Identifying galaxies, quasars, and stars with machine learning: A new catalogue of classifications for 111 million SDSS sources without spectra

A. O. Clarke, A. M. M. Scaife, R. Greenhalgh, and V. Griguta

Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester M13 9PL, UK
e-mail: alex.clarke-3@manchester.ac.uk

Received 24 September 2019 / Accepted 28 April 2020

ABSTRACT

We used 3.1 million spectroscopically labelled sources from the Sloan Digital Sky Survey (SDSS) to train an optimised random forest classifier using photometry from the SDSS and the Widefield Infrared Survey Explorer. We applied this machine learning model to 111 million previously unlabelled sources from the SDSS photometric catalogue which did not have existing spectroscopic observations. Our new catalogue contains 50.4 million galaxies, 2.1 million quasars, and 58.8 million stars. We provide individual classification probabilities for each source, with 6.7 million galaxies (13%), 0.33 million quasars (15%), and 41.3 million stars (70%) having classification probabilities greater than 0.99; and 35.1 million galaxies (70%), 0.72 million quasars (34%), and 54.7 million stars (93%) having classification probabilities greater than 0.9. Precision, Recall, and F_1 score were determined as a function of selected features and magnitude error. We investigate the effect of class imbalance on our machine learning model and discuss the implications of transfer learning for populations of sources at fainter magnitudes than the training set. We used a non-linear dimension reduction technique, Uniform Manifold Approximation and Projection, in unsupervised, semi-supervised, and fully-supervised schemes to visualise the separation of galaxies, quasars, and stars in a two-dimensional space. When applying this algorithm to the 111 million sources without spectra, it is in strong agreement with the class labels applied by our random forest model.

Key words. galaxies: statistics – quasars: general – stars: statistics – catalogs – methods: statistical – surveys

1. Introduction

The classification scheme of galaxies, quasars, and stars is one of the most fundamental in astronomy. The early cataloguing of stars and their distribution in the sky has led to the understanding that they make up our own galaxy (Herschel 1789) and, following the distinction that Andromeda was a separate galaxy to our own (Opik 1922; Hubble 1929), numerous galaxies began to be surveyed as more powerful telescopes were built. The designation of quasars arose after radio emission was detected from unresolved star-like sources with high redshifts (e.g. 3C48 and 3C273; Smith & Hoeffleit 1961; Greenstein & Matthews 1963; Matthews & Sandage 1963; Schmidt 1963; Greenstein & Schmidt 1964). This emission was later demonstrated to have been produced by accretion disks surrounding super-massive black holes at the centre of some galaxies (Burbidge et al. 1963; Rees 1984; Begelman et al. 1984). For quasars, the emission from this central region on scales less than a light year, known as an active galactic nucleus (AGN; Urry & Padovani 1995), is expected to dominate over the light from the host galaxy. Large samples of quasars (e.g. Pâris et al. 2018) are now routinely selected through the identification of characteristic high-ionisation emission lines in their optical spectrum (e.g. C IV, Mg II; Francis et al. 1991; Vanden Berk et al. 2001), as well as via spectroscopic follow-up of optical sources where a radio counterpart is also present (Gürkan et al. 2019). Whilst quasars are typically located at redshifts high enough to be unresolved with optical telescopes, some nearby resolved galaxies (Seyfert galaxies;

Weedman 1977; Antonucci 2012) are also labelled as quasars, having bright and compact cores with associated emission lines from AGN, although they are comparatively less bright than high-redshift quasars.

Providing classification labels for astronomical catalogues containing large numbers of sources has a wide range of benefits, both for studies of individual systems and for statistical population analyses. In particular, a significant range of science goals are dependent on large samples of quasars, which are still the minority class. This consideration has been important in motivating the construction of new facilities, such as the Square Kilometre Array (SKA; Jarvis et al. 2015) and the Large Synoptic Survey Telescope (LSST; LSST Science Collaboration 2009; LSST Dark Energy Science Collaboration 2012; Ivezić et al. 2019). Science objectives reliant on quasar samples include Lyman- α forest surveys (Rauch 1998; McDonald & Eisenstein 2007), cosmic magnetism studies (Scranton et al. 2005), general cosmology (Leistedt & Peiris 2014; Hutsemékers et al. 2005), and the evolution of galaxies (Schmidt & Green 1983; Sanders et al. 1988; Kauffmann & Haehnelt 2000), amongst others.

Millions of sources have already been catalogued from telescopes such as the Sloan Digital Sky Survey (SDSS; Aguado et al. 2019), the Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010) and the LOw Frequency ARray (LOFAR; van Haarlem et al. 2013; Shimwell et al. 2019), amongst others. The next generation of telescopes are predicted to significantly increase the size of source catalogues. The LSST is expected to catalogue approximately 20 billion galaxies and a

similar number of stars (Ivezić et al. 2019). Source count predictions for the SKA indicate a source density of around ten galaxies per square arcminute for phase one of the array, and up to 75 galaxies per square arcminute for phase two (assuming a detection threshold of $S_{1\text{GHz}} \geq 100$ nJy). This results in totals of 1 and 8 billion sources, respectively, for a survey area of 3π steradians (Jarvis et al. 2015). Consequently, it is becoming unfeasible for astronomers to manually verify and label individual sources and, whilst efforts such as Galaxy Zoo (Lintott et al. 2008, 2011) bring in many more people to help sift through data, this effort alone is not expected to be able to keep up with the source counts anticipated for the next generation of telescopes. For such large datasets, machine learning algorithms are becoming an increasingly valuable tool for analysis and data exploration. The development of such algorithms in computer science fields has accelerated rapidly in the last decade, focusing on processing large datasets in high performance computing workflows and cloud computing systems (Jones 2014; Wu et al. 2016).

Although distinguishing astronomical source type is normally straightforward where detailed data are available, such as spectroscopy and multi-wavelength observations, the complexity of obtaining detailed observations for millions of individual sources is time consuming and generally impractical for the largest samples of sources given the survey speeds of the current generation of telescopes. In contrast, classifying sources using only photometry in multiple wavebands, and labelling them based on their colours (differences between pairs of photometry measurements), is comparatively fast. With three orders of magnitude less wavelength coverage than spectroscopy, photometry cannot capture the same detail as spectra, however it can capture the overall shape of the spectrum that distinguishes different types of sources. In the optical, stars show a black-body spectrum, galaxies show the superposition of many black-body spectra, and quasars show a comparatively flatter spectrum due to AGN emission. The most widely used colours for source classification are $u - g$ and $g - r$ in SDSS data, and $w1 - w2$ and $w2 - w3$ in WISE data (Nikutta et al. 2014; Peters et al. 2015). Consequently, photometry data have been demonstrated to be useful as machine learning features in source-type classification by a number of studies (e.g. Carrasco et al. 2015; Schindler et al. 2019; Kang et al. 2019; Nakoneczny et al. 2019; Bai et al. 2019). Furthermore, testing whether a source is resolved or unresolved can help distinguish the extended profiles of galaxies from stars and quasars (Aguado et al. 2019; Baldry et al. 2010; Morice-Atkinson et al. 2018) and serve as a useful machine learning feature.

In this work, we apply supervised, unsupervised and semi-supervised machine learning algorithms to classify galaxies, quasars, and stars using SDSS and WISE data. In Sect. 2, we introduce the SDSS and WISE data and describe our feature set. In Sect. 3, we introduce the machine learning algorithms, describe how the models were optimised using a spectroscopically selected training dataset, and give an in-depth evaluation of the model performance as a function of source magnitude and classification. In Sect. 4, we apply our optimised model to an SDSS photometric catalogue of 111 million previously unclassified sources. In Sect. 5, we discuss the results, and in Sect. 6 we draw our conclusions. All of the code used in this paper to gather, process, and analyse the data is available on our Github repository¹.

2. Data

Two datasets were constructed from the SDSS Data Release 15 (DR15 Aguado et al. 2019) catalogue. These comprise (i) a labelled dataset of spectroscopically observed sources and (ii) an unlabelled dataset of photometrically observed sources with no associated spectroscopic observation. In both cases we selected sources that have WISE counterparts.

For SDSS, photometric measurements are provided in five optical bands: u ($\lambda = 0.355 \mu\text{m}$), g ($\lambda = 0.477 \mu\text{m}$), r ($\lambda = 0.623 \mu\text{m}$), i ($\lambda = 0.762 \mu\text{m}$) and z ($\lambda = 0.913 \mu\text{m}$), with associated errors. WISE provides photometric measurements in four infrared bands: $w1$ ($\lambda = 3.4 \mu\text{m}$), $w2$ ($\lambda = 4.6 \mu\text{m}$), $w3$ ($\lambda = 12 \mu\text{m}$) and $w4$ ($\lambda = 22 \mu\text{m}$) with their associated errors. SDSS photometric measurements, as described in Stoughton et al. (2002), are optimised for different types of source. For unresolved point sources (e.g. stars and quasars) that are well-described by the point spread function (PSF), the best measure of the total flux is determined by fitting a PSF model to the source, referred to as *psfMag*. However, for resolved sources such as galaxies this is not the case, and a better measure of total flux comes from a model fitted to a source's radial profile. There are several of these model magnitudes associated with each catalog source. The *devMag* and *expMag* magnitudes are associated with de Vaucouleurs and exponential model fits, respectively. These magnitudes are calculated from independent models in each of the five bands. In addition, there is the *modelMag*, which uses the better of the two fits in the r -band as a matched aperture to calculate the flux in all bands. For extended sources, this option provides the best measurement of colours due to the flux being measured over equivalent apertures across all bands. A composite model magnitude *cmodelMag* is also defined, taking the best fit from either the de Vaucouleurs or the exponential model in each band and obtaining a linear combination of the two that best fits the image. In practice, *cmodelMag* is the optimum total flux indicator, agreeing well with *psfMag* and *modelMag*. However, it does not result in as high a signal-to-noise measurement for colours compared to *modelMag*.

For each source we determine if it is resolved or point-like by calculating the difference between the *cmodel* magnitude and the PSF magnitude in the r -band:

$$resolved_r = |psf_r - cmod_r|. \quad (1)$$

2.1. Spectroscopic data

The SDSS data release 15 includes 3 238 003 unique sources, spectroscopically observed and cross-matched with detections in the WISE catalogue. Each source in the catalogue is labelled as “STAR”, “GALAXY” or “QSO” (quasar) depending on the outcome of fitting various pre-defined models to the spectrum of each source (Bolton et al. 2012). We obtained the DR15 data by submitting Structured Query Language (SQL) jobs to the CasJobs component of SkyServer API². We then applied various conditions upon retrieving the data to ensure the dataset is clean, with minimal contamination from sources which could have incorrect classifications or problematic photometry.

Firstly, 21 086 sources have more than one WISE match within five arc-seconds. We removed the duplicate entries, only keeping the entry with the closest WISE match. Next, we removed 130 sources that have -9999 for the *cmodel* magnitudes where the fit failed. Sources with spectra that are known to have problems are indicated in the DR15 dataset by the *zwarning* flag.

¹ <https://www.github.com/informationcake/SDSS-ML/>

² <http://skyserver.sdss.org/casjobs/>

Fig. 1. Histogram of 3.1 million spectroscopically observed galaxies, quasars, and stars from both the SDSS and BOSS instruments (after removing unclean data). There is a double peaked distribution for galaxies and quasars due to each instrument having different target magnitude selection criteria.

We only selected sources where this flag has a value of either 0 or 16. If the flag is zero, this indicates the spectra has no problems. If the flag is 16, this indicates the “MANY_OUTLIERS” warning which is only present for data taken with the SDSS spectrograph and not with the BOSS spectrograph, and usually indicates a high signal-to-noise spectrum or broad emission lines in a galaxy. Consequently, it rarely signifies a true error. These *zwarning* conditions removed 116 026 sources, 92 349 of which are due to the “SMALL_DELTA_CHI2” warning (*zwarning* = 4), which indicates that the chi-squared value of the best fit spectrum is too close to that of the second best. Finally, we removed 1304 sources where the value of any of the WISE band magnitudes were set to 9999, where extracting the magnitude failed.

In total the cleaning process described above removed 138 546 sources: 90 969 galaxies, 38 890 quasars, and 8687 stars. This left a total of 3 099 457 sources, with 2 209 333 galaxies, 377 888 quasars, and 512 236 stars. The class ratio is approximately 5.8:1.0:1.4. Within the resulting dataset, 1 789 194 sources were observed using the BOSS spectrograph (58%) and 1 310 263 were observed using the SDSS spectrograph (42%). Due to the different survey goals for each of these spectrographs, there is a difference in the magnitude distribution for the sources measured in each case. This is illustrated in Fig. 1, where a double peaked distribution for both galaxies and quasars is evident from the dotted and dashed lines of the BOSS and SDSS spectrographs respectively. We note that due to WISE having a slightly lower resolution than SDSS, 13 904 sources in our spectroscopic dataset have a duplicate WISE match due to multiple SDSS sources being in close proximity. For these SDSS sources in close proximity, their mean separation is 2.36 arcsec with a standard deviation of 1.38 arcsec.

2.2. Photometric data

The complete SDSS photometric catalogue contains 1 231 051 050 entries, including repeat observations. To create the catalogue of previously unclassified photometrically observed sources, we used sources from the SDSS *PhotoPrimary* table with WISE matches that have no associated spectroscopic observation and clean photometry (where the flag *clean* = 1). We limited the dataset to have SDSS (and WISE) magnitudes between 0 and 35, which prevented contamination from spurious nonphysical values. This resulted in a catalogue of 111 447 786 sources. 52 318 SDSS sources were found to be associated with more than one WISE source. We removed duplicate WISE associ-

ations by selecting the WISE source closest in angular separation to the SDSS position. Following these selection steps, the remaining dataset contained 111 395 468 unclassified photometrically observed sources. We note that 14 768 549 of these sources have a duplicate WISE match due to multiple SDSS sources being in close proximity. For these SDSS sources, their mean separation is 2.40 arcsec, with a standard deviation of 1.38 arcsec.

2.3. Feature Set

We used SDSS and WISE photometry, and the *resolved_r* parameter, as features in our machine learning models. We did not explicitly calculate colours, and instead gave all five SDSS bands, and all four WISE bands as individual features. We used the dataset of spectroscopically observed sources to train, validate, and test a model to predict the class labels which are spectroscopically confirmed (we do not use the spectra itself as a feature or in the analysis). We used this model to predict the class labels on the second dataset of unlabelled photometrically observed sources that do not have spectra in SDSS.

3. The machine learning model

In this work we use a random forest (thoroughly reviewed in Louppe 2014) as a supervised learning algorithm to classify galaxies, quasars, and stars. A random forest is an ensemble of independent decision trees where each individual tree is trained on a random subset of both features and data samples. The predicted classification for a new data sample comes from a majority consensus classification across the full set of approximately uncorrelated machine learning models from all the trees in the forest. These principles make a random forest robust to over-fitting, and minimise the variance and bias in its predictions. In addition, random forests are a commonly used algorithm for supervised learning problems due to a number of other strengths: they can deal with numerical and categorical features over different scales, are effective at multi-class problems, and naturally return classification probabilities and feature importance rankings.

We divide our spectroscopically-classified dataset into a training dataset and a testing dataset. The fraction of the complete training set used for each category is discussed in Sect. 3.1. The training set is used to train the random forest classifier and fit the machine learning model. It must be large enough for the random forest to extract a sufficiency of information on the expected types of galaxies, quasars, and stars and their relation to the features we are using. The test dataset is used to derive a variety of performance metrics which assess how the machine learning model would perform on unseen data. These metrics will be used to assess the confidence of the classifications when applying the model to unlabelled sources without spectra. When creating these two subsets of data, we ensure that each contains the same ratio (5.8:1.0:1.4) of the three different classes, and that they have the same distribution in feature space, in order not to bias training, validation or testing.

Instead of splitting out a fixed validation dataset we implement a cross-validation scheme during training to tune the hyperparameters, described in more detail in Sect. 3.2. Validation of this type is used to ensure that the fitted machine learning model does not over- or under-fit the training data.

We use precision, recall, and F_1 score as metrics to assess the performance of the model:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (2)$$

Fig. 2. Precision (*top*), recall (*middle*), and F_1 score (*bottom*) per class as a function of the fraction of the training dataset (1.55 million sources) used to train the random forest. Balancing the classes was done by taking 20% of the galaxies in the training set. All models were evaluated on the test dataset of 1.55 million spectroscopically confirmed sources, without balancing the classes. Class weights inversely proportional to the class frequency were used in the training in all cases.

indicating how good the classifier is at identifying true positives (TP), which are the correctly identified sources. A low precision for an individual class would indicate a low fraction of positive identifications.

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (3)$$

indicates how good the classifier is at minimising false negatives. A low recall for an individual class would indicate it is often misclassified as another class.

$$F_1 = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}} \quad (4)$$

is the harmonic mean of precision and recall and is used as an overall performance metric. In this multi-class scheme these metrics are calculated per class to show the relative performance of each class.

3.1. Training data volume

Given the large number of sources available to us in the spectroscopic catalogue, we initially split the full dataset into two halves, using one half as a preliminary training set, and the other as a preliminary test set. By taking increasing fractions of the full training set to train a random forest model, we can use these

models to predict the class labels of the sources in a consistent test dataset and assess how the volume of data in the training set affects the model performance.

During training we set a weight per class which is inversely proportional to the frequency of that class in the dataset. This is used to prevent the imbalance between the different classes, primarily caused by an excess of galaxies, influencing the fit. Figure 2 shows how the F_1 score per class varies as a function of the fraction of the full training set used to fit the model. The performance increase becomes linear for all classes when using more than 60% of the of the full training dataset, equivalent to 30% of all sources with spectra. When using fewer than 60% of the samples in the full training set the classes show a non-linear increase in performance that is more significant for the minority classes (quasars and stars). This effect is due to the training data subset not capturing a broad enough range of examples for all classes and indicates that a minimum of 60% of the training dataset, equivalent to 30% of the full spectroscopic catalogue, should be used to fit the machine learning model in order to avoid biased results.

Given the class imbalance in our dataset that arises from there being approximately 4 times more galaxies than stars, and 6 times more than quasars, we repeat the procedure described above including only 20% of the galaxies in the training set. Balancing the three classes in the training dataset whilst maintaining the same unbalanced test dataset gives the result shown by a dashed line in Fig. 2. The resulting F_1 scores for the model are lower for all classes, decreasing by the same fraction for all fractions of the training dataset. There is a small increase in precision for galaxies, but a larger decrease in precision for quasars and stars. There is a decrease in recall for galaxies, but an increase in recall for quasars and stars. By limiting the number of galaxies available to train on, the model is hindered on all classes, with a particularly large drop in the precision of quasars (a large increase in false positives in this class from galaxies misclassified as quasars). Overall our approach of including weights in the fitting of the random forest is robust to the galaxy class imbalance without requiring us to sub-sample this class. Reducing the number of galaxies in the training set results in a poorer model for all classes.

3.2. Hyper-parameter optimisation

Random forest algorithms are inherently designed to prevent over-fitting through their selection of random subsets of data and features during training. However, tuning hyper-parameters in the random forest model can lead to over- or under-fitting if set incorrectly, and so we use a validation dataset when optimising the model. To do this we use cross-validation scheme (Mosteller et al. 1968) with a five-fold split. This means that our training dataset is randomly split up into five parts, and each part is used to validate the model trained on the other four parts combined. Thus, the model is trained and validated on five different subsets of the training dataset, where each one is referred to as a fold. 80% of our full training dataset (40% of the full spectroscopic catalogue) is used to fit the machine learning model in each fold, and 20% of our training set (10% of the full spectroscopic catalogue) is used for validating each fold. Given the results described in Sect. 3.1, this reduction in the full training dataset will not have an adverse effect on the performance of our model applied to each fold.

We optimize over three main hyper-parameters that influence the fitting of the random forest model: the number of trees in the forest ($n_{\text{estimators}}$), the maximum number of features per

Table 1. Performance metrics derived from applying the random forest model to the test dataset of 1.55 million spectroscopically confirmed sources.

Features	Precision			Recall			F_1 score		
	Galaxy	Quasar	Star	Galaxy	Quasar	Star	Galaxy	Quasar	Star
SDSS + WISE + <i>resolved_r</i>	0.987	0.961	0.991	0.995	0.944	0.965	0.991	0.952	0.978
SDSS + WISE	0.981	0.957	0.989	0.993	0.929	0.958	0.987	0.943	0.973
SDSS	0.956	0.905	0.958	0.980	0.865	0.888	0.968	0.885	0.922
SDSS + <i>resolved_r</i>	0.981	0.948	0.979	0.993	0.924	0.945	0.987	0.936	0.962
WISE	0.880	0.886	0.711	0.954	0.862	0.466	0.915	0.874	0.563

Notes. SDSS PSF magnitudes are in bands u, g, r, i, z , WISE magnitudes are in bands $w1, w2, w3, w4$.

decision tree (*max_features*), and the minimum number of samples required for a decision tree to split (*min_samples_leaf*). All other hyper-parameters are left at their default values whilst tuning these parameters one at a time.

We found that using 200 estimators in the random forest gave the best performance whilst minimising run-time. Using 1000 trees did not provide a notable increase in performance over 200, but required five times longer to train the model. This is consistent with the expected computational complexity for random forests, which scales linearly with the number of estimators (Louppe 2014). Increasing the *min_samples_leaf* parameter in the range from one to 500 always results in a large drop in the F_1 score, and so we keep it set to one. Adjusting the *max_features* parameter in the range from two to six results in a very small increase in F_1 score. However, given that we only have ten features, substantially increasing this parameter will lead to over-fitting. We therefore use the default value which is the square root of the number of features, which is rounded down to three in our case.

After optimising the hyper-parameters individually, we verify the optimisation by re-fitting the random forest using all three optimised hyper-parameters together, whilst also implementing five-fold cross-validation. This demonstrated that the model had not over- or under-fitted the training data, and would generalise to an unseen dataset.

Finally, to maximise the volume of data available for training, we re-fit the machine learning model without implementing a cross-validation scheme. This uses the complete training dataset (50% of the full spectroscopic catalogue) and the random forest with optimised hyper-parameters. The resulting machine learning model is our production model.

3.3. Feature optimisation

A number of other measurements are available for each source in the SDSS catalogue, which might alternatively be used as machine learning features. In order to establish that there was no advantage in using these alternative features, we evaluated the model performance for a variety of feature iterations using the test dataset. For each of these feature iterations, we also repeated the procedures in Sect. 3.2 to establish that the changes to the feature set were not significant enough to affect the optimal hyper-parameter values. The precision, recall, and F_1 scores per class are shown in Table 1 for each different set of features.

Including the *resolved_r* parameter (Eq. (1)) as a feature gives an improved classification performance for all classes, mostly for the recall of quasars (Table 1). This is illustrated in Fig. A.1 where we see a large improvement in the classification of unresolved stars and quasars, and resolved galaxies. We use the

r -band magnitude for this measure as it has the highest signal to noise. We see the performance decrease when using either u - or z -bands for this purpose due to the poorer signal to noise in these bands, agreeing with similar work by Morice-Atkinson et al. (2018). Using the three high signal to noise bands, g, r, i , to calculate the resolved nature of a source as three separate features does not increase performance as this does not provide the classifier with any additional information. Any increase in signal to noise from doing this is insignificant when also using the nine other photometric features. Consequently we use only the r -band value in order to minimise the number of features.

Excluding WISE magnitudes from the features gives worse results for all classes. This affects quasars the most with the F_1 score dropping from 0.952 to 0.936, which is due to an increase in false negatives (quasars classified as galaxies or stars) with recall being lower by 0.02, and an increase in false positives (stars and galaxies classified as quasars), with precision lower by 0.014. Only using WISE magnitudes as features significantly lowers the F_1 scores for all classes, particularly so for stars.

Whilst a correction for total Galactic extinction is provided per band for SDSS sources (Schlegel et al. 1998), with a median and standard deviation of 0.13 ± 0.30 , 0.10 ± 0.24 , 0.07 ± 0.16 , 0.05 ± 0.12 , 0.04 ± 0.09 for the u, g, r, i, z bands, respectively, applying this correction does not alter the performance of the model for any class. However, this correction is only appropriate for extra-galactic sources and when we apply our model to the unlabelled SDSS photometric catalogue we will not know if sources are galactic or extra-galactic. Therefore we do not include SDSS extinction-corrected magnitudes in our feature set.

Using *cmodel* magnitudes in place of PSF magnitudes for the SDSS bands gives a poorer performance across all target classes, with the most significant effect evident for stars. Using a combination of *cmodel* and PSF magnitudes provides no improvement over using PSF magnitudes alone.

Overall, the optimal combination of features includes the SDSS PSF magnitudes, the WISE magnitudes and the *resolved_r* parameter. The relative weight assigned to each feature from the resulting model is shown in Fig. 3.

3.4. Performance evaluation

The random forest model derived in Sect. 3.2 was applied to our test dataset of 1.55 million spectroscopically confirmed sources (50% of our full spectroscopic catalogue) using the optimal combination of features described in Sect. 3.3. Figure 4 shows the how the features vary for correctly identified (correct) and missed sources per class. There are 1 098 880 correct galaxies, 178 269 correct quasars, and 247 313 correct stars. There are 3942 galaxies that were missed as quasars, 9715 quasars missed

Fig. 3. Relative feature importances and their one standard deviation returned by the random forest classifier.

as galaxies, 5933 stars missed as galaxies, 1845 galaxies missed as stars, 2872 stars missed as quasars, and 960 quasars missed as stars. Figure 4 shows the distribution of correctly identified and missed sources as a function of feature space. The average magnitude and one standard deviation range is shown for the SDSS and WISE bands on the left side in Fig. 4, and a histogram of the *resolved_r* parameter is shown on the right side. This is shown pairwise for all classes (top: quasar/star, middle: galaxy/star, bottom: galaxy/quasar) indicating which feature ranges contain particular sources that have been misclassified as others. Furthermore, a histogram of each of the SDSS and WISE magnitude features is shown in Fig. A.2, per class, for correct and missed sources. In summary, it can be seen that correctly identified galaxies and stars show a similarly shaped spectrum across the SDSS bands, with galaxies being brighter in the WISE bands. Correctly identified quasars show a flatter spectrum across the SDSS bands compared to correctly identified galaxies and stars. Correctly identified quasars also have a much larger $w1 - w2$ difference than galaxies and stars. Correctly identified quasars and stars are mostly unresolved, with similar distributions, whilst correctly identified galaxies are mostly resolved.

Quasars missed as galaxies (9715 sources) appear similar to galaxies across both the SDSS and WISE bands, and are more often resolved (bottom row of Figs. 4 and A.2). Most of these sources have redshifts less than 1 (see Fig. 5). They can appear as resolved galaxies where the light from the galaxy dominates over the AGN, giving a galaxy-like spectrum across the SDSS and WISE bands whilst there is still an AGN present (e.g. the top row of Fig. 6). Alternatively, they can be more unresolved galaxies that are much redder (e.g. the bottom row of Fig. 6), although their *resolved_r* parameter is still much larger on average than point source quasars and stars (shown by the dashed pink line in the bottom right panel of Fig. 4). Their distribution per magnitude feature in SDSS and WISE is shown in Fig. A.2. In particular, their distribution has a shape that follows that of the correctly identified galaxies (shown by the dashed pink line in the bottom row of Fig. A.2). Overall, quasars missed as galaxies tend to be resolved objects (at low redshifts) which are redder in colour.

Galaxies missed as quasars (3942 sources) appear similar to quasars in the WISE bands, and have a spectrum across the SDSS bands of a similar shape to quasars but are fainter than quasars on average. Compared to correctly identified galaxies, they tend to be unresolved (shown by the dashed green line in the bottom right panel of Fig. 4), and have a similar distribution

in redshift from $z = 0$ to $z = 1$ (Fig. 5). Whilst number density of galaxies drops significantly above $z = 1$, higher redshift galaxies are more often misclassified as quasars than stars. Whilst we do not use redshift as a feature, we expect the photometry to be correlated with redshift, and therefore for this effect to be seen in our classification scheme. In the SDSS bands galaxies missed as quasars appear bluer on average than correctly identified galaxies, with a flatter spectrum. Figure A.4 shows a selection of examples of galaxies missed as quasars. Their distribution per magnitude feature in SDSS and WISE is shown in the bottom row of Fig. A.2.

Galaxies missed as stars (1845 sources) look similar to correctly identified galaxies in the SDSS bands, but are fainter in the WISE bands (middle row of Figs. 4 and A.2). Figure A.5 shows a selection of examples of these sources. The top-left image in Fig. A.5 is a case where the SDSS classification pipeline is incorrect, and our model has correctly identified it as a star. There may be a small number other cases like this, although without manually inspecting all of them (which we have not done), it is not possible to find these cases automatically with this dataset. Their distribution in redshift and the *resolved_r* parameter is similar to correctly identified galaxies (middle row of Figs. 4 and 5).

Stars missed as galaxies (5933 sources) have a similarly shaped spectrum to stars and galaxies within the SDSS and WISE bands, but are much fainter in the SDSS bands, and mostly resolved (middle row of Figs. 4 and A.2). The top row of Fig. A.6 shows examples of these resolved sources with *resolved_r* > 1. The first image shows a galaxy with a foreground star at its centre. The second image shows a galaxy with a foreground star just off-centre, which is close and bright enough to dominate the spectrograph, hence why the SDSS pipeline has incorrectly labelled it as a star whilst our algorithm has correctly labelled it as a galaxy. The fifth image shows a foreground star on top of a galaxy where the SDSS pipeline has correctly labelled both the star and the background galaxy (which was in our training set) from their spectra. However our algorithm has missed this star as a galaxy, likely due to contamination from the background galaxy, which is 0.35 magnitudes brighter in the PSF *r*-band, and the large value of the *resolved_r* parameter. The bottom row of Fig. A.6 shows the rarer examples of more unresolved sources (*resolved_r* < 0.3), where the reason for the misclassification is predominantly due to these stars having a spectrum across the SDSS and WISE bands that is very similar to galaxies.

Quasars missed as stars (960 sources) tend to be fainter in the $w2$ -band than correctly identified quasars, which results in a smaller $w1 - w2$ difference, and have a spectrum across the SDSS bands similar to stars (top row of Figs. 4 and A.2). They are much fainter in the *u*-band than correct quasars making them appear redder on average. They have a similar distribution in the *resolved_r* parameter to correct quasars, correct stars and missed stars (top right panel of Fig. 4). They tend to be at higher redshifts since lower redshift quasars are more likely to be resolved and misclassified as galaxies (shown in Fig. 5). The main reason for their misclassification is due to the shape of their spectrum within the SDSS and WISE bands, and being very faint in the *u*-band. Some examples of these are shown in Fig. A.7.

Stars missed as quasars (2872 sources) tend to be much fainter across all bands than correctly identified stars and have a larger $w1 - w2$ difference (top row of Figs. 4 and A.2). They are significantly fainter in the *z*-band than correct stars, giving them a flatter spectrum across the SDSS bands. They have a similar distribution in the *resolved_r* parameter to correct stars, correct quasars and missed quasars. The main reason for their misclassification is due to the shape of their spectrum across the SDSS

Fig. 4. *Left:* average magnitude and one standard deviation error bar for each feature (waveband) for correct and missed sources, per class. Error bars and lines are offset in the x-axis per feature for clarity. The resolved feature was also used in the model. The F_1 score is 0.990 for galaxies, 0.953 for quasars and 0.977 for stars. *Right:* histogram of the $resolved_r$ parameter per class, for correct and missed sources.

and WISE bands, and being very faint sources where there are much fewer stars in the training set. Some examples of these are shown in Fig. A.7.

3.4.1. Precision, recall, and F_1 score

The average precision, recall, and F_1 scores per class are shown in Table 1. These metrics represent the overall performance of the model on the test dataset of 1.55 million spectroscopically

confirmed sources (50% of our full spectroscopic catalogue). Furthermore, we calculate precision, recall, and F_1 score for sources binned along various variables to assess the model performance throughout the 10-dimensional feature space.

Figure 7 shows precision, recall, and F_1 score calculated as a function of the binned PSF r -band magnitude, PSF r -band error, one dimensional feature, and $resolved_r$ parameter. The shaded area shows a one standard deviation confidence interval calculated per bin from the Wilson score interval (Wilson 1927). The

Fig. 5. Histogram of spectroscopic redshifts for correct and missed sources per class from the random forest model applied to the test dataset of 1.55 million spectroscopically confirmed sources.

Wilson interval score behaves well for probability distributions with many values close to zero or one (unlike a Normal distribution) or with small sample sizes. A histogram per class is also shown to guide where the source counts drop, normalised relative to the galaxy class which peaks at a half. Furthermore, Fig. A.3 shows histograms of precision, recall, and F_1 score over each of the individual SDSS and WISE magnitude bands.

Stars are classified correctly with very high precision as a function of the PSF r -band magnitude. The precision has values of 0.99–1 between PSF r -band magnitudes of 14–18, dropping to 0.95 at a magnitude of 21, before rising again and then tailing off as the source density of stars decreases. The recall for stars is higher than 0.95 in the magnitude range 14–19, but drops to 0.73 at magnitude 22 before rising again. This drop in recall is due to an increase in false negatives. These stars missed as galaxies (4071 sources) are mostly stars super-posed along the line of sight towards background galaxies, increasing their resolved parameter. The precision, recall, and F_1 score as a function of each magnitude feature in SDSS and WISE is shown in Fig. A.3. The other source of false negatives, the stars missed as quasars (2328 sources), are mostly faint unresolved sources with a distribution per magnitude feature in SDSS and WISE shown in Fig. A.3. Overall, the F_1 score for stars is greater than 0.8 for the majority of sources, excluding the very high and low magnitude limits where the number of sources is low. The F_1 score drops at magnitudes 20–22.5 due to an increase in false negatives. As a function of the $resolved_r$ parameter (bottom row in Fig. 7) the precision stays above 0.9, dropping when $resolved_r$ is greater than 4. The recall drops significantly for resolved stars.

Galaxies are classified correctly with high precision as a function of the PSF r -band magnitude. This is higher than 0.95 from magnitude 15.5 to 23, and reduces to 0.8 from magnitude 23 to 25 as the source density falls. The recall for galaxies is also very high, being above 0.98 from magnitude 16 to 23.5, only dropping either side of this as the source density significantly decreases. Overall the F_1 score is very high over the entire magnitude range, only dropping when the source density significantly falls. As a function of the $resolved_r$ parameter, the precision and recall for galaxies drops below 0.95 when $resolved_r$ is less than 0.2 or greater than 6, as the source density significantly drops.

Quasars are classified correctly with high precision as a function of the PSF r -band magnitude, and is greater than 0.95 between magnitude 15.7 and 20.7. Precision falls either side of this as the source density drops. The recall for quasars is above 0.95 between magnitude 18.6 and 20.7. Outside of this range the recall drops due to an increased number of false negatives. These

are mostly due to quasars missed as galaxies (9715 sources) while only a small number (960 sources) are due to quasars missed as stars. The bottom row in Fig. 7 shows how the precision and recall drop below 0.95 when the $resolved_r$ parameter is greater than 0.8 and 1.1, respectively. Whilst the precision stays high over the majority of the source population, the recall drops significantly for resolved sources due to resolved galaxies being misclassified as quasars.

Whilst fainter sources generally have larger PSF r magnitude errors we do not see evidence that this affects the performance of the classifier (second row of Fig. 7). For example, the recall for stars falls as the PSF magnitude error increases, however the precision remains high. This drop in recall for stars (due to an increase in misclassified stars) is seen at fainter magnitudes from 20 to 22, where the classifier confuses them with galaxies and quasars. In general for all three classes, the F_1 scores only drop as a function of PSF r error when the source density falls significantly.

3.4.2. Classification probabilities

The random forest algorithm naturally returns the probabilities per class of a classification, which sum to one over the number of possible classes. The class with the largest probability is the predicted class assigned by the algorithm. The left of Fig. 8 shows a histogram per classification possibility over the probabilities for the assigned class. For example, the dashed pink line shows quasars that the random forest predicted were stars, and is a histogram of the probability that the random forest assigned to it being a quasar. Incorrect classes must have probabilities less than one half, whilst correct classes must have probabilities greater than one third. The right side of Fig. 8 shows the same but as a cumulative normalised histogram.

The 0.99–1 bin in Fig. 8 contains 71%, 52% and 81% of the correctly classified galaxies, quasars, and stars, respectively. Furthermore, 96% of galaxies, 84% of quasars and 94% of stars have classification probabilities above 0.9. Quasars are the weakest performing class in this regard, where in general the correctly classified quasars have lower classification probabilities than correctly classified stars or galaxies. This indicates that this class is the most difficult to classify.

The last column in Fig. 7 has the classification probabilities (dashed lines) and one standard deviation (dotted lines) plotted along with the F_1 scores. Overall they follow the trend in the F_1 scores for all classes. They only start to decrease in areas with very low numbers of sources (for example at the edges of the plots).

Figure 9 shows a 2-D histogram of the PSF r magnitude and classification probability. Misclassified sources are also plotted individually and a normalised 1-D histogram of the PSF r magnitude is overlaid. Overall the random forest probabilities show a similar distribution over the whole PSF r magnitude range. Misclassified sources are mostly found with r magnitudes from 20 to 22.

3.4.3. Limiting the training set

One of the primary intended uses for machine learning algorithms is to classify sources detected in new astronomical surveys based on models trained from existing data, a form of transfer learning (e.g. Pratt et al. 1991; Tang et al. 2019). As telescopes become more powerful, new surveys typically become progressively deeper in sensitivity and recover fainter populations of sources. One possible limitation for this form of transfer learning is the introduction of biases when deploying a model on a

Fig. 6. Examples of quasars missed as galaxies. The spectroscopically observed target is at the centre in a red box, and photometrically observed sources are circled in blue. *From top left to bottom right* in order of how resolved they are, their SpecObjIDs, *resolved*, parameters and redshifts are: 5331182333197262848, 2.544 ($z = 0.137$); 706011020694415360, 1.655 ($z = 0.058$); 996574564462913536, 1.603 ($z = 0.170$); 6097903957849382912, 1.347 ($z = 0.527$); 148177585556364800, 1.163 ($z = 0.100$); 8848470304915169280, 1.047 ($z = 0.238$); 6729619569105674240, 0.97151 ($z = 0.142$); 668808218473949184, 0.968 ($z = 0.098$); 7523363890043789312, 0.899 ($z = 0.592$); 9072765177980493824, 0.841 ($z = 0.802$); 1612303285509711872, 0.782 ($z = 0.107$) and 9897866743979679744, 0.631 ($z = 0.337$).

population of sources that are fainter than the population used for training. Here we consider how the performance of our machine learning model changes as a function of source brightness.

Figure 10 shows the precision, recall, and F_1 score when upper limits were set on the r magnitude in the training set. For a given limit, 50% of all sources brighter than the limit were used to train the model, and all remaining sources were used to test the model. Figure 11 shows the relationship between the r magnitude training limit and the r magnitude at which the F_1 score drops below a certain fraction of the original value. These figures show that the model maintains its performance up to the training magnitude limit as expected, however, beyond this limit the model weakens at various rates for the three classes of source. For galaxies and quasars, the recall is significantly affected, due to an increase in false negatives. For stars, the precision is significantly affected, due to an increase in false positives.

For galaxies, the precision is not strongly affected by imposing any magnitude limit in the training until above magnitude 23. The recall (and therefore F_1 score) is immediately affected above the training limit, but improves significantly as the training magnitude limit increases beyond 19.5. Overall, the F_1 score (dominated by the poor recall) shows a jump in performance above magnitude 19.5 (shown in Fig. 11).

For quasars, the precision is not strongly affected by any magnitude limit in the training until above magnitude 21. However, the recall is immediately affected above the training limits imposed. Overall, the F_1 score (dominated by the poor recall) improves linearly as the training magnitude limit increases (shown in Fig. 11).

For stars, the recall is not strongly affected by any magnitude limit in the training until above magnitude 21. However, the precision is immediately affected above the training limits imposed. Overall, the F_1 score (dominated by the poor precision) improves linearly as the training magnitude limit increases (shown in Fig. 11).

When imposing the r magnitude training limit at its lowest value of 18.5, the overall fraction of each class used to fit the model is 6.5%, 3.3% and 34.6% for galaxies, quasars, and stars, respectively. Comparing Figs. 2 and 10 indicates that on average the F_1 scores should still be much higher despite this reduction

in the number of training samples. Furthermore, the F_1 score as a function of the PSF r magnitude when training on 1% of sources shows the same distribution as that shown in the first row of Fig. 7, which used 50% of the data for training. In other words, training on 1% of each class gives a distribution in r magnitude of F_1 scores that is similar to training on 50% of each class. Overall this demonstrates that when limiting the training set to sources below magnitude 19, the training set does not contain enough examples of each class that are representative out to fainter magnitudes. This degraded performance is alleviated significantly once the training set goes up to and beyond magnitude 19.5, but overall the models cannot be applied to sources at much fainter magnitudes than were included in the training set without resulting in a significantly poorer performance. Furthermore, this decrease in performance will not be quantifiable for a distribution of sources at magnitudes fainter than those in the training dataset. Consequently, we consider that the metrics used to quantify the performance of the random forest model are only relevant within the magnitude range explored by our training dataset, as shown in Fig. 7.

3.4.4. Confidence of classification for individual sources

In the previous section, we make approximations in order to specify how the model performs as a function of individual features, reducing the full 10-D space used by the performance metrics. However these feature-specific metrics are still only correct on average within a specific range of feature space, and cannot be used to quantify performance for individual sources. In order to narrow further the range of feature space that a performance metric represents, we implement a nearest neighbour approach for a subset of sources in one and ten dimensions, and we investigate the relation between these metrics and the classification probabilities returned by the random forest for individual sources.

For each source, the ten features are transformed into one single feature by using PCA (as seen in the third row of Fig. 7). In this single feature, 75% of the variance is retained from the set of 10 features. Given the computationally intensive nature of finding large numbers of nearest neighbours for many sources,

Fig. 7. Precision, recall, and F_1 score as a function of: PSF r magnitude (*first row*), PSF r magnitude error (*second row*), features transformed into one dimension (*third row*), and the *resolved_r* parameter (*fourth row*). These metrics are calculated from applying the random forest model to the test dataset of 1.55 million spectroscopically labelled sources. The shaded area shows a one standard deviation confidence interval calculated per bin from the Wilson score interval. A histogram per class is shown normalised relative to galaxies which peaks at a half. In the *last column* the dashed lines shows the mean classification probability per bin for the predicted classes, and the dotted lines shows one standard deviation below this mean. Galaxies are in green, quasars in pink and stars in blue.

Fig. 8. Histogram (*left*) and cumulative normalised histogram (*right*) of the random forest classification probabilities using a bin size of 0.005. 71%, 52% and 81% of the correctly classified galaxies, quasars, and stars have classification probabilities greater than 0.99. Correct classifications must have a probability greater than one third and incorrect classifications must have a probability less than one half. In general correctly classified quasars have lower classification probabilities than correctly classified stars or galaxies.

Fig. 9. Random forest classification probabilities plotted as a function of PSF r magnitude. A normalised histogram of the PSF r magnitude is also overlaid, though note that the missed galaxies histogram is multiplied by a factor of ten to be visible. 71%, 52% and 81% of the correctly classified galaxies, quasars, and stars have classification probabilities greater than 0.99 (as detailed in Fig. 8).

we find nearest neighbours for only 100 000 sources from the test dataset. For each of these sources we find the 10 000 nearest neighbours in the 1-D feature space, calculating the F_1 score from these nearest neighbours. The top left panel in Fig. 12 shows the F_1 score as a function of the 1-D feature for these 100 000 sources. Whilst in Fig. 7 the bin size is constant, this nearest neighbours approach uses a variable bin size, defined by the requirement to have 10 000 nearest neighbours within the bin. This is represented by the horizontal error bars in the top left panel of Fig. 12, calculated as one standard deviation from the 10 000 nearest neighbours. These error bars are small where source density is high, and only begin to become visible where source density drops at high and low values of the 1-D feature. The error in the F_1 score is calculated using the Wilson score interval, and is only significant when the source density

drops at the edges of the plot. In the top right panel of Fig. 12 we show the classification probabilities for the same sources. As shown in Fig. 8, quasars generally have lower classification probabilities, whilst galaxies and stars have higher values. The probabilities for all classes are positively correlated with the F_1 score calculated from the 10 000 nearest neighbours. Whilst this nearest neighbour method does not fully explore the 1-D feature space (only using 6.5% of the test dataset), the F_1 scores and probabilities are in agreement with those shown in Fig. 7 (the final panel of the third row).

We repeat this process using all of the features, finding the 10 000 nearest neighbours for each of 100 000 sources in the original 10-D feature space (bottom row of Fig. 12). The error bars on the 1-D feature are now much larger, as a sample of 10 000 nearest neighbours selected in 10-D will have a wider

Fig. 10. Precision, recall, and F_1 score plotted as a function of PSF r magnitude for galaxies (*left*), quasars (*middle*), and stars (*right*). Each of the coloured lines depicts the upper limit of the r magnitude for the random forest training. Below each of these limits 50% of the dataset was used for training.

Fig. 11. r magnitude training upper limit plotted against r magnitude at which the F_1 score drops to a fraction of the F_1 score (shown in the legend) without a r magnitude training limit. Below this limit 50% of the data were used for training. A linear relationship is maintained for stars and quasars, while galaxies show a non-linear change between r magnitude of 19 and 20 for F_1 score fractions less than 0.9 (also seen in Fig. 10).

Fig. 12. 100 000 sources from our spectroscopic test dataset, the F_1 score is calculated from 10 000 nearest neighbours in 1-D (*top*) and 10-D (*bottom*) feature spaces, plotted against the mean 1-D feature (*left*) and mean classification probability (*right*) for those neighbours. Sources are binned per pixel and colours combined proportional to how many of each class are in that pixel bin. The brightness corresponds to the total source count in that pixel on a logarithmic scale. A small sample of individual data points with their error bars are also overlaid.

distribution than in 1-D. A broadly similar trend in F_1 scores is seen when compared with that for the 1-D case, see the top left panel of Figs. 12 and 7. However, we now see a much larger scatter in the F_1 scores for each class as a result of exploring more sparsely populated areas of the 10-D feature space. In particular, there are regions in the 10-D feature space where the F_1 scores are significantly below the average, most notably for stars and quasars. This occurs because the source density in the 10-D space can drop significantly, even when the corresponding 1-D feature is densely populated with a high F_1 score. In the bottom right panel of Fig. 12 we show the mean classification probability per source calculated from its nearest neighbours in 10-D. There are significantly more galaxies, quasars, and stars with lower probabilities due to the lower source density in some regions of the 10-D feature space. Despite this, the classification probabilities are still positively correlated with the F_1 score. In other words, areas of the 10-D feature space which had low F_1 scores calculated from their nearest neighbours in 10-D also have low classification probabilities.

Overall, Fig. 12 shows that the random forest probabilities are positively correlated with the calculated F_1 scores across the full 10-D feature space. This means that even in sparsely populated regions of the 10-D feature space, the random forest probabilities reflect this uncertainty in the classification for individual sources.

Calculating F_1 scores in the 10-D space like this is highly computationally intensive. For each of our 111 million photometrically observed sources that we classify in Sect. 4, finding the 10 000 nearest neighbours from the 1.55 million sources in our spectroscopic test dataset and calculating F_1 scores localised to each newly classified source would require significant computational resources. Whilst this would be beneficial when using

classification algorithms in higher dimensions that do not return a classification confidence per source, the classification probabilities returned by the random forest provide a more efficient solution. Furthermore, through testing a subset of 100 000 sources we have demonstrated that the random forest probabilities are in agreement with the localised F_1 scores. Overall, the random forest probabilities are an effective method to estimate the reliability of a classification, representing the confidence of classifications in that particular area of 10-D feature space.

3.5. Unsupervised clustering with UMAP

In order to interrogate the performance of our machine learning model further, we use a form of unsupervised machine learning. Unsupervised machine learning algorithms have no prior knowledge of target class labels, and attempt to group sources by inferring their similarities in feature space. Here we use the UMAP (Uniform Manifold Approximation and Projection: McInnes et al. 2018) non-linear dimension reduction algorithm to explore potential patterns or clustering within our dataset. UMAP is a non-linear dimension reduction algorithm based on manifold learning techniques and topological data analysis. It is designed to preserve structure information on local scales, but is also effective at preserving structure on global scales, particularly as compared to alternative non-linear algorithms such as t -distributed Stochastic Neighbour Embedding (t -SNE: van der Maaten & Hinton 2008). It also has significantly superior scaling performance to t -SNE (McInnes et al. 2018, Figs. 5–7), allowing for much larger samples sizes and number of dimensions. UMAP can also be semi-supervised or fully-supervised, and allows for metric learning where a model can be applied to unseen data. Furthermore, unlike t -SNE it returns meaningful distances between clusters of points. However, the

Fig. 13. Test dataset of 1.55 million spectroscopically observed sources processed with UMAP in unsupervised (*top left*) and supervised (*top right*) schemes to reduce from ten features to two. The resulting two dimensions are plotted, with sources binned per pixel and colours combined proportional to how many of each class are in that pixel bin. Galaxies are shown in green, quasars in pink and stars in blue. The brightness corresponds to the total source count in that pixel on a logarithmic scale. Whilst the axis have no physical meaning in this 2-D space, UMAP returns proportionate distances between points and clusters, effectively displaying local and global structures in both supervised and unsupervised schemes. The *bottom row* shows the mean and standard deviation of the random forest classification probabilities.

dimensions returned in embedded space have no specific meaning, unlike principal component analysis, for example, where the returned dimensions are the direction of greatest variance in the original data.

In an unsupervised scheme, we applied UMAP to our test dataset (1.55 million spectroscopically observed sources) without supplying the class labels to the algorithm. UMAP transforms the 10-D data into 2-D, and the resulting embedding is shown in the top left panel of Fig. 13, where each point has been colour coded by its class label, and where multiple classes are present in a single pixel bin those colours are combined proportionally. The brightness of the pixel represents the number of sources within that bin, and is on a logarithmic scale.

In a supervised scheme, we used our training dataset (1.55 million spectroscopically observed sources), to train a UMAP model embedding the 10-D data into two dimensions. In this way, UMAP sees the labels of these sources whilst fitting the embedding which maps them into the 2-D space. Once trained, we applied the model to the unseen test dataset of 1.55 million sources. This is shown in the top right panel of Fig. 13. The same

result is achieved in a semi-supervised scheme where only half of the class labels in the training dataset are given to UMAP.

In both supervised and unsupervised scenarios, UMAP has effectively separated the data into distinct classes in this 2-D space. This is achieved much clearer in the supervised scheme. Furthermore, UMAP has distributed the classes in a way that represent more than just their class label. The structures present in Fig. 13 represent their magnitude, the shape of the spectrum across the SDSS and WISE bands, plus how resolved the sources are. Figure 14 is coloured by these parameters as a guide. Galaxies that are predominantly star forming (termed the blue cloud in colour-magnitude diagrams) are distinguished as a semi-separate cluster with smaller $|\text{PSF}_u - \text{PSF}_z|$ values (bluer colours). These also tend to be more resolved galaxies. From these bluer galaxies, the source density drops as their colours become greener. This is referred to as the green valley (Salim 2014; Angthopo et al. 2019), which contains galaxies where star-formation has been quenched. Galaxies then become redder in the left of the structure which contains old passive galaxies (termed red sequence galaxies). Stars are arranged in

Fig. 14. Test dataset of 1.55 million spectroscopically observed sources processed with UMAP to reduce from ten features to two in supervised (*top row*) and unsupervised (*bottom row*) schemes. The resulting two dimensions are plotted, with sources binned per pixel and coloured by the *resolvedr* parameter (*left column*), an SDSS colour $[\text{PSF } u - \text{PSF } z]$ (*middle column*) and a WISE colour $|w1 - w2|$ (*right column*). Whilst the axis have no physical meaning in this two dimensional space, UMAP returns proportionate distances between points and clusters, effectively displaying local and global structures in both supervised and unsupervised schemes. Galaxies can be seen distributed by how resolved they are, separating nearby extended galaxies from distant point sources. Furthermore, red galaxies can be distinguished from blue ones and green valley galaxies can be seen connecting the two regions. Stars are arranged in an alternate representation of a Hertzsprung Russell diagram, distinguishing their colour and luminosity. Quasars are predominantly unresolved bluer sources, though redder quasars are seen on the outskirts of the cluster.

an alternate representation of a Hertzsprung Russell diagram, distinguishing their colour and luminosity. Quasars are predominantly unresolved bluer sources, though redder quasars are seen on the outskirts of this cluster.

When taking a closer view of the clusters in Fig. 13 it is noticeable that some regions are contaminated at a low level by colours from other classes. For example the region dominated by galaxies is contaminated at a low level by stars and quasars (shown by pink and blue colours within the green). This contamination is most noticeably so for the galaxy region, less so for the quasar region, and barely noticeably so in the star region. This qualitatively reflects the precision and recall for each class presented by the random forest. Stars have the highest precision (few false positives in the blue region of Fig. 13), followed by galaxies then quasars. Galaxies have the best recall (few false negatives in the pink and blue region of Fig. 13), however stars and quasars have a poorer recall, shown by blue and pink appearing outside of their respective classes.

The bottom row of Fig. 13 shows the mean and standard deviation of the random forest classification probabilities in the supervised and unsupervised metric-learning schemes for the spectroscopic test dataset. The mean classification probabilities

are very high (also seen in Fig. 8) throughout this 2-D space, and only decrease in regions where classes overlap in the unsupervised case, and on the outskirts of clusters in the supervised case. The standard deviation of the classification probabilities is very low, varying slightly throughout the galaxy and quasar clusters, but remains low for most of the cluster of stars. Overall, a source on the outskirts of a cluster in these UMAP projections does not always correlate with it having a lower random forest classification probability.

Whilst we have used UMAP as a qualitative assessment of the class labels, it affirms that using UMAP in an unsupervised, semi-supervised, or fully-supervised scheme is an effective tool in understanding the distribution of sources where labels are not available. In Sect. 4.2 we explore using UMAP on the 111 million un-labelled photometrically observed sources.

4. Classifying new sources

4.1. Applying the random forest model

Using the machine learning model described in Sect. 3, we classify the 111 395 468 previously unlabelled SDSS photometric

Fig. 15. Histogram of the PSF r magnitude for all 111 million newly classified sources. Spectroscopically observed sources are also plotted.

sources described in Sect. 2.2 as either a galaxy, quasar or star. This returns 50 417 547 galaxies, 2 137 839 quasars, and 58 840 082 stars, with their distribution shown in Fig. 15.

Figure 16 shows that the average shape of a spectrum across the SDSS and WISE bands for these newly labelled photometrically observed sources is similar to that of the spectroscopically observed sources. Newly labelled galaxies and quasars are fainter across all bands (apart from $w3$ and $w4$), whilst newly labelled stars are only significantly fainter in the shorter wavelength SDSS bands. The $resolved_r$ feature effectively distinguishes stars from galaxies for these new sources, shown in the right plot of Fig. 16, although there are significantly more newly labelled galaxies that are unresolved point sources than were in the spectroscopically observed dataset.

Figure 17 shows the distribution of classification probabilities for the newly labelled sources. The 0.99–1 probability bin is the most populated bin, containing 6 683 526 galaxies (13%), 330 666 quasars (15%), and 41 279 349 stars (70%). There are 35 075 918 galaxies (70%), 722 159 quasars (34%), and 54 673 689 stars (93%) with probabilities greater than 0.9, which is a lower fraction than those in our spectroscopically observed test dataset (96%, 84%, 94%). The cumulative histogram in the right of Fig. 17 shows that the probabilities of newly labelled stars have a similar distribution to those with spectra. However, newly labelled galaxies and quasars have significantly more sources with lower classification probabilities than those with spectra. This indicates that stars are easier to classify than the other two classes, and quasars are the most difficult, having lower classification probabilities on average.

Figure 18 shows how the classification probabilities vary as a function of the PSF r magnitude. For stars fainter than a PSF r magnitude of 18, the classification probabilities are very high. However, above magnitudes of 18, there are many more stars with lower classification probabilities. This is likely because stars at fainter magnitudes are more easily confused with the higher number density of galaxies and quasars at fainter magnitudes. Most of the newly labelled quasars with lower classification probabilities have magnitudes above 20. Above a magnitude of 23, new sources have low classification probabilities mainly split between galaxies and stars.

Figure 20 shows how the classification probabilities vary as a function of the $resolved_r$ parameter for these newly classified sources. The number density of galaxies drops at

smaller $resolved_r$ values, with 8.1 million galaxies having $resolved_r < 0.2$, and 4.2 million with $resolved_r < 0.1$ (also seen in Fig. 16). Whilst 70% of galaxies have probabilities greater than 0.9 and the peak in the normalised histogram for galaxies overlaid on Fig. 20 is at $resolved_r = 0.7$, for unresolved galaxies the classification probabilities are much lower, plateauing at 0.66 when $resolved_r < 0.01$. This is also reflected in the training and testing datasets where there are significantly less unresolved galaxies to train on, and therefore the classification probabilities are much lower (Fig. 7).

There are photometric redshifts available for 47 million sources in our catalogue (Beck et al. 2016). Figure 19 shows their distribution, with quasars and galaxies having distributions peaked above $z = 0.1$, and stars having a broader distribution cover lower redshifts. When looking at the classification probabilities for these stars with photometric redshifts, we see no difference to that of the overall distribution of probabilities. In other words, it is unlikely that these stars are all misclassified by our model. It is likely that the SDSS photometric redshift pipeline has fitted incorrect photometric redshifts to these sources which they have misclassified as stars.

4.2. Clustering with UMAP

We applied UMAP to 10% of the 111 395 468 million photometrically observed sources in an unsupervised scheme, shown in the left plot of Fig. 21, where sources are coloured by the classification labels assigned by the random forest model. UMAP effectively separates the galaxies and stars as labelled by our random forest model, but does not pick out quasars as a separate cluster due to the significantly lower number of sources. However, quasars are one continuous structure mainly mixed in with the cluster of galaxies. Stars are clustered into various groups with complex structure. Galaxies are clustered as a single group which broadly resembles the global structures displayed in the spectroscopic sources, seen in Fig. 13. The middle plot in Fig. 21 includes the 3.1 million spectroscopically observed sources without their labels, combined with 10% of the unlabelled photometric sources. This unsupervised approach guided by the inclusion of unlabelled spectroscopically observed sources improves the clustering by increasing the number in the minority class. It particularly helps to separate out quasars which now have much less overlap with other classes. The right hand plot in Fig. 21 includes the same sources as the middle plot, but in this case the labels for the spectroscopically observed sources were explicitly used by the UMAP clustering algorithm. This semi-supervised approach helps UMAP achieve a clearer clustering as defined by the galaxy, quasar, star classification scheme. The galaxy-star separation is very distinct in this case, although there is still some overlap with quasars. Even when classes overlap in the embedded 2-D space shown in Fig. 21, the structures visualised in overlapping regions are coherent and consistent with the labels from the random forest. This can be interpreted as non-overlapping clusters of sources being distinguished in the higher 10-D space by the random forest, but overlapping in this 2-D space due to the embedding derived by UMAP.

We also implement supervised dimension reduction by deriving UMAP models from the spectroscopic training dataset in supervised and unsupervised schemes, and using these fitted models to transform all unlabelled photometric sources into the resulting 2-D space. This form of metric-learning allows us to efficiently apply the UMAP transformation derived from 1.55 million spectroscopically observed sources (as seen in Fig. 13) to all 111 million photometrically observed sources. This is computationally advantageous compared to running UMAP on 111 million photometric sources in an un-

Fig. 16. *Left:* average magnitude and one standard deviation error bar for each feature (waveband) for 111 million newly classified sources. Error bars and plot lines are offset in the x -axis for clarity. *Right:* histogram per class over a measure of how resolved the source is for the newly classified sources (solid line), and for our dataset of 3.1 million sources with spectra (dashed line).

Fig. 17. Histogram (*left*) and cumulative normalised histogram (*right*) of the random forest classification probabilities using a bin size of 0.005. 13%, 15% and 70% of the 111 million classified galaxies, quasars, and stars have classification probabilities greater than 0.99. 70%, 34% and 93% of galaxies, quasars, and stars have classification possibilities greater than 0.9.

semi-supervised scheme. The results are shown in Fig. 22, using the unsupervised model in the top left plot, and the supervised model in the top right plot, where sources are colour coded post-projection with the labels assigned by the random forest algorithm. As expected in this metric learning scheme, the global and local clustering patterns appear the same as those from the spectroscopically observed sources shown in Fig. 13. There is a clear separation of the three classes as designated by the random forest algorithm, similar to that seen for the spectroscopically observed sources. Given the significant increase in source density in this figure, regions that were previously unpopulated in Fig. 13 are now filled. There is significant overlap in the unsupervised case where each of the three classes intersect, and many more stars labelled by the random forest appearing in the quasar and galaxies clusters. In the supervised scheme there is significantly less overlap of the three classes, with UMAP and the random forest in a much stronger agreement. Nevertheless in the supervised scheme there are still regions where stars labelled by the random forest are co-located with the main cluster of galaxies and on the edge of the cluster of quasars. We do not see quasars labelled

by the random forest noticeably in the main cluster of galaxies or stars, although there are only 2.1 million quasars compared to the 50 million galaxies and 59 million stars. Galaxies labelled by the random forest overlap the edge of the cluster of quasars in the supervised scheme where their colour is much redder (see the top middle plot of Fig. 14), indicating a disagreement between UMAP and the random forest for a small number of these unresolved red galaxies or quasars. Although these objects have much lower random forest classification probabilities. Galaxies labelled by the random forest do not overlap the cluster of stars with any significance.

The bottom row of Fig. 22 shows the mean and standard deviation of the random forest classification probabilities in the supervised and unsupervised metric-learning schemes for the 111 million photometric sources. Whilst the mean classification probabilities are lower than in the test dataset (as seen in Figs. 13 and 17), the mean remains high throughout this 2-D space. It decreases in regions where classes overlap in the unsupervised case, however there are small clusters of stars

Fig. 18. Random forest classification probabilities plotted as a function of the PSF r magnitude for the 111 million newly classified sources. A histogram of the PSF r magnitude is also overlaid normalised for each class. 13%, 15% and 70% of the galaxies, quasars, and stars have classification probabilities greater than 0.99. 70%, 34% and 93% of galaxies, quasars, and stars have classification possibilities greater than 0.9 (as detailed in Fig. 17).

Fig. 19. Histogram of photometric redshifts for newly classified sources. 47 million (42%) have available photometric redshifts. We note the large portion of stars with large photometric redshifts, even when limiting to those with the most confident fits (error class = 1). The distribution of classification probabilities for these sources is not shifted towards lower values, and we suggest it is likely due to the photometric redshift pipeline incorrectly identifying galaxies as stars.

between the quasar and galaxy clusters which still have high classification probabilities. This suggests that this unsupervised run of the metric-learning UMAP scheme is not the optimal separation in this 2-D space. In the supervised case the probabilities mainly only decrease in the region between the three clusters of sources, though there are noticeable regions within the cluster of galaxies where the probabilities are lower. The standard deviation of the classification probabilities varies significantly throughout the galaxy and quasar clusters, but remains low for most of the cluster of stars in both supervised and unsupervised

Fig. 20. Classification probabilities for 111 million newly classified sources as a function of the *resolved_r* parameter. One standard deviation is shaded, though note that the distribution is not Gaussian (as shown in Fig. 17) and as such one standard deviation can exceed a probability of 1. Dashed lines show normalised histograms per class to highlight the source density as a function of the *resolved_r* parameter.

cases. Figure 23 shows the same UMAP projections coloured by features, as in Fig. 14, but with two orders of magnitude more sources much more of the 2-D space is filled in.

The metric-learning schemes (Fig. 22) and direct application of UMAP to the photometric data in un-supervised or semi-supervised schemes (Fig. 21) achieve similar results, but appear to have biases towards different classes. When applying UMAP directly to the photometric data in a semi-supervised scheme including photometric data (the right plot in Fig. 21), UMAP effectively separate stars from galaxies. However, in the unsupervised metric-learning scheme (Fig. 22) there is significant contamination from stars in the galaxy and quasar clusters. Quasars

Fig. 21. UMAP applied to various subsets of data, reducing the original ten features to the two which are plotted. Sources are binned per pixel, with colours combined proportional to how many of each class are in that pixel bin. The brightness corresponds to the total source count in that pixel on a logarithmic scale. Source labels for galaxies (green), quasars (pink), and stars (blue) were derived from our random forest model. *Left:* 11 million photometrically observed sources without spectra run in an unsupervised scheme. *Middle:* all spectroscopically observed sources in our dataset, plus 11 million photometrically observed sources, run in an unsupervised scheme. *Right:* same as the middle plot but with the class labels of spectroscopically observed sources passed to the UMAP algorithm in a semi-supervised scheme. Including the spectroscopically observed sources in unsupervised (middle plot), and semi-supervised schemes (right plot) UMAP helps UMAP separate out the classes without them overlapping (particularly the minority class of quasars). Furthermore, our labels assigned from the random forest model show that even when classes overlap in this 2-D space, the structures picked out by UMAP are consistent with our class labels.

are most effectively separated in the supervised metric-learning schemes. Whilst the un-supervised and semi-supervised implementations only use 10% of the photometric sources and the metric-learning scheme uses all photometric sources, the same result is achieved when limiting the metric-learning method to the same subset of sources.

Overall, the labels assigned to photometrically observed sources by the random forest model are consistent with the clusters derived from UMAP. The metric-learning scheme gives the most consistent and interpretable results with the best separation of the classes. Whilst much slower, applying UMAP directly to a subset of the photometric dataset can be effective, and its effectiveness at separating clusters increases when including spectroscopically observed sources in an unsupervised scheme, and further increases in a semi-supervised scheme. The metric-learning scheme has the advantage that it can be efficiently scaled to embed any number of previously unseen sources using a parallel or distributed computing system, making it a viable algorithm to use with the next generation of surveys from telescopes such as LOFAR, SKA and LSST. While we could use UMAP as a pre-processing tool in anticipation of a supervised clustering scheme in two dimensions, for example a nearest neighbour approach, this will always be less accurate than a supervised learning scheme in the higher 10-D feature space. Hence, our use of UMAP is primarily as a diagnostic tool for this dataset, in particular for visualising the effectiveness of the classification labels assigned to photometrically observed sources by the random forest.

4.3. Catalogue description

The catalogue resulting from this work contains 111 395 468 labelled sources. These sources all have associated WISE matches but do not have associated spectra in SDSS. We provide the catalogue under a digital object identifier³. The whole catalogue is in a single Pandas (McKinney 2010) data frame, saved as

a Pickle file (.pkl). We also provide it as Pandas data frames per class, which are sorted by classification probability in descending order. The column names are described in Table 2. An extract from the full catalogue is shown in Appendix B for illustration. There are 50 417 547 galaxies, 2 137 839 quasars, and 58 840 082 stars. Of these, 6 683 526 galaxies (13%), 330 666 quasars (15%) and 41 279 349 stars (70%) have classification probabilities greater than 0.99. At the same time, 35 075 918 galaxies (70%), 722 159 quasars (34%), and 54 673 689 stars (93%) with classification probabilities greater than 0.9.

5. Discussion

5.1. Feature selection

We have used a set of ten features that are readily available in order to predict class labels for a significant fraction of SDSS sources. As random forests compare random sets of features in each decision tree estimator, they effectively interpret colours in different bands naturally throughout the forest. This method helps remove bias and over-fitting, which could arise when only using a few colours to identify objects. Furthermore, utilising all SDSS and WISE wavebands, and the *resolved_r* parameter together in this way provides a greater confidence in the classifications over using selected colours alone. We observe that removing absolute magnitude dependence from the model, i.e. taking the difference between the photometry from each band and the SDSS *r*-band photometry, resulting in a set of nine total features, did not alter the results.

The feature rankings returned from the random forest show that the *resolved_r* parameter separates the classes very well. This is also shown by the histograms in the right panel of Fig. 4. The next best features are the SDSS magnitudes, followed by the WISE magnitudes, although the W4 band provides very little additional information in this classification scheme. Whilst higher signal to noise colours can be calculated using *model-Mag*, we found that this made no difference to the performance of our classifier. The higher signal to noise colours calculated

³ <https://www.doi.org/10.5281/zenodo.3459293>

Fig. 22. UMAP applied to all 111 million photometrically observed sources in a metric-learning scheme, reducing the number of dimensions (features) from ten to two. UMAP models were obtained in unsupervised (*top left*) and supervised (*top right*) schemes using the spectroscopic training dataset (as in Fig. 13), and then all 111 million photometrically observed sources were embedded into this 2-D space. Colours were added afterwards for galaxies (green), quasars (pink), and stars (blue), derived from the random forest model trained on the same spectroscopically observed sources. Source are binned per pixel, and colours are combined proportional to how many of each class are in that pixel bin. The brightness corresponds to the total source count in that pixel on a logarithmic scale. UMAP effectively separates the classes in the same way as with the spectroscopic dataset (Fig. 13) and is in a strong agreement with the random forest. The bottom row shows the mean and standard deviation of the random forest classification probabilities.

from *modelMag* are therefore likely not significant compared to the ensemble approach of biased estimators from the random forest.

In our classification scheme, where we do not know if a source is galactic (a star) or extra-galactic (a galaxy or quasar), we cannot apply galactic extinction corrections to sources with unknown class labels, such as our photometric catalogue. Whilst this was a small correction for our dataset and made no difference to the results, we envisage that accounting for galactic extinction will be an important consideration for surveys finding quasars near the galactic plane where extinction would be significant (Masci 1998). Furthermore, the same problem will occur for quasars at high redshifts due to extinction from the intergalactic medium. Such dust-obscured quasars will appear much redder in the SDSS bands (or not detected altogether) although they will be rare in our dataset of optically selected quasars. Furthermore, the inclusion of the WISE bands helps significantly

when selecting quasars (shown by Table 1), and such longer wavelength observations will be important for classifying dust-obscured quasars.

SDSS provide a binary classification of star or galaxy depending on their *resolved_r* parameter being greater or less than 0.145 (Aguado et al. 2019). Due to its simplicity, in many cases this classification is incorrect (for example the middle right panel of Fig. 4) and will mainly incorrectly identify unresolved galaxies as stars. Overall, our method of leveraging information from the photometry alongside the *resolved_r* parameter provides a more accurate classification label, along with a quantified estimation of its probability.

5.2. Resolution differences

The lower resolution of the WISE data compared to the SDSS data allows for multiple SDSS sources to be matched to the

Fig. 23. Photometric dataset of 111 million sources processed with UMAP to reduce from ten features to two in supervised (*top row*) and unsupervised (*bottom row*) schemes, using models trained on the spectroscopic test dataset. The resulting two dimensions are plotted, with sources binned per pixel and coloured by the *resolvedr* parameter (*left column*), an SDSS colour $[\text{PSF } u - \text{PSF } z]$ (*middle column*) and a WISE colour $|w1 - w2|$ (*right plots*). Structures are as in Fig. 14, but the plot is more densely populated.

same WISE counterpart. When two or more SDSS sources are matched with the same WISE counterpart, they are never more than 8 arcsec apart, with a mean separation of 2.36 arcsec in the spectroscopic dataset and 2.40 arcsec in the photometric dataset. In our spectroscopic dataset this problem is minor, with 13 904 sources having a WISE counterpart that is matched to more than one source. This occurs for a very low fraction of the training data and is not significant enough to bias the model. Table 1 shows that including the WISE data as a feature improves the classification of quasars and stars (with an increase in the F1 score of 0.016) more so than galaxies (with an increase in F1 score of 0.004). In our photometric dataset there are 14 768 549 sources that have a WISE counterpart matched to more than one SDSS source, with 1.4 million of these having WISE counterpart matched to more than two sources. In such cases the WISE emission could be from any one of the SDSS sources, or a combination of all of them. The mean *resolvedr* parameter for these sources is 0.18, indicating they are slightly resolved sources, potentially more likely to be interacting if at a mean separation of 2.40 arcsec. We make no attempt to interpret if such sources are associated or not in our models.

In the datasets we have used, the resolution difference between SDSS and WISE is relatively small. Our model is dominated by information from the higher resolution SDSS data (see

Fig. 3), and in cases where there is incorrect WISE data that suggests an alternate class label, the overall classification probability will be lower. When using data from various surveys with much larger resolution differences and implementing more complex classification schemes, if the information from a lower resolution survey is the only way to make a confident source classification then this would lead to a substantial number of false positives. Additional features that quantify the likelihood of a correct match would help in such scenarios. Finally we note that we do not account for the chance alignment of a WISE source on top of a single SDSS source where the two are not of the same origin. Such a scenario becomes more problematic as the resolution difference between surveys increases, and the source density increases.

5.3. Class imbalance

In our new labelled photometric catalogue there are 23 times more galaxies, 6 times more quasars and 115 times more stars than in the spectroscopic dataset we used for training and testing. It is notable that the ratio between the three classes is different in the spectroscopic catalogue used for training compared to the resulting photometric catalogue. This bias is expected as a result of spectroscopic surveys prioritising particular sources at different magnitude depths. For example there is greater

Table 2. Description of catalogue columns for our 111 million classified objects.

Column name	Description
<i>objid</i>	SDSS object ID from the <i>PhotoPrimary</i> table
<i>ra</i>	Right Ascension
<i>dec</i>	Declination
<i>psf_u</i>	SDSS PSF magnitude in the <i>u</i> -band
<i>psf_g</i>	SDSS PSF magnitude in the <i>g</i> -band
<i>psf_r</i>	SDSS PSF magnitude in the <i>r</i> -band
<i>psf_i</i>	SDSS PSF magnitude in the <i>i</i> -band
<i>psf_z</i>	SDSS PSF magnitude in the <i>z</i> -band
<i>w1</i>	WISE band 1 magnitude
<i>w2</i>	WISE band 2 magnitude
<i>w3</i>	WISE band 3 magnitude
<i>w4</i>	WISE band 4 magnitude
<i>resolvedr</i>	A measure of how extended a source is: $ psf_r - cmod_r $
<i>class_pred</i>	Predicted class
<i>class_prob_galaxy</i>	Probability of prediction as a galaxy
<i>class_prob_quasar</i>	Probability of prediction as a quasar
<i>class_prob_star</i>	Probability of prediction as a star

scientific demand for galaxy spectra over those of stars, and a rare occurrence of bright quasars to target. It is therefore expected that the ratio of the three classes will be different in our new photometric catalogue compared to the ratio in the spectroscopic sources training data. Figure 2 demonstrates that the model still performs well when trained on several orders of magnitude fewer sources, and that sub-sampling the galaxy class hinders the model. However, under-sampling any classes in such a way that restricts the distribution in magnitude space does affect the results, as is described in Sect. 3.4.3. We note that training the random forest model on only 1% or 10% of the spectroscopically observed sources still results in the same ratio of source classifications in the resulting catalogue, though the classification probabilities are slightly lower.

5.4. Spectroscopic follow-up of quasars

Our new catalogue contains 2.1 million quasars, which is four times the 526 356 spectroscopically and visually confirmed quasars previously identified from SDSS (Pâris et al. 2018). Spectroscopic follow-up observations of these new quasars could be done most efficiently by prioritising those with high classification probabilities. In particular there are 330 666 sources that have classification probabilities greater than 0.99. As an increasing number of quasars are spectroscopically confirmed in future observations, these could be progressively included into the training set to improve the model's performance. As discussed in Sect. 3.4.4, increasing the density of sources in the 10-D feature space by spectroscopically confirming more quasars will increase the random forest classification probabilities for other quasars in similar areas of the feature space. This would be particularly effective at fainter magnitudes where the training set has lower source counts. We note one example seen in the literature where a spectroscopic survey of hot white dwarf stars confirmed a quasar classification in our new catalogue (Fix et al. 2015), with an *r*-band magnitude of 17.32 ± 0.01 .

5.5. Future surveys using machine learning

The next generation of telescopes performing large surveys are expected to take observations over many years before

completion, using increasingly more efficient and sensitive instruments. Given the cadence of data acquisition and processing, waiting until each survey is complete before building a classification pipeline will not maximise the scientific output from the instrument to the community. Instead, an iterative scheme of re-training and updating an existing model as more data become available will be very important for the telescopes which aim to survey large portions of the sky over many years. Furthermore, the efficiency of such an approach would be improved were models developed specifically with the aim of increasing a training set gradually over the course of the survey. In such schemes, it is very important to have metrics that can assess the reliability of classifications as a function of multiple variables that change over the lifetime of the survey. In this paper, we demonstrate some of these aspects, providing classification metrics as a function of feature space, as well as variables not included as features, such as magnitude error and redshift.

6. Conclusions

We have trained a random forest machine learning model on 1.55 million spectroscopically confirmed sources from SDSS in order to classify sources as galaxies, quasars, and stars. As features we have used photometry from both SDSS and WISE bands, plus a measure source extension. We have used cross-validation to tune the hyper-parameters in the model, and ensure that the class imbalance in the training data does not affect the performance of the model. Using a test dataset of a further 1.55 million spectroscopically confirmed sources we determine that the random forest achieves F_1 scores of 0.990, 0.953, and 0.977 for galaxies, quasars, and stars, respectively. Precision, recall, and F_1 score are also derived as a function of individual features as well as other parameters, in order to illustrate how the performance of the classifier varies across the three classes. We use the classification probabilities from the random forest as a measure of the likelihood of an individual classification, and we show that this value is in agreement with the F_1 score derived from a nearest neighbour search around each source in one and ten dimensions.

We applied the random forest model to 111 395 468 previously unlabelled photometrically observed sources from SDSS. Our model returns 50 417 547 galaxies, 2 137 839 quasars, and 58 840 082 stars. Each source has an associated classification probability per class. 6 683 526 galaxies (13%), 330 666 quasars (15%), and 41 279 349 stars (70%) have classification probabilities greater than 0.99. At the same time, 35 075 918 galaxies (70%), 722 159 quasars (34%), and 54 673 689 stars (93%) have classification probabilities greater than 0.9.

As a way to validate our model we use a non-linear dimension reduction technique (UMAP) in supervised, unsupervised, and semi-supervised schemes to reduce the number of features from ten to two, in order that the result can be plotted and the class labels visualised. In all three of these schemes UMAP clearly separates the galaxy, quasar, and star classes in two dimensions for the spectroscopic sources.

We applied UMAP to 10% of our new photometric catalogue of labelled sources (11 million sources) in three different schemes: (i) we used UMAP in an unsupervised scheme, where it is in a strong agreement with the star and galaxy class labels assigned by the random forest, but is poor at separating quasars; (ii) we used UMAP in an unsupervised scheme, but also including 3.1 million spectroscopically observed sources (without labels) in the dataset. This gives a superior performance, separating quasars from stars and galaxies more effectively; (iii) we

use UMAP in a semi-supervised scheme, including 3.1 million spectroscopically labelled sources in the dataset. This gives a further improvement in the result, clearly separating the classes. When the three classes overlap in any of these three scenarios, the labels from the random forest show coherent structures per class, demonstrating that in the original 10-D space there is a clearer separation of the classes than in the 2-D space. However, these three schemes are very computationally expensive.

We further utilised the ability of UMAP to perform metric-learning by training it on 1.55 million spectroscopically labelled sources in supervised and unsupervised schemes to embed the dataset into two dimensions. We used the resulting models to embed the test dataset of 1.55 million spectroscopically labelled sources, and all 111 million unlabelled photometric sources into this 2-D space. The supervised metric-learning scheme gave the optimal performance, clearly separating out all galaxies, quasars, and stars in a 2-D space. Furthermore, this clustering was in strong agreement with the random forest class labels. A significant advantage of the metric-learning schemes is that it can be efficiently scaled in a parallel or distributed computing system to embed any number of unlabelled sources using a model built from a significantly smaller, partially labelled training dataset. This makes it a viable algorithm to run on the next generation of surveys (e.g. LSST, SKA, LOFAR) for both classification and data visualisation.

Acknowledgements. The authors gratefully acknowledge support from the UK Research & Innovation Science & Technology Facilities Council (UKRI-STFC). We thank Justin Bray for numerous discussions helping to guide this work. We thank the anonymous referee for providing feedback that has improved this paper. We thank the language editors at the journal for providing corrections. We made use of the Python packages Numpy (Oliphant 2006), Scikit-learn (Pedregosa et al. 2011), Matplotlib (Hunter 2007), Pandas (McKinney 2010) and Dataloader: <https://datashader.org/>.

References

- Aguado, D. S., Ahumada, R., Almeida, A., et al. 2019, *ApJS*, **240**, 23
- Angthopo, J., Ferreras, I., & Silk, J. 2019, *MNRAS*, **488**, L99
- Antonucci, R. 2012, *Astron. Astrophys. Trans.*, **27**, 557
- Bai, Y., Liu, J., Wang, S., & Yang, F. 2019, *AJ*, **157**, 9
- Baldry, I. K., Robotham, A. S. G., Hill, D. T., et al. 2010, *MNRAS*, **404**, 86
- Beck, R., Dobos, L., Budavári, T., Szalay, A. S., & Csabai, I. 2016, *MNRAS*, **460**, 1371
- Begelman, M. C., Blandford, R. D., & Rees, M. J. 1984, *Rev. Mod. Phys.*, **56**, 255
- Bolton, A. S., Schlegel, D. J., Aubourg, É., et al. 2012, *AJ*, **144**, 144
- Burbidge, G. R., Burbidge, E. M., & Sandage, A. R. 1963, *Rev. Mod. Phys.*, **35**, 947
- Carrasco, D., Barrientos, L. F., Pichara, K., et al. 2015, *A&A*, **584**, A44
- Fix, M. B., Smith, J. A., Tucker, D. L., Wester, W., & Annis, J. 2015, *Astron. Nachr.*, **336**, 614
- Francis, P. J., Hewett, P. C., Foltz, C. B., et al. 1991, *ApJ*, **373**, 465
- Greenstein, J. L., & Matthews, T. A. 1963, *AJ*, **68**, 279
- Greenstein, J. L., & Schmidt, M. 1964, *ApJ*, **140**, 1
- Gürkan, G., Hardcastle, M. J., Best, P. N., et al. 2019, *A&A*, **622**, A11
- Herschel, W. 1789, *Phil. Trans. R. Soc. London Ser. I*, **79**, 212
- Hubble, E. P. 1929, *ApJ*, **69**, 103
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Hutsemékers, D., Cabanac, R., Lamy, H., & Sluse, D. 2005, *A&A*, **441**, 915
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- Jarvis, M., Bacon, D., Blake, C., et al. 2015, *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 18
- Jones, N. 2014, *Nature*, **505**, 146
- Kang, S.-J., Fan, J.-H., Mao, W., et al. 2019, *ApJ*, **872**, 189
- Kauffmann, G., & Haehnelt, M. 2000, *MNRAS*, **311**, 576
- Leistedt, B., & Peiris, H. V. 2014, *MNRAS*, **444**, 2
- Lintott, C. J., Schawinski, K., Slosar, A., et al. 2008, *MNRAS*, **389**, 1179
- Lintott, C., Schawinski, K., Bamford, S., et al. 2011, *MNRAS*, **410**, 166
- Louppe, G. 2014, ArXiv e-prints [arXiv:1407.7502]
- LSST Dark Energy Science Collaboration 2012, ArXiv e-prints [arXiv:1211.0310]
- LSST Science Collaboration (Abell, P. A., et al.) 2009, ArXiv e-prints [arXiv:0912.0201]
- Masci, F. J. 1998, PhD thesis [arXiv:astro-ph/9801181]
- Matthews, T. A., & Sandage, A. R. 1963, *ApJ*, **138**, 30
- McDonald, P., & Eisenstein, D. J. 2007, *Phys. Rev. D*, **76**, 063009
- McInnes, L., Healy, J., & Melville, J. 2018, ArXiv e-prints [arXiv:1802.03426]
- McKinney, W., et al. 2010, *Proceedings of the 9th Python in Science Conference*, 445, 51 Austin, TX
- Morice-Atkinson, X., Hoyle, B., & Bacon, D. 2018, *MNRAS*, **481**, 4194
- Mosteller, F., & Tukey, J. W. 1968, in *Handbook of Social Psychology*, eds. G. Lindzey, & E. Aronson (Addison-Wesley), 2
- Nakoneczny, S., Bilicki, M., Solarz, A., et al. 2019, *A&A*, **624**, A13
- Nikutta, R., Hunt-Walker, N., Nenková, M., Ivezić, Ž., & Elitzur, M. 2014, *MNRAS*, **442**, 3361
- Oliphant, T. E. 2006, *A Guide to NumPy* (USA: Trelgol Publishing), 1
- Opik, E. 1922, *ApJ*, **55**, 406
- Pâris, I., Petitjean, P., Aubourg, É., et al. 2018, *A&A*, **613**, A51
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, **12**, 2825
- Peters, C. M., Richards, G. T., Myers, A. D., et al. 2015, *ApJ*, **811**, 95
- Pratt, L. Y., Mostow, J., Kamm, C. A., & Kamm, A. A. 1991, *AAAI*, **91**, 584
- Rauch, M. 1998, *ARA&A*, **36**, 267
- Rees, M. J. 1984, *ARA&A*, **22**, 471
- Salim, S. 2014, *Serb. Astron. J.*, **189**, 1
- Sanders, D. B., Soifer, B. T., Elias, J. H., et al. 1988, *ApJ*, **325**, 74
- Schindler, J.-T., Fan, X., Huang, Y.-H., et al. 2019, *ApJS*, **243**, 5
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, **500**, 525
- Schmidt, M. 1963, *Nature*, **197**, 1040
- Schmidt, M., & Green, R. F. 1983, *ApJ*, **269**, 352
- Scranton, R., Ménard, B., Richards, G. T., et al. 2005, *ApJ*, **633**, 589
- Shimwell, T. W., Tasse, C., Hardcastle, M. J., et al. 2019, *A&A*, **622**, A1
- Smith, H. J., & Hoeffleit, D. 1961, *AJ*, **70**, 295
- Stoughton, C., Lupton, R. H., Bernardi, M., et al. 2002, *AJ*, **123**, 485
- Tang, H., Scaife, A. M. M., & Leahy, J. P. 2019, *MNRAS*, **488**, 3358
- Urry, C. M., & Padovani, P. 1995, *PASP*, **107**, 803
- Vanden Berk, D. E., Richards, G. T., Bauer, A., et al. 2001, *AJ*, **122**, 549
- van der Maaten, L., & Hinton, G. 2008, *J. Mach. Learn. Res.*, **9**, 2579
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, **556**, A2
- Weedman, D. W. 1977, *ARA&A*, **15**, 69
- Wilson, E. B. 1927, *J. Am. Stat. Assoc.*, **22**, 209
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, **140**, 1868
- Wu, C., Buyya, R., & Ramamohanarao, K. 2016, ArXiv e-prints [arXiv:1601.03115]

Appendix A: Supplementary plots

Fig. A.1. Correct and missed spectroscopically observed sources per class when including the *resolved_r* parameter as a feature (*top*) and when excluding it as a feature (*bottom*). Both use SDSS PSF magnitudes and WISE magnitudes as features. The F_1 score per class increases from 0.988 to 0.991 (galaxy), 0.942 to 0.952 (quasar), 0.974 to 0.978 (star).

Fig. A.2. Histogram of each magnitude feature, per class, for correct and misclassified sources from the random forest model applied to the test dataset of 1.55 million spectroscopically confirmed sources. Galaxies are shown in green, quasars in pink and stars in blue.

Fig. A.3. Precision, recall, and F_1 score plotted as a function of each magnitude feature per class, as derived from the random forest model applied to the test dataset of 1.55 million spectroscopically confirmed sources. In the *top row*, a histogram of each magnitude feature is also shown as a dotted line per class (normalised to one third) to show the source density. In the *bottom row* the classification probabilities from the random forest model for the assigned classes are shown by dashed lines. Galaxies are shown in green, quasars in pink, and stars in blue.

Fig. A.4. Examples of galaxies missed as quasars. The spectroscopically observed target is at the centre in a red box, and photometrically observed sources are circled in blue. From top left to bottom right in order of how resolved they are, their SpecObjIDs, $resolved_r$, parameters and redshifts are: 527093262050158592, 1.627 ($z = 0.162$); 337809312191637504, 1.298 ($z = 0.186$); 3087234932888594432, 1.137 ($z = 0.094$); 658795790411524096, 1.129 ($z = 0.199$); 2976929452103067648, 1.064 ($z = 0.303$); 1928687220470867968, 0.919 ($z = 0.162$); 5492327033231876096, 0.742 ($z = 0.882$); 1780082440165943296, 0.296 ($z = 0.280$); 4946128968019320832, 0.090 ($z = 0.640$); 7990778382286561280, 0.083 ($z = 0.460$); 9339371187510157312, 0.030 ($z = 1.078$) and 1621402570222233600, 0.022 ($z = 0.051$)

Fig. A.5. Examples of galaxies missed as stars. The spectroscopically observed target is at the centre in a red box, and photometrically observed sources are circled in blue. From top left to bottom right in order of how resolved they are, their SpecObjIDs, $resolved_r$, parameters and redshifts are: 892985450690537472, 4.685 ($z = 0.003$); 3383360912827195392, 2.357 ($z = 0.0$), 5584568912597458944, 1.858 ($z = 0.724$); 3340604199258843136, 0.386 ($z = 0.038$); 2402769422277175296, 0.120 ($z = 0.089$) and 686843782766815232, 0.004 ($z = 0.656$). Note that the second source is actually a real star, incorrectly labelled by the SDSS classification pipeline.

Fig. A.6. Examples of stars missed as galaxies. The spectroscopically observed target is at the centre in a red box, and photometrically observed sources are circled in blue. The top row are resolved sources with SpecObjIDs and $resolved_r$ parameters of: 1142953103180457984, 2.323; 1190170526966900736, 1.609; 2430897686078056448, 1.133; 2660651076583188480, 5.279; 3141329535987902464, 1.498 and 5382020452850835456, 1.897. Top bottom row are unresolved sources with SpecObjIDs and $resolved_r$ parameters of: 9568057781371510784, 0.268; 9929352360244518912, 0.063; 4543105995102986240, 0.294; 4725488312604663808, 0.270; 4728959746027003904, 0.286 and 9197691961909157888, 0.097.

Fig. A.7. Examples of quasars missed as stars. The spectroscopically observed target is at the centre in a red box, and photometrically observed sources are circled in blue. Their SpecObjIDs and redshifts from left to right are: 1578634038246139904 ($z = 4.450$), 4316783906617532416 ($z = 2.257$), 6092406673766326272 ($z = 3.087$), 6434584599746027520 ($z = 3.998$), 7539033580141256704 ($z = 2.995$), 7595591358561099776 ($z = 1.818$).

Fig. A.8. Histogram of each magnitude band for all 111 million newly classified sources. Galaxies are plotted in green, quasars in pink, and stars in blue.

Appendix B: Catalogue extract

Table B.1. Sample of 20 galaxies, 20 quasars, and 20 stars from our new catalogue of classified sources.

Objid	RA	Dec	psf_u	psf_g	psf_r	psf_i	psf_z	w1	w2	w3	w4	resolvdrr	class_pred	class_prob_galaxy	class_prob_quasar	class_prob_star
1237663782600311620	30.97881	-1.1961	24.05239	22.94274	21.72948	20.89259	20.18513	17.004	16.64	13.011	9.35	0.17654	GALAXY	0.85	0.035	0.115
1237662224094987749	253.04789	16.46796	23.38386	23.02948	21.74218	20.1817	20.92808	17.121	16.255	12.591	9.03	0.10353	GALAXY	0.995	0.005	0.0
1237667212670927178	169.1246	29.96615	23.87665	23.06762	21.51385	20.906	20.21203	17.263	15.922	12.425	8.955	0.13124	GALAXY	0.845	0.055	0.1
123765706633045647	19.49697	-2.5862	24.27969	23.02265	21.38321	20.85167	20.26984	15.531	15.433	12.627	9.001	1.10773	GALAXY	0.995	0.0	0.005
1237679439352693622	37.94041	-1.12846	24.72696	23.59927	22.73186	20.93435	21.46056	17.329	16.524	12.602	9.573	0.9788	GALAXY	0.99	0.02	0.0
1237663916799493064	114.34315	45.38042	22.7852	21.79741	21.06585	20.85956	20.17542	17.015	16.844	12.216	8.763	1.0378	GALAXY	0.98	0.0	0.01
123766865898258502	263.24811	3.58866	25.02485	21.90447	20.5326	20.00715	19.70523	17.191	17.006	12.799	9.292	0.57036	GALAXY	0.63	0.015	0.355
1237676440385880795	2.63491	30.80671	22.69767	23.41565	22.46992	20.44775	17.79413	17.484	16.698	11.819	8.791	0.03125	GALAXY	0.605	0.355	0.04
1237659324960998467	255.46213	29.25622	22.68677	22.64796	21.78658	20.11527	20.73625	16.631	16.39	12.614	9.12	0.8028	GALAXY	1.0	0.0	0.0
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366	16.746	12.97	9.464	0.25505	GALAXY	0.895	0.07	0.035
1237665566072701160	220.53691	20.68538	22.41847	20.729	19.7746	19.33793	19.02539	16.111	15.915	13.398	9.211	0.53491	GALAXY	0.98	0.0	0.02
1237661069789889582	160.9454	13.80896	23.55087	23.38648	22.26917	20.52199	20.7818	16.131	15.653	11.989	8.35	0.39083	GALAXY	0.96	0.035	0.005
1237654398616732586	172.03828	63.61214	23.22376	23.29302	22.82913	20.72159	20.72159	17.037	16.951	13.063	9.291	0.82094	GALAXY	0.99	0.01	0.0
1237661950270506022	222.08511	10.39907	24.16395	24.06727	22.9681	22.29608	21.48033	17.512	17.092	13.105	9.481	0.36943	GALAXY	0.885	0.1	0.015
1237648721250354405	224.83936	-0.17729	25.11269	24.19453	22.86048	21.46819	20.51785	16.366								