis of DESI targets observed during the SV phase (see DESI Collaboration et al. 2024b, for more details). The conditions to define a reliable redshift are detailed in Table 2. Note that `COADD_FIBERSTATUS==0` represents the spectra that are free from any hardware issue, while the `ZWARN==0` condition implies that redrock fits are

<sup>54</sup> [https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.lsq\\_linear.html](https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.lsq_linear.html)



**Table 2**

Redshift Success Rate Criteria ( $Q_o$ ) for Spectra Having Nominal Exposure Time (i.e., for Bright Tiles,  $T_{\text{eff}} > 180$  s and for Dark Tiles,  $T_{\text{eff}} > 1000$  s)

Target	Redshift Selection Criteria Common Condition: COADD_FIBERSTATUS=0	References
BGS	ZWARN=0, $\Delta\chi^2 > 40$	(a)
LRG	ZWARN=0, $\Delta\chi^2 > 15$ , $z < 1.5$	(b)
ELG	ZWARN=0, $\log_{10}(\text{O II}_{\text{SNR}}) > f(\Delta\chi^2)$	(c)
QSO <sup>a</sup>	SPECTYPE=QSO	(d)

**Notes.**  $\text{O II}_{\text{SNR}}$  is the signal-to-noise ratio of [O II] flux, which also requires that  $[\text{O II}]_{\text{flux}} > 0$  and corresponding  $\sigma_{[\text{O II}]_{\text{flux}}} > 0$ , and  $f(\Delta\chi^2) = 0.9 - 0.2 \cdot \log_{10}(\Delta\chi^2)$ . ZWARN=0 is the bit warning for a redshift without any obvious issue.

<sup>a</sup> SPECTYPE is from Redrock. This is a very preliminary selection criterion to define the redshift success rate for QSOs to compare with our archetype results. Our goal here is to just understand if the archetype recovers the QSO targets as QSOs. For cosmology purposes, the good redshift for QSOs is defined based on a combination of redshifts obtained with afterburners that use redrock outputs as priors to estimate refined redshifts using a CNN and emission-line modeling. More details on this can be found in Chaussidon et al. (2023).

**References.** (a) Hahn et al. (2023), (b) Zhou et al. (2023), (c) Raichoor et al. (2023), (d) Chaussidon et al. (2023).

good (see DESI Collaboration et al. 2024a, for more details). We use the same criteria to assess the performance of our archetype-based redshift fitter.

### 3.2. Priors on Polynomial Coefficients

This Section discusses the rationale for using priors on polynomial coefficients in the context of our current archetype-based redshift fitting method. The reason that we want to use priors is to avoid the correction for “multiplicative” broadband variations due to an incorrect broadband color of the model spectra, for instance, for QSOs. The challenge here arises from the fact that quasar spectra are often described by a power-law continuum, i.e.,  $F_c \propto \lambda^{-\beta}$ , with  $\beta > 0$ , leading to a spectral slope. Unfortunately, this spectral slope can sometimes be incorrectly attributed to calibration issues within our proposed model. As a result, quasars can be misclassified as galaxies due to extreme freedom to our Legendre coefficients, which impacts spectral classification. We will highlight this more in Section 4.4.2 when we calculate redshift success rates for all target classes. The use of priors aids in reducing the misclassification of genuine quasars as galaxies.

We want polynomial correction to only correct for the “additive” terms caused by the CCD bias issue, which we can understand better with sky fibers. By providing a user-defined prior ( $\sigma_a$ ), the algorithm adds a prior to the  $\chi^2$  of Equation 3, but only to the polynomial terms, before solving for their coefficients (see Figure 2 for the method schematic). This is similar to the Gaussian regularization that quantifies the prior by which the coefficients should be close to zero. With the prior on polynomial terms, the modified  $\chi^2_{\text{new}}$  takes the following form:

$$\chi^2_{\text{new}} = \chi^2 + \sum_j \sum_i \left( \frac{a_{ij}}{\sigma_a} \right)^2. \quad (3)$$

It is important to note that the unit of prior is the same as the coefficients, which is proportional to the calibrated flux.

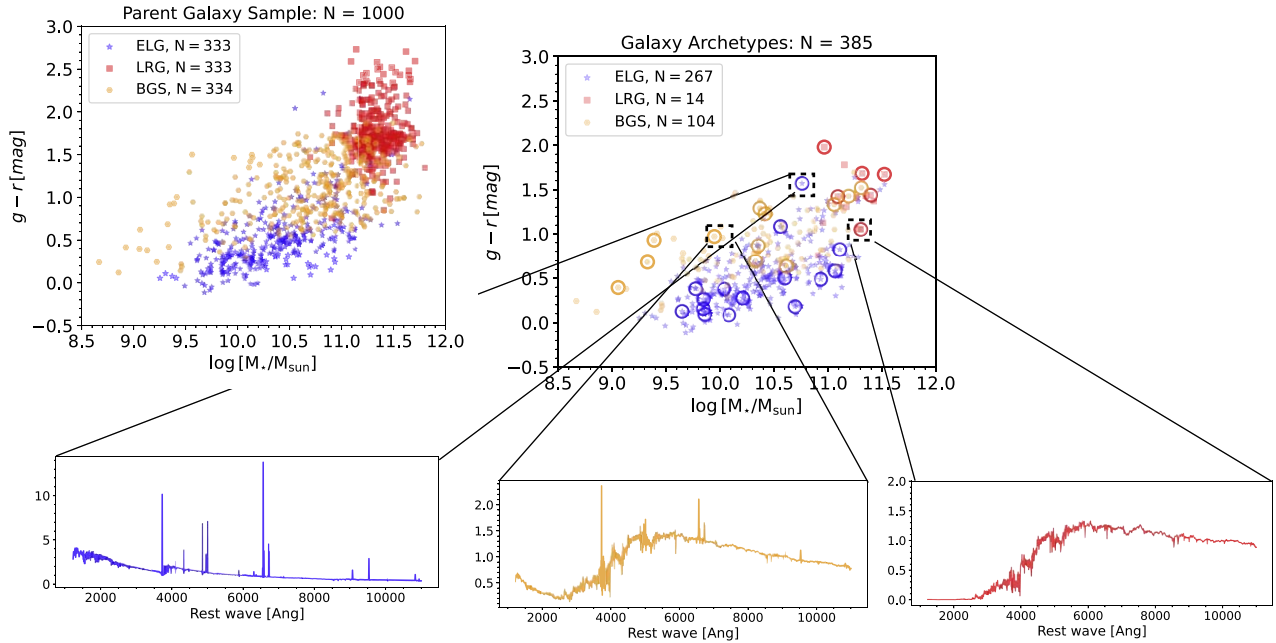
Another key observation is that the default PCA galaxy templates result in artificially high  $\Delta\chi^2$  values when fitting to sky spectra. While the priors avoid correcting quasar continuum variations (which are physical), they also help reduce plausible redshift estimates for sky fibers. These fibers, typically targeted to observe blank regions of the sky, should ideally yield negligible signal and thus low  $\Delta\chi^2$ . A small value of  $\sigma_a$  adds a larger correction to the final  $\chi^2$ .

For this study, we set  $\sigma_a = 0.1$ , a value determined empirically based on our analysis of a subset of 339,712 sky spectra from the Y1 data set of DESI. We discuss this choice in more detail in Appendix B, where we present the distribution of Legendre coefficients for sky fibers for both dark and bright tiles. Notably, the spread of the coefficient distribution varies between 0.1 and 0.5. Our choice of  $\sigma_a = 0.1$  closely aligns with this spread, ensuring that the application of priors also constrains the false-positive redshift estimates for sky fibers. In summary, the priors on the Legendre coefficient are a valuable tool for reducing the misclassification of quasars as galaxies and for not yielding reliable redshift estimates for sky fibers. We show the results for both with and without priors in all Tables and Figures in Section 4 and quantify the performance of our new method compared with Redrock (without archetypes). Finally, we show a detailed flowchart of all of the steps of our method in Figure 2.

### 3.3. Synthetic Galaxy Spectra

Prior to constructing a set of galaxy archetypes, it is important to carefully generate a set of synthetic galaxy spectra that closely align with DESI spectra in terms of wavelength coverage, resolution, and physical properties. The method to generate synthetic spectra is based on a combination of observation and simulation, which is described in detail in S. J. Bailey et al. (2024, in preparation). However, we provide a short summary of the approach below. We compile an extensive sample of galaxies at redshifts  $0 < z < 2$  observed in different photometric and spectroscopic surveys. Then, we use physically motivated emission-line fluxes and stellar population synthesis models to generate synthetic spectra free of any instrumental effects of those surveys. Using these models, we generate very high-resolution spectra of galaxies in a rest-frame wavelength range so that the important emission and absorption lines fall within the DESI wavelength range. We describe the basic properties of these synthetic spectra below.

To generate synthetic spectra for ELGs, we select galaxies (at  $z \sim 1$ ) from the Deep Extragalactic Evolutionary Probe 2 (DEEP2) Galaxy Redshift Survey (Newman et al. 2013) Data Release 4 (DR4; Matthews et al. (2013)). The DEEP2 DR4 consists of  $\sim 50,000$  high-resolution spectra of emission-line galaxies that have measured [O II]  $\lambda\lambda 3727$ , 29 doublet lines (Newman et al. 2013) at an observed wavelength of 6500–9300 Å (typical [O II] doublet region in DESI spectra for ELGs). ELGs are typically star-forming galaxies and appear blue in the image. The interaction of stellar activity (particularly the UV radiation from hot young stars) with the neighboring gas clouds can photoionize metals, which, in turn, can produce narrow forbidden metal lines that are visible on the top of the intrinsic stellar continuum in the galaxy spectra (see Conroy 2013; Kewley et al. 2019, for review). The emission lines were generated based on theoretical models as described in Stasińska & Izotov (2003) that can reproduce the observed



**Figure 3.**  $g-r$  vs. stellar mass for parent galaxy sample (left) and archetypes (right). In both panels, we show the properties of ELGs (blue), LRGs (red), and BGS (orange). We see that galaxies lie in different regions on this plane, and archetypes also span a similar range in the color space. In the right panel, the open circles denote the archetype that we selected randomly to show their spectra in Figure 14 in Appendix A.

trends in synthetic spectra. One caveat is that our emission-line galaxy templates do not include any active galactic nuclei (AGNs), as they can show a combination of narrow and broad emission lines produced by complex physical processes due to central black holes and ongoing star formation. We will look into such subtypes in the future.

The BGS galaxy synthetic spectra were produced by using a flux-limited spectrophotometric sample of 11,000 galaxies ( $0.1 < z < 0.8$ ) observed in AGN and Galaxy Evolution Survey (AGES; Kochanek et al. (2012)). The spectral resolution and wavelength coverage are similar to DESI spectrographs.

Finally, we use a parent sample of 111,114 LRGs observed with DECaLS/DR7 *grzW1W2* photometry for the LRGs. At the same time, the redshifts were compiled from previous large spectroscopic surveys such as BOSS, AGES, and DEEP2 (Zhou et al. 2023). We aimed to simulate a physically motivated representative set of LRGs for targeting purposes and other DESI-related galaxy science. With this selected sample of LRGs, we generated synthetic spectra LRGs using `iSEDfit` package as described in Moustakas et al. (2013) and Moustakas (2017). The method can fairly reproduce the observed SEDs of the LRGs. Note that our LRG modeling does not include emission lines,<sup>55</sup> which may be important for a small fraction of DESI LRGs. A small fraction of the LRG targets can be AGNs that exhibit narrow metal lines in their spectra.

Finally, with the public DESI code,<sup>56</sup> we can generate synthetic galaxy spectra whenever necessary. Using this package, we generated a sample of 1000 (approximately 333 each) rest-frame ( $\lambda = 1228-11000 \text{ \AA}$ ,  $\Delta\lambda = 0.1 \text{ \AA}$ ) synthetic spectra of LRGs, BGS, and ELGs. We present certain physical properties, specifically  $g-r$  (color) versus stellar mass, in the

top-left panel of Figure 3. The colors blue, orange, and red correspond to ELGs, BGS, and LRGs, respectively. The galaxy population dichotomy (Kauffmann et al. 2003) is clearly visible where ELGs predominantly exhibit star-forming characteristics with a typical stellar mass of  $\sim 10^{10.5} M_{\odot}$ , while LRGs are primarily passive and exhibit higher stellar masses ( $\sim 10^{11.5} M_{\odot}$ ). At the same time, the BGS sample occupies an intermediate position between ELGs and LRGs.

### 3.4. Galaxy Archetypes


After generating these synthetic spectra, we use the classification technique, `SetCoverPy`<sup>57</sup> (Zhu 2016) to identify a subsample of spectra (i.e., *archetypes*) that can sufficiently span the physical parameters of the parent sample. In a nutshell, it solves the Set Cover Problem (SCP), i.e., finding an optimal subsample that can represent the given parent sample using similarity (or distance metric) between the input spectra. A detailed description of both `SetCoverPy` and SCP is beyond the scope of this paper; however, we encourage readers to refer to Zhu (2016) for a detailed discussion. It uses the Lagrangian relaxation algorithm (Held & Karp 1971; Geoffrion 1974; Caprara et al. 1999; Fisher 2004) to find the minimum number of *archetypes* that span the parent sample. One example of such use of generating archetypes using `SetCoverPy` is Brodzeller & Dawson (2022), where it was used to identify quasar archetypes for the purpose of creating physical models of quasar spectra.

This method has two free parameters: the  $s^2$  threshold (below which two instances or objects will be considered similar) on distance metric and the choice of weights as a function of wavelength. Once these two parameters are defined, `SetCoverPy` finds the minimum number of *archetypes*. We follow the same approach as described in Bolton et al. (2012) to measure

<sup>55</sup> LRGs are usually passive and have no or very low star formation activity; therefore, they lack many emission lines.

<sup>56</sup> <https://github.com/desihub/desimim>

<sup>57</sup> <https://github.com/guangtunbenzhu/SetCoverPy>



**Figure 4.** Empirical determination of optimal threshold for archetype generation using the *SetCoverPy* method. We show the  $s_{\text{red}}^2$  as a function of the number of galaxy archetypes for each galaxy subtype. Our optimal value for  $s_{\text{red}}^2 \approx 0.003$  lies where the number of LRG archetypes significantly decreases (as shown in the inset). Beyond this threshold, the number of archetypes begins to saturate at higher  $s_{\text{red}}^2$  values.

the similarity matrix of input spectra. We first calculate the distance metric,  $s_{ij}^2$ , which measures the similarity between spectra  $f_i$  and its model  $am_j$ , where  $a$  is some scaling factor

$$s_{ij}^2 = \sum_{l=1}^{N_{\text{pix}}} [f_i(\lambda_l) - am_j(\lambda_l)]^2. \quad (4)$$

$f_i$  corresponds to the normalized flux of  $i$ th spectra. We normalize the flux in such a way that  $\sum_{l=1}^{N_{\text{pix}}} f_i^2(\lambda_l) = N_{\text{pix}}$ . Note that we have used a constant uniform weight for all wavelengths (i.e., weights = 1); therefore, they do not appear in the equation. As described in Bolton et al. (2012), we can perform a least-squares minimization on  $s_{ij}^2$  (differentiating w.r.t  $a$ ), which gives  $a_{\text{best}}(ij) = N_{\text{pix}}^{-1} \sum_{l=1}^{N_{\text{pix}}} f_i(\lambda_l) m_j(\lambda_l)$ . With this, we finally get the minimum  $s_{ij_{\text{min}}}^2$

$$s_{ij_{\text{min}}}^2 = N_{\text{pix}}(1 - a_{\text{best}}^2(ij)); \quad s_{\text{red}}^2 = s_{ij_{\text{min}}}^2 / N_{\text{pix}}. \quad (5)$$

The rest-frame wavelength range we use in our  $s_{ij_{\text{min}}}^2$  calculation is 1228 Å to 11000 Å with  $\Delta\lambda = 0.1$  Å, which gives us a total of  $N_{\text{pix}} = 97,720$  wavelength pixels.<sup>58</sup>

A key point to note is that there is no optimal way to decide the best threshold ( $s_{ij_{\text{min}}}^2$ ) one should choose; it has to be decided empirically (see Zhu 2016, for more discussion). However, it can be understood from Equation (5) that a large threshold will only generate a small number of archetypes, which may not be a fair representation of the parent sample. In contrast, a small threshold will include similar archetypes and overrepresent the parent sample. It is important to run *SetCoverPy* on all galaxies simultaneously without employing distinct thresholds for each galaxy type. Galaxy spectra occupy a large dimensional space lacking sharp boundaries in properties. Consequently, employing different thresholds for different types may not generate a set of diverse archetypes that span the galaxy properties space well.

<sup>58</sup> This is another dimension to explore in future. We chose this resolution in the paper as we want archetypes to have the same resolution as Redrock’s PCA templates. However, this increases the file size.

Hence, we employ an empirical approach to determine our optimal selection. For our purpose, we used the  $s_{\text{red}}^2$  value to decide the optimal threshold. We varied the  $s_{\text{red}}^2$  threshold from 0.001 to 0.1 and observed the number of archetypes associated with each galaxy subtype (ELGs, LRGs, and BGS) and the computational time. For a small  $s_{\text{red}}^2$ , the algorithm selects many similar archetypes as expected. We identify the optimal threshold parameter by noting when the number of archetypes for any subtype significantly changes to a very small number. In our methodology, we optimize the  $s_{\text{red}}^2$  value with respect to the number of LRG archetypes. This is a choice informed by the predominantly passive nature of LRGs, which results in fewer distinct spectral features compared to ELGs or BGS. We observe a significant reduction in the number of LRG archetypes, decreasing from 14 to 2, as the  $s_{\text{red}}^2$  increases from approximately 0.003 to  $s_{\text{red}}^2 \approx 0.005$ . Moreover, the number of LRG archetypes is reduced to 0 if we increase  $s_{\text{red}}^2$  to 0.1. We also note that the 1 or 2 archetypes cannot effectively model DESI LRG’s observed features. We show the variation of the number of archetypes for each galaxy class as a function of  $s_{\text{red}}^2$  in Figure 4. Finally, we select the  $s_{\text{red}}^2 \approx 0.003$  (somewhere between a very low to very high number of LRG archetypes) and generate our final set of galaxy archetypes.

After running the *SetCoverPy* with this choice, we get 385 (14 LRGs, 104 BGS, and 267 ELGs) archetypes from an input set of 1000 spectra. Note that there are more archetypes for ELGs than other subtypes. This is due to the substantial variations in the emission-line properties of ELGs, which are usually associated with their color and stellar properties. In the top-right panel of Figure 3, we show the color ( $g - r$ ) versus stellar mass for our final set of archetypes. One can see the archetypes represent the parent sample (left panel) distribution well in this plane, which shows that *SetCoverPy* performs fairly well in solving the set covering problem for our purpose. Additionally, we also present representative spectra of ELGs (in blue), BGS (in orange), and LRGs (in red) in the bottom panel, showing distinctive spectral features characteristic of each galaxy subclass. We also observe that there are few LRGs with  $g - r > 2$  in the parent galaxy sample while there are none in the archetype sample. The reason is possibly due to very limited distinct features intrinsic to LRGs that slightly differ only in their continuum amplitudes rather than shape. In fact, we see that a few ( $\sim 4$ ) LRG archetypes can actually model the majority of the parent LRG sample. We show those LRGs in Appendix A.

Finally, we also compare some more parent and archetype galaxy properties in Appendix A to understand how well our archetypes span parent galaxies in those dimensions. We plan to explore ways to construct optimal samples of archetypes using a combination of galaxy properties and clustering methods (also see Section 5.2).

### 3.5. Code Implementation

Redrock is written in `python` and updated and maintained as a part of the DESI public GitHub repository.<sup>59</sup> It is a high-performance parallel code that uses MPIs and multiprocessing to run on CPUs and GPUs. The code can fit thousands of spectra, classify their spectral type, and measure redshifts within a few minutes. This is crucial to analyzing the ongoing

<sup>59</sup> <https://github.com/desihub/redrock>



DESI survey as it will observe more than 40 million objects in the next 5 yr. Our updated algorithm is now part of that repository, where we describe the newly added functional arguments in detail. The galaxy archetypes can be downloaded by following the instructions in the Redrock repository.

The archetype code can also run on both CPUs and GPUs; however, the only remaining step that is not GPU accelerated yet is the `scipy`'s BVLS method, which we plan to explore in the future. GPUs significantly improve the runtime for our archetype method, which we illustrate below. These spectral data sets are processed on NERSC's<sup>60</sup> Perlmutter super-computer GPU nodes, each equipped with four GPU cores (each with 40 GB memory) and 64 CPU cores (with a total of 256 GB memory). For each of the tiles we analyzed, there were 10 coadded spectra files (one for each DESI spectrograph), each containing 500 targets (including real and sky fibers). Our archetype processing (including modeling in each camera) exhibited an average processing time of approximately  $\sim 70$  s for 500 spectra when utilizing GPU and CPU resources on NERSC's GPU nodes. We also note that BVLS step (described in Section 3.1) is one of the slowest steps, as it is an iterative method to estimate coefficients. In contrast, processing the same spectral files solely on CPU nodes (each having 64 cores with a total of 512 GB memory) required an average of  $\sim 90$  s ( $\sim 30\%$  slower than CPU+GPU runtime). For reference, when employing the current version (0.19.0) of Redrock (without archetypes), the processing time is 40 s ( $\sim 2$  times faster than archetype) to analyze 500 spectra on NERSC's CPU-only nodes, while the computation time reduces to  $\sim 30$  s ( $\sim 3$  times faster) on NERSC's CPU+GPU nodes. This is expected, given that the Redrock (without archetypes) employs `numpy`'s (or `cupy`'s) linear algebra solvers to solve for PCA coefficients, which is significantly faster compared to the iterative BVLS method.

We emphasize that the new approach is significantly slower (by a factor of  $\sim 2-3$ ) than the Redrock (without archetypes) implementation. However, it is still practical for analyzing large data sets containing millions of spectra. This trade-off in processing time allows for improved physical modeling of spectra and better redshift measurement, and spectral classification and holds significant promise for extending it to other large spectroscopic surveys in the future.

#### 4. Tests and Performance

In this Section, we illustrate the comprehensive tests conducted to assess the performance and efficiency of our archetype method applied to the DESI data set. The analysis involved measuring the method's performance across various data sets characterized by distinct observing conditions, exposure times, SNR, and visual inspection criteria. We compare our test results against the existing DESI redshifts from PCA-only redrock, i.e., without archetypes (version 0.19.0). To provide a comprehensive comparative study, we run our archetype method on three distinct data sets: (a) SV tiles that were observed on several nights to quantify product reproducibility of the pipeline, (b) SV deep coadded spectra that were also visually inspected, and (c) an extensive test run encompassing millions of targets selected from main survey data of Y1. For all of the tests in the subsequent Sections, we use the nine best redshifts ( $N_{z\text{best}} = 9$ , refer Figure 2) of PCA-

template-based redshift scan. We performed some tests exploring  $N_{z\text{best}}$  (see Section 4.5) where we changed it between 3 and 15, and found 9 to be the optimal value. We find that this value is neither very small (which might miss the true redshift) nor very large (computationally expensive). We show all of the comparison statistics for our archetype method with ( $\sigma_a = 0.1$ , Equation (3)) and without prior with redrock (without archetypes) in the sections that follow.

##### 4.1. Archetype Fitting of DESI Spectra

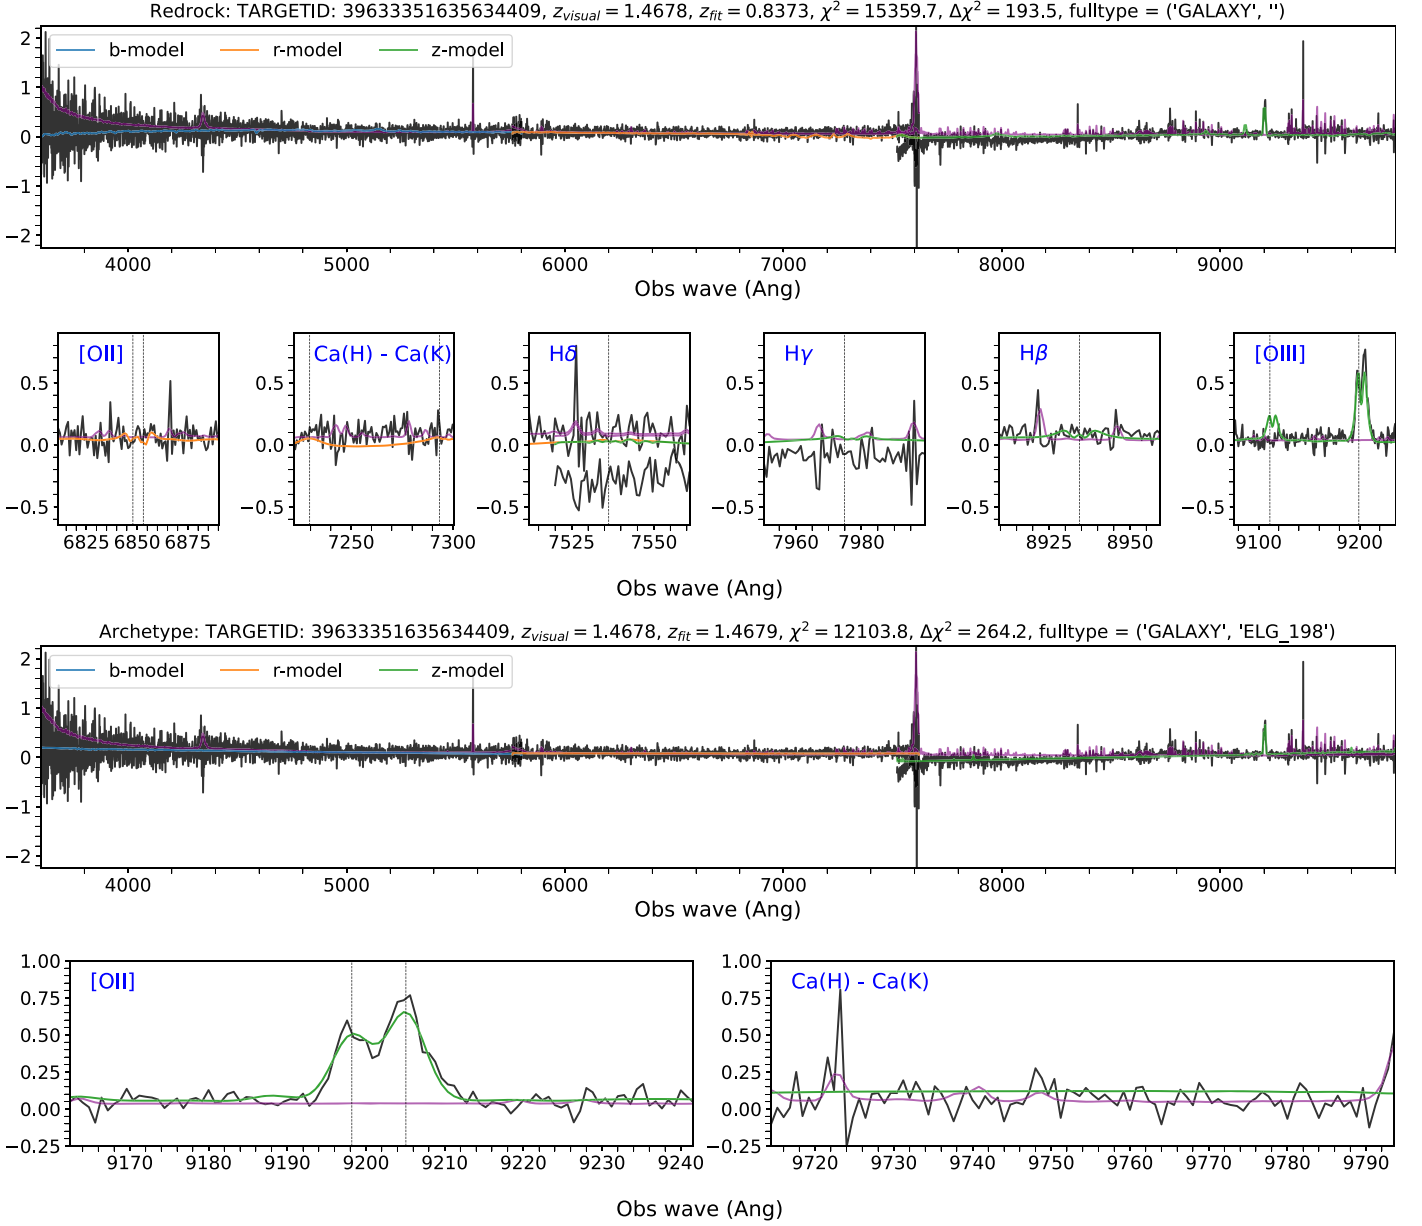
As described in Section 3, our model is more physical than the PCA-based approach and can absorb the pipeline defects more accurately. In Figures 5 and 6, we present two example spectra to illustrate this comparison between the redrock without archetypes and the archetype-based model. Figure 5 shows a spectrum where the CCD discontinuity is clearly visible, and the spectrum has shifted vertically in  $z$ -camera ( $\lambda > 7500$  Å). The spectrum was visually inspected and assigned a redshift of 1.4678. However, the current PCA model (redrock without archetypes, first panel) fails to find the correct redshift. In fact, it misidentifies the [O III] doublet as [O II] (see  $\lambda \sim 9200$  Å). At the same time, our archetype approach (third panel) finds the correct redshift (see the nice [O II] doublet at  $\lambda \sim 9200$  Å), fits the data better, and absorbs the vertical CCD discontinuities in the spectrum more accurately. This is evident not only by eye but also in the final  $\chi^2$  values, where the archetype  $\chi^2 = 12104$  is notably lower than the PCA-only (without archetypes) fit ( $\chi^2 = 15360$ ).

Next, in Figure 6, we demonstrate another example spectrum where redrock fails to find the correct redshift due to an unphysical negative emission-line model of [O II] at  $\lambda \sim 6200$  Å and H $\beta$  model at  $\lambda \sim 8080$  Å (top panel). At the same time, our archetype model (third panel) shows improved physical modeling of the spectrum, yielding a more accurate redshift estimate. This is evidenced by the location of the [O II] ( $\lambda \sim 8320$  Å) line corresponding to the newly determined redshift, where we can see the emission features (also confirmed by small noise in the wavelength region). This spectrum was also visually inspected by DESI collaborators, who assigned a redshift of 1.2317, and our archetype redshift estimate matches that. Notably, the  $\chi^2$  is smaller while  $\Delta\chi^2$  is larger in our archetype model than that of the PCA-only (without archetypes) run. The higher  $\Delta\chi^2$  indicates that the next best-fit model is far less probable. The wrong redshift estimate in the PCA-only (without archetypes) model is likely due to the flexibility of PCA to fit the negative feature as [O II] line, which possibly drives the  $\chi^2$  to be smaller than the model at the right redshift. To test this hypothesis in Figure 7, we show the redrock model (without archetypes) at the right redshift ( $z = 1.2317$ ) and the archetype model at the wrong redshift ( $z = 0.6615$ ).

The top panel shows the PCA model for the spectrum estimated at the “right” redshift ( $z = 1.2317$ ). It is clear that despite the fact that the PCA model well fits the physical features of the spectrum (see zoomed version), the estimated  $\chi^2 = 11038$  is larger than the best-fit redshift ( $\chi^2 = 11003$ ; compare top panel of Figure 6). This confirms our hypothesis that the PCA model is driven by fitting the unphysical negative dip at  $\lambda \sim 6200$  Å. In the bottom panel, we show the archetype model estimated at the “wrong” redshift ( $z = 0.6615$ ). As expected, we find that  $\chi^2 = 10911$  is larger than the best-fit  $\chi^2 = 10833$  (bottom panel of Figure 6). Furthermore, this

<sup>60</sup> National Energy Research Scientific Computing Center, [www.nersc.gov](http://www.nersc.gov).





**Figure 5.** First: PCA-template-based best-fit model (shown in colored solid lines) of a galaxy spectrum. We can see the flux discontinuity in the  $z$  camera ( $\lambda \sim 7550 \text{ \AA}$ ). Second: zoomed version of PCA model of the same spectrum, showing the expected location of emission lines. The continuum is not well-fitted with PCA at the camera boundary. [O III] doublet is also not well fitted. Third: archetype-based per-camera best-fit model of the same spectrum. The fit looks better visually and absorbs the flux discontinuity in the  $z$  camera more accurately, also evident by smaller  $\chi^2$ . Fourth: zoomed version of archetype-based per-camera model of the same spectrum, showing the expected location of emission lines. [O II] emission line ( $\lambda \sim 9200 \text{ \AA}$ ) is more accurately modeled. The black curve is the observed flux, the purple curve shows the error spectra in all panels, and the fitted model is shown in solid color lines.

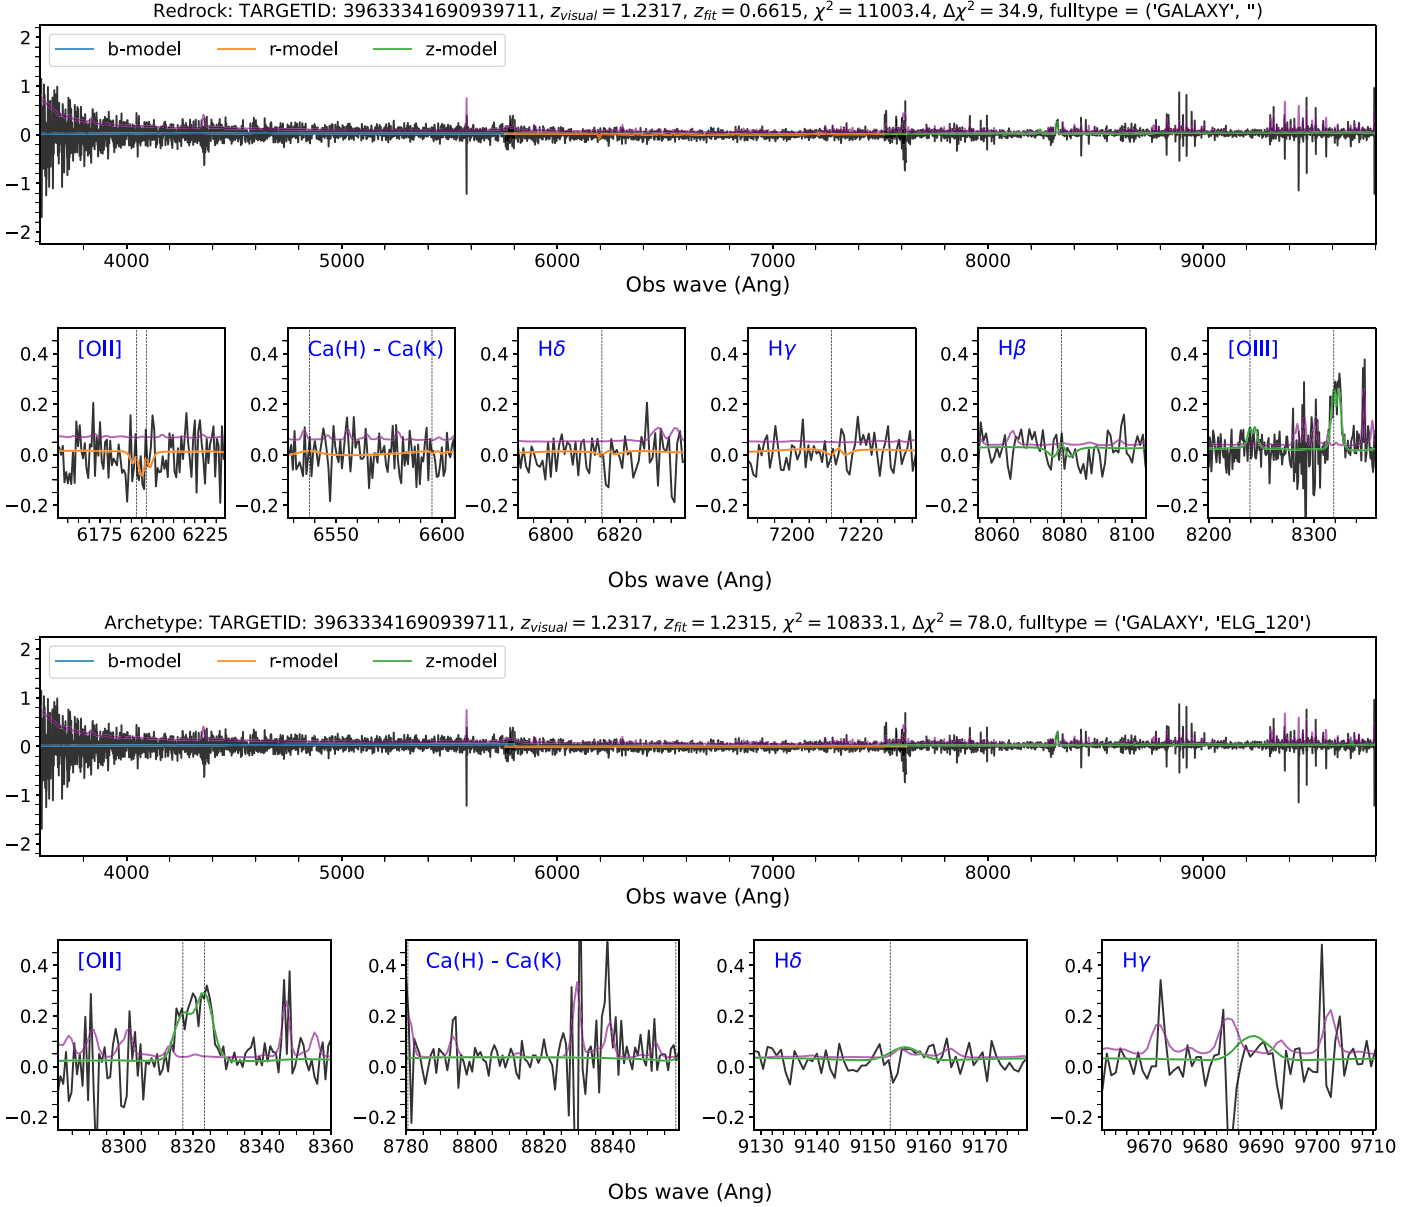
model does not fit the unphysical negative dip (see zoomed version around  $\lambda \sim 6200 \text{ \AA}$ ), even at the wrong redshift. This is because our archetypes are physical galaxy models, and their coefficients are always positive in our model, which naturally prohibits any such unphysical fitting.

#### 4.2. Survey Validation Repeat Observations

A total of six unique “tiles” were specifically designed to target distinct galaxy target classes in the SV phase. Three of these tiles were designed to observe ELGs (TILEIDs 80606, 80608, and 80610; see Raichoor et al. 2023 for more details). As discussed earlier, ELGs were divided into ELG\_HIP,

ELG\_LOP, and ELG\_VLO. Two additional tiles (TILEIDs 80605 and 80609; see Zhou et al. 2023) were designed to target LRGs, and one tile (TILEID 80613; see Hahn et al. 2023) was designed for BGS. As described in Section 2.1, the BGS targets were further subdivided into BGS\_BRIGHT and BGS\_FAINT.

These tiles were systematically observed in the SV phase on different nights between 2020 December and 2021 May. The nominal exposure times (same for the main survey) varied depending on the target class, ranging from approximately  $\sim 180 \text{ s}$  for BGS targets to around  $\sim 1000 \text{ s}$  for LRG and ELG targets. In total, 8910 unique galaxy targets were observed on these tiles, with 1153 spectra attributed to BGS\_BRIGHT, 792 to BGS\_FAINT, 2522 to ELG\_LOP, 644 to ELG\_VLO, 300 to



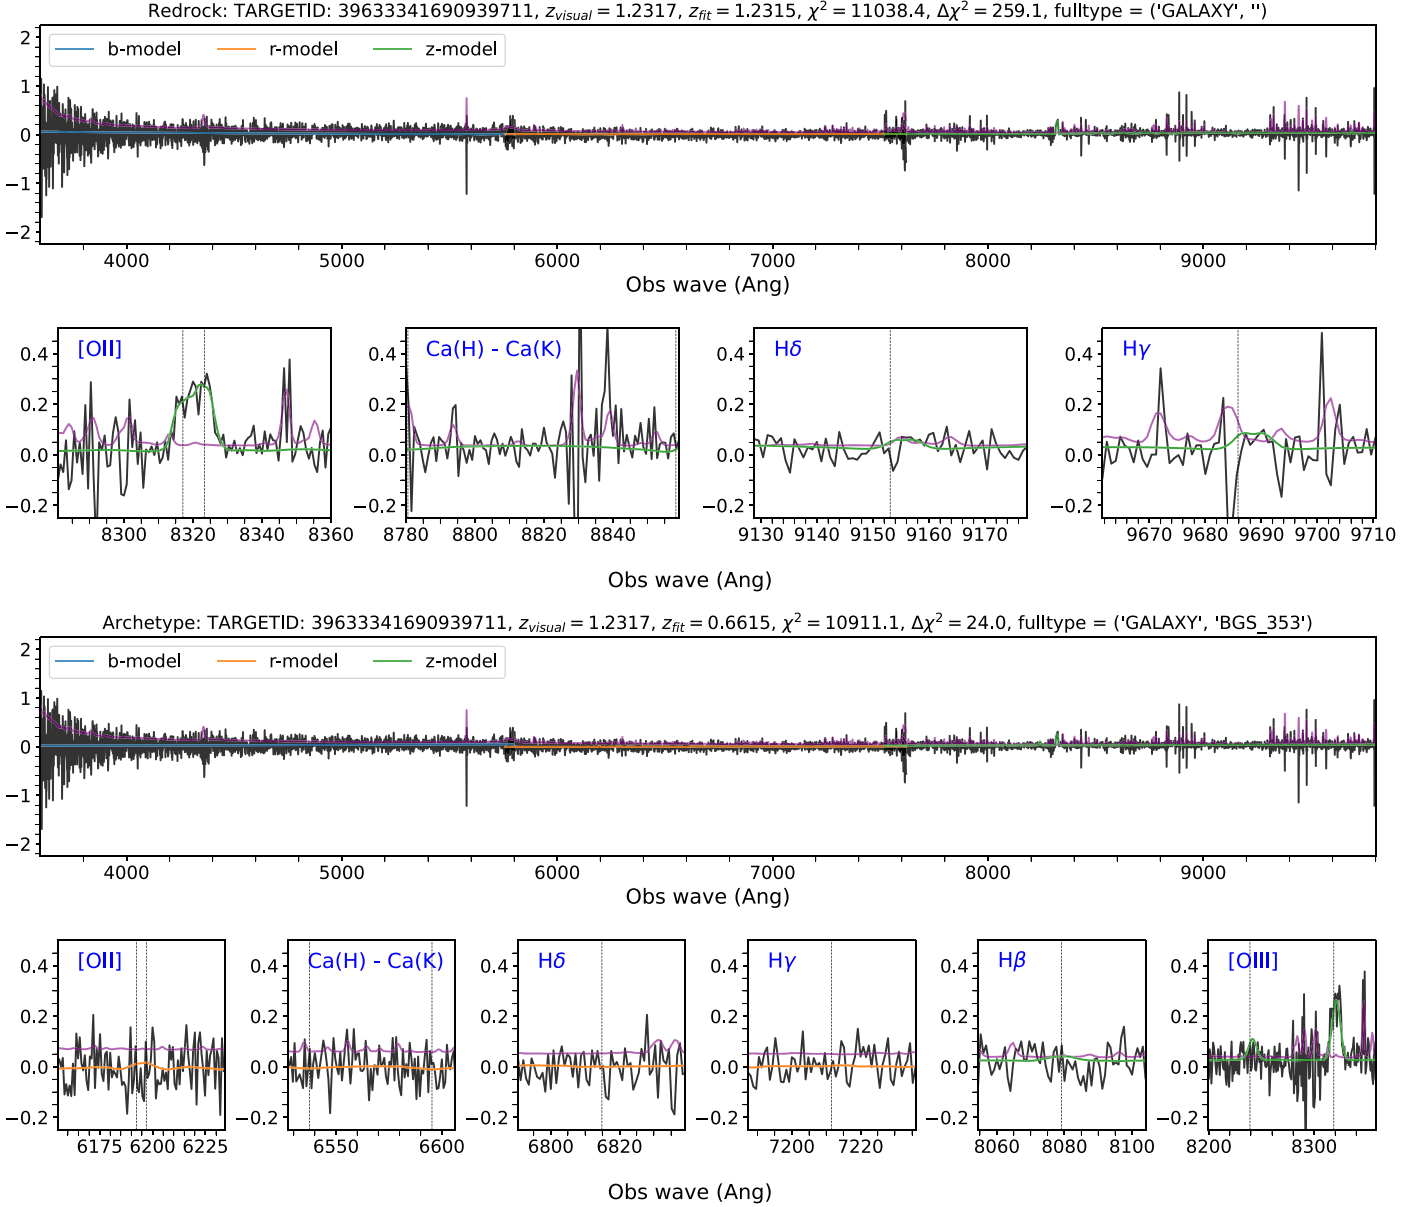
**Figure 6.** First: PCA-template-based best-fit model (shown in colored solid lines) of a galaxy spectrum. Second: zoomed version of PCA model of the same spectrum. We can see the unphysical negative absorption line fitting for [O II] ( $\lambda \sim 6200 \text{ \AA}$ ) and  $H\beta$  ( $\lambda \sim 8080 \text{ \AA}$ ) lines. In fact, Redrock (without archetypes) misidentifies [O II] doublet as [O III]. Third: archetype-based best-fit model of the same spectrum. Fourth: zoomed version of archetype-based per-camera model of the same spectrum. The new model is more physical, and the redshift is more accurate, as confirmed by visual inspection of [O II] doublet. The black curve is the observed flux, the purple curve shows the error spectra in both panels, and the model is shown in solid color lines.

ELG\_HIP, and 3499 to LRGs. The individual exposures of these targets were observed multiple times, ranging from four to ten nights per target class spread across several nights. We ran the redshift fitter for these targets and compared their redshift estimates with those obtained from their deep coadded spectra (discussed in Section 4.3) using redrock (without archetypes) and archetype-based per-camera approach and compared their performance.

This test is crucial to quantify the performance of redshift fitter software, as it shows its ability to analyze the spectra with varying observing conditions and SNR; therefore, understanding the software performance on single-epoch spectra is important. In addition, the test also allows us to empirically define “good redshift” criteria (defined in Table 2) for different target classes based on a combination of the rate of

“catastrophic failure” and  $\Delta\chi^2$  of the best-fit model. The catastrophic redshift failure is defined as the difference between redshift ( $z_{\text{epoch}}$ ) measured with low-SNR single-epoch spectrum and redshift ( $z_{\text{coadd}}$ ) measured with deep coadded high-SNR spectrum of the same target. If  $|\Delta v| = \left| \frac{z_{\text{epoch}} - z_{\text{coadd}}}{1 + z_{\text{coadd}}} \right| \cdot c > 1000 \text{ km s}^{-1}$ , then the measurement is termed a “catastrophic failure” (DESI Collaboration et al. 2024b).

The detailed comparison statistics of redshift success and catastrophic failure can be found in Table 3. We present the redshift differences (expressed as  $\Delta v$ ) between single-epoch observations and the deep coadds for the same DESI targets, as a function of the  $\Delta\chi^2$  values obtained from single-epoch observations for BGS\_FAINT, BGS\_BRIGHT, ELG\_LOP, ELG\_VLO, and LRGs for PCA-only (without archetypes)



**Figure 7.** First: PCA-template-based model of the same galaxy spectrum (shown in Figure 6) at right redshift. Second: zoomed version of the model. We can see that  $\chi^2$  is higher than the best-fit model (see Figure 6). Third: archetype-based fit of the same spectrum at wrong redshift. Fourth: zoomed version of the archetype model. We see that our archetype model does not fit any unphysical feature, and  $\chi^2$  is higher than the best-fit model (see third panel of Figure 6). The black curve is the observed flux, the purple curve shows the error spectra in both panels, and the model is shown in solid color lines.

model (left panels) and our archetype model (right panels) in Figures 8, 9, and 10, respectively. To guide the readers, we define the meaning of each quadrant for these Figures: bottom right is good, meaning individual nights agree with the deep coadd, and they have a highly confident quality metric; bottom left are missed opportunities: they agreed with deep coadd but were not flagged as confident, upper left are failures, but at least they were not flagged as confident, upper right are the worst: catastrophic failures with redshift discrepancies but flagged as confident. In all Figures, the blue dots refer to the accepted redshifts ( $N_a$ ), while red crosses refer to the rejected redshifts ( $N_r$ ). Note that the sum of blue and red points in all samples is equal to the number of targets shown in Table 3, except for ELG targets, where the sum of plotted points is smaller than the total number of targets in the Table. This is because we select

spectra having positive [O II] flux and their errors before we calculate [O II] SNR, which is plotted in Figure 9.

Next, we define additional statistics for these data sets. Utilizing the truth tables for redshifts (i.e., the visual redshifts) of these targets, we calculate the precision, recall, and F1 scores for both methods. The results are presented in Table 7 in Appendix C. Overall, these metrics are very high for both methods, indicating their excellent performance on the DESI data sets. However, the metrics are marginally higher for the archetype method (with  $\sigma_a = 0.1$ ). Here, we also point out that we cannot perform such analysis for main survey targets, as there are no truth tables for them. However, we expect very similar metrics for them as well, as many of these targets were also selected for the main survey and have comparable exposure times.

**Table 3**  
Survey Validation Repeat Tiles Redshift Success Comparison—Redrock (without Archetypes) vs. Archetype Method

Target Class	$N_{\text{target}}$	$N$ (selected by Quality Cut, $Q_o$ )			$N(Q_o,  \Delta v  > 1000 \text{ km s}^{-1})$			Redshift Purity (%)		
		Redrock	Archetype (no prior)	Archetype ( $\sigma_a = 0.1$ )	Catastrophic Failure			(Defined in the Text)		
					Redrock	Archetype (no prior)	Archetype ( $\sigma_a = 0.1$ )	Redrock	Archetype (no prior)	Archetype ( $\sigma_a = 0.1$ )
BGS_FAINT	2267	2178	2183	2184	9	8	8	99.58	99.63	99.63
BGS_BRIGHT	3256	3178	3184	3187	11	13	11	99.65	99.60	99.66
LRG	9981	9556	9625	9633	80	62	62	99.16	99.36	99.36
ELG_LOP	9806	7647	7658	7647	31	25	19	99.59	99.67	99.75
ELG_VLO	2417	2327	2326	2332	6	2	3	99.74	99.91	99.88
ELG_HIP	1136	944	953	951	1	1	2	99.89	99.89	99.81

**Note.**  $X$  and  $X_o$  are the good redshift criteria parameters for the corresponding target classes with nominal effective exposure time as described in Table 2. The catastrophic failure and redshift purity are defined in the text. We also apply the following quality cuts ( $Q_o$ ) on single-epoch data. (a) BGS: ZWARN = 0,  $X > X_o$ , (b) LRG: ZWARN = 0,  $z_{\text{pernight}} < 1.5$ ,  $X > X_o$ , (c) ELG:  $X > X_o$  and  $[\text{O II}]_{\text{flux}}, \sigma_{[\text{O II}]_{\text{flux}}} > 0$ , i.e., valid measurements of single-epoch [O II] fluxes and their errors.  $z_{\text{pernight}}$  and ZWARN denote the redshift and its warning bits obtained for single-epoch spectra.

We observe that both redrock (without archetypes) and archetype models (with prior  $\sigma_a = 0.1$ ) perform equally well for almost all target classes. We find that the overall catastrophic redshift failure rate is relatively very small ( $\gtrsim 0.5\%$ – $1\%$  of the full sample) in both models. This shows that the current PCA-based Redrock is already very efficient. For the BGS\_BRIGHT targets (Figure 8), our archetype-based approach (with  $\sigma_a = 0.1$ ) yields 3187 accurate redshifts ( $\Delta\chi^2 > 40$ ) compared to 3178 in the Redrock (without archetypes) run, while similar success is seen for BGS\_FAINT target classes (2178 versus 2184 in redrock (without archetypes) versus archetype with  $\sigma_a = 0.1$ ), though slightly fewer catastrophic failures in archetype mode (eight versus nine in Redrock (without archetypes)). For ELG\_LOP, ELG\_VLO and LRGs, we see that in the archetype-based method, the relative improvement is  $\sim 10\%$ – $40\%$  in reducing catastrophic redshift failure ( $|\Delta v| > 1000 \text{ km s}^{-1}$ ; compare the blue dots in the upper-right quadrant in each panel). For ELG\_LOP targets, the archetype-based approach (with  $\sigma_a = 0.1$ ) yields 19 catastrophic failures (25 without prior case), a  $\sim 40\%$  improvement compared to the 31 failures in the Redrock (without archetypes) run (see Table 3 and the top panel of Figure 9, for detailed comparison). Next, for ELG\_VLO, there are three catastrophic failures in the archetype approach (with  $\sigma_a = 0.1$ ) as opposed to six in Redrock (without archetypes), showing a significant improvement (see bottom panel of Figure 9). The performance is similar for both approaches for ELG\_HIP. Similarly, for LRGs, the archetype method delivers  $\sim 25\%$  (62 compared to 80 in Redrock (without archetypes)) fewer catastrophic redshift failures than redrock (see Figure 10). Also, the redshift success is slightly lower in Redrock (without archetype) mode (9556 versus 9625). Furthermore, our archetype-based modeling approach consistently delivers marginally higher redshift purity, i.e., the fraction of accurate redshifts that satisfy redshift success criteria,  $\frac{N(Q_o, |\Delta v| < 1000 \text{ km s}^{-1})}{N(Q_o)}$ .  $Q_o$  is the good redshift criteria in terms of  $\Delta\chi^2$  and  $[\text{O II}]_{\text{SNR}}$  for the corresponding target classes as described in Tables 2 and 3. The redshift purity for our archetype method is slightly higher (0.1%–0.2%) than the current Redrock. The relatively higher purity and lower catastrophic failures in the archetype approach underscore its improved performance in achieving precise redshift measurements. However, seeing more significant changes in these statistics would be our goal in the future.

It is important to emphasize that the overall performance of both Redrock (without archetypes) and archetype approach on DESI data well exceeds the quality thresholds necessary for cosmological analysis. For instance, the rates of catastrophic failure are significantly lower ( $< 1\%$ ) than the acceptable limit of  $< 5\%$  (DESI Collaboration et al. 2024b). Furthermore, the typical redshift errors achieved with the DESI pipeline meet the precision requirements for large-scale structure analysis. Specifically, the typical redshift errors for the BGS, ELGs, and LRGs are approximately 10, 20, and  $40 \text{ km s}^{-1}$ , respectively, which are well below the required redshift precision of approximately  $200 \text{ km s}^{-1}$  (DESI Collaboration et al. 2024b; Lan et al. 2023).

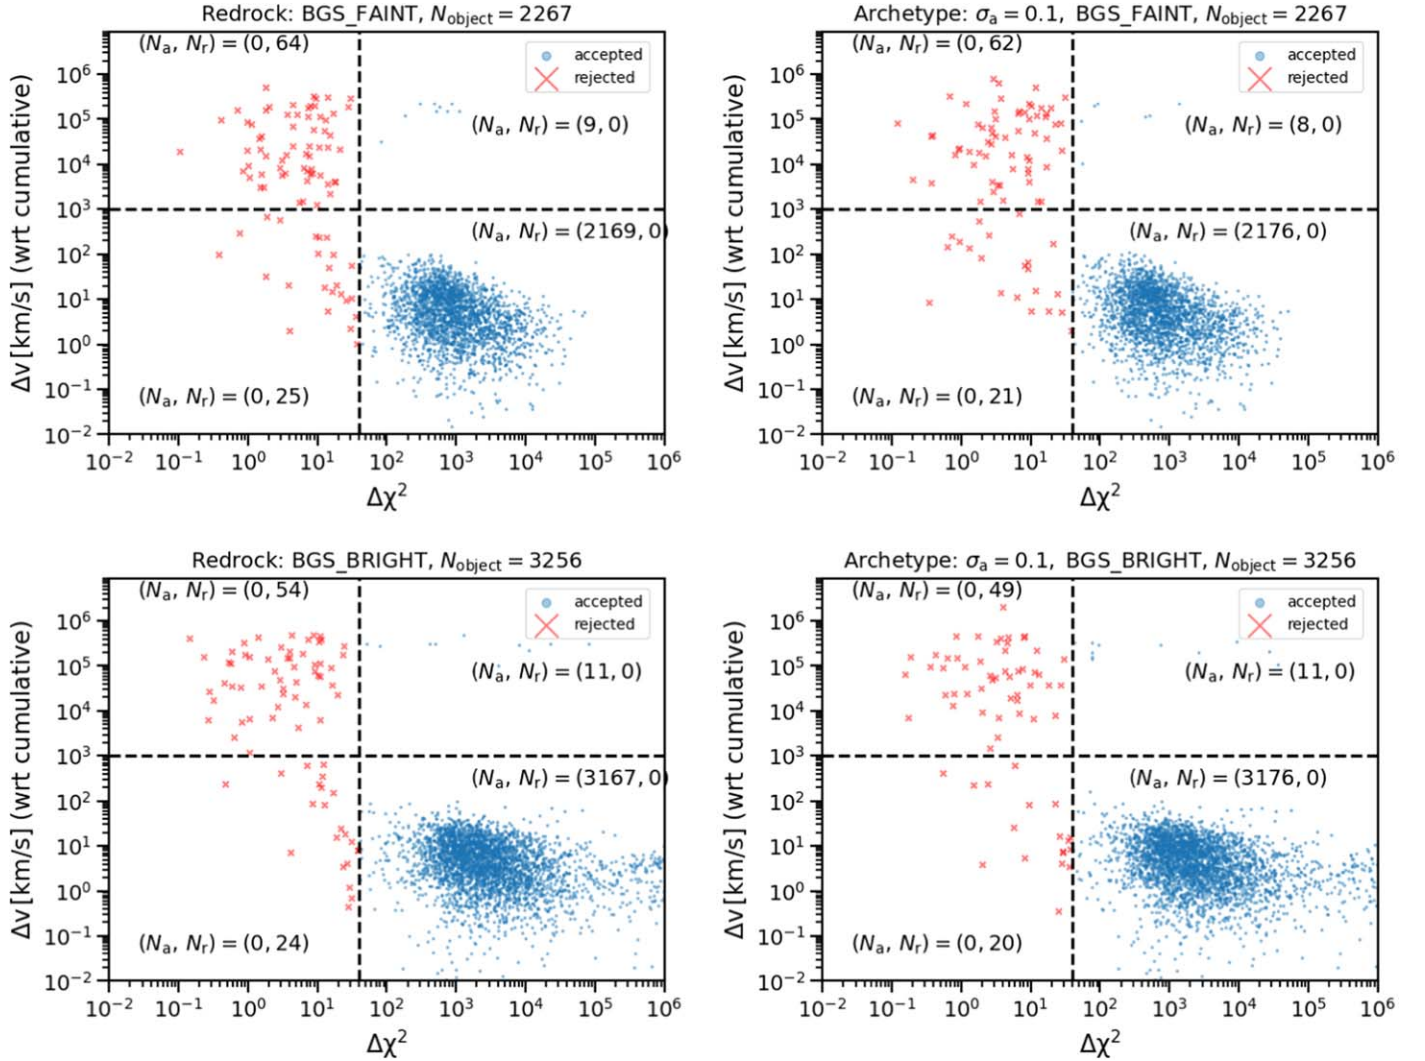
### 4.3. Survey Validation Visually Inspected Deep Tiles

In this Section, we describe the efficiency of the archetype method in measuring the redshifts for targets that DESI collaborators visually inspected during the SV phase. As described earlier, the single-epoch observations of these tiles were combined across exposures to construct higher-SNR coadded spectra. The cumulative effective exposure times vary from  $\sim 1500 \text{ s}$  to around  $\sim 5000 \text{ s}$ . These spectra were visually inspected by at least two DESI collaborators, following a well-defined, uniform approach, which is detailed in Lan et al. (2023). In summary, visual inspectors look for the most prominent emission lines (e.g.,  $\text{H}\alpha$ ,  $\text{H}\beta$ , [O II], [O III]),  $4000 \text{ \AA}$  break, and the shape of the continuum to verify the redshift and spectral class of DESI spectra estimated by redshift fitter. A set of predefined metrics, as described in Lan et al. (2023), is employed to maintain consistency in the inspection process. These metrics aid in assigning quality labels (QA\_VI)<sup>61</sup> to each coadded spectrum based on spectral features to ascertain the reliability of the redshift estimates and spectral classes. We ran our archetype method on these targets and then compared our redshifts obtained from the pipeline and those determined through visual inspection.

The comparison results are presented in Table 4, where we compare our results (with and without priors on Legendre terms) with redrock (without archetypes) measurements. Our approach shows a 10%–30% improvement in reducing

<sup>61</sup> The members assign a quality index between 0 (not a robust redshift) and 4 (most robust redshift) depending upon the spectral features visible in the spectra.





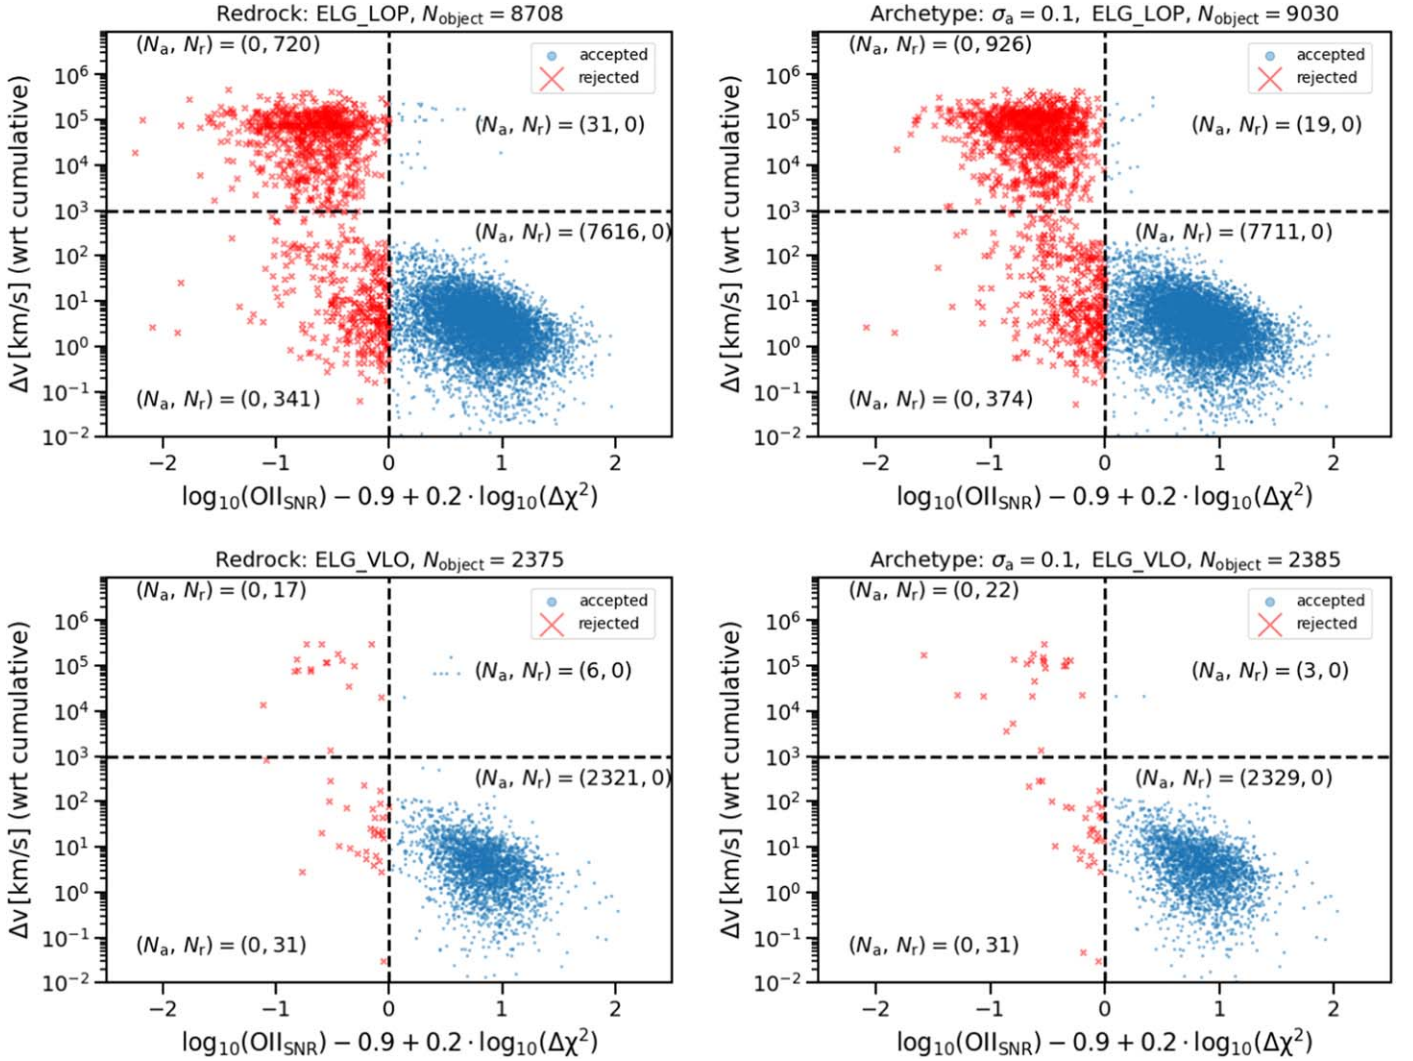
**Figure 8.** The redshift difference between single-exposure and deep coadded spectra of same BGS\_FAINT (top) and BGS\_BRIGHT (bottom) targets. The  $\Delta\chi^2$  values are taken from single-epoch observations. The vertical dashed line ( $\Delta\chi^2 > 40$ ) shows the boundary to define confident redshift, and the horizontal dashed line defines the catastrophic redshift failure, i.e.,  $|\Delta v| > 1000 \text{ km s}^{-1}$ . The left panel shows the results for Redrock (without archetype), and the right panel shows the results for our archetype approach (with a prior of  $\sigma_a = 0.1$  on polynomial coefficients). In both panels, the blue dots show the accepted redshifts ( $N_a$ ) while the red crosses show the rejected redshifts ( $N_r$ ). To guide the readers, the bottom-right quadrant in each subpanel represents good performance, showing that individual night redshifts match with deep coadded redshifts of the same spectrum, while the upper right shows the redshifts that differ from the deep coadded spectrum and also are flagged as very robust redshift in the pipeline.

catastrophic redshift failure for all target classes. For example, in the case of BGS\_BRIGHT, the number of catastrophic redshift failures is nine in redrock (without archetypes) compared to six in archetype (with prior  $\sigma_a = 0.1$ ) case. Similarly, ELG\_LOP targets have 21 catastrophic failure in the archetype (with prior  $\sigma_a = 0.1$ ) mode, compared to 25 in redrock (without archetype) mode. Finally, for LRG, there are only 26 redshift failures as opposed to 39 in redrock (without archetype). Next, we also compare the redshift estimates, color-coded by the  $\chi^2$  difference between the PCA-only model (i.e., Redrock without archetypes) versus archetype model (with  $\sigma_a = 0.1$ ) in Figure 11. We see a spurious clustering at  $z_{\text{redrock}} \sim 1.6$  in the redrock method (left), which goes away in the archetype method (right). The archetype method consistently yields lower  $\chi^2$ , indicating better model fits for the input spectra. These results clearly show that our method

performs better than the existing method on deep coadded spectra.

#### 4.4. Large Test Runs on the Y1 data set

As described above, our archetype approach has improved overall redshift estimates for targets observed with SV tiles. However, these are very small data sets (a few thousand only) relative to what DESI will provide in the coming years. Subsequently, a comprehensive large-scale test run was conducted on a large suite of tiles extracted from the Y1 data set of DESI. Our test data set included observations from 36 nights spread across 2021 May to 2022 June, averaging three nights per month, encompassing a diverse range of observing and instrumental conditions. Within this data set, 553 tiles were selected, comprising 233 dark and 320 bright tiles. This subset of data encompassed a total of 2,281,969 target spectra and 339,712 sky spectra selected from fibers pointing to targets and



**Figure 9.** The redshift difference between single-epoch observation and deep coadded spectra of same ELG\_LOP targets. The good redshift selection criteria are defined in terms of [O II] SNR and  $\Delta\chi^2$  where both quantities are taken from single-epoch observations (see Table 2). We only include spectra that have valid [O II] line strength (i.e.,  $[\text{O II}]_{\text{flux}} > 0$  and  $\sigma_{[\text{O II}]_{\text{flux}}} > 0$ ). The vertical dashed line ( $X > 0$ , see Table 2) shows the boundary to define good redshift, and the horizontal dashed line defines the catastrophic redshift failure, i.e.,  $|\Delta v| > 1000 \text{ km s}^{-1}$ . The left panel shows the results for Redrock (without archetypes), and the right panel shows the results for our archetype approach (with a prior of  $\sigma_a = 0.1$  on polynomial coefficients). In both panels, the blue dots show the accepted redshifts ( $N_a$ ) while the red crosses show the rejected redshifts ( $N_r$ ). To guide the readers, the bottom-right quadrant in each subpanel represents good performance, showing that individual night redshifts match with deep coadded redshifts of the same spectrum, while the upper right shows the redshifts that differ from the deep coadded spectrum and also are flagged as very robust redshift in the pipeline. Archetype performs  $\sim 50\%$  better in reducing catastrophic redshift failure.

blank sky region,<sup>62</sup> respectively. We used the archetype-based per-camera spectral fitting method and compared redshifts obtained with redrock (without archetypes).

#### 4.4.1. Performance on Sky Fibers

If the pipeline was perfect and the subtracted sky spectra were pure zero-mean uncorrelated noise, all templates would have the best solution of coefficients = 0 and  $\Delta\chi^2 = 0$ . In this run, we first demonstrate our analysis by focusing exclusively on sky fibers to evaluate the efficacy of our methodology in the context of objects specifically targeted toward the blank sky. Note that the sky fibers are excluded from the final target redshift catalog. However, this is an important exercise to mitigate the false-positive cases in our archetype method, as

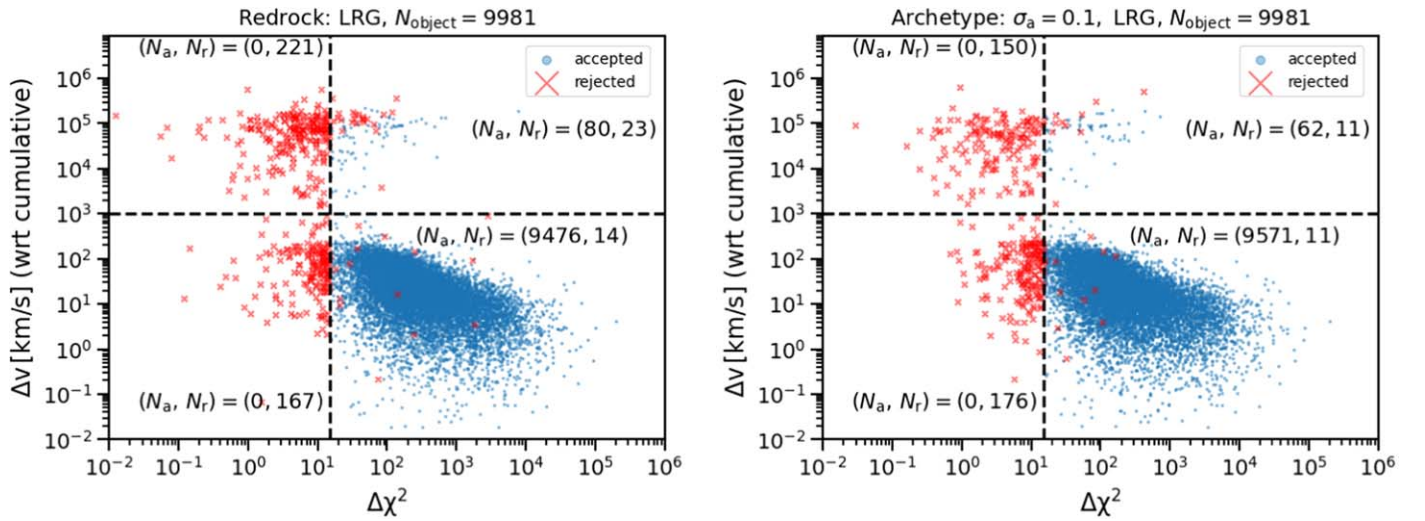
these spectra should not yield robust redshift. It also helps reduce false redshifts of real targets in sky-dominated wavelength regions. We also emphasize that some confident sky redshifts may be real serendipitous objects falling into sky fibers, though such occurrence is rare, and we do not focus on that here.

For this purpose, we use the same good redshift criteria (defined in Table 2) to understand the statistics of plausible redshift measurements for these sky fibers. Our expectation is to reduce the number of false positives for these fibers.

We show the best-fit redshifts for these sky spectra as a function of fiber number in Figure 12. These best-fit redshifts point to the “attractor solutions” due to upstream pipeline imperfections, which risk pulling Redrock toward incorrect fits on very-low-SNR data. As elucidated in Table 5 and illustrated in Figure 12, our archetype approach yields a notably lower number of sky fibers with confident redshifts than the default PCA-only (without archetypes) run. This outcome aligns with

<sup>62</sup> Each DESI coadded fits file includes a header storing OBJTYPE of observed TARGETIDS. A reader can apply the conditions of OBJTYPE==TGT and OBJTYPE==SKY to select target and sky fibers.





**Figure 10.** The redshift difference between single-epoch and deep coadded spectra of the same LRG targets. The  $\Delta\chi^2$  values are taken from single-epoch observations. The vertical dashed line ( $\Delta\chi^2 > 15$ ) shows the boundary to define good redshift, and the horizontal dashed line defines the catastrophic redshift failure, i.e.,  $|\Delta v| > 1000 \text{ km s}^{-1}$ . The left panel shows the results for Redrock (without archetypes), and the right panel shows the results for our archetype approach (with a prior of  $\sigma_a = 0.1$  on polynomial coefficients). In both panels, the blue dots show the accepted redshifts ( $N_a$ ) while the red crosses show the rejected redshifts ( $N_r$ ). To guide the readers, the bottom-right quadrant in each subpanel represents good performance, showing that individual night redshifts match with deep coadded redshifts of the same spectrum, while the upper right shows the redshifts that differ from the deep coadded spectrum and also are flagged as very robust redshift in the pipeline. There are some rejected redshifts in the bottom-right quadrant. That is because we also reject LRGs with  $z > 1.5$ . Archetype performs  $\sim 20\%$  better in reducing catastrophic redshift failure.

**Table 4**

Visually Inspected Deep Coadded Tiles Redshift Comparison for Redrock vs. Archetype

Target class	$N_{\text{target}}$ ( $\text{QA\_VI} \geq 2.5$ )	$ \Delta v  > 1000 \text{ km s}^{-1}$ (Catastrophic failure)		
		Redrock	Archetype (no prior)	Archetype ( $\sigma_a = 0.1$ )
BGS_FAINT	754	9	6	6
BGS_BRIGHT	1091	5	3	3
ELG_LOP	2132	25	21	21
ELG_VLO	639	2	1	1
ELG_HIP	263	2	2	2
LRG	3440	39	26	26

**Note.** Catastrophic redshift failure is defined as,  $|\Delta v| = \frac{|z_{\text{ref}} - z_{\text{visual}}|}{1 + z_{\text{visual}}} \cdot c$ , where  $z_{\text{ref}}$  and  $z_{\text{visual}}$  are pipeline and visual redshifts, respectively. We select spectra that are visually classified as “good”; i.e.,  $\text{QA\_VI} \geq 2.5$  (see Lan et al. 2023, for details).

expectations, as our method relies on physical galaxy models, and sky fibers typically do not correspond to any physical galaxy, being spectra of blank sky. In contrast, PCA templates exhibit a high degree of flexibility, allowing them to confidently fit such spectra, as evident by the data (see the “Redrock (without archetypes)” column) in Table 5 and left panel of Figure 12. In the daily quality assessment, dedicated DESI collaborators look at these plots to understand the redshift performance.

In summary, our archetype method demonstrates a 6%–40% improvement in mitigating erroneous redshift fits for sky fibers. This is a substantial improvement over the current redrock. For instance, when applying redshift success criteria akin to LRG on sky fibers (bottom panel of Figure 10), the Redrock (without archetypes) run identifies 5808 objects with robust redshifts. In contrast, our archetype method identifies only 4498 without priors and 4275 ( $\sim 27$  percent improvement) with  $\sigma_a = 0.1$  on


polynomial coefficients. This improvement in reducing false-positive redshifts is consistently observed in our archetype approach (see Table 5 and the right panel of Figure 12) with redshift selection akin to BGS (top-right panel, 3184 in archetype versus 1902 in Redrock (without archetypes)) and ELG (middle right panel, 1969 in archetype versus 2075 in Redrock (without archetypes)) target classes. This outcome aligns with our rationale for utilizing priors to reduce such misidentifications in the first place.

Another improvement we observe with our archetype approach is that it reduces the accumulation of spurious redshifts around any particular redshift as a function of fiber number. In the bottom panel of Figure 12, we show results for sky fibers selected with redshift selection criteria akin to LRGs (see Table 2). We find that the Redrock (without archetypes) redshifts (bottom-left panel) show spurious redshift accumulation between fibers 1400 and 2000. However, the archetype case has no such accumulation (Figure 12, bottom-right panel). This can be attributed to the fact that archetypes represent physical galaxy spectra; consequently, they do not cluster around any fibers as the final fit can be any redshift with small  $\Delta\chi^2$ .

#### 4.4.2. Redshift Success Rate

Next, we estimate the redshift success rate for all DESI targets using the same redshift criterion defined in Table 2 and compare our archetype results with success rates from the Redrock (without archetypes) run. We also note that we do not need to define a new criterion<sup>63</sup> to define the redshift success rate in our archetype approach while also improving the overall performance of the spectral fitting and redshift estimation pipeline. We present the comparison in Table 6. We find that our archetype approach (with  $\sigma_a = 0.1$ ) consistently yields a

<sup>63</sup> Although this can be explored in the future when we run the archetype mode for all spectral types and redshifts.



**Figure 11.** The redshift comparison between the PCA-based Redrock method (left) and the archetype method (right), color-coded by their  $\chi^2$  difference for visually inspected galaxy spectra, reveals notable differences. The archetype method demonstrates reduced spurious clustering of redshifts and fewer catastrophic failures around  $z \sim 1, 1.6$ . Furthermore, the archetype method generally yields lower  $\chi^2$  values, indicating better model fits.

higher (by 0.1%–0.8%) redshift success rate for almost all galaxy target classes. For example, the redshift success rate for ELG\_LOP is 70.90% in the Redrock (without archetypes) run while it increases to 71.56% in our archetype (with prior,  $\sigma_a = 0.1$ ) approach ( $\sim 0.7\%$  improvement). The lower success rate in the PCA-only model can, in part, be attributed to the misclassification of strong emission lines (e.g., see Figure 6). However, since the archetypes are derived from actual galaxy spectra, which include all of the expected emission lines, they can help find the correct redshift corresponding to their location in observed spectra. A similar relatively higher-redshift success rate is seen for ELG\_VLO and ELG\_HIP; compare the redrock versus archetype columns in Table 6. Similarly, the success rate for LRGs is 98.81% in the archetype method compared to 98.79% in Redrock (without archetypes). Finally, for the QSO targets, Redrock (without archetypes) delivers a 67.15% redshift success rate, while it reduces to 66.84% in archetype mode without prior. However, when we use prior  $\sigma_a = 0.1$ , the success rate increases to 67.11%. This justifies our proposition that adding priors on Legendre polynomials in our model reduces quasar misclassification and does improve the overall performance for them (see Section 3.2).

We also observe that the average redshift success rate is significantly lower for the ELG\_LOP class compared to other target classes. This discrepancy arises from a combination of the target selection algorithm and the redshift selection criterion. First, ELG targets are selected using a combination of  $grz$  colors, which serve as a proxy for the [O II] flux and star formation. These targets extend nearly to the imaging depth, and the color-based cuts are imperfect, leading to the inclusion of many faint targets with undetectable [O II] flux for reliable redshift estimation.

Additionally, the wavelength coverage of DESI targets poses another issue. For any target with  $z > 1.62$ , the [O II] line falls outside the spectrograph’s range, resulting in no emission line being detected by DESI. Consequently, this results in redshift failure, as redshift efficiency is also dependent on the SNR of the detected [O II] flux.

It is important to point out that in terms of relative percentage, these improvements are low; however, DESI will collect more than  $\sim 30$  million galaxy spectra in 5 yr, and even such a marginal increase will translate to very large numbers in

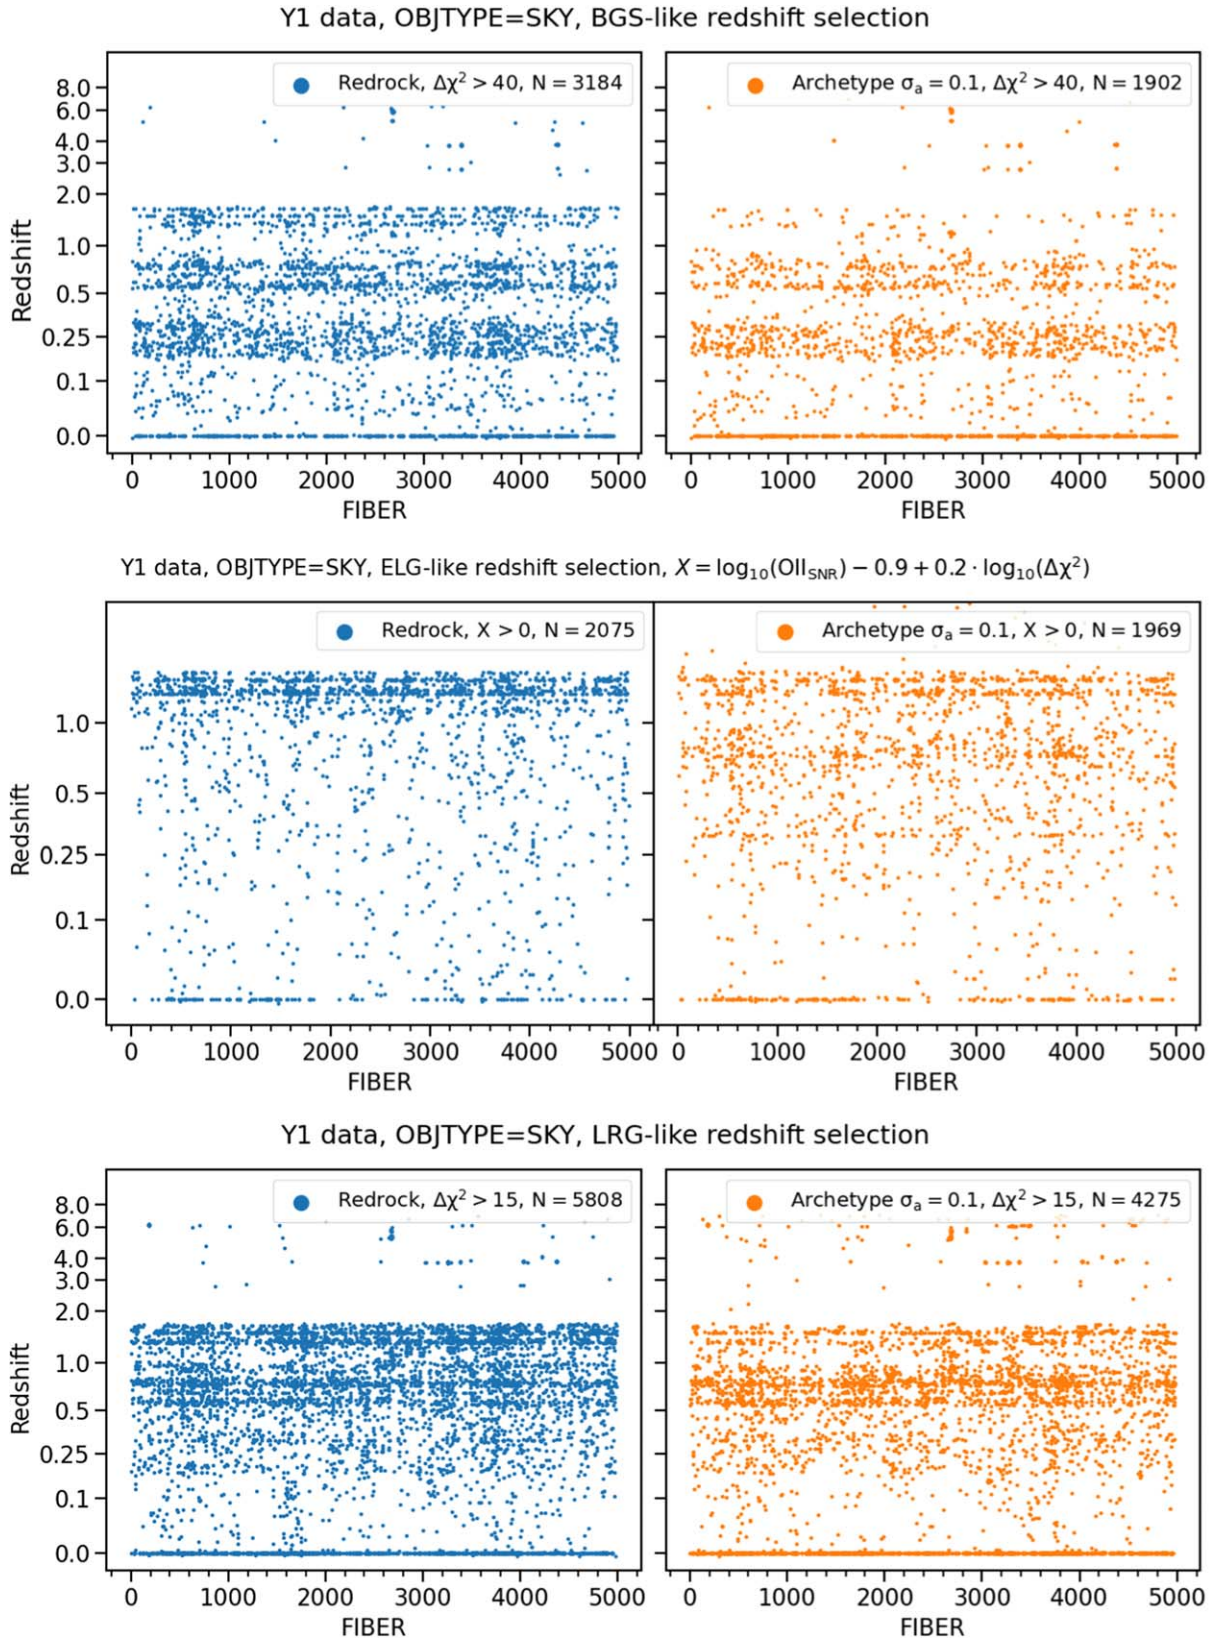
absolute terms. At the same time, the method also resolves the inherent issues with the current PCA-based method we discussed in Section 3.1. Furthermore, within our archetype approach, the redshift success rate is consistently higher for runs with prior (rightmost column of Table 6) than without prior (middle column of Table 6) on polynomial coefficients. This should not surprise us, as putting priors on the polynomial coefficients allows us to restrict their freedom, which reduces the unphysical modeling of input spectra.

#### 4.5. Further Algorithmic Tests

We also explored various algorithmic adaptations by changing the model parameters and including additional complexities within our archetype-based galaxy redshift fitting method. The first test involved varying the number of best redshifts ( $N_{z_{\text{best}}}$ ; see Figure 2) on which the archetype-based model should be run. We varied the  $N_{z_{\text{best}}}$  redshifts from 3 to 15. We observe that the redshift performance got worse for  $N_{z_{\text{best}}} \leq 6$  compared to the Redrock (without archetypes). This outcome is likely attributable to the possibility that a small number of PCA-based best redshifts may not include the correct redshifts, particularly if the initial fit was unphysical. Next, increasing  $N_{z_{\text{best}}}$  to 9 significantly improved the performance compared to Redrock (without archetypes) without compromising ( $\sim 2$  times higher than default run) much on runtime, and we use this value for our final analysis. On the other hand, increasing  $N_{z_{\text{best}}}$  to 12 and 15 did not improve the  $\chi^2$  values, and the best redshifts remained unchanged while significantly ( $\sim 3$  times more than default run) increasing the runtime.

Next, we also explored including additional nearest neighbors in our method to use multiple nearby archetypes in superposition so as to achieve a “sub-grid” interpolation. For that purpose, after getting the best archetype fit for a given spectrum (i.e., the archetype corresponding to the minimum  $\chi^2$ ), we consider  $N$ -nearest archetype neighbors around that in  $\chi^2$  space. Then, we construct a model that is a linear combination of these nearest archetypes and the Legendre polynomials and estimate the new  $\chi^2$ . This is to test if increasing the number of archetypes improves the overall spectral fitting for galaxies further. However, we observe very





**Figure 12.** Redshift vs. fibers for sky fibers (selected with OBJTYPE=SKY criteria) for sky fibers from DESI Y1. In each panel, we also show the number of false positives selected using good redshift criteria (see Table 2) like true galaxy targets. The top panel shows results for BGS-like selection; the middle panel shows results for ELG-like selection, and the bottom panel presents the results for LRG-like redshift selection. The left panel (blue) shows the result for Redrock (without archetypes), and the right panels show the results of our archetype method (with a prior  $\sigma_a = 0.1$  on polynomial coefficients, shown in orange) in all panels. It is evident that the archetype performs 6%–40% better in rejecting false-positive redshifts for sky fibers.

**Table 5**  
Redshift Statistics of Sky Fibers from Y1 Data Set

Program	N (skyfibers)	Selection Criteria	N (False Positives)		
			(Confident Redshifts)		
			Redrock	Archetype (no prior)	Archetype ( $\sigma_a = 0.1$ )
dark	136, 134	ELG	2075	1980	1969
	136, 134	LRG	5808	4498	4275
bright	203,578	BGS	3184	1936	1902

**Table 6**  
Test Run on DESI Targets from Y1 Data Set

Target	N	Redshift Success Rate (%)		
		Redrock	Archetype (no prior)	Archetype ( $\sigma_a = 0.1$ )
BGS_FAINT	227,784	98.78	98.34	98.78
BGS_BRIGHT	562,848	98.64	98.53	98.72
ELG_LOP	340,064	70.90	71.55	71.56
ELG_VLO	44,170	94.63	94.68	94.68
ELG_HIP	94,046	75.76	76.29	76.29
LRG	233,784	98.79	98.77	98.81
QSO <sup>a</sup>	238,775	67.15	66.84	67.11

**Note.**

<sup>a</sup> Also see the footnote for Table 2.

slight improvements in  $\chi^2$  in some cases, albeit at the expense of significantly higher runtime ( $\sim 3$  times higher than a single archetype run), as it entails an extra computational step. The slight reduction in  $\chi^2$  is expected, as adding more nearest neighbors introduces additional degrees of freedom into the equation. Consequently, considering the insignificant gain in best fit at the expense of increased runtime, we decided to run the algorithm with just one archetype.

We also ran our method only with archetypes without Legendre polynomials, which performed worse than the current approach. This is expected, as archetypes alone can only model the shape and features of the spectrum and not the pipeline defects, because they do not have the same flexibility as the Legendre polynomials in absorbing those defects accurately.

Another test was to include higher-order (quadratic and above terms; see Equation (1) and Figure 2) polynomial terms in the model. They do improve the  $\chi^2$  in some cases, and thus, the best redshifts also change. However, this also yields some unphysical negative flux fittings, as the combination of negative coefficients and higher degree polynomials can overfit the data. This also comes from the fact that the linear combination polynomial basis vectors can fit almost any curve. We plan to explore this in more detail using the NNLS-like method in the future.

## 5. Discussion

This is one of the first comprehensive modifications in the redshift fitting algorithm of DESI spectral data and extensive testing on DESI data to quantitatively assess its performance. This effort is of particular significance in light of the substantial scale of DESI observations anticipated over the next 3–4 yr, addressing crucial issues inherent to the current algorithm.

### 5.1. Improved Galaxy Modeling and Redshift Estimation with archetypes

In Section 3.1, we provide the details of our new archetype model coupled with a per-camera polynomial fitting (see Equation (1)) for modeling DESI spectra. Our method is significantly different from previous studies (Cool et al. 2013; Hutchinson et al. 2016) that have used archetypes to fit the low-resolution spectra and estimate their redshifts. They allow the archetype coefficients to take negative values, which are susceptible to yielding unphysical fits occasionally. They also do not include polynomial modeling of spectra in individual cameras and fail to accurately absorb the CCD and throughput offsets. Moreover, we also add priors (Section 3.2) on our polynomial coefficients that help us reduce false positives in sky fibers and the misclassification of quasar spectra. Since the archetypes represent physical galaxy spectra with emission (such as [O II], [O III], H $\alpha$ , H $\beta$ ) and absorption feature (4000 Å break, Ca–K, H lines) characteristics to their physical properties, they seem to be better at fitting those features in observed spectra and provide less flexibility in fitting unphysical features (see Figure 6). However, the reduced flexibility lacks the ability to capture some peculiar galaxy spectra that might also be detected serendipitously, which we will discuss in the next Section. We ran our redshift fitter with the archetype-based approach on various small and large data sets and compared its performance against Redrock (without archetypes), as elucidated in Section 4.1.

To quantify the improvement in redshift estimates with our archetype model, we performed a large test on DESI tiles across diverse data sets. Notably, these enhancements are achieved without changing the existing good redshift criteria, though this can be further explored in the future. An extensive analysis of our method on tiles observed on multiple nights during the survey validation phase (see Section 4.2) shows that the new method yields fewer catastrophic failures (10%–30%; see Table 6 and Figures 8, 9, and 10) while simultaneously increasing the redshift purity for all galaxy subclasses, as it can deliver better models for the CCD discontinuities (see Figure 5). The other noteworthy improvement is in sky emission-line regions. In cases where PCA-only (without archetypes) templates might erroneously fit the noise as an actual spectral feature, the archetype method inherently does not find physical features and, therefore, reduces false-positive redshifts. Next, we also find that the new method is more effective (10%–20% in a relative sense) in reducing catastrophic redshift failure for visually inspected galaxy spectra (see Section 4.3 and Table 4).

Finally, when run on millions of DESI targets from the main survey (Section 4.4), our method shows an increased redshift success rate (see Table 6) by 0.5%–0.8% for all galaxy target classes while yielding a similar success rate for QSOs as compared to the Redrock (without archetypes). Although the



relative changes are small, the numbers are much larger in the absolute sense. The new method is relatively fast to process such a large number of targets within acceptable time frames, as described in Section 3.5. Furthermore, we find that adding a prior while solving for polynomial coefficients in our model further enhances the ability to reject false-positive redshift measurements for sky fibers (i.e., the spectrum taken for blank sky regions; see Table 5) while not reducing the redshift success rate for any galaxy subclass (Sections 4.4.1 and 4.4.2). Based on our extensive tests, we conclude that the new archetype-based per-camera polynomial fitting method shows significant improvement over Redrock (without archetypes). Additionally, the method is generic enough that its functionalities can be easily extended to other similar large spectroscopic surveys in the near future.

It is also important to compare the redshift success rate of the DESI pipeline with that of previous large surveys. For example, GAMA achieved approximately 97% correct redshifts down to an apparent magnitude of 21 in the SDSS  $r$  band (Loveday et al. 2012). Similarly, LRGs in the BOSS program had a redshift success rate of approximately 93% for  $i_{\text{fiber}} < 21.5$  (Bolton et al. 2012), while the DEEP2 survey delivered an overall success rate of  $\geq 70\%$  (Cooper et al. 2006). In comparison, the DESI pipeline, Redrock, consistently achieves higher success rates across all target classes. This improved performance can be attributed to enhanced imaging, a refined target selection algorithm, robust PCA templates, better sky modeling, and also better resolution than BOSS.

### 5.2. Future Improvements and Extensions

It is also important to understand the shortcomings of our method. An important step would be to test our method independently of PCA-based template fitting. As described earlier, due to the runtime complexity, we still rely on a list of redshifts derived from the initial PCA-based redshift scan; therefore, if the true redshift is not present in that list, our method will fail. In the near future, we plan to optimize the BVLS method and extend our method to all redshifts to perform an independent comparison with current DESI redshifts. Another important improvement can be brought to the synthetic spectra data set (Section 3) that was generated using spectrophotometry from the data sets that were in place before DESI (like DEEP2 and AGES). In the future, however, we will use DESI observations themselves. This may impact the overall quality and diversity of galaxies in these data sets, as they are all flux-limited surveys and have restricted wavelength coverage and theoretical modeling of SEDs. This is an important point to remember while interpreting the current results.

With the ongoing DESI survey, we have collected millions of galaxy spectra with updated and more robust spectrophotometry and SED modeling. Notably, these spectra originate from the same instrument, facilitating the construction of a more consistent synthetic galaxy spectral data set that aligns more accurately with observations. Concerted efforts are underway within the DESI collaboration to refine the galaxy templates based on the wealth of data accumulated from DESI. The initial value-added catalogs for these DESI galaxies, incorporating new state-of-the-art spectrophotometry (i.e., `fastspecfit`<sup>64</sup>; Moustakas et al. 2023) and derived physical properties, are set to be released in the near future

(J. Moustakas et al. 2024, in preparation). This will be an important step in enhancing the accuracy and applicability of galaxy archetypes in redshift fitting.

Another important issue is that our current ELG archetypes lack the broad diversity in line ratios associated with different metal transitions, which is usually associated with their stellar and gas properties. This may be critical for identifying and modeling different line ratios within the ELG subclass. Additionally, the current set of archetypes also lacks AGNs and LINER-like galaxies and high redshifts ( $z > 1.5$ ) Lyman break galaxies or Ly $\alpha$  emitters. It is important to understand that the intricate theoretical modeling of such galaxies is extremely challenging; however, they are found in DESI data sets. Addressing this limitation represents another improvement that will be integrated into upcoming archetypes based on `fastspecfit` galaxy properties and line ratios in the future. Therefore, our archetypes may not comprehensively cover the vast diversity of galaxy spectra. Consequently, if we encounter galaxy spectra that are fairly different from our archetypes, we might miss them. However, running `SetCoverPy` with appropriate weights while generating archetypes can be very useful for handling such outliers. Even the default PCA-based templates were generated from a galaxy sample that excluded these special cases. However, the inherent flexibility of PCA allows adaptability in identifying diverse objects; it is not inherently designed for such specialized classification. Currently, a combination of visually inspecting the spectral features and redshift measurements from previous surveys is used to identify them.

However, as described in Section 4.4.2, the overall low redshift success rate for ELGs is intrinsically related to their selection algorithm and stringent redshift criterion. The redshift criterion was developed to ensure a very low rate of catastrophic failures, which is crucial for clustering analysis. While relaxing this criterion might increase the success rate, it would also lead to an increased rate of catastrophic failures (Raichoor et al. 2023). To achieve significantly higher efficiency, the target selection algorithm must be refined, and the definition of reliable redshifts needs to be reconsidered.

Our future strategy involves incorporating rare galaxy classes when generating PCA templates and galaxy archetypes. This step is deemed crucial for enhancing the precision of our redshift fitter, especially in light of upcoming surveys like DESI-II, which are being designed to target these peculiar galaxies (Ruhlmann-Kleider et al. 2024).

In the current analysis, we are using only galaxy archetypes, though it is run on all redshifts regardless of spectral type. So, there still can be unphysical fitting for QSOs and stars, which have different spectral features than galaxy archetypes, and PCA-based  $\chi^2$  can still be smaller. The challenge lies in constructing a set of archetypes that span the extensive diversity inherent to quasars. Currently, their unified model is not well understood and modeled (see Netzer 2015, for a recent review). On the other hand, constructing archetypes for stars is also challenging, given the diversity of spectra from young to old stars. Another complexity is in modeling the spectra of CVs and WDs, which show very distinct features from main-sequence stars.

In the immediate future, a potential avenue for future enhancement would be redshift fitting in each amplifier within the DESI camera. Notably, the pipeline occasionally presents calibration challenges in amplifiers within the camera. This manifests again as an offset or gradient-like polynomial feature in regions within a given camera. This per-amplifier modeling

<sup>64</sup> <https://fastspecfit.readthedocs.io/en/latest/>



is challenging due to wavelength overlap, and defining a method to split the spectra in each amplifier is not straightforward. This refinement can show further improvement in error absorption than what we have gained now. Ongoing efforts are directed toward incorporating such intricacies in future releases of Redrock.

Additionally, we can explore the possible algorithms to construct archetypes from a given input set. Currently, we are using a generic classification algorithm (*SetCoverPy*) for this, which may not be very efficient sometimes, as the optimal reduced  $\chi^2$  value for the distance matrix is not well defined. Therefore, developing an approach that includes clustering based on physical properties rather than relying solely on machine-learning-based techniques can be more useful. However, this is quite challenging, given the difficulty in defining an  $N$ -dimensional parametric space that effectively splits galaxies or quasars based on their properties in that hyperplane. In any case, exploring this avenue is quite exciting and can offer potential insights into a more detailed understanding of galaxy and quasar classification. This also has implications for measuring their redshifts while also offering useful information on physical properties such as stellar mass or star formation rate.

Finally, as described earlier, our method is generic enough to quickly extend to other large surveys similar to DESI, such as WEAVE,<sup>65</sup> WAVES (planned in the 4MOST survey), Prime Focus Spectrograph on Subaru (Tamura et al. 2016), and MOONS<sup>66</sup> on Very Large Telescope (Cirasuolo et al. 2020). For instance, the upcoming WEAVE instrument (Dalton et al. 2014; Jin et al. 2024) will have two cameras to obtain high-resolution spectra of galaxies and quasars. Given that our method can extract information on the number of cameras from the input spectrum, it can automatically fit the WEAVE spectrum in each of those two cameras and estimate the redshift and spectral class. We aim to explore this avenue and test its applicability to the WEAVE spectra in the near future.

## 6. Summary of Conclusions

In this paper, we have presented a modification to the existing redshift estimation algorithm for DESI spectra. The archetype-based per-camera polynomial fitting approach fits the spectra of DESI galaxies and measures their redshifts and spectral classes. Our comprehensive tests demonstrate the efficacy of our method in mitigating unphysical model fits to galaxies and adeptly addressing pipeline defects introduced during the spectral reduction process. The main conclusions are as follows:

1. The primary DESI redshift fitter Redrock occasionally suffers from the issue of yielding unphysical fits and inaccurate redshift and spectral type of galaxies. It also struggles to account for pipeline defects that are visible in the extracted spectra as vertical shifts or gradient-like polynomial features in the spectrum within one of the cameras due to CCD bias and camera throughput issues.
2. To develop a more physical galaxy model fit and improve the redshift estimates, we employ an archetype-based approach to model DESI galaxy targets. Furthermore, we introduce a per-camera polynomial spectral fitting to robustly absorb pipeline defects.

3. Our extensive tests show that the new method performs 10%–40% better (e.g., 31 versus 19 in redrock versus archetype method for ELGs) in reducing catastrophic redshift failure than redrock (without archetypes) on tiles observed on several nights during the survey validation phase. Additionally, the redshift purity is consistently higher than the current redrock for these tiles in our archetype method. Furthermore, our method shows a notable ability to find precise redshifts for visually inspected high-SNR coadded spectra of DESI targets.
4. We find that the archetype method performs marginally better on survey validation targets, which also confirms the overall robustness of the Redrock pipeline and, hence, the downstream DESI science results.
5. Additionally, our method performs better on millions of DESI targets selected from the Y1 data set. We observe an increased redshift success rate (by 0.5%–0.8%) for all galaxy target classes (compare 70.90% versus 71.55% for ELGs in the redrock versus archetype method, respectively) while maintaining similar success for QSO targets. At the same time, we also find that our method reduces the number of false-positive redshifts (by 5%–40%) for sky fibers.
6. Finally, we explain the shortcomings and future improvements to our method and discuss the generic nature of our method and its applicability to other large upcoming surveys.

## Acknowledgments

We thank the anonymous referee for providing insightful comments that have helped improve the clarity of the paper. All of the computations were performed at the DOE’s super-computer facility NERSC, located at Berkeley Lab. This material 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 and Innovation of Spain (MICINN), 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.

*Software:* Matplotlib (Hunter 2007), NumPy (Harris et al. 2020), Scipy (Virtanen et al. 2020), SetCoverPy (Zhu 2016).

<sup>65</sup> William Herschel Telescope Enhanced Area Velocity Explorer.

<sup>66</sup> Multi-Object Optical and Near-infrared Spectrograph.

### Data Availability


The data plotted in the results are available at <https://doi.org/10.5281/zenodo.12007781>.

### Appendix A

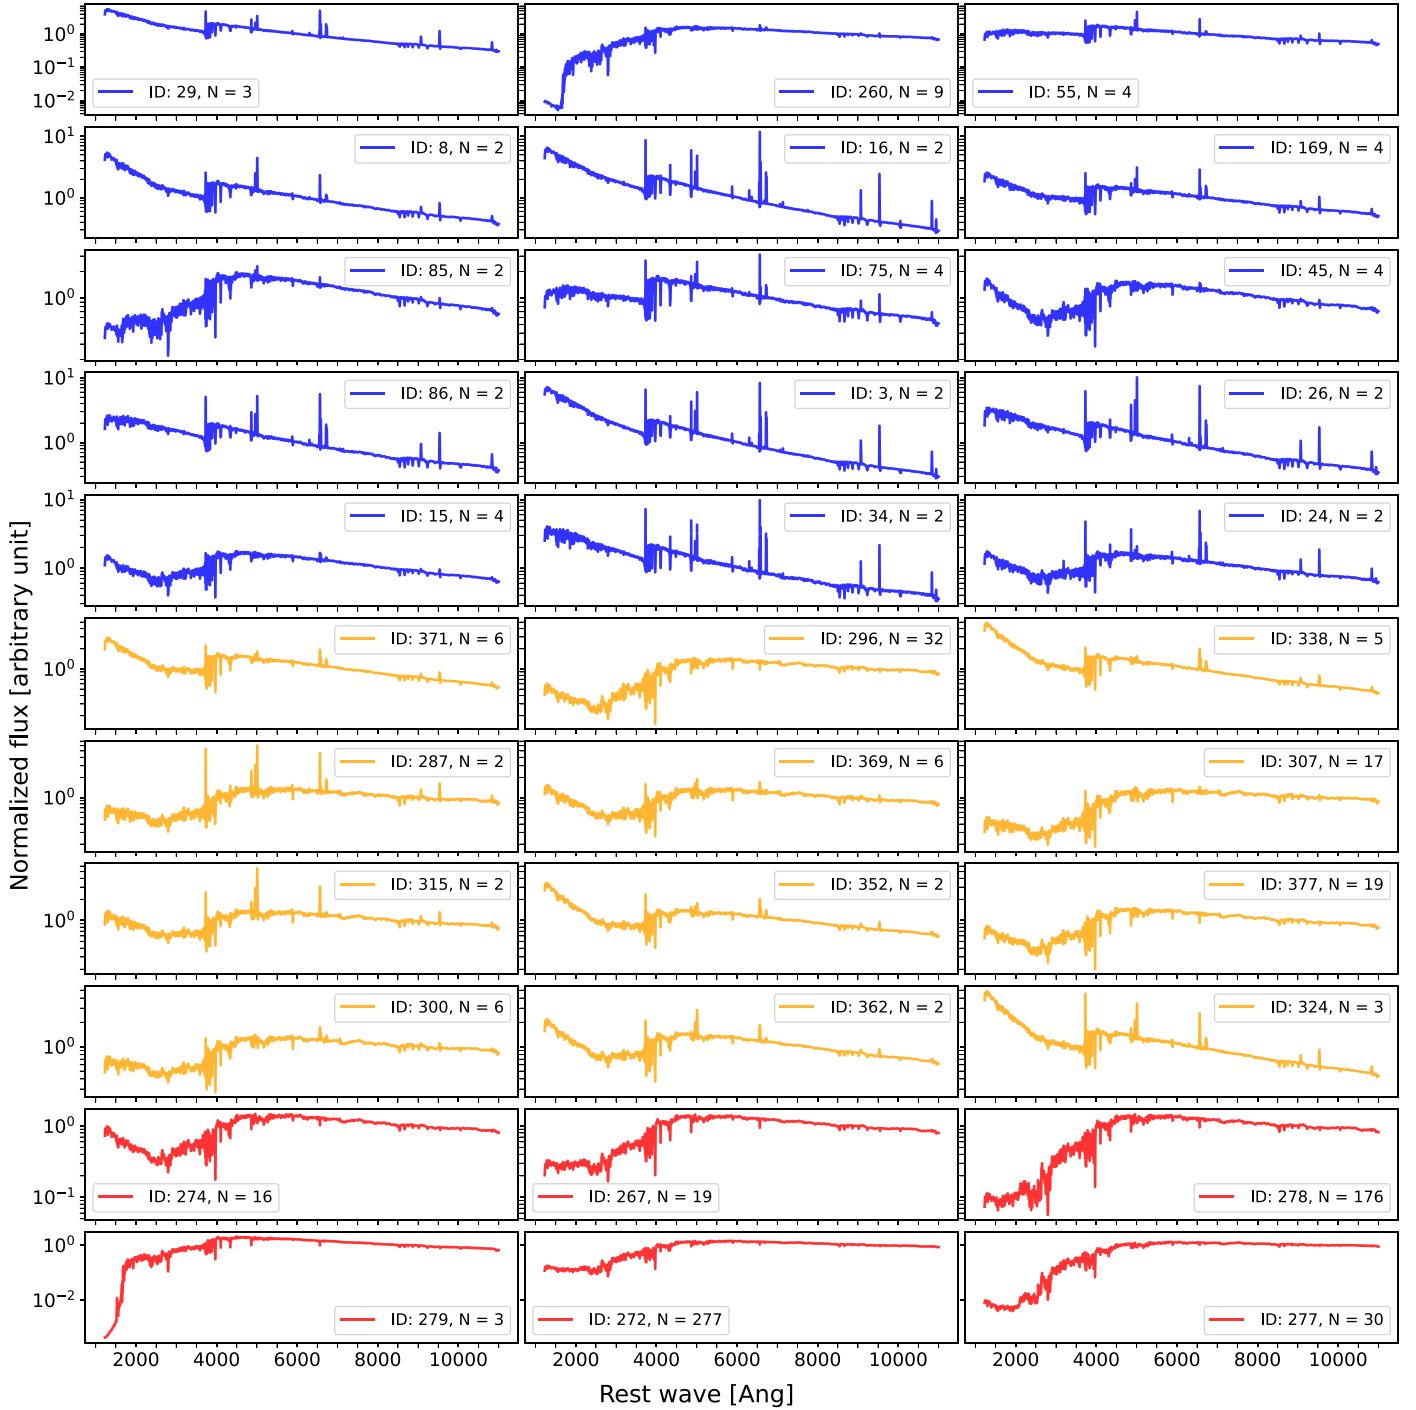
#### Archetypes versus Parent Galaxies Comparison in Color–Color Space

To qualitatively assess the performance of *SetCoverPy* in finding an optimal set of archetypes, we compared their distribution with the parent sample in physical space. We present one such comparison in Figure 13, where we show the distribution of parent (left) and archetype (right) galaxies in  $r - z$  versus  $g - r$  color–color space. The blue, orange, and red colors represent the ELGs, BGS, and LRGs, respectively. In the parent sample, we see that ELGs and LRGs occupy very different regions in the color–color space. A similar distribution is visible in archetypes, though we see far fewer LRGs in the archetype, because (1) their spectra usually do not show emission-line features, and (2) continuum shape does not vary

much. In fact, we see that there are no LRG archetypes with  $g - r > 2$ ; this is evident when we show below the spectra of a few LRG archetypes that span most of the parent LRGs. At the same time, the BGS sample lies somewhat in the intermediate region in the color–color diagram. They are low-redshift, very bright galaxies and can range from highly star-forming to passive galaxies. This observation underscores the reliability of *SetCoverPy* in finding an optimally representative subset of galaxies. As pointed out in Zhu (2016), since the fluxes are normalized, stellar mass versus star formation rate may likely not be the most informative dimension to compare parent and archetype galaxy samples. Subsequently, we select 33 random galaxies from our archetypes (shown with open circles in Figure 3) and show their spectra in Figure 14. One can see the diverse spectral features in ELGs and BGS, while LRGs do not show such diversity. Most of the LRGs are spanned by the three archetypes (ID: 272, 278, and 267). In each subpanel, we also show the number of spectra that the particular archetype can represent in the parent sample. We can see that a few LRGs can span almost all of the LRGs in the parent sample.



**Figure 13.**  $g - r$  vs.  $r - z$  color–color space comparison between parent (left) and archetype (right) galaxies. The archetypes obtained with *SetCoverPy* span the properties of the parent sample well in this space. The colors are not corrected for interstellar extinction.



**Figure 14.** The rest-frame spectra of archetypes (shown as open circles in Figure 3) that can represent more sources than themselves in the parent galaxy sample. The number  $N$  shows how many galaxies in the parent sample the individual archetype can represent, i.e., with distances shorter than the minimum distances defined by the  $s^2$  threshold in the *SetCoverPy* method. We can see a larger diversity in ELGs than in LRGs.

### Appendix B Optimal Prior Value for Polynomial Coefficients in Archetype Approach

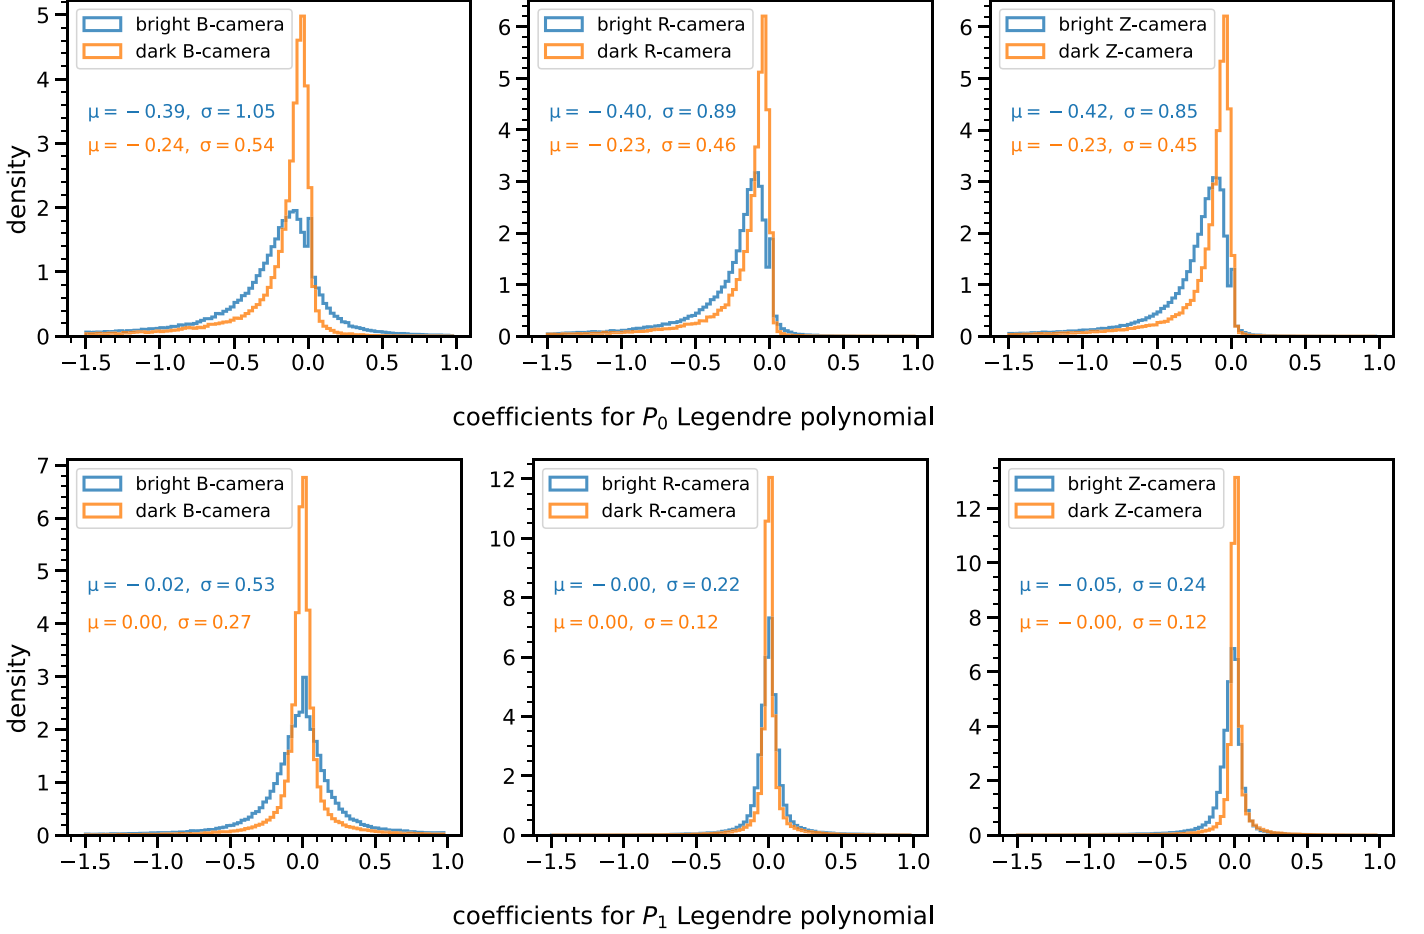
In order to select an optimal value for prior on coefficients of Legendre polynomials, we employ results of fibers targeting the blank sky. These sky spectra are used to construct robust sky models for the spectral reduction process. As illustrated in Section 3.2, we ran our method on numerous tiles within the DESI Y1 data set. This extensive analysis helped us understand the distribution of Legendre polynomials on hundreds of

thousands of sky fibers. This comprehensive analysis enabled us to select an optimal value of the prior value for these polynomial coefficients.

The sky fibers are expected only to have sky emissions and no emission lines associated with actual astronomical objects. Consequently, modeling spectra using Legendre polynomials and archetypes for these fibers should demonstrate minimal variability compared to the diversity observed in the spectra of galaxies. Therefore, the spread in the coefficient distribution allows us to constrain the flexibility of coefficients of the



Year 1 data, OBJTYPE=SKY, SPECTYPE=GALAXY



**Figure 15.** Distribution of Legendre polynomial coefficients (in each camera, the constant term is in the top panel, and the slope term is in bottom panel) for sky fibers (for both dark and bright tiles) from the Y1 data set. The spread of distribution varies between 0.1 and 0.5 for dark tiles and between 0.2 and 1 for bright tiles. Based on this, we choose a prior value of 0.1 (lowest) on polynomial coefficients for our current archetype method. More details are presented in Section 3.2.

Legendre polynomial while fitting real spectra in the archetype method. As demonstrated in Section 3.2, the physical slope of the quasar continuum may occasionally be confused with slope features associated with pipeline defects due to excessive freedom in polynomial coefficients if no prior is used. The prior determined based on sky fibers restricts such freedom, thereby mitigating the misidentification of quasar spectra as galaxies or vice versa within the archetype approach.

We present the distribution of fitted coefficients of Legendre polynomials for these sky fibers in all three cameras in Figure 15. The observed distributions look approximately Gaussian, with a spread between 0.1 and 1. However, notable asymmetry is observed, potentially due to other pipeline defects introduced during the spectral reduction. Therefore, we have opted for a prior smaller ( $\sigma_a = 0.1$ ) than these variances. This choice ensures that the Legendre coefficients take smaller values for the sky fibers, as they are featureless spectra. In fact, we checked the distribution of Legendre coefficients for sky

fibers after applying the priors, and we found that the average value is close to zero while the spread reduces to  $\lesssim 0.05$ . This is evidence that our regularization method (see also Section 3.2) has worked fairly well.

### Appendix C More Performance Metrics for SV Phase Data

The truth catalogs available for the spectra from the SV phase provide us the opportunity to define additional performance metrics such as precision, recall, and F1 scores for the galaxy data sets. We utilize these tables to compare the true redshifts with the redshifts measured by the pipeline. The redshift difference is calculated as  $|\Delta v| = \left| \frac{z_{\text{visual}} - z_{\text{pernight}}}{1 + z_{\text{visual}}} \right| \cdot c$ , where  $z_{\text{visual}}$  and  $z_{\text{pernight}}$  are the visual and measured redshifts of the targets observed on several nights during the SV phase (see Section 4.2). The good redshift criteria is defined in Table 2. We use these two definitions to calculate the precision, recall, and F1 scores for all galaxy targets for both the Redrock

**Table 7**  
Recall and Precision Statistics for Redrock vs. Archetype for Targets Observed during SV Phase

Target	Redrock			Archetype			Archetype ( $\sigma_a = 0.1$ )		
	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score
BGS_FAINT	0.980	0.993	0.987	0.981	0.996	0.988	0.981	0.996	0.988
BGS_BRIGHT	0.977	0.995	0.986	0.977	0.996	0.987	0.977	0.996	0.986
LRG	0.987	0.984	0.986	0.988	0.983	0.986	0.988	0.984	0.986
ELG_LOP	0.991	0.968	0.979	0.989	0.962	0.975	0.989	0.966	0.977
ELG_VLO	0.997	0.988	0.992	0.995	0.987	0.991	0.997	0.987	0.992
ELG_HIP	0.995	0.970	0.982	0.994	0.972	0.983	0.994	0.975	0.984

(without archetypes) and archetype methods. We define the following quantities:

True Positive,  $TP = N(|\Delta v| < 1000 \text{ km s}^{-1}, X > X_0)$

False Positive,  $FP = N(|\Delta v| > 1000 \text{ km s}^{-1}, X > X_0)$ ; i.e., the catastrophic failures

True Negative,  $TN = N(|\Delta v| > 1000 \text{ km s}^{-1}, X < X_0)$

False Negative,  $FN = N(|\Delta v| < 1000 \text{ km s}^{-1}, X < X_0)$
















$$\text{precision} = \frac{TP}{TP + FP}; \quad \text{recall} = \frac{TP}{TP + FN}$$

$$\text{F1 score} = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

where  $X$  and  $X_0$  are defined in Table 2. We calculate these quantities for both methods and compile the values in Table 7 for all galaxy targets. As one of our key objectives is to minimize the number of catastrophic failures (i.e., false positives), precision is more valuable than recall. Our results indicate that the overall precision of both the Redrock method (without archetypes) and the archetype method is quite high and very similar across all galaxy target classes. This finding again demonstrates that Redrock is already highly efficient, while our archetype-based model also achieves comparable high performance for DESI spectra, along with improved modeling of spectral defects.

#### ORCID iDs

Abhijeet Anand  <https://orcid.org/0000-0003-2923-1585>  
 Julien Guy  <https://orcid.org/0000-0001-9822-6793>  
 Stephen Bailey  <https://orcid.org/0000-0003-4162-6619>  
 John Moustakas  <https://orcid.org/0000-0002-2733-4559>  
 S. Ahlen  <https://orcid.org/0000-0001-6098-7247>  
 A. S. Bolton  <https://orcid.org/0000-0002-9836-603X>  
 A. Brodzeller  <https://orcid.org/0000-0002-8934-0954>  
 S. Cole  <https://orcid.org/0000-0002-5954-7903>  
 A. de la Macorra  <https://orcid.org/0000-0002-1769-1640>  
 Biprateep Dey  <https://orcid.org/0000-0002-5665-7912>  
 K. Fanning  <https://orcid.org/0000-0003-2371-3356>  
 J. E. Forero-Romero  <https://orcid.org/0000-0002-2890-3725>  
 S. Gontcho A Gontcho  <https://orcid.org/0000-0003-3142-233X>  
 C. Howlett  <https://orcid.org/0000-0002-1081-9410>  
 D. Kirkby  <https://orcid.org/0000-0002-8828-5463>  
 T. Kisner  <https://orcid.org/0000-0003-3510-7134>  
 A. Kremin  <https://orcid.org/0000-0001-6356-7424>

M. Landriau  <https://orcid.org/0000-0003-1838-8528>  
 L. Le Guillou  <https://orcid.org/0000-0001-7178-8868>  
 M. Manera  <https://orcid.org/0000-0003-4962-8934>  
 A. Meisner  <https://orcid.org/0000-0002-1125-7384>  
 G. Niz  <https://orcid.org/0000-0002-1544-8946>  
 N. Palanque-Delabrouille  <https://orcid.org/0000-0003-3188-784X>  
 W. J. Percival  <https://orcid.org/0000-0002-0644-5727>  
 F. Prada  <https://orcid.org/0000-0001-7145-8674>  
 A. Raichoor  <https://orcid.org/0000-0001-5999-7923>  
 M. Rezaie  <https://orcid.org/0000-0001-5589-7116>  
 E. Sanchez  <https://orcid.org/0000-0002-9646-8198>  
 E. F. Schlafly  <https://orcid.org/0000-0002-3569-7421>  
 G. Tarlé  <https://orcid.org/0000-0003-1704-0781>  
 R. Zhou  <https://orcid.org/0000-0001-5381-4372>  
 H. Zou  <https://orcid.org/0000-0002-6684-3997>

#### References

Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, *ApJS*, 249, 3  
 Allende Prieto, C., Koesterke, L., Hubeny, I., et al. 2018, *A&A*, 618, A25  
 Anand, A., Kauffmann, G., & Nelson, D. 2022, *MNRAS*, 513, 3210  
 Anand, A., Nelson, D., & Kauffmann, G. 2021, *MNRAS*, 504, 65  
 Anand, A., Roy, N., & Gopal-Krishna 2019, *RAA*, 19, 083  
 Bailey, S. 2012, *PASP*, 124, 1015  
 Baldry, I. K., Alpaslan, M., Bauer, A. E., et al. 2014, *MNRAS*, 441, 2440  
 Baron, D. 2019, arXiv:1904.07248  
 Bolton, A. S., Burles, S., Koopmans, L. V. E., et al. 2008, *ApJ*, 682, 964  
 Bolton, A. S., Schlegel, D. J., Aubourg, É., et al. 2012, *AJ*, 144, 144  
 Brodzeller, A., & Dawson, K. 2022, *AJ*, 163, 110  
 Brodzeller, A., Dawson, K., Bailey, S., et al. 2023, *AJ*, 166, 66  
 Busca, N., & Balland, C. 2018, arXiv:1808.09955  
 Caprara, A., Fischetti, M., & Toth, P. 1999, *Oper. Res.*, 47, 730  
 Chaussidon, E., Yèche, C., Palanque-Delabrouille, N., et al. 2023, *ApJ*, 944, 107  
 Chen, Y.-M., Kauffmann, G., Tremonti, C. A., et al. 2012, *MNRAS*, 421, 314

- Cirasuolo, M., Fairley, A., Rees, P., et al. 2020, *Msngr*, **180**, 10
- Conroy, C. 2013, *ARA&A*, **51**, 393
- Cool, R. J., Moustakas, J., Blanton, M. R., et al. 2013, *ApJ*, **767**, 118
- Cooper, A. P., Kposov, S. E., Allende Prieto, C., et al. 2023, *ApJ*, **947**, 37
- Cooper, M. C., Newman, J. A., Croton, D. J., et al. 2006, *MNRAS*, **370**, 198
- Dalton, G., Trager, S., Abrams, D. C., et al. 2014, *Proc. SPIE*, **9147**, 91470L
- Davies, L. J. M., Thorne, J. E., Robotham, A. S. G., et al. 2021, *MNRAS*, **506**, 256
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036
- DESI Collaboration, Abareshi, B., Aguilar, J., et al. 2022, *AJ*, **164**, 207
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024a, *AJ*, **168**, 58
- DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024b, *AJ*, **167**, 62
- Driver, S. P., Davies, L. J., Meyer, M., et al. 2016, in *The Universe of Digital Sky Surveys*, ed. N. R. Napolitano et al. (Springer: Cham), 205
- Fisher, M. L. 2004, *Manage. Sci.*, **50**, 1861
- Geoffrion, A. M. 1974, *Approaches to Integer Programming* (Berlin: Springer), 82
- Green, D., & Bailey, S. 2023, arXiv:2311.04855
- Guy, J., Bailey, S., Kremin, A., et al. 2023, *AJ*, **165**, 144
- Hahn, C., Wilson, M. J., Ruiz-Macias, O., et al. 2023, *AJ*, **165**, 253
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Held, M., & Karp, R. M. 1971, *Math. Program.*, **1**, 6
- Holwerda, B. W., Knabel, S., Thorne, J. E., et al. 2022, *MNRAS*, **510**, 2305
- Holwerda, B. W., Baldry, I. K., Alpaslan, M., et al. 2015, *MNRAS*, **449**, 4277
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Hutchinson, T. A., Bolton, A. S., Dawson, K. S., et al. 2016, *AJ*, **152**, 205
- Jin, S., Trager, S. C., Dalton, G. B., et al. 2024, *MNRAS*, **530**, 2688
- Kauffmann, G., Heckman, T. M., White, S. D. M., et al. 2003, *MNRAS*, **341**, 33
- Kewley, L. J., Nicholls, D. C., & Sutherland, R. S. 2019, *ARA&A*, **57**, 511
- Kochanek, C. S., Eisenstein, D. J., Cool, R. J., et al. 2012, *ApJS*, **200**, 8
- Krolewski, A., Yu, J., Ross, A. J., et al. 2024, arXiv:2405.17208
- Lan, T.-W., Tojeiro, R., Armengaud, E., et al. 2023, *ApJ*, **943**, 68
- Lawson, C. L., & Hanson, R. J. 1995, *Classics in Applied Mathematics* (University City, PA: SIAM)
- Lee, D. D., & Seung, H. S. 1999, *Natur*, **401**, 788
- Levi, M., Bebek, C., Beers, T., et al. 2013, arXiv:1308.0847
- Loveday, J., Norberg, P., Baldry, I. K., et al. 2012, *MNRAS*, **420**, 1239
- Maddox, N., Hess, K. M., Blyth, S. L., & Jarvis, M. J. 2013, *MNRAS*, **433**, 2613
- Mathews, D. J., Newman, J. A., Coil, A. L., Cooper, M. C., & Gwyn, S. D. J. 2013, *ApJS*, **204**, 21
- Miller, T. N., Doel, P., Gutierrez, G., et al. 2024, *AJ*, **168**, 95
- Moustakas, J., 2017 iSEDfit: Bayesian Spectral Energy Distribution Modeling of Galaxies, Astrophysics Source Code Library, ascl:1708.029
- Moustakas, J., Scholte, D., Dey, B., & Khederlarian, A., 2023 FastSpecFit: Fast Spectral Synthesis and Emission-line Fitting of DESI Spectra, Astrophysics Source Code Library, ascl:2308.005
- Moustakas, J., Coil, A. L., Aird, J., et al. 2013, *ApJ*, **767**, 50
- Myers, A. D., Moustakas, J., Bailey, S., et al. 2023, *AJ*, **165**, 50
- Napolitano, L., Pandey, A., Myers, A. D., et al. 2023, *AJ*, **166**, 99
- Netzer, H. 2015, *ARA&A*, **53**, 365
- Newman, J. A., Cooper, M. C., Davis, M., et al. 2013, *ApJS*, **208**, 5
- Raichoor, A., Moustakas, J., Newman, J. A., et al. 2023, *AJ*, **165**, 126
- Ruhlmann-Kleider, V., Yèche, C., Magneville, C., et al. 2024, arXiv:2404.03569
- Schlafly, E. F., Kirkby, D., Schlegel, D. J., et al. 2023, *AJ*, **166**, 259
- Silber, J. H., Fagrelus, P., Fanning, K., et al. 2023, *AJ*, **165**, 9
- Stark, P. B., & Parker, R. L. 1995, *Comput. Stat.*, **10**, 129
- Stasińska, G., & Izotov, Y. 2003, *A&A*, **397**, 71
- Tamura, N., Takato, N., Shimono, A., et al. 2016, *Proc. SPIE*, **9908**, 99081M
- Tsalmantza, P., & Hogg, D. W. 2012, *ApJ*, **753**, 122
- Tumlinson, J., Peebles, M. S., & Werk, J. K. 2017, *ARA&A*, **55**, 389
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Wang, K., Mo, H. J., Li, C., Meng, J., & Chen, Y. 2020, *MNRAS*, **499**, 89
- Yu, J., Ross, A. J., Rocher, A., et al. 2024, arXiv:2405.16657
- Zhou, R., Dey, B., Newman, J. A., et al. 2023, *AJ*, **165**, 58
- Zhu, G. 2016, arXiv:1606.07156
- Zhu, G., & Ménard, B. 2013, *ApJ*, **770**, 130