



HAL
open science

Spectroscopic QUasar extractor and redshift (z) estimator squeeze – II. Universality of the results

Ignasi Pérez-Ràfols, Matthew M. Pieri

► To cite this version:

Ignasi Pérez-Ràfols, Matthew M. Pieri. Spectroscopic QUasar extractor and redshift (z) estimator squeeze – II. Universality of the results. *Monthly Notices of the Royal Astronomical Society*, 2020, 496 (4), pp.4941-4950. 10.1093/mnras/staa1786 . hal-02922228

HAL Id: hal-02922228

<https://hal.science/hal-02922228>

Submitted on 23 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Spectroscopic QUasar extractor and redshift (z) estimator SQUEZE – II. Universality of the results

Ignasi Pérez-Ràfols ^{1,2}★ and Matthew M. Pieri¹

¹CNRS, CNES, LAM, Aix Marseille Univ, F-13388 Marseille cedex 13, France

²CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies, LPNHE, Sorbonne Université, Université Paris Diderot, 4 Place Jussieu, F-75252 Paris, France

Accepted 2020 June 17. Received 2020 June 12; in original form 2019 November 1

ABSTRACT

In this paper, we study the universality of the results of SQUEZE, a software package to classify quasar spectra and estimate their redshifts. The code is presented in Pérez-Ràfols et al. (2019). We test the results against changes in signal to noise, spectral resolution, wavelength coverage, and quasar brightness. We find that SQUEZE levels of performance (quantified with purity and completeness) are stable to spectra that have a noise dispersion four times that of our standard test sample, Baryon Oscillation Spectroscopic Survey (BOSS). We also find that the performance remains unchanged if pixels of width 25 Å are considered, and decreases by ~ 2 per cent for pixels of width 100 Å. We see no effect when analysing subsets of different quasar brightness, and we establish that the blue part (up to 7000 Å) of the spectra is sufficient for the classification. Finally, we compare our suite of tests with samples of spectra expected from WEAVE-QSO and DESI, and narrow-band imaging from J-PAS. We conclude that SQUEZE will perform similarly when confronted with the demands of these future surveys as when applied to current BOSS data.

Key words: quasars: absorption lines – quasars: emission lines – cosmology: observations.

1 INTRODUCTION

Current spectroscopic surveys are generating hundreds of thousands of spectra of objects targeted as quasars. These spectra need to be inspected to determine whether the observed object is indeed a quasar, and to estimate their redshift when that is the case. One of such surveys is the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013, part of the Sloan Digital Sky Survey – III survey Eisenstein et al. 2011), with 546 856 spectra targeted as quasars that were visually inspected (Pâris et al. 2017). In light of next generation of surveys such as WEAVE-QSO (Pieri et al. 2016; as part of WEAVE, Dalton et al. 2016), DESI (DESI Collaboration 2016), *Euclid* (Laureijs et al. 2010), and J-PAS¹ (Benitez et al. 2014), it has become clear that these surveys will provide potential quasar (pseudo-)spectra in numbers too large for a survey-wide visual inspection to be viable.

In Pérez-Ràfols et al. (2019), hereafter [Paper I](#), we presented a code, SQUEZE,² to automatically inspect all these spectra. SQUEZE works by measuring the presence and relative strength of potential quasar emission lines in coarse bands of spectrum using a series of metrics. The performance of the algorithm is tested using BOSS data and yields a purity of 97.40 ± 0.47 per cent (99.59 ± 0.06 per cent for quasars with $z \geq 2.1$) and a completeness of 97.46 ± 0.33 per cent (98.81 ± 0.13 per cent for quasars with $z \geq 2.1$) when a confidence threshold of $p_{\min} = 0.32$ is used. Here (and throughout this paper),

purity is defined as the number of quasars that are correctly classified divided by the total number of objects called quasars, and completeness as the number of quasars correctly classified divided by the total number of quasars that were fed to SQUEZE. Note that for a quasar to be correctly classified we require its redshift to be within 0.15 of the ‘true’ redshift (defined currently as that which is given by visual inspection). Other codes, such as QUASARNET (Busca & Balland 2018) or RedRock (an unpublished DESI code that develops the methods used by the BOSS pipeline; see Bolton et al. 2012), appear to have roughly similar performance to SQUEZE (a detailed comparison between the performance of SQUEZE and these codes will be given in follow-up papers in the series), but access the entire spectra. The critical decision-making algorithms of SQUEZE, on the other hand, use only high-level metrics. The difference is noteworthy because all current training and performance tests have made based on BOSS data (the only large enough visually inspected data set available), while these codes will be applied to *other* surveys with different resolution, noise properties, and pipeline reduction software. The minimal direct access of SQUEZE to the spectra makes it more resistant to these potential peculiarities.

SQUEZE sets only minimal requirements on the data: a spectral resolution sufficient to marginally resolve quasar emission lines and a wavelength range wide enough to access a small number of detectable emission lines simultaneously at the redshifts of interest. This versatility is a key element of SQUEZE, as the main goal of these classifiers is to operate on data from new surveys, rather than from the already complete BOSS survey. We test the limits of this versatility with low signal to noise (increasing the noise dispersion by a factor of up to 4), low resolution (down to pixels of 100 Å; in width), and changes in quasar properties with brightness.

* E-mail: iperezra@lpnhe.in2p3.fr

¹J-PAS is an imaging survey survey with 56 narrow-band filters as can be treated as effectively providing ‘pseudo-spectra’.

²Publicly available at <https://github.com/iprafols/SQUEZE>.

Our tests are not arbitrarily chosen. They refer in broad terms to the next-generation surveys stated above. All these quasar samples are dominated by their faint limits and identification is limited by the success at these faint limits. Both WEAVE-QSO and DESI will acquire spectra of $z > 2.1$ quasars with signal to noise $\gtrsim 0.4$ at full depth, and this is the approximate minimum signal to noise of BOSS quasar spectra. A further challenge is presented by DESI since it must identify and perform a redshift estimate on essentially all quasars targeted with an exposure time four times shorter, and this corresponds to our four times noise test limit. Finally, J-PAS will acquire narrow-band imaging data of quasars with filter centres separated by 100 \AA ; broadly equivalent 100 \AA ; binning, and will reach depth necessary to identify WEAVE-QSO targets ($r < 23.2$), which corresponds to our four times noise test limit. Finally, some target selection from J-PAS data will be acquired at times by only partial wavelength coverage, addressed by our other tests.

We therefore estimate the performance of the algorithm when applied to data sets with different properties, and determine the conditions where the trained algorithm is no longer applicable. Using only high-level metrics, SQUEZE is resistant to the changes mentioned above. The object of this paper is to quantify the extent to which these changes affect the performance of SQUEZE. We will first address how performance is affected by changes related to the usage of different instrumentation: changes in the signal to noise, changes in spectral binning (and effectively resolution), and changes in wavelength coverage. We also explore the impact on performance when studying samples of fainter quasars. Finally, we explore the challenge of acquiring faint quasar samples from J-PAS data by combining both increased noise and decreased resolution.

As in Paper I, in this paper we assume that the visual inspection is always correct. Based on a preliminary visual re-inspection of a subset of the apparent failures, we believe that the visual inspection catalogue suffers from low levels (~ 1 per cent) of impurity and incompleteness. Therefore, a more detailed study of the validity of this statement will be addressed in a follow-up paper on the series. The data sets used here are the same used in Paper I and are described in section 3 of Paper I, but we summarize its properties here. We use eight independent pairs of training and validation samples. Each of the samples consists of 64 plates of BOSS (~ 6800 quasars and $\sim 11\,520$ contaminants), and all 16 samples are independent of each other. The spectra in these samples are modified for the different tests as explained in the corresponding sections.

This paper is organized as follows. We start by giving an overview of SQUEZE in Section 2. Then, we move to testing changes in signal to noise in Section 3, changes in spectral binning in Section 4 and changes in wavelength coverage in Section 5. Tests on the effect in changes on quasar properties are performed in Section 6. Finally, we discuss the performance of SQUEZE on upcoming surveys in Section 7 and summarize our results in Section 8.

2 SQUEZE OVERVIEW

In this section, we review the SQUEZE algorithm. The purpose of this section is not to provide a detailed description of the code (which is done in Paper I), but to summarize its main features, and introduce the key concepts used throughout the paper.

To reproduce the essential elements of visual inspection, SQUEZE follows a three-step procedure. In the first step, the peak finder locates all significant peaks on a given spectrum. Each of these peaks is assigned a trial emission line classification (see Table 1). Every trial classification for every peak, implies a trial redshift, z_{try} , and an

Table 1. Line’s name and nominal wavelength for the selected lines. All wavelengths are given in \AA ; at the restframe.

Line	Wavelength (\AA)
Ly β	1033.03
Ly α	1215.67
Si IV	1396.76
C IV	1549.06
C III	1908.73
Ne IV	2423.83

accompanying set of high-level metrics are computed. We refer the reader to Paper I for a detailed definition of these metrics. Note that spectra without any significant peaks are rejected at this stage and never get processed by the random forest classifier.

After this first step, the list of trial redshifts and accompanying metrics is fed into a random forest classifier that will determine whether the trial quasar and redshift classification does indeed correspond to a quasar at that redshift, and assigns a confidence level for this classification. Note that Ly α forest quasar candidates (i.e. those with $z_{\text{try}} \geq 2.1$) are classified separately from lower redshift candidates. We treat a classification as correct if the redshift is within 0.15 of the true redshift, z_{true} (given by the visual inspection).

SQUEZE third and final step consists of choosing for each spectrum, the preferred classification and constructing the final catalogue. To construct the final catalogue, a confidence threshold, p_{min} , is set. Objects with lower confidence values will not be included in the catalogue. Naturally, purity and completeness are affected by the choice of p_{min} . Higher values of p_{min} correspond to samples that are purer and less complete. We follow the standard we set in Paper I and focus on the case where the confidence value gives approximately equal purity and completeness for the entire quasar sample (hence the $z > 2.1$ subset is not deliberately balanced here). Other choices are possible based on the needs of the user.

At this point, we want to introduce the concept of line confusion (which will be relevant in Section 4). This phenomenon occurs when the wrong emission line is assigned to a peak. In these cases, the trial redshift error will be too large, and we expect it not to be classified as a quasar. However, because the wavelength ratio of lines is sometimes similar, some trial redshifts originating from assigning the wrong emission line are still classified as quasars.

3 PERFORMANCE VERSUS SIGNAL TO NOISE

As stated above, we want to estimate the performance of SQUEZE when applied to data sets of different properties. The first property we change is the signal to noise. We modify the signal to noise in BOSS data by adding Gaussian noise to the spectra. The noise is added at the pixel level as

$$f'_i = f_i + (N_{\text{noise}} - 1) \sigma_i G(0, 1), \quad (1)$$

$$\sigma'_i = \sigma_i \sqrt{N_{\text{noise}}}, \quad (2)$$

where f_i and σ_i are the flux and the standard deviation, respectively, in a given pixel, f'_i and σ'_i are the modified fluxes and the standard deviations, respectively, and $G(0, 1)$ is a random number drawn from a Gaussian distribution with mean 0 and standard deviation 1.

The case $N_{\text{noise}} = 1$ corresponds to the original data, and we explore the cases $N_{\text{noise}} = 2, 3, 4$, corresponding to having two, three, and four times the original noise. We name these cases noise2, noise3, and noise4. Fig. 1 illustrates the noise addition in a randomly selected

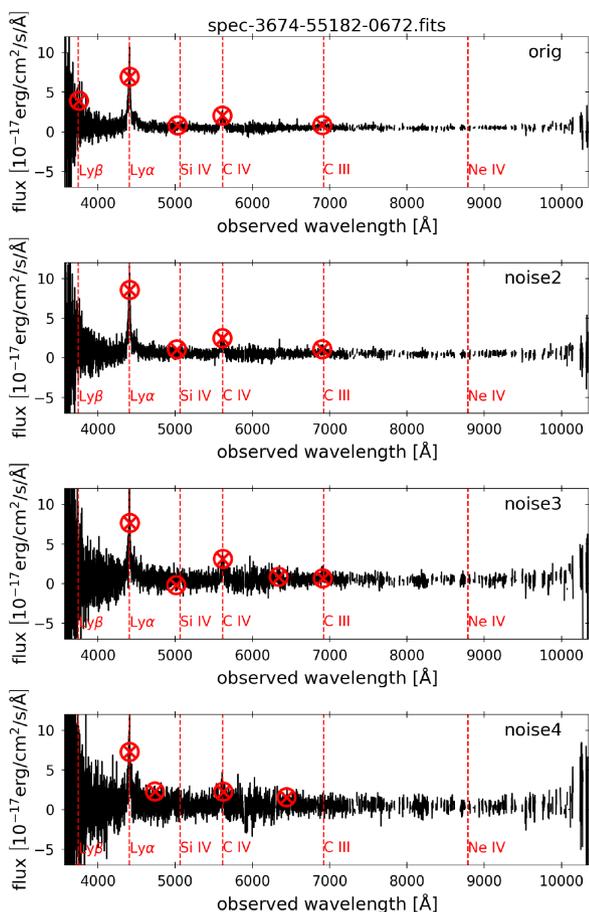


Figure 1. Example of the performance of the peak finder with spectra with different signal to noise. The noise is added to the spectrum following the prescription in equation (1). The cases noise2, noise3, and noise4 correspond to increasing the noise by a factor of 2, 3, and 4 respectively. The dashed lines show the emission lines in Table 1 redshifted to the observed wavelength using the visual inspection reshift of the quasar. The gaps are present due to missing or unreliable data (for example, due to sky lines).

spectrum and the behaviour of the peak finder in each of the cases. We can see that in the original spectrum we successfully identify several emission lines: Ly β , Ly α , Si IV, C IV, and C III (see the rest-frame wavelength for these lines in Table 1. When we increase the noise we maintain the detection of Ly α and C IV, whereas we lose the weaker Ly β and Si IV lines. As noise increases, the peak finder fails to detect the less prominent peaks and begins to select noise spikes. However, since spectra are smoothed before the peak search is performed, some real peaks are still found even as the noise is increased.

We now explore the performance of SQUEZE on these noise-augmented spectra. We ran SQUEZE twice, first retraining the models on the modified data, and then using the models obtained in the original training. Fig. 2 shows the results of this test taking values of p_{\min} such that purity is approximately equal to completeness for the entire sample (also the standard choice in Paper I). In the top panel, using the retrained models, we see that the performance decreases by ~ 1 per cent for case noise2, by $\sim 3 - 4$ per cent for case noise3, and by ~ 7 per cent for case noise4. While the decrease is expected, we note that SQUEZE is obtaining purity and completeness of around 91 per cent even when the noise dispersion is four times that of the original sample. For Ly α quasars (with $z \geq 2.1$) the results are even

more stable: purity decreases by only ~ 1 per cent for case noise4, while completeness decreases by ~ 3 per cent. In the bottom panel, we show the difference in performance when using the original or the retrained models, where we see that there is essentially no need to retrain if small changes on the noise levels are expected. If the sample is significantly noisier, as is the case for case noise4, then retraining the models minimizes losses.

The above assessment of additional noise refers only to the choice of confidence threshold where the purity and completeness are approximately equal. We may broaden the exploration to the different choices of confidence threshold. Fig. 3 shows this focusing on the difference between unmodified noise and our largest noise test, noise4. Again we see that additional noise has a weak effect even when we retain the training model generated from the unmodified data, and that weak effect is made even smaller by retraining.

4 PERFORMANCE VERSUS BINNING/RESOLUTION

Now that we have seen that SQUEZE is insensitive to changes in the signal to noise, we shift our attention to changes in the resolution. Now, we modify the original data by rebinning the data into wider bins. The rebinned flux is computed by averaging fluxes into wider bins, and the rebinned error estimates are computed by standard error propagation. The new bins are created so that there is a bin centred at 4000 Å and have widths 3.125, 6.25, 12.5, 25, 50, and 100 Å. We name these cases rebin3.125, rebin6.25, rebin12.5, rebin25, rebin50, and rebin100, respectively. For reference, the size of the original pixels is ~ 1 Å. Fig. 4 illustrates this rebinning and the behaviour of the peak finder for the same example spectrum as used in Fig. 1. Again, we can see that in the original spectrum five emission line peaks are successfully identified. As the bins become broader, this example shows that the peak finder no longer retains the Ly β line in the rebin100 test as the peak height is diminished through the dilution of local averaging. Also, we see that there is some instability in the peak detection for peak 25, 50, and 100 Å for our example spectrum, but the emission lines in question were not needed for classification. Note that the scale of smoothing performed by the peak finder before looking for peaks is decreased as the pixel size increases (see Table 2).

Once we rebinned the spectra, we ran SQUEZE on them to assess its performance. As before we ran SQUEZE twice, both retraining the model on the modified data, and using the original training. We show the results of this test in Fig. 5, where again we choose values of p_{\min} such that purity and completeness are similar. We see that there is no change in the performance for bins of sizes up to 25 Å. Up to bins of this scale, retraining the model does not improve the performance. The performance starts decreasing for bin size of 50 Å, but is only ~ 2 per cent lower even for bin size of 100 Å. Using the original training has no effect for the case rebin50 but decreases the performance of SQUEZE by an extra ~ 2 per cent for the case rebin100. For Ly α quasars, the decrease in the performance is less prominent: purity is not modified and completeness drops by $\lesssim 1.5$ per cent. We conclude that SQUEZE is resistant to changes in the resolution, provided that the pixel size is smaller than 25 Å. When binning further to 100 Å, retraining based on equivalent data helps alleviate the decrease in performance.

We now analyse in more detail the case rebin100. As already seen in Fig. 5, the performance of SQUEZE is decreased in this case study, and the decrease is significantly worse if the original training is used when the confidence threshold is set such that purity and completeness are similar. In Fig. 6, we explore whether this

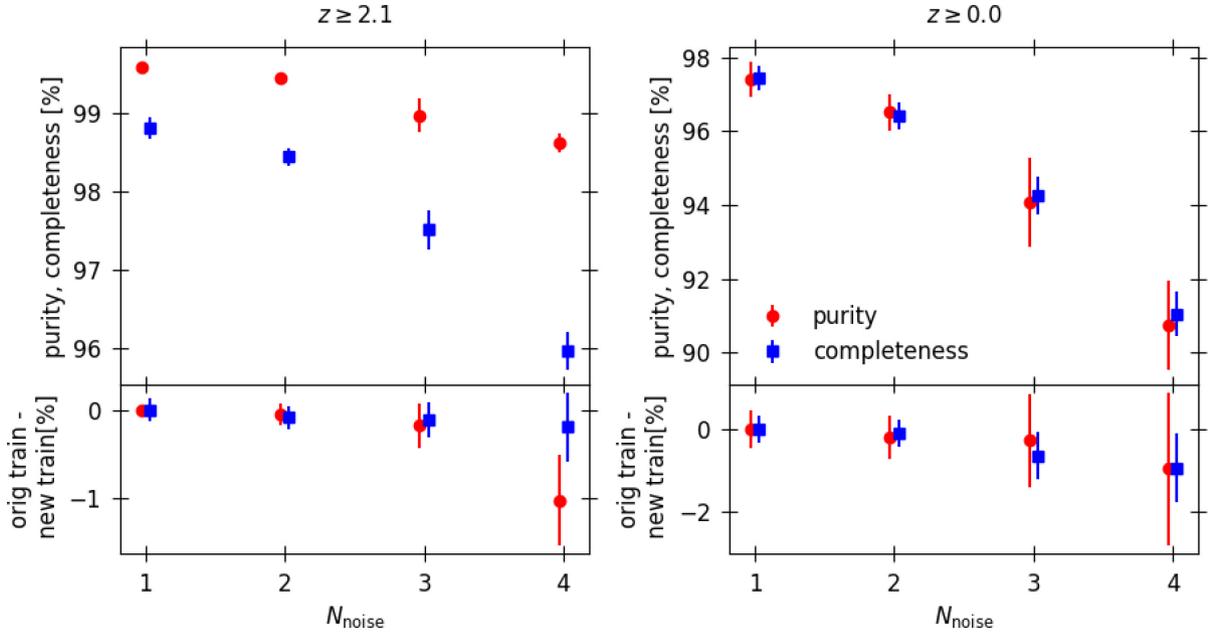


Figure 2. Top panels: Purity and completeness as a function of the noise level (see equation 1). For each sample, a new model is retrained on the modified data, and we select a p_{min} such that purity and completeness are similar for the entire sample and the right-hand panels show the performance for these samples. The left-hand panels show the performance limiting ourselves to quasars with $z \geq 2.1$. The points are horizontally shifted to avoid overlap. Bottom panels: difference in the performance between using the original or the retrained models. Negative (positive) values indicate that the retrained (original) models are better. $N = 1$ correspond to the original data in Paper I.

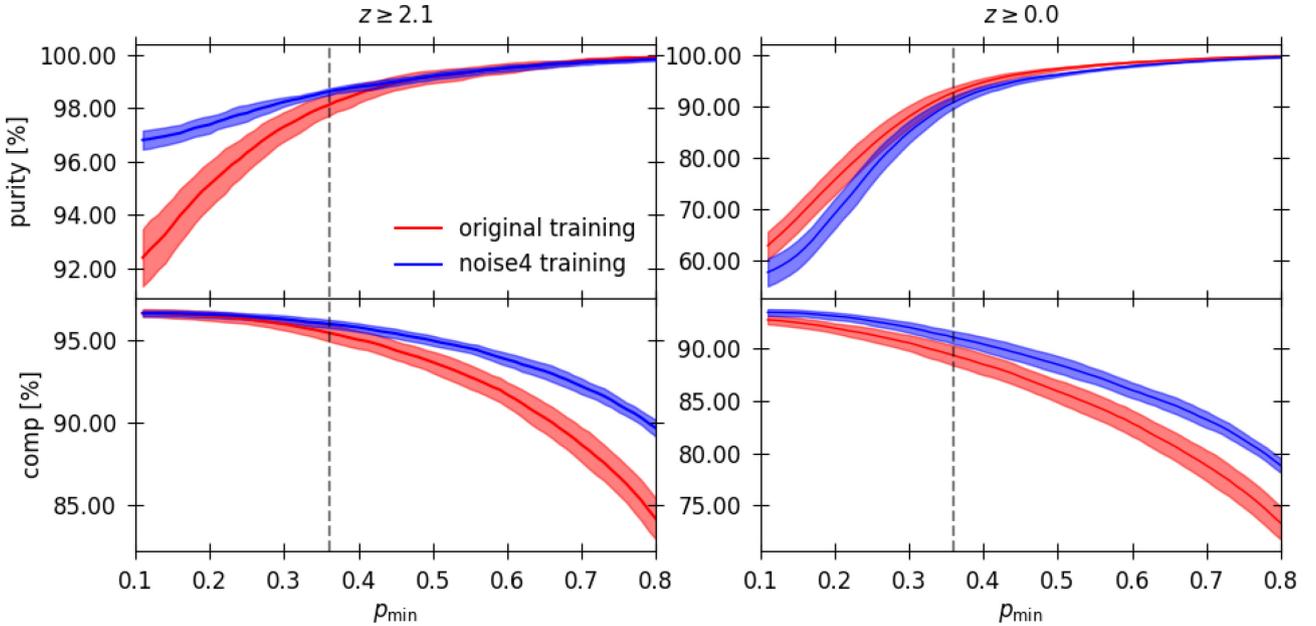


Figure 3. The dependence of the noise4 results using the original or the retrained models for varying probability thresholds. The plot shows the purity (top panels) and the completeness (bottom panels) measured on the test sample with the noise variance increased by a factor of 4. The red and blue lines show the results when the classifier is trained with the original data and with the noise-augmented data, respectively. Left-hand panels show the results for $z_{\text{tr}} \geq 2.1$, whilst right-hand panels show the results at all redshifts. The dashed vertical lines show the p_{min} used in Fig. 2.

statement holds for varying confidence threshold p_{min} . We see that retraining improves the completeness for the all redshifts quasar sample, particularly for high values of p_{min} . We note that purity significantly increases when using the retrained model (with respect to the purity obtained using the original model) for $z \geq 2.1$ at low values of p_{min} .

A possible (naive) explanation for the decrease in the performance of SQUEZE in the case rebin100 would be that the rebinning is moving the centre of the emission peaks, and therefore the estimated quasar redshift. In this scenario, quasar redshift errors would be larger than the tolerance redshift, $\Delta z_r = 0.15$, discarding some quasars that are correctly identified. The validity of this scenario can be

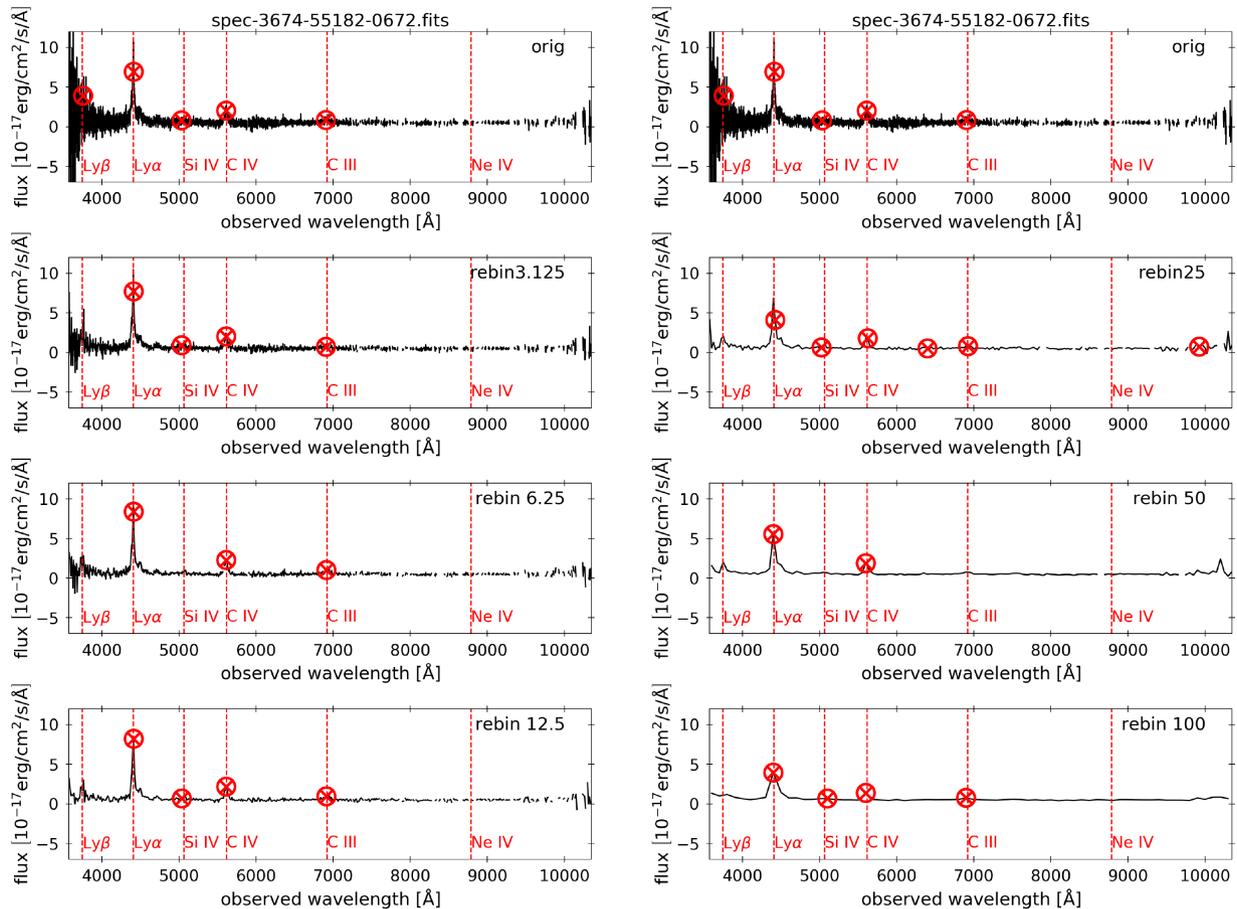


Figure 4. Example of the performance of the peak finder with spectra with different resolution.

Table 2. Bin size and kernel of smoothing applied by the peak finder before looking for peaks for the different case studies involving changes in resolution. New bins are made so that there is a bin centred at 4000 Å.

Case study	Pixel size (Å)	Smoothing (num bins)
Original	~1	70
rebin3.125	3.125	22
rebin6.25	6.25	11
rebin12.5	12.5	6
rebin25	25	3
rebin50	50	1
rebin100	100	0 (no smoothing)

easily seen in a so-called line confusion plot. A line confusion plot shows z_{try} as a function of z_{true} and highlights the presence of line confusion, seen as straight lines trends in the failure cases far from the $z_{\text{try}} = z_{\text{true}}$ line and with slope different from unity. If the naive explanation were correct, we would see quasar contaminants (i.e. actual quasars classified as quasars but not meeting our redshift requirements) very close to the $z_{\text{try}} = z_{\text{true}}$ line in the line confusion plot (see Fig. 7). Since we do not see such an effect, we can rule this out. Thus, it appears that at this level of binning the calculation of the metrics themselves becomes somewhat compromised since the bins widths approach the size of bands used in their calculations.

5 PERFORMANCE WITH LIMITED WAVELENGTH COVERAGE

We now evaluate the performance of SQUEZE as a function of wavelength coverage. We explore this by returning to the original data and removing 1/4 and 1/2 of the total coverage. We analyse five cases: red1, red2, blue1, blue2, and mid. Their wavelength coverage is specified in Table 3.

Fig. 8 summarizes the result of this exercise. SQUEZE performance is strongly reduced for samples mid, red1, and particularly red2. The samples blue1 and blue2, on the other hand, present a much milder decrease in performance. This suggests that the blue end of the spectra are driving the classification. To better understand this behaviour, we analyse the line confusion plots for samples blue1 and red1 (Figs 9 and 10, respectively). Note that p_{min} is allowed to vary between wavelength coverage case studies such that each one balances purity and completeness for the sample of all quasars. We see that the decrease on purity and completeness for sample red1 is mostly due to line confusion (note the multiple lines of quasar contaminants with slope far from unity). Losing the blue part of the spectrum results in SQUEZE misidentify quasar emission lines. This is most likely due to the fact that the blue portion of a quasar spectrum (above the Lyman limit in the quasar restframe) is more crowded with major emission lines (see in fig. 1 of Pérez-Ràfols et al. 2019), which typically fall in the blue portion of the observed wavelength range in our sample. The strongest line is of course Ly α , which motivates our treatment $z > 2.1$ quasar targets as a special case.

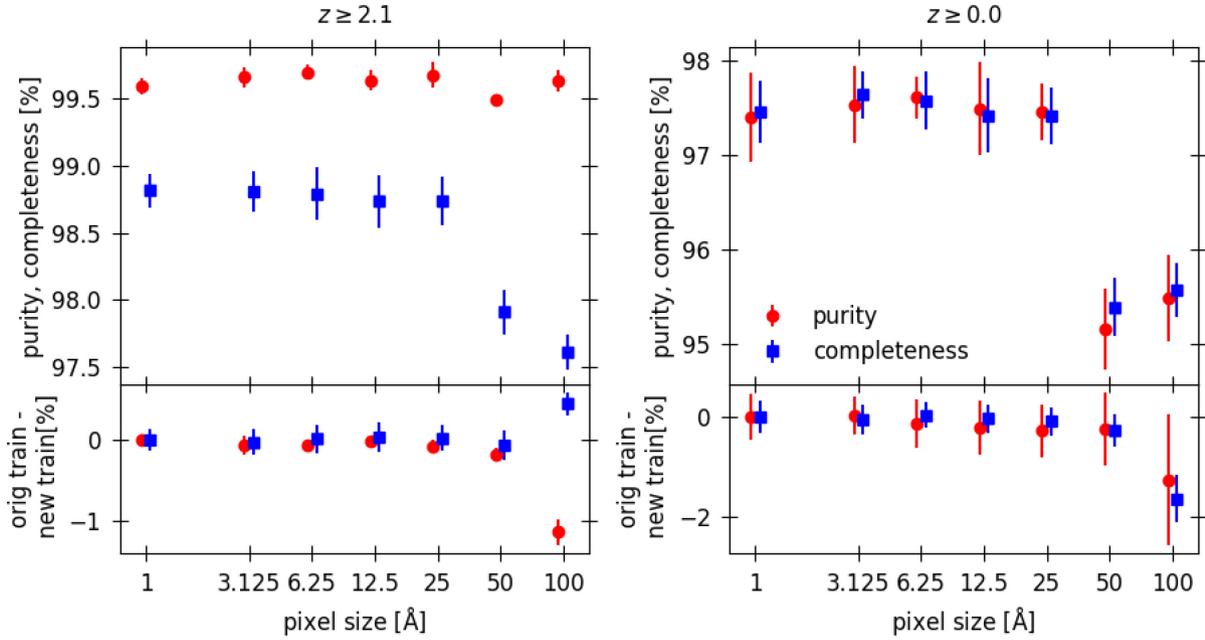


Figure 5. Similar to Fig. 2 but as a function of the resolution. Pixel size of 1 Å; correspond to the original data.

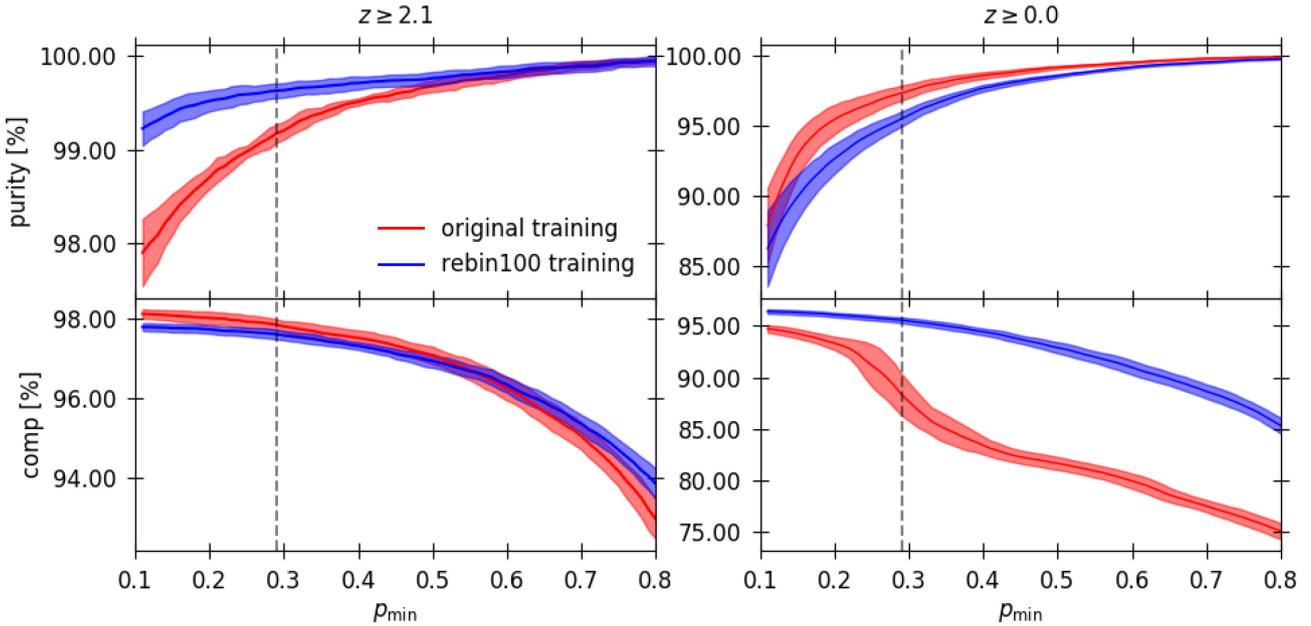


Figure 6. The dependence of the rebin100 results using the original or the retrained models for varying probability thresholds. The plot shows the purity (top panels) and the completeness (bottom panels) measured on the test sample with pixels rebinned to 100 Å-wide pixels. The red and blue lines show the results when the classifier is trained with the original data and with the rebinned data, respectively. Left-hand panels shows the results for $z_{\text{try}} \geq 2.1$, whilst right-hand panels show the results at all redshifts. The dashed vertical lines show the p_{\min} used in Fig. 5.

The lower panels of Fig. 8 shows the performance decrease if the original training from full wavelength is used instead of retraining on the limited coverage. This shows that SQUEZE performs significantly better upon retraining. This change is easily understood by the behaviour of the code. When SQUEZE attempts to compute the metrics of a line that is outside the covered range, then it assigns it a NaN value. This means that the values of some metrics in the restricted wavelength case will have a NaN value, whereas it will have a non-NaN value in the original training data set (without the wavelength coverage restriction). This substantial difference in the values of

metrics with and without retraining has a significant impact. Hence, retraining to deal with markedly limited wavelength coverage is recommended.

6 PERFORMANCE VERSUS QUASAR BRIGHTNESS

Thus far, we have presented the kind of sample differences that might arise for a survey that observes the same quasar target sample as BOSS but with different observatories, instrumentation, and set-ups.

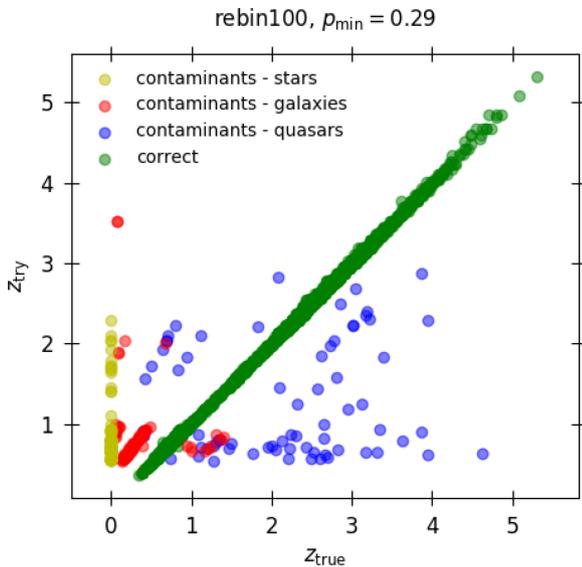


Figure 7. Line confusion plot for the case study rebin100, where the model has been retrained. For all the objects in the catalogue, z_{try} against z_{true} . In this plot, the green circles correspond to correct classifications, the yellow circles to stellar contaminants, the red circles to galactic contaminants, and the blue circles to quasar contaminants (i.e. actual quasars classified as quasars but not meeting our redshift requirements). Note that this plot singles out the false positives.

Table 3. Wavelength coverage of the original data and the five case studies.

Sample	Wavelength coverage (Å)
orig	3600–10 400
red1	5300–10 400
red2	7000–10 400
blue1	3600–8700
blue2	3600–7000
mid	5300–8700

Important performance modifications may also arise due to changes in the quasar target sample. We explore this point here by testing the impact of quasar brightness over the dynamic range provided by the BOSS sample.

One of the major assumptions of SQUEZE is that quasars are self-similar, but in reality this is just an approximation: we know that there is a quasar-to-quasar variation. This variation arises from a number of things: presence of broad absorption lines, damped Ly α absorbers, difference in continuum slope(s) and differences emission line strengths. All this variation exists in our test samples, but if the survey on which SQUEZE is implemented differs significantly in its realization of these various effects, SQUEZE may not perform as we have described. It is clear that one way future surveys will differ from BOSS is in their luminosity distributions, targeting to a fainter limiting magnitude. Going to fainter magnitudes will most likely change the distribution of quasar properties. An example of such a change is the well-known Baldwin effect (Baldwin 1977), which states that the relative strength of the emission lines depends on the luminosity.

While we cannot yet test SQUEZE on significant numbers of quasars fainter than the BOSS limit ($r \lesssim 22.2$), we can test the sensitivity to quasar faintness by taking brightness-dependent sample subsets. To this end, we split the quasar target sample into two bright/faint subsamples of roughly equivalent signal to noise by obtaining

approximately equal values for the ensemble quantity

$$\sum \left(\frac{r_i}{\delta r_i} \right)^2, \quad (3)$$

where r_i and δr_i are the r -band magnitude and its error³ for the i th spectrum of that subsample.

The boundary that conserves r -band signal to noise between bright and faint samples is $r \sim 16$. We apply models trained using the bright samples (bright models) to both the bright and the faint validation samples. Similarly, we apply models trained using the faint samples (faint models) to both bright and faint validation samples.

Results of this exercise are shown in Fig. 11. We observe that SQUEZE performance on $z > 2.1$ quasars is not significantly affected by these sample changes. It is particularly striking that SQUEZE performs well on faint $z > 2.1$ quasars, including quasars as faint as $r \sim 22$, even when it is trained using quasars brighter than $r \sim 16$.

This does not remain the case when all redshifts are included. It would appear that the performance for faint quasars does indeed benefit from retraining on similar quasars but only at the ~ 2 – 3 per cent level. It appears that this is driven by a dependence on the weaker emission lines with wavelengths longwards of Ly α .

7 PERFORMANCE ON UPCOMING SURVEYS

The tests we have presented in the previous sections are presented as survey-neutral stress tests that explore how SQUEZE performs given a variety of differences of observing conditions or quasar population. However, they do have direct relevance for real-world expected use-cases for up-coming surveys. In this section, we place those tests in context of the upcoming surveys where SQUEZE will be used.

In the near future, the two most relevant surveys are WEAVE-QSO and DESI. In both cases, the minimum signal to noise full depth for quasars with $z > 2.1$ will be similar to that of BOSS and therefore the results shown in Paper I are already a reasonable guide to quasar catalogue making performance with regards to noise. However, the quasar populations may themselves differ from that of BOSS since these surveys reach a limiting magnitude $\Delta r \sim 1$ fainter than BOSS. While we do not assess the performance for this fainter population, we do test the importance of quasar brightness. We find that within the BOSS magnitude distribution, SQUEZE performs satisfactorily upon faint $z \geq 2.1$ quasars even when training using bright $z \geq 2.1$ quasars. When including $z < 2.1$ quasars retraining on quasars with representative magnitudes improves performance by ~ 2 – 3 per cent. Both DESI and WEAVE-QSO provide higher spectral resolution than BOSS but our binning tests show that a factor of ~ 2 or ~ 3 in the resolution, respectively, will have no impact on SQUEZE performance.

For both DESI and WEAVE-QSO, significant challenges are faced with regards to building the survey. In the case of DESI, only $z > 2.1$ quasars will be observed to full depth in four layers (or passes) over the survey footprint. Quasars must be efficiently identified (and assigned approximate redshifts) based on a single layer (or a quarter of the exposure time) to either determine whether to continue in subsequent observing layers or to refine the redshift with the data as-is. Single layer data is equivalent to our noise4 case (see tests in Section 3). We find that SQUEZE will be able to identify $z \geq 2.1$ quasars on the first pass to high purity but that completeness may be reduced by a few percent (depending on the choice of probability threshold). Including quasars with $z < 2.1$ presents more

³As reported in the superset catalogue file, available at https://data.sdss.org/sas/dr12/boos/qso/DR12Q/Superset_DR12Q.fits.

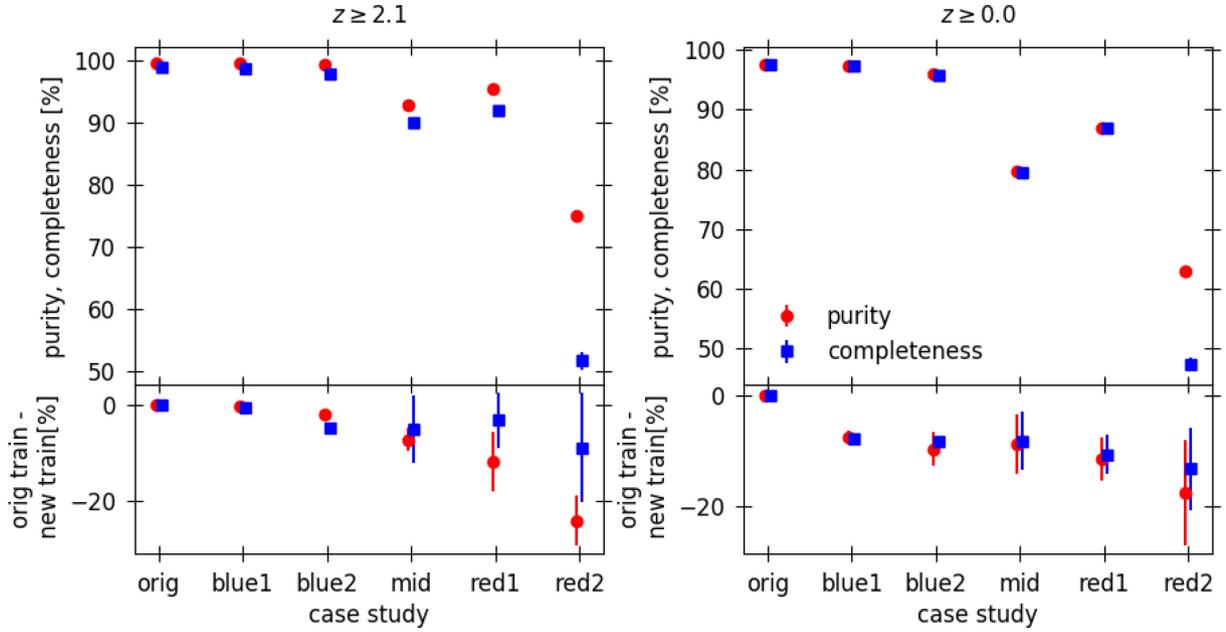


Figure 8. Similar to Fig. 2 but for the case studies with different wavelength coverage (see Table 3 for details).

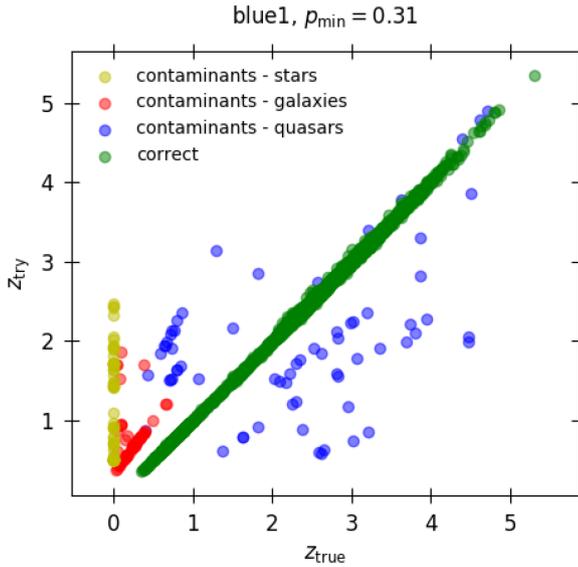


Figure 9. Similar to Fig. 7 but for the case study blue1, where the model has been retrained.

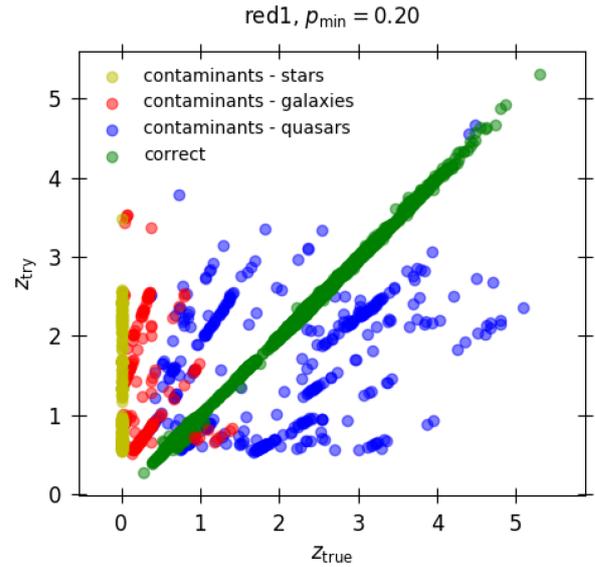


Figure 10. Similar to Fig. 7 but for the case study red1, where the model has been retrained.

of a challenge with a larger impact on both purity and completeness (of approximately 7 per cent) that cannot be recovered trivially by tuning the probability threshold.

In the case of WEAVE-QSO, SQUEZE is also expected to perform well for catalogue production and as a redshift prior for additional redshift refinement step. Applying SQUEZE to the unmodified BOSS sample for $z > 2.1$ we see that, assuming there are no dramatic changes to the quasar or contaminant population, we expect almost no losses or impurities for this task. The main challenge for WEAVE-QSO is at the target selection stage, since *gtr* 90 per cent completeness and purity of $z > 2.1$ quasars is needed from the quasar targets. The WEAVE-QSO plan is to achieve this goal using data from the survey J-PAS. J-PAS is a narrow-band photometric survey with 56 narrow-band filters effectively providing spectra of very

low resolution. These pseudo-spectra have an effective bin width of $\sim 100 \text{ \AA}$. Our results of the impact of binning show SQUEZE performs well even at the level of 100 \AA ; bins (case rebin100). The purity of $z > 2.1$ sample seems to be unaffected and the completeness shows only an additional 1 per cent loss compared to the unmodified sample. The WEAVE-QSO survey is entirely focused on Lyman α forest quasars, but it is noteworthy that the performance of SQUEZE on J-PAS data at all redshifts remains strong.

An additional challenge for WEAVE-QSO is reflected in the tight schedule of J-PAS targeting and WEAVE observing. As a result it is likely that in some instances only some of the filters will be observed in time for propagation to WEAVE-QSO fiber assignment. In this regard, SQUEZE performance tests with limited wavelength coverage could be critical. The success of the blue2 test (using only the blue

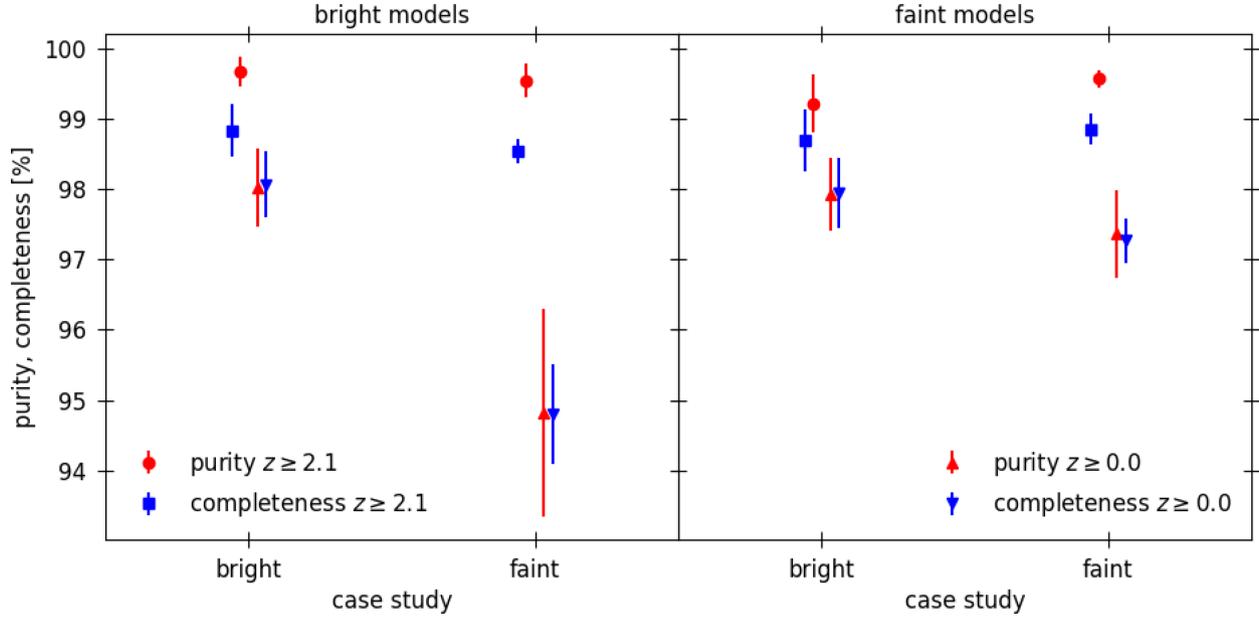


Figure 11. Performance of SQUEZE when applied to samples of different brightness. Left-hand (right-hand) panel shows the result applying the bright (faint) models. See text for details.

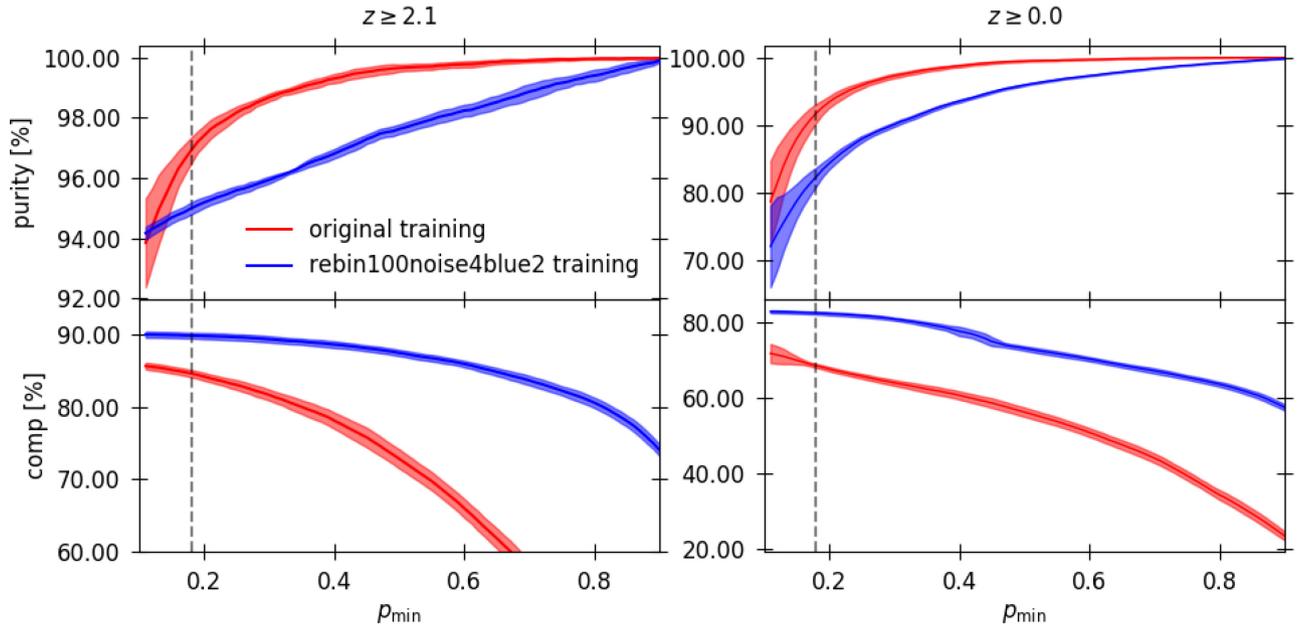


Figure 12. The dependence of the rebin100+noise4+blue2 results using the original or the retrained models for varying probability thresholds. The plot shows the purity (top panels) and the completeness (bottom panels) measured on the test sample with the noise variance increased by a factor of 4, with the pixels rebinned to 100 Å-wide pixels, and the wavelength range limited to only the blue half of the spectra (see Table 3). The red and blue lines show the results when the classifier is trained with the original data and with the rebinned data, respectively. Left-hand panels shows the results for $z_{\text{try}} \geq 2.1$, whilst right-hand panels show the results at all redshifts. The dashed vertical lines show the value of p_{min} for which purity and completeness are equal for the retrained model and at all redshifts.

half of the optical range) indicates a potential way forward for the planning for J-PAS and WEAVE-QSO scheduling using SQUEZE.

This is a simplification, however, since these are isolated tests. True quasar identification in J-PAS will involve a combination of increased noise (since the limiting magnitude required by WEAVE-QSO will be 1 magnitude fainter than BOSS), ~ 100 Å binning), and (at times) limited wavelengths coverage. Realistic J-PAS mocks are needed to fully address this challenge. To take initial steps towards

a more realistic test including all these effects simultaneously, we perform the analysis on a modified sample combining the cases rebin100, noise4, and blue2. Our noise4 modification should produce spectra of roughly the expected signal to noise of J-PAS pseudo-spectra (Carolina Queiroz, private communication). We stress that here we are considering the worst-case scenario in which all the data have limited wavelength coverage. The results of this exercise are given in Fig. 12. We can see that the performance obtained

in this case is lower than the performance obtained in the cases noise4, rebin100, and blue2. Purity and completeness drop below 90 per cent for all redshift, but we recover a 90 per cent completeness and 95 per cent purity for the Ly α quasars. We note that all these changes combine in a non-trivial way. Interestingly, the purity seems to be higher when the original training is used, but we believe this is a direct consequence of the reduced completeness. For example for $p_{\min} = 0.3$ and for Ly α quasars purity is better by ~ 3 per cent, but completeness drops by ~ 10 per cent, and this is even worse for higher values of $p_{\min} = 0.3$. A high purity sample can be achieved using the retrained sample with a sufficiently high choice for p_{\min} with less cost in terms of completeness compared to using the original training. The results presented here suggest that SQUEZE will be able to adapt to this challenge, but we leave such detailed tests for future studies.

Other surveys such as *Euclid* or WEAVE-LOFAR (another survey part of the WEAVE Collaboration) may also benefit from the use of SQUEZE, but their specific requirements have not yet been explored.

8 SUMMARY

In this work, we have explored the universality of the performance of SQUEZE (presented in Paper I), a software packaged designed to identify quasar spectra, among a set of contaminating spectra, and estimate their redshifts. We define this universality two ways. First, we confront SQUEZE with the more straightforward challenge of training on and processing data with different spectroscopic properties (compared to BOSS data in Paper I) and explore how these changes affect performance. Secondly, we assess the universality of SQUEZE combined with the training established in Paper I by confronting BOSS-trained SQUEZE with different data. This latter test allows us to assess the sensitivity to the realism of the training set, and even exploring whether SQUEZE needs retraining on the different data at all.

We have addressed the effect of changing the signal to noise, the spectral binning (and resolution), the wavelength coverage, and the brightness of the sample. After all these tests, we are confident that the performance of SQUEZE is largely survey independent. However, for optimal performance we do recommend altering BOSS spectra to resemble the characteristics of the survey to be analysed and retraining the model on the modified data. If more precise and reliable tests are required, a more realistic contamination sample matched to the details of the given survey must be explored (such a sample could be obtained, for example, in the survey validation phase). We do not recommend retraining using synthetic spectra, unless full modelling of the contaminants is also taken into account.

All the tests performed here assume that the visual inspection is always correct, but this is not necessarily the case. We will explore this in more detail in a subsequent paper of the series. However, we remark that the tests are fair since what we analyse is the difference in the results on the original and modified spectra, and whatever

misclassifications are present in truth table for the original samples are also present in the modified samples.

Finally, we make initial tests of SQUEZE performance on upcoming surveys. We conclude that SQUEZE will perform satisfactorily on spectra from DESI (including single layer spectra from the first pass), and WEAVE-QSO. We also explore the case of J-PAS pseudo-spectra, which will be used for the target selection of WEAVE-QSO, reaching the same conclusion. Other surveys such as *Euclid* or WEAVE-LOFAR may also benefit from the use of SQUEZE, but their specific requirements have not yet been explored.

ACKNOWLEDGEMENTS

We thank Carolina Queiroz for very helpful discussions on the expected quality of J-PAS data.

This work was partly supported by the A*MIDEX project (ANR-11-IDEX-0001-02) funded by the ‘Investissements d’Avenir’ French Government program, managed by the French National Research Agency (ANR), and by ANR under contract ANR-14-ACHN-0021.

DATA AVAILABILITY

The data underlying this article is available in SDSS repository at https://data.sdss.org/sas/dr12/boos/spectro/redux/v5_7_0/spectra/ite/. The code and the standard training originally presented in Paper I is publicly available at <https://github.com/iprafols/SQUEZE>.

REFERENCES

- Baldwin J. A., 1977, *ApJ*, 214, 679
 Benitez N. et al., 2014, preprint (arXiv:1403.5237)
 Bolton A. S. et al., 2012, *AJ*, 144, 144
 Busca N., Balland C., 2018, preprint (arXiv:1808.09955)
 Dalton G. et al., 2016, in Evans J., Simard L., Takami H., eds, in Proc. SPIE Conf. Ser. Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI. SPIE, Bellingham, p. 99081G
 Dawson K. S. et al., 2013, *AJ*, 145, 10
 DESI Collaboration, 2016, preprint (arXiv:1611.00036)
 Eisenstein D. J. et al., 2011, *AJ*, 142, 72
 Laureijs R. J. et al., 2010, in Oschmann J. M. Jr., Clampin M. C., MacEwen H. A., eds, Proc. SPIE, Conf. Ser. Vol. 7731, Space Telescopes and Instrumentation 2010: Optical, Infrared, and Millimeter Wave. SPIE, Bellingham, p. 77311H
 Pâris I. et al., 2017, *A&A*, 597, A79
 Pérez-Ràfols I., Pieri M. M., Blomqvist M., Morrison S., Som D., 2019, *MNRAS*, 496, 4931 (Paper I)
 Pieri M. M., 2016, in Reylé C., Richard J., Cambréys L., Deleuil M., Pécontal E., Tresse L., Vauglin I., eds, SF2A-2016: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics. p. 259

This paper has been typeset from a \LaTeX file prepared by the author.