



LJMU Research Online

Cook, TL, Bandi, B, Philipsborn, S, Loveday, J, Bellstedt, S, Driver, SP, Robotham, ASG, Bilicki, M, Kaur, G, Tempel, E, Baldry, I, Gruen, D, Longhetti, M, Iovino, A, Holwerda, BW and Demarco, R

Wide Area VISTA Extra-galactic Survey (WAVES): unsupervised star-galaxy separation on the WAVES-Wide photometric input catalogue using UMAP and HDBSCAN

<http://researchonline.ljmu.ac.uk/id/eprint/25074/>

Article

Citation (please note it is advisable to refer to the publisher's version if you intend to cite from this work)

Cook, TL, Bandi, B, Philipsborn, S, Loveday, J, Bellstedt, S, Driver, SP, Robotham, ASG, Bilicki, M, Kaur, G, Tempel, E, Baldry, I, Gruen, D, Longhetti, M, Iovino, A, Holwerda, BW and Demarco, R (2024) Wide Area VISTA Extra-galactic Survey (WAVES): unsupervised star-galaxy separation

















LJMU has developed [LJMU Research Online](#) for users to access the research output of the University more effectively. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LJMU Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain.

The version presented here may differ from the published version or from the version of the record. Please see the repository URL above for details on accessing the published version and note that access may require a subscription.

For more information please contact researchonline@ljmu.ac.uk

<http://researchonline.ljmu.ac.uk/>

Wide Area VISTA Extra-galactic Survey (WAVES): unsupervised star-galaxy separation on the WAVES-Wide photometric input catalogue using UMAP and HDBSCAN

Todd L. Cook ¹, ¹★ Behnood Bandi ¹, Sam Philipsborn,¹ Jon Loveday ¹, Sabine Bellstedt ², Simon P. Driver ², Aaron S. G. Robotham ², Maciej Bilicki ³, Gursharanjit Kaur ³, Elmo Tempel ^{4,5}, Ivan Baldry ⁶, Daniel Gruen ^{7,8}, Marcella Longhetti ⁹, Angela Iovino ⁹, Benne W. Holwerda ¹⁰ and Ricardo Demarco ¹¹

¹Astronomy Centre, University of Sussex, Falmer, Brighton BN1 9QH, UK

²International Centre for Radio Astronomy Research (ICRAR), M468, University of Western Australia, Crawley, WA 6009, Australia

³Center for Theoretical Physics, Polish Academy of Sciences, Aleja Lotników 32/46, PL-02-668 Warsaw, Poland

⁴Tartu Observatory, University of Tartu, Observatooriumi 1, 61602 Tõravere, Estonia

⁵Estonian Academy of Sciences, Kohtu 6, 10130 Tallinn, Estonia

⁶Astrophysics Research Institute, Liverpool John Moores University, IC2, Liverpool Science Park, 146 Brownlow Hill, Liverpool L3 5RF, UK

⁷University Observatory, Faculty of Physics, Ludwig-Maximilians-Universität München, Scheinerstr 1, D-81679 Munich, Germany

⁸Excellence Cluster ORIGINS, Boltzmannstr 2, D-85748 Garching, Germany

⁹INAF – Osservatorio Astronomico di Brera, via Brera 28, I-20121 Milano, Italy

¹⁰Department of Physics and Astronomy, University of Louisville, Natural Science Building 102, Louisville, KY 40292, USA

¹¹Institute of Astrophysics, Facultad de Ciencias Exactas, Universidad Andrés Bello, Sede Concepción, Talcahuano, Casilla 160-C, Chile

Accepted 2024 October 15. Received 2024 September 2; in original form 2024 May 30

ABSTRACT

Star-galaxy separation is a crucial step in creating target catalogues for extragalactic spectroscopic surveys. A classifier biased towards inclusivity risks including high numbers of stars, wasting fibre hours, while a more conservative classifier might overlook galaxies, compromising completeness and hence survey objectives. To avoid bias introduced by a training set in supervised methods, we employ an unsupervised machine learning approach. Using photometry from the Wide Area VISTA Extragalactic Survey (WAVES)-Wide catalogue comprising nine-band $u - K_s$ data, we create a feature space with colours, fluxes, and apparent size information extracted by PROFOUND. We apply the non-linear dimensionality reduction method UMAP (Uniform Manifold Approximation and Projection) combined with the classifier HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise) to classify stars and galaxies. Our method is verified against a baseline colour and morphological method using a truth catalogue from *Gaia*, SDSS (Sloan Digital Sky Survey), GAMA (Galaxy And Mass Assembly), and DESI (Dark Energy Spectroscopic Instrument). We correctly identify 99.75 per cent of galaxies within the AB magnitude limit of $Z = 21.2$, with an F1 score of 0.9971 ± 0.0018 across the entire ground truth sample, compared to 0.9879 ± 0.0088 from the baseline method. Our method's higher purity (0.9967 ± 0.0021) compared to the baseline (0.9795 ± 0.0172) increases efficiency, identifying 11 per cent fewer galaxy or ambiguous sources, saving approximately 70 000 fibre hours on the 4MOST (4-m Multi-Object Spectroscopic Telescope) instrument. We achieve reliable classification statistics for challenging sources including quasars, compact galaxies, and low surface brightness galaxies, retrieving 92.7 per cent, 84.6 per cent, and 99.5 per cent of them, respectively. Angular clustering analysis validates our classifications, showing consistency with expected galaxy clustering, regardless of the baseline classification.

Key words: methods: data analysis – catalogues – surveys – galaxies: photometry – large-scale structure of Universe.

1 INTRODUCTION

The classification of astronomical objects through their imaging is a key tool for astronomy. Spectroscopic surveys such as Sloan Digital

Sky Survey (SDSS),¹ Galaxy And Mass Assembly (GAMA),² and the upcoming surveys using the 4-m Multi-Object Spectroscopic

* E-mail: t.cook@sussex.ac.uk

¹Sloan Digital Sky Survey

²Galaxy and Mass Assembly survey

Telescope (4MOST)³ instrument (York et al. 2000; Driver et al. 2011; de Jong et al. 2019, respectively) all require an input catalogue of selected targets. These targets are generated through the analysis of prior imaging, with the star-galaxy classification of the targets being a crucial step. For an extragalactic spectroscopic survey, a target catalogue based on a more liberal classifier will result in a high number of stars, leading to wasted fibre-hours. A more conservative classifier on the other hand will cause the omission of more galaxies in the spectroscopic observations. Star-galaxy separation can be conducted by multiple different methods. Modern-day star-galaxy separation techniques can be split into colour, morphological, and machine learning methods. The DEVILS survey (Davies et al. 2018) utilizes NIR colours and surface brightness to filter stars from their target catalogue. The GAMA input catalogue (Baldry et al. 2010) utilizes SDSS imaging, and classifies sources into stars versus galaxies using a combination of profile fitting and colour separation. Morphological star-galaxy separation techniques are able to differentiate between extended sources (galaxies) and point sources. They cannot however differentiate between the nature of the point source, as they could be stars or quasi-stellar objects (QSOs). Morphological techniques such as Slater, Ivezić & Lupton (2020) and Soumagnac et al. (2015) show promise but show limited reliability at faint magnitudes. This is due to astronomical seeing ‘smearing’ psf-sized point sources such as stars and quasars, which is apparent even in space-based imaging (Holwerda et al. 2024).

For decades, star-galaxy separation has been seen as an exemplary classification task for supervised machine learning (Odewahn et al. 1993; Weir, Fayyad & Djorgovski 1995; Bertin & Arnouts 1996; Bailer-Jones, Fouesneau & Andrae 2019; Clarke et al. 2020; Baqui et al. 2021), with numerous models being used (neural networks, random forest, support vector machines, etc.). However, supervised machine learning requires prior training data, with classifications being most effective when the training data are abundant and representative of the test data. With that said, there are ways of retrieving unbiased results for classifications or parameter estimation given a small amount of training data (e.g. using active learning; Lochner & Bassett 2021; Stevens et al. 2021) or using a biased training set (e.g. Gruen & Brimiouille 2017).

We utilize unsupervised machine learning for our star-galaxy separation, without the use of any training data. Unsupervised machine learning (including dimensionality reduction and clustering) has the advantage of making use of the entire given catalogue to find underlying clusters. It also has the advantage of not being biased by a given training set, and is instead only dependent on the characteristics of the sample it is given. In this case, these include the depth of the survey, the elimination of artefacts, etc. Instead of training data, we use data in which stars and galaxies are identified as such to high significance (which we term ‘ground truth’ henceforth) purely for the purpose of validation. As we are only classifying stars and galaxies, it is easy to assign star or galaxy labels to the two clusters identified. This is as opposed to finding several distinct populations of sources such as in Siudek et al. (2018).

Unsupervised machine learning has been used in astronomy in the past. (A comprehensive recent review of unsupervised machine learning in astronomy can be found in Fotopoulou 2024.) In the past, it has also been used specifically for star-galaxy separation. Logan & Fotopoulou (2020) also use Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) for star-galaxy-QSO separation with a different pre-processing stage. Instead

of performing star-galaxy-QSO classification, we focus solely on distinguishing between stars and galaxies. This decision is based on WAVES’ avoidance of quasars in their target selection process. If quasars are mistakenly classified as stars, they will not be in the target catalogue. However, if they are correctly classified as galaxies, they should be filtered out by WAVES’ photometric redshift selection criteria and will not be part of the part of the target catalogue regardless. We investigate this in Section 7.1.

In this work, we aim to provide star-galaxy separation for the input catalogue for the Wide region of the WAVES survey, part of the 4MOST suite of surveys. Our goal is to enable WAVES to achieve the 95 per cent completeness required for its science goals with as few fibre hours as possible. First in Section 2, we summarize the photometric catalogue of the WAVES target catalogue. We then describe how we compiled a sample of ground truth data we use for the validation of our method, and analyse how it compares to the overall target catalogue. In Section 3, we describe a baseline star-galaxy separation method currently used in WAVES as a point of comparison for our method. In Section 4, we outline the pre-processing part of the method, including data cleaning, feature formation, feature scaling, and the dimensionality reduction which performs most of the ‘heavy lifting’ of the method. In Section 5, we explain the actual clustering method used to separate stars and galaxies in the reduced data. In Section 6, we measure the fidelity of our method using the ground truth sample we compiled using confusion matrices and analysing F1 score as a function of magnitude. In Section 7, we measure the effectiveness of our method at classifying galaxies that star-galaxy separation methods would find challenging. Finally in Section 8, we examine the two-point angular correlation function of our sources.

2 DATA

The data utilized in this work is comprised of the multiband photometric data on which we classify objects, and a compiled sample of ground truth data which we use for verification. The photometric data are described in Section 2.1, while the ground truth data are described in Section 2.2.

2.1 Photometric catalogue

For our photometric data, we use the parent photometric catalogue of the upcoming WAVES-Wide survey (Driver et al. 2019), a survey of the local Universe to be performed with the 4MOST spectroscopic instrument⁴ WAVES will use approximately 1.75 million low resolution fibre hours providing spectra at $R = 4000 - 7000$ covering the wavelength range from 370 - 950 nm. WAVES-Wide covers an area of $\sim 1170 \text{ deg}^2$ over its north and south regions, and contains $\sim 14\,800\,000$ sources within the $Z < 21.2$ magnitude limit of the parent catalogue.⁵ The north region lies on the equatorial plane and spans 157.25 to 225.0 deg in right ascension. The south region sits at -30 deg declination and spans -30 to 52.5 deg right ascension.

Catalogue construction will fully be outlined in Bellstedt et al. (in preparation), and is derived from deep optical and NIR imaging from VST KiDS⁶ (de Jong et al. 2013; Kuijken et al. 2019) (with the

³4-metre Multi Object Spectroscopic Telescope

⁴4MOST is observing different types of targets (e.g. stars, galaxies, and quasars) simultaneously as if they were one single survey (Tempel et al. 2020a, b)

⁵The actual WAVES magnitude limit will likely be brighter in both the WAVES and Deep regions.

⁶VLT Survey Telescope Kilo-Degree Survey

u , g , r , and i bands) and the VISTA VIKING⁷ survey (Edge et al. 2013) (with the Z , Y , J , H , and K_s bands), respectively. We note that far- and mid-infrared photometry also provide useful information for star-galaxy separation (Kovács & Szapudi 2015; Krakowski et al. 2016; Kurcz et al. 2016). Logan & Fotopoulou (2020) note that the use of $W1$ and $W2$ bands from *Wide-field Infrared Survey Explorer* (*WISE*⁸; Wright et al. 2010) in their unsupervised star-galaxy-quasar classifier are critical for their precision and accuracy. The Random Forest analysis of the features identifies $W1$ and $W2$ as the most important bands. However, Nakoneczny et al. (2021) using KiDS and VIKING can effectively classify quasars even without *WISE* data. We do not incorporate *WISE* data in this work because the corresponding observations are not deep enough compared to the KiDS and VIKING bands. After background subtraction, a large fraction of sources have negative *WISE* fluxes which are clearly unphysical, and cannot be converted to magnitudes. As demonstrated in further sections, we are confident this work has the capability to reliably classify faint targets without *WISE* data.

Source detection and characterization are carried out by the PROFOUND package (Robotham et al. 2018). An inverse variance weighted stack of $r + i + Z + Y$ bands is utilized for initial source detection. PROFOUND works by detecting and maintaining the original isophote (or segment) for a source, which can vary in shape from regular to irregular. Then, segment dilation occurs to derive pseudo-total fluxes. This involves progressively adding layers of pixels surrounding each segment until the flux reaches convergence. Segments for each source are generated and large, fragmented sources (see fig. 11 of Bellstedt et al. 2020) are manually regrouped after being flagged by the PROFOUND pipeline. The radius and ellipticity of each source is determined depending on the size and shape of each segment. The resulting flux and flux error are estimated for each band within each segment.

The Planck $E(B - V)$ extinction map is applied to the sources (Planck Collaboration IX 2013), correcting their flux for Galactic dust absorption. Finally, stars brighter than a G -band magnitude of 16.0 are removed, and all sources within a radius of $10^{1.6-0.15G}$ arcmin of these bright stars are masked out because their flux can affect the estimate of other sources' fluxes in the photometry, where G is the *Gaia* G -band magnitude. This results in brighter stars having a larger exclusion radius. After star masking, we are left with 14 802 032 sources within the $Z < 21.2$ magnitude limit that need to be classified.

The WAVES target catalogue will also be limited by a photometric redshift upper limit of $z < 0.2$ in the Wide fields and $z < 0.8$ in the Deep fields.⁹ Each source will have an associated photometric redshift, which are found by combining multiple photo- z methods (detailed in Bellstedt et al. in preparation) using the same KiDS + VIKING photometry. All star-galaxy separation conducted in this work is prior to the photo- z selection cuts.

2.2 Ground truth catalogue

To assess the performance of our classifier, we compile a catalogue of sources with known classification as a test set. As we cannot confidently infer ground truth classification from photometry, we rely

⁷Visible and Infrared Survey Telescope for Astronomy VISTA Kilo-degree Infrared Galaxy

⁸Wide-field Infrared Survey Explorer

⁹We do not address the Deep fields in this work, but our method is used for the Deep fields in the WAVES target selection.




Figure 1. Histogram of the radial difference between sources in the WAVES-Wide photometric input catalogue and our ground truth catalogue, with a line showing our 0.6 arcsec cross-match radius. The upper-right panel shows the difference in RA (WAVES – ground truth) and Dec., with a circle indicating the 0.6 arcsec cross-match radius.

on pre-existing spectroscopic data and *Gaia* parallax measurements in the WAVES-Wide fields. When matching sources between the WAVES input catalogue and the ground truth data, we use a 0.6 arcsec cross-match radius. This can be visualized in Fig. 1, which shows a histogram of the radial difference between sources in the WAVES-Wide photometric catalogue and sources in the ground truth catalogue. The upper right panel shows the differences in RA and Dec., showing a small systematic offset. We choose 0.6 arcsec as a conservative cross-match radius, preventing many potential spurious matches, as we prioritize purity in the ground truth data set. For WAVES equatorial coordinates, we use R_{max} and Dec_{max} generated from the PROFOUND package, which correspond to the position of the pixel which contains the greatest flux in the segment.

2.2.1 GAMA

The GAMA survey is an extragalactic spectroscopic survey conducted with the AAOmega wide-field facility on the Anglo Australian Telescope. GAMA consists of four $\sim 50 \text{ deg}^2$ equatorial fields located at 2^{h} (G02), 9^{h} (G09), 12^{h} (G12), 15^{h} (G15), and one field at 23^{h} and -32.5° (G23). Two of these fields lie within the WAVES-Wide North equatorial region, and G23 lies within the WAVES-Wide South region. GAMA is a strictly extragalactic survey. However, due to the imperfect star-galaxy separation used to define the target catalogue, a number of stars were observed. We use the fourth data release of GAMA (Driver et al. 2022), which also utilizes KiDS and VIKING photometry. We filter `gkvScienceCatv02`¹⁰ by $NQ > 2$, meaning a reliable spectroscopic redshift was measured. After a cross-match with our sample, we retrieve 180 831 galaxies, 17 082 stars, and 7 ambiguous sources that have been classified through GAMA's analysis of spectral features and redshift (Hopkins et al. 2013). We do not apply any additional redshift cuts in distinguishing stars and galaxies. We discard the seven ambiguous sources.

Due to the magnitude limit of GAMA ($r \lesssim 19.65$), the majority of the sample lies well below the $Z < 21.2$ magnitude limit of this work.

¹⁰<https://www.gama-survey.org/dr4/schema/>




Figure 2. Histograms of the ground truth data sourced for this analysis compared to the total number of sources needed to be classified in the WAVES-Wide regions (light grey) as a function of their Z-band magnitude. Red and blue regions (top and bottom rows) notate galaxies and stars, respectively. The different hatchings notate the different sources of ground truth data (GAMA, SDSS, DESI, and *Gaia*). The dark grey region indicates the total number of ground-truth sources. The ratio of ground truth data to the total number of WAVES sources drops off quickly at faint magnitudes.

This can be seen in the first column of plots in Fig. 2. GAMA provides a plethora of galaxies up to $Z < 19$, but the distribution drops off quickly, and no GAMA galaxies or stars lie at the magnitude limit of this work. However, GAMA has the advantage of being highly complete and is not biased against any particular galaxy type.

2.2.2 Dark Energy Spectroscopic Instrument

For sources fainter than the magnitude limit of GAMA, we utilize the Early Data Release (EDR) from the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration 2023). DESI is mounted on the 4-m Mayall telescope at Kitt Peak National Observatory in Arizona and will take spectra of roughly 40 million galaxies and quasars over its 5 yr programme over an area of $14\,000\text{ deg}^2$ (DESI Collaboration 2022). While DESI is primarily targeting the Northern sky, its footprint does overlap with the equatorial WAVES-Wide North region. We find 132 333 galaxies and 71 025 stars cross-matched to sources in the WAVES-Wide catalogue from DESI’s ‘Target Selection Validation’ (SV1) and ‘One-Percent Survey’ (SV3). We filter `zall-tilecumulative-edr-vac` by `OBJTYPE=TGT` which excludes faulty and sky fibres, and remove any sources with problematic spectra using `ZWARN==0`. We also remove galaxies with a redshift lower than $z < 0.0015$, and remove stars with a redshift higher than $z > 0.0015$ to prevent contamination.

The DESI galaxy sample contains three tracers: bright galaxies (Hahn et al. 2023) that are magnitude limited down to $r < 19.5$ and colour-selected in the magnitude range $19.5 < r < 20.175$, luminous red galaxies (Zhou et al. 2023) which are colour-selected down to Z -band fibre magnitude of $Z < 21.6$, and emission-line galaxies

(Raichoor et al. 2023) which are colour selected in the magnitude range $g > 20$ and g -band fibre magnitude of $g_{\text{fibre}} < 24.1$. The DESI Milky Way Survey (Prieto et al. 2020; Cooper et al. 2023) select *Gaia* colour-selected stars in the magnitude range $16 < r < 19$, and provide the majority of stars used for our sample. The magnitude distribution of DESI sources can be seen in the second column of Fig. 2. Visual inspection of DESI galaxies (Lan et al. 2023) and quasars (Alexander et al. 2023) reveal good classification of galaxies (~ 99 per cent purity), and reasonable classification of quasars (~ 86 per cent purity), although the spectroscopic classification pipeline for the EDR has been updated since. The relative depth of DESI is critical for the validation of the star-galaxy separation, as it provides spectra for sources all the way down to the magnitude limit of this analysis.

2.2.3 SDSS

We also utilize data from the 17th data release of SDSS (Abdurro’uf et al. 2022) for our validation. SDSS has a plethora of spectroscopically confirmed Galactic and extragalactic sources in the WAVES-Wide North equatorial region, with some reaching out to the magnitude limit of $Z < 21.2$. We filter our SDSS sample by `ZWARNING==0`, implying that a reliable redshift measurement has been found. We are able to cross-match 123 730 sources with SDSS DR17. Of these sources, 41.6 per cent are galaxies from the SDSS main survey (Strauss et al. 2002) and 52.0 per cent are galaxies from the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013). This sample also contains stars from the stellar survey, the Sloan Extension for Galactic Understanding and Exploration

(SEGUE; Yanny et al. 2009), although the vast majority of faint stars come from the BOSS survey. In fact, of the stars that are sourced from SDSS, 93.4 per cent of those between the magnitudes $20.0 < Z < 21.2$ are misclassified quasars from BOSS (Ross et al. 2012).

2.2.4 Gaia

Gaia is a space-based telescope, designed to create an all-sky optical map, mainly focusing on the stars in our Galaxy (Gaia Collaboration 2016) but also including extragalactic sources such as quasars (Gaia Collaboration 2023). We make use of *Gaia*'s third data release (Brown et al. 2021) that contains astrometry and photometry for 1.8 billion sources. We use the *Gaia* archive query described in appendix B of Gaia Collaboration (2018) to extract a significant number of stars in the WAVES-Wide regions. This query utilizes *Gaia*'s parallax measurements and ensures each source has a *R*, *G*, and *B* flux with a signal-to-noise ratio of at least 10. Ultimately, we are able to generate a catalogue of 281 522 sources that we are confident are stars in the WAVES-Wide regions. The resulting magnitude distribution of these stars can be seen in the fourth plot on the bottom row of Fig. 2. As parallax measurements require a high signal-to-noise ratio, the stars we retrieve from *Gaia* are significantly biased towards brighter magnitudes.

2.2.5 Compilation

We find that 102 904 sources overlap across different surveys, most of these being between GAMA, DESI, and SDSS observations in the WAVES north region. We find 218 sources with contradictory labels between surveys, which we remove. Overall, our ground-truth catalogue consists of 367 888 stars and 335 295 galaxies, with a ratio of stars to galaxies roughly what we expect from the WAVES catalogue at fainter magnitudes. Fig. 2 shows the number counts of sources with ground truth labels compared to the total sources in the WAVES-Wide regions as a function of *Z*-band magnitude. This plot demonstrates the advantage of unsupervised over supervised machine learning in this problem. At the 21.2 *Z*-band magnitude limit, only 0.27 per cent of sources have a ground truth label, which provides a challenge regardless of the method of star-galaxy separation. Supervised machine learning methods are, however, known to be heavily biased by the training data they receive. Our ground-truth labels do not constitute a significant fraction of the total number of sources, as a result, we make no use of the ground-truth labels in the classification of objects, and use them purely for the purpose of verification.

3 BASELINE METHOD

One of the ways in which we assess the performance of our star-galaxy separation is by comparing against a baseline algorithm. We choose the classification algorithm used by Galaxy And Mass Assembly Data Release 4 (GAMA DR4; Driver et al. 2022) outlined in section 2.9.1 of Bellstedt et al. (2020), which uses very similar photometry. This algorithm utilizes a combination of colour and size criteria.

The colours are derived from the ‘total’ magnitudes in different bands. This involves adding the total flux within the segment of each source, estimated on the detection bands, and then converting the flux to magnitudes. This is different from ‘colour’ magnitudes,




Figure 3. The baseline star-galaxy separation algorithm outlined in Section 3 for all 14 802 032 sources in the WAVES-Wide fields within the $Z < 21.2$ magnitude limit. The upper panel shows $J - K_s$ colour as a function of *r*-band magnitude. The lower panel shows the log of the half-light radius, R_{50} as a function of *r*-band magnitude. The solid lines indicate the galaxy, star and ambiguous regions of the plots, with the colour indicating the final baseline classification. The sharp cut in sources brighter than $r = 16$ is due to the *Gaia* star mask, in which all stars brighter than a *G*-band magnitude of 16.0 are removed. The discretization of the smallest $\log(R_{50})$ values is due to the discrete seeing values across tiles.

which are derived using a fixed aperture across the multiple bands. For size information, the angular half-light radius R_{50} is used. This is the radius in arcseconds that contains half of the detection band flux within the segment.

The classification method can be visualized in Fig. 3. Essentially, sources are plotted first in $(J - K_s)$ colour versus *r*-band magnitude space and then in $\log R_{50}$ versus *r*-band magnitude space. By doing this, stars and galaxies appear to occupy separate regions on the plots: galaxies tend to have larger $(J - K_s)$ colours and increase in apparent radius as a function of brightness unlike stars. From this, lines are drawn to classify the sources into galaxy, star, and ambiguous regions through the equations

$$\begin{aligned}
 (J - K_s) &= 0.025, & \text{if } r < 19.5, \\
 (J - K_s) &= 0.025 + 0.025(r - 19.5), & \text{if } r > 19.5, \\
 (J - K_s) &= 0.025 - 0.1(r - 19.5)^2, & \text{if } r > 19.5,
 \end{aligned} \tag{1}$$

where the ambiguous region lies between the two lines beyond $r > 19.5$ and

$$\begin{aligned} \log(R_{50}) &= \Gamma + 0.05 - 0.075(r - 20.5), & \text{any } r \\ \log(R_{50}) &= \Gamma + 0.05, & \text{if } r > 20.5, \end{aligned} \quad (2)$$

where Γ is the median LOG10SEEING value, the log of the seeing in arcseconds in the detection band ($r + i + Z + Y$). This value varies between -0.3 and -0.1 $\log(\text{seeing}/\text{arcsec})$ depending on the tile.

Initially, all sources are labelled ambiguous. If a source is in the same region for both planes, then that region label is adopted. If a source is in differing star/galaxy regions on each plane, then it keeps its ambiguous label. And if it is ambiguous in one plane but star/galaxy in the other plane, then it adopts that class. One of the obvious drawbacks of this method is that it is likely to misclassify quasars, which are near point sources, especially if they have star-like colours.

4 PRE-PROCESSING AND DIMENSIONALITY REDUCTION

The pre-processing stage of machine learning is critical, and plays a significant role in the final classification. In this work, the pre-processing and dimensionality reduction means that the final classification algorithm HDBSCAN has no trouble in separating the star and galaxy clusters in the lower dimension space. This process improves the accuracy of the final classifier, and also reduces the computing resources necessary to make the classification.

4.1 Data cleaning

First, all artefacts are removed from the catalogue. Artefacts are identified as sources with extremely unusual colours as outlined in Bellstedt et al. (in preparation). We then remove sources with any missing photometry from any band, as our dimensionality reduction method does not work with missing data. This equates to 1.07 per cent of the catalogue within the adopted magnitude limit of $Z < 21.2$. We then remove sources that have a negative flux after sky subtraction, as magnitudes cannot be calculated with negative flux. This removes a further 0.73 per cent of sources from the catalogue. Investigations were conducted using raw fluxes as the features instead of magnitudes, but this did not prove nearly as effective as using magnitudes.

One of the features used in star/galaxy separation is the KiDS u -band magnitude. However, a significant fraction, 10.57 per cent, of sources have a negative u -band flux after sky subtraction. We found that the u -band does not significantly improve the quality of the star-galaxy separation overall, but does help to distinguish stars from quasars due to quasars' UV excess (Malkan 1983). We therefore run our classification algorithm twice, first with the u -band and then without, and combine the classifications, prioritizing the label given by the run including the u -band. This means we only have to discard 1.80 per cent of sources in total.

We note the omission of sources with missing data is a drawback of this method, and the baseline method would be required for sources with missing data for the formation of a target catalogue. Alternatively, data imputation could be utilized such as in Miller et al. (2017), although this is computationally expensive and requires further study. SED fitting could also be used to fill in missing data, as some photometric redshift template methods do not need all bands to function. This has the advantage of being physically motivated.

4.2 Feature formation

We input the magnitudes from all nine available fluxes as features, using the total flux output from PROFOUND. We also use as all 36 possible colour combinations of the photometric bands using the colour flux output from PROFOUND. The colour flux is used to calculate colours as the segment is the same size in each band, unlike with total flux. Our colours will include some redundant information [e.g. $g - i$ colour is the same as $(g - r) + (r - i)$]. However, we find that the dimensionality reduction method works best when all possible information is given to it, so that the most important features can be extracted. We also include as features the log of R_{50} , the effective half-light radius, and the axial ratio of each source. Finally, we include the log of the astronomical seeing of each source which varies from tile to tile. This leaves us with a total of 48 features for each source or 34 without u -band. PROFOUND does output flux errors however we do not incorporate them as features.

4.3 Scaling

We first use the STANDARDSCALER package from SCIKIT-LEARN (Pedregosa et al. 2011) to appropriately scale our data. This involves scaling each feature to have a mean of 0 and unit variance (Z-score transformation), meaning that no feature is prioritized over another, thus avoiding bias.

4.4 Dimensionality reduction using UMAP

The majority of the 'heavy lifting' of our star-galaxy separation is performed using Uniform Manifold Approximation and Projection (UMAP; McInnes, Healy & Melville 2020), a dimensionality reduction technique. We use dimensionality reduction due to the expensive cost of running a classifier on all 48 features, and we also find that using UMAP improves classification performance. UMAP works by first constructing a graph of nearest neighbours in high-dimensional space, where each data point is connected to its n -nearest neighbours where n is set by the user. Each vertex between points in this graph is weighted by the probability that these points are connected based on their distance, creating a 'fuzzy' structure. This high-dimensional space is then projected to lower dimensional embeddings, where the layout of the graph is maintained as much as possible. UMAP is able to preserve both the global and local structure of the data. In the context of astronomical imaging, the global structure would be the separation between stars, galaxies, and quasars, and the local structure could be the properties of each source such as stellar classification, galaxy redshift and galaxy morphology.

UMAP works in a similar way to t-distributed Stochastic Neighbor Embedding (t-SNE) that has been implemented in astronomy research previously (Traven et al. 2017; Anders et al. 2018; Reis et al. 2018; Nakoneczny et al. 2021; Queiroz et al. 2023; Guiglian et al. 2024). However, UMAP has been demonstrated to be more scalable than t-SNE and better able to preserve the global structure of the data (Becht et al. 2019; McInnes et al. 2020). UMAP is also advantageous to Principal Component Analysis (PCA, a review of which can be found in Jolliffe & Cadima 2016) despite the latter's prolific use in data science, as UMAP is able to capture complex non-linear relationships within the data unlike PCA which is a linear dimensionality reduction technique. Self-organizing maps are also a good candidate for unsupervised star-galaxy separation. They have been used widely in astronomy for classification purposes such as object classification (Geach 2012) and identifying galaxy populations (Holwerda et al. 2022), but also regression problems

such as photometric redshifts (Masters et al. 2015; Wright et al. 2020).

Several star-galaxy separation methods have used UMAP for data visualization and validation of their algorithm (Clarke et al. 2020; Stoppa et al. 2023), but fall short of using it for the main classification. Whilst there are some concerns that UMAP finds ‘mirages’ of spurious local structure (Chari & Pachter 2023), we are confident in its ability to find the global structure of our data in this context: distinguishing stars and galaxies, due to their differences in apparent features.

For our UMAP hyperparameters, we use 200 nearest neighbours and a minimum distance of 0. After running UMAP several times, these hyperparameters most often result in two clearly visible clusters in the UMAP feature space. We choose to reduce our 48-dimensional data down to 10 dimensions. This significantly reduces the computational power required for the classifier whilst maintaining some higher dimensional correlations. Using a similar number of dimensions (e.g. 8, 9, 11, and 12) gives very similar results.

The first two projected dimensions of UMAP applied to the 6593 731 sources of the WAVES-Wide South region can be seen in Fig. 4. We run UMAP on the North and South regions separately in case there are systematic differences between the two regions, and for computational reasons. The plot is coloured by each source’s label given by the baseline algorithm described in Section 3. Clearly, sources appear to be clustered into two distinct nodes, and the baseline algorithm classification indicates that the left-hand node contains galaxies, and the right-hand node contains stars. Ambiguous sources appear to be distributed between both nodes, but are also densely populated in an area attached to the main galaxy node. These sources are primarily QSOs (explained in further sections). The sources ‘connecting’ the two nodes are primarily blended sources with a mixture of flux from both galaxies and stars. This can happen when the source-finding algorithm mistakes multiple sources for a single source. In this case, the flux from foreground stars are contaminating the flux from background galaxies. We show all ten projected dimensions of UMAP in Fig. A1 in Appendix A, visualizing the entire 10-dimensional space which we use for clustering. Fig. 4 is a positive indication of the fidelity of the algorithm, showing a general agreement between the UMAP embeddings of the data and the baseline algorithm.

5 CLASSIFICATION WITH HDBSCAN

Despite the UMAP embeddings being very distinctly clustered in two dimensions, we still employ an unsupervised cluster algorithm to precisely split the data. We use HDBSCAN (McInnes, Healy & Astels 2017) for this purpose, and a robust cluster-finding algorithm built on Density-Based Spatial Clustering of Applications with Noise (DBSCAN: Ester, Kriegel & Xu 1996). These algorithms work differently from other clustering algorithms such as k -means because they do not require a user-set number of clusters to work towards, and instead simply find areas of high density. k -means also require clusters to be ‘spherical’ in feature space which is not always the case for features derived from these surveys (Turner et al. 2019; Holwerda et al. 2022).

They work by first establishing a set of ‘key’ points, those which lie in groups of at least k points (set by the `min_samples` argument in HDBSCAN), where points are linked if they lie within a set value, ϵ . HDBSCAN improved upon DBSCAN by automatically finding an optimum ϵ instead of being user-set, making it more generalizable. HDBSCAN also uses a user-set `min_cluster_size` argument, a

minimum threshold value below which clusters are discarded and the most sensitive parameter in this work.

HDBSCAN is used in astronomy for both extragalactic and Galactic science. Logan & Fotopoulou (2020) use HDBSCAN also for star-galaxy-QSO separation for SDSS using PCA and Random Forest as their feature selection. Queiroz et al. (2023) use HDBSCAN in combination with t-SNE to group *Gaia* stars according to their spectra.

After several iterations of parameters, we settle on `min_samples=1100` and `min_cluster_size=10000`, to reduce the chance of any fragmentation of the star and galaxy clusters. Any fragmented clusters are manually merged into one of their parent clusters until two clusters remain. HDBSCAN also generates a probability for each source of belonging to its associated cluster.

We run HDBSCAN on the 10 dimensions of the data generated by UMAP, which results in each source being given a 0 or 1 label. We analyse the baseline classification for sources labelled 0 and 1, and change to ‘star’ or ‘galaxy’ accordingly. The resulting classification can be seen in the upper left panel of Fig. 5. HDBSCAN has successfully distinguished the two main clusters, which we classify as stars and galaxies by comparing to the Baseline method in Fig. 4. Running HDBSCAN on the entire feature set is much more computationally expensive. We found that running HDBSCAN on a small patch of sky with the full feature set (without dimensionality reduction) is not as effective as applying UMAP first and then clustering.

For the sources with a u -band measurement, we find a small number of contradictory classifications running our method with and without the u band. Of the 9118 285 sources classified as galaxies using u band, 9099 596 of them (99.8 per cent) of them are also classified as galaxies *without* using u band. Similarly, of the 4309 681 sources classified as stars using u -band, 4289 080 of them (99.5 per cent) were also classified as stars *without* using u -band. As stated before, for the sources which have contradictory labels, the classification using u -band takes priority.

We use HDBSCAN’s outlier detection (`outlier_scores`) to give each source a likelihood that it is an outlier to the two main clusters generated by UMAP. HDBSCAN finds very few number of outliers in the UMAP feature space, with just 0.25 per cent of sources having an outlier score greater than 0.5. The majority of these are found in between the two UMAP clusters, connecting them.

6 CLASSIFICATION PERFORMANCE USING GROUND TRUTH LABELS

The overall classifications of the 14 802 032 sources is summarized in Table 1. If using the baseline classification scheme, we would select sources classified as galaxy or ambiguous in order to ensure the required completeness. With UMAP/HDBSCAN (hereafter ‘cluster’) classification, we will simply select targets classified as galaxies. Moving to this new classifier will result in 1672 758 fewer sources identified as galaxy or ambiguous, 11.3 per cent of the catalogue. We believe that this is due to the ambiguous class in the baseline method containing primarily stars. This is associated with a significant reduction in the size of the WAVES target catalogue. We are able to use the 4MOST exposure time calculator (4FS_ETC) to estimate the reduction in fibre hours. 4FS_ETC estimates exposure times based on the magnitude and assumed template for each source, as well as the required signal-to-noise ratio to recover a redshift measurement. We can compare the exposure times associated with the WAVES target catalogues using the baseline star-galaxy classifier




Figure 4. The first two UMAP projected dimensions generated from the 6593 731 sources of the WAVES-Wide South region within the $Z < 21.2$ magnitude limit. The colours indicate the class label determined by the baseline method described in Section 3. Clearly, the sources have been separated into stars and galaxy ‘nodes’ as indicated by separation in baseline class label. Ambiguous sources appear distributed in both nodes.

and our new method to estimate the number of fibre hours saved. After the photometric redshift cuts have been made, we find that moving from the baseline to the cluster classification results in a reduction of 69 504 fibre hours for the WAVES wide region, and an additional 69 903 fibre hours reduction for the WAVES deep region.

We can tentatively gauge the fidelity of the classifier by looking at the properties of the stars and galaxies we have classified. Even though we are only looking at observed and not intrinsic properties of these sources, we can still use them for informative analysis. Fig. 6 shows some observable properties of our classified stars and galaxies. The left panel shows a colour–colour plot, namely $J - K_s$ versus $g - i$: a plot classically used in colour-based star-galaxy separation (e.g. Ivezić et al. 2002; Baldry et al. 2010). The distinctive stellar locus can be seen, where stars have a roughly constant $J - K_s$ colour beyond $g - i > 2$, and then reduces at smaller $g - i$ colours. There is a significant amount of overlap at smaller $g - i$ colours. This indicates that a straight dividing line through $J - K_s = 0.025$ to separate stars and galaxies as is done in the Baseline method would not be optimal. Similarly in the right panel of Fig. 6, we plot our sources’ log of half-light radius as a function

of Z-band magnitude. We see that in the bright regime, stars and galaxies can easily be classified according to their apparent radius. However, at the faintest magnitudes, the distinction becomes less clear. We see a slight uptick in apparent radius of stars at the faintest magnitudes, which could indicate some stars being misclassified as galaxies.

6.1 Comparison to ground truth

The most intuitive way to measure the fidelity of the new classification method is to compare with the ground truth classifications. Fig. 7 shows the confusion matrix between these labels. Of the 335 295 true galaxies, 334 407 of them have been correctly identified, an accuracy of 99.75 per cent compared to the baseline’s accuracy of 93.94 per cent. Similarly, of the 367 888 true stars, 99.70 per cent of them have been correctly identified, compared to the baseline accuracy of 98.10 per cent. Most of the losses of the baseline algorithm are due to the ambiguous class, which contains faint, harder to classify sources. Despite this, more galaxies are missed by the baseline algorithm, with 1186 galaxies being misclassified as stars, compared to 846 from our method. This implies that our method




Figure 5. The same as Fig. 4, but coloured by different observed source properties. The top left shows the labels generated by HDBSCAN, clearly distinguishing the two clusters into stars and galaxies. The upper right plot is coloured by Z-band magnitude. The lower left plot is coloured by the log of the half-light radius. The lower right plot is coloured by the $J - K_s$ colour of each source.

Table 1. The overall classifications of the WAVES-Wide sample made by our method and the baseline method.

Method	Total WAVES-Wide sample			Total
	Galaxy	Star	Ambiguous	
Cluster label	66.5%	33.5%	–	100.0%
	9840 496	4961 236	–	14 802 032
Baseline	66.8%	22.2%	11.0%	100.0%
	9890 177	3288 478	1623 377	14 802 032

has been able to retrieve more galaxies (higher completeness) and includes fewer stars (higher purity) than the baseline method.

To quantify the fidelity of the method, we use the F1 score metric, which uses a harmonic mean of completeness and purity. In the context of identifying galaxies, we define a true positive (TP) as a true galaxy correctly identified, a false positive (FP) as a true star misclassified as a galaxy and a false negative (FN) as a true galaxy misclassified as a star. Purity is then defined as

$$P = \frac{TP}{TP + FP}, \quad (3)$$

the fraction of true positives to total positives. Completeness is defined as

$$C = \frac{TP}{TP + FN}, \quad (4)$$

the fraction of positive prediction to the total number of positives in the sample. These are combined to form the F1 score

$$F1 = 2 \frac{P \cdot C}{P + C}, \quad (5)$$

the harmonic mean of purity and completeness.

In the context of an extragalactic target catalogue and taking true positives as correctly identified galaxies, our method achieves a purity of 0.9967 ± 0.0021 , a completeness of 0.9975 ± 0.0018 and an F1 score of 0.9971 ± 0.0018 . We derive the uncertainties of our metrics by taking the standard deviation of the metric across magnitude bins. The metrics' variability as a function of magnitude is explored later.

Using the baseline algorithm, it can be assumed that all ambiguous objects would be part of the target catalogue for an extragalactic survey. This is done to maximize the completeness, a metric critical for the construction of group catalogues (Robotham et al. 2011; Tempel et al. 2014; Tully 2015). Assuming this, the baseline algorithm achieves a purity of 0.9795 ± 0.0172 , a completeness of 0.9965 ± 0.0026 , and an F1 score of 0.9879 ± 0.0088 .

We observe that whilst the completeness stays the same between the two methods (i.e. roughly the same number of galaxies are retrieved), the purity increases from 0.9795 ± 0.0172 to 0.9967 ± 0.0021 and consequently the F1 score increases too from 0.9879 ± 0.0088 to 0.9971 ± 0.0018 . In the case of a spectroscopic extragalactic survey, this would imply fewer stars in the targeting




Figure 6. Observable properties of the sources labelled stars and galaxies by our classifier. The left panel shows $J - K_s$ colour versus $g - i$ colour, and the right panel shows the log of half-light radius as a function of Z-band magnitude. Red and blue contours/points indicate galaxies and stars, respectively. Galaxies occupy the upper regions of both plots. Contours are scaled logarithmically and the points show a random 10 per cent of each population.




Figure 7. The confusion matrices generated by the ground-truth data set compiled in Section 2.2. The left confusion matrix uses the labels generated by HDBSCAN, and the right confusion matrix uses the labels generated by the baseline algorithm described in Section 3.

catalogue and therefore less wasted fibre hours. A summary of the purity, completeness and F1 metric for the cluster and baseline regime can be found in Table 2.

In addition, we find that of the $z < 0.2$ galaxies in the ground truth sample (the photo- z limit of WAVES-Wide), we identify 99.75 per cent of them as galaxies, compared to 88.50 per cent of them being classified as galaxies by the baseline algorithm. This highlights a major improvement in the star-galaxy separation in the context of WAVES-Wide. For galaxies with redshifts $z < 0.8$ (the photo- z limit of WAVES-Deep), we correctly identify 99.82 per cent, compared to 94.85 per cent from the baseline algorithm.

To demonstrate the importance of UMAP in our method, we attempt star-galaxy classification with a number of pre-processing and clustering methods. We create a random subset of WAVES-Wide which is a 10th of the size but fully representative of the entire catalogue in all dimensions. For our pre-processing, we use UMAP, PCA, and also attempt clustering without any pre-processing. We use k -means clustering and HDBSCAN as our clustering methods.

The results can be seen in Table 3, in which we show the F1 scores achieved from the ground truth sample by combining the various

Table 2. Overall purity, completeness, and F1 scores of the two methods, using all of the ground-truth labels.

Method	Purity	Completeness	F1
Cluster label	0.9967 ± 0.0021	0.9975 ± 0.0018	0.9971 ± 0.0018
Baseline ^a	0.9795 ± 0.0172	0.9965 ± 0.0026	0.9879 ± 0.0088

^aTaking all ambiguous sources to be galaxies.

Table 3. The F1 scores achieved by a selection of pre-processing and clustering methods, run on a one tenth representative sample of WAVES-Wide.

Clustering	Pre-processing		
	None	PCA	UMAP
k -means	0.621 ± 0.194	0.623 ± 0.193	0.997 ± 0.004
HDBSCAN	0.935 ± 0.046	0.947 ± 0.037	0.997 ± 0.004

pre-processing and clustering methods. We find that when using UMAP as our pre-processing method, both k -means and HDBSCAN both achieve the same F1 score, as UMAP makes the clustering trivial (shown in Fig. 4). They achieve F1 scores of 0.997 ± 0.004 ,




Figure 8. The F1 score (equation 5), purity and completeness for identifying galaxies as a function of Z-band magnitude. The solid line shows the performance of our algorithm, compared to the baseline algorithm in grey. Errors are based on Poisson statistics from the number of ground truth sources we have, shown in the bottom plot.

compared to the F1 scores of 0.947 ± 0.037 and 0.935 ± 0.046 achieved by HDBSCAN when using no pre-processing and PCA respectively. This indicates that it is indeed UMAP that is driving the performance of our star-galaxy classification, and not HDBSCAN.

6.2 As a function of magnitude

We can also measure the fidelity of our method as a function of different source properties. For example, Fig. 8 shows how the F1 score (equation 5), purity, and completeness for classified galaxies varies as a function of Z-band magnitude compared to the baseline algorithm. Across all magnitudes, the purity produced by our method exceeds the baseline algorithm, suggesting our sample is far less contaminated, especially at very faint ($Z > 20$) and very bright ($Z < 17$) magnitudes. The purity from our method never dips below 0.99, whereas the purity from the baseline method is 0.970 at the faintest magnitudes. The completeness of both our method and the baseline algorithm hover around 0.995, with a dip at about $Z \sim 20$, which we believe is caused by a shift from the GAMA to DESI regime in the ground truth sample. This results in the total F1 score never reaching below 0.990, and consistently above the baseline method at all magnitudes.




Figure 9. The mean probability of classifications in bins of Z-band magnitude, calculated by perturbing the flux values of sources by their uncertainties and measuring the number of contradictory classifications. The bottom panel shows the number of sources per magnitude bin. In this analysis, we use a representative sample of WAVES-Wide, one hundredth of the size.

6.3 Classification probabilities

We also investigate how the uncertainties in the flux measurements of each source can change the classification. The flux errors are calculated in PROFOUND from sky subtraction and sky rms, with errors being higher in areas of great noise, e.g. near bright objects. We create a representative sample of WAVES-Wide, one hundredth the size, for ease of computation. We assume Gaussian uncertainties, and perturb the fluxes of each source by their flux uncertainty 100 times, and classify the sources as stars and galaxies as before, keeping track of the classification of each source. We then assign each source a classification probability, based on the number of contradictory classifications after the 100 perturbations.

The results can be visualized in Fig. 9, in which we show the mean of the classification probabilities as a function of Z-band magnitude. As expected, the probabilities decrease towards fainter magnitudes, with an average probability of 0.9757 ± 0.0007 in the faintest magnitude bin. This implies that on average, 2.4 out of every 100 perturbations will result in a contradictory classification at $Z = 21.2$. The mean probability of the overall sample is 0.9898 ± 0.0002 . When comparing to our ground truth catalogue, we find we find a standard deviation of 0.0006 in F1 score based on the 100 perturbations. This however is likely to be an underestimate due to the ground truth catalogue being biased towards bright objects. With enough computing power, we could use this method for the entire WAVES-Wide sample to give every source a classification probability. This would give another parameter for the 4MOST target catalogue, with fibre allocation being prioritized for galaxies with high probabilities.

7 CLASSIFICATION OF CHALLENGING GALAXIES

In addition to assessing the ground truth sample as a whole, we also look at galaxies that could be challenging for star-galaxy classification. These include quasars, compact galaxies and low-surface brightness galaxies. We do this to ensure that the completeness of the star-galaxy classifier is effective even with these fringe cases. Fig. C1 shows nine of these challenging galaxies using KiDS and VIKING photometry. All have been correctly classified as galaxies by our algorithm, but some are misclassified by the baseline algorithm.

Table 4. The classification of DESI EDR quasars achieved by the method described in this paper and the baseline method.

Method	DESI EDR quasars			Total
	Galaxy	Star	Ambiguous	
Cluster label	92.7%	7.3%	–	100.0%
	5494	431	–	5929
Baseline	23.9%	12.3%	63.8%	100.0%
	1418	729	3782	5929

7.1 Quasars

First, we assess how well our star-galaxy separation performs at classifying quasars. Quasars or QSOs are galaxies with an extremely luminous active galactic nucleus, meaning they can be observed even at high redshifts. Like stars, they appear almost as point sources, meaning their classification from stars is a difficult task, and careful analysis is required. Quasars can be identified using infrared and optical photometry (Assef et al. 2018; Chaussidon et al. 2023), UV excess (Richards et al. 2002), photometric variability (MacLeod et al. 2011), and X-ray excess (Maccacaro, Gioia & Stocke 1984).

For our true quasar sample, we use those identified spectroscopically in the DESI EDR. Quasars are identified in DESI through their spectral classification pipeline using Redrock template fitting software (Guy et al. 2023). Redrock fits a set of PCA templates to each source at every redshift, and the classification and redshift is returned with the smallest χ^2 fit, regardless of the target selection type. Many galaxies with active galactic nuclei are classified as quasars in the EDR (DESI Collaboration 2023), producing a ‘bump’ of sources at $z < 0.5$. We filter these out by selecting only sources with a PSF morphological type, increasing the median redshift from $z = 1.42$ to $z = 1.61$.

We are able to cross-match 5929 quasars in the North region. The results of the two classifiers can be seen in Table 4. Of the 5929 quasars, our method was able to correctly identify 5494 or 92.7 per cent of them as galaxies. This is compared to the baseline algorithm, which was only able to identify 1418 or 23.9 per cent of them as galaxies, with the majority being classified as ambiguous sources. Even if all ambiguous sources were observed (an additional 1.2 million sources), our method would have still identified more quasars in these data.

Further study could be conducted in using HDBSCAN to create an additional quasar cluster¹¹ separated from the galaxy and star clusters. However, for the purpose of this work, we are only interested in separating them from stars. We mention earlier that any quasars correctly identified as galaxies should be filtered out of the WAVES-Wide target catalogue due to the $z < 0.2$ photo- z selection. We find that of the DESI EDR quasars with spectroscopic redshifts of $z > 0.2$, ~ 95 per cent of them are filtered out using the photo- z selection.

7.2 Compact galaxies

Next, we measure the effectiveness of our star-galaxy separation in identifying compact galaxies. Compact galaxies are characterized as having a high concentration of their mass and luminosity in a small, central area of the galaxy. Their compact profile makes them difficult to distinguish from stars using only morphological data, so colour information is necessary.

¹¹ ‘Clusters’ in the sense of the classification algorithm.

Table 5. The classification of compact galaxies described in Section 7.2 achieved by the method described in this paper and the baseline method.

Method	Compact GAMA galaxies			Total
	Galaxy	Star	Ambiguous	
Cluster label	84.6%	15.4%	–	100.0%
	743	135	–	878
Baseline	81.8%	12.1%	6.2%	100.0%
	718	106	54	878

We utilize the metric $\Sigma_{1.5} = \log(M_*/M_\odot) - 1.5\log(r_{50}/\text{kpc})$ to quantify compactness, first outlined in Barro et al. (2013) for the CANDELS survey and later in Baldry et al. (2021) for SDSS, where M_* is stellar mass and r_{50} is the physical half-light radius of the galaxy. The metric essentially measures the deviation from the size-mass relation of high-mass, quiescent galaxies, and is measured in units of $M_\odot \text{kpc}^{-1.5}$. A greater $\Sigma_{1.5}$ value indicates a more compact galaxy.

We use stellar masses from GAMA DR4 (Driver et al. 2022), which uses code first outlined in Taylor et al. (2011), but has been developed into DR4 using Source Extractor photometry from Driver et al. (2016), matched aperture photometry from LAMBDAR (Wright et al. 2016) and PROFOUND photometry from Bellstedt et al. (2020). It also uses GAMA spectroscopic redshifts combined with the apparent half-light radius to calculate the physical half-light radius. We limit the GAMA sample to redshifts less than 0.6 to avoid quasars, and define our compact sample as galaxies with the highest 0.5 per cent of $\Sigma_{1.5}$ values. This leaves us with a sample of 878 galaxies with a mean compactness value of $\Sigma_{1.5} = 10.28 M_\odot \text{kpc}^{-1.5}$, compared to $\Sigma_{1.5} = 9.34 M_\odot \text{kpc}^{-1.5}$ for the whole GAMA sample.

Table 5 shows the performance of our star-galaxy separation and the baseline algorithm on the compact GAMA sample. We correctly identify 84.6 per cent of the compact galaxy sample as galaxies, and misclassify 15.4 per cent of them as stars. This is compared to the baseline algorithm classifying 81.8 per cent, 21.2 per cent, and 6.2 per cent of the compact galaxies as galaxies, stars and ambiguous sources, respectively. This is a slight improvement in the retrieval of compact galaxies, unless all ambiguous sources are also considered. We do make the caveat that GAMA is magnitude limited to $r < 19.8$, and its input catalogue likely omitted more compact galaxies in its own star-galaxy separation (Baldry et al. 2010). We also note that if all ambiguous objects were observed, this would result in more compact galaxies being retrieved (93.2 per cent compared to 84.5 per cent), although this increase in completeness comes at the cost of observing an additional 1.2 million sources.

7.3 Low surface brightness galaxies

Finally, we look at low surface brightness galaxies. These are galaxies with a faint overall brightness over their area, making them difficult to detect in photometric data. They span a wide range of colours (Greene et al. 2022), with their bluer members actively forming stars. There is a concern that some of these low surface brightness galaxies have extreme colours, and thus could be missed by simple colour-based star-galaxy separation methods. They may also have unreliable photometric redshifts which one may also use in a method for star-galaxy separation.

We calculate the effective surface brightness $\mu_{z,50}$ for each DESI and GAMA galaxy using their area calculated from their half-light radius R_{50} and Z-band magnitude. We limit the sample to have a redshift of $z < 0.2$ and define the low surface brightness

Table 6. The classification of low surface brightness galaxies described in Section 7.3 achieved by the method described in this paper and the baseline method.

Method	Low surface brightness galaxies			Total
	Galaxy	Star	Ambiguous	
Cluster label	99.5%	0.5%	–	100.0%
	442	2	–	444
Baseline	58.8%	0.0%	41.2%	100.0%
	261	0	183	444

sample as having the smallest 0.5 per cent $\mu_{z,50}$ values. This sample comprises 444 galaxies, with a mean $\mu_{z,50}$ surface brightness value of 22.93 mag arcsec⁻², compared to the mean surface brightness of 20.28 mag arcsec⁻² of the rest of the $z < 0.2$ GAMA and DESI samples.

The results can be seen in Table 6. Of the 444 low surface brightness galaxies in the sample, we correctly identify 442 of them as galaxies. This is compared to the baseline algorithm which classifies just 261 as galaxies and the rest as ambiguous sources. Their large apparent radii mean they lie in the galaxy region in the lower plot of Fig. 3, but they exhibit a wide range of colours. The 25th and 75th percentile of their $J - K_s$ colours are -0.31 and 0.28 , respectively, compared to that of the ground truth galaxy sample of 0.12 and 0.32 . 50.2 per cent of the low-surface brightness galaxies have a $J - K_s$ colour less than the baseline cut-off of 0.025 , placing them in the star region of the upper panel (colours versus magnitude) of Fig. 3, thus classing them as ambiguous despite their large apparent radii.

8 CLASSIFICATION PERFORMANCE USING ANGULAR TWO-POINT CORRELATION FUNCTION

This section utilizes the angular two-point correlation function, $\omega(\theta)$, for assessing star-galaxy separation in the population as a whole. This tool can be useful in differentiating stars from galaxies, given their distinct clustering patterns. Galaxies typically cluster following a power-law on small angular scales,

$$\omega(\theta) = A_\omega \theta^\delta, \quad (6)$$

with $\delta \approx -0.8$ (Groth & Peebles 1977; Coil 2012). In contrast to galaxies, stars exhibit a different clustering pattern, as evidenced by their angular two-point correlation function. Stars typically do not demonstrate the large-scale clustering patterns that are observable in galaxies. Unlike galaxies, stars are essentially randomly distributed apart from a large-scale gradient towards the Galactic Centre and plane. This difference in clustering pattern allows an assessment of star/galaxy classification, independent of the ground-truth data set.

8.1 Methodology

In this test, we use the Landy & Szalay (LS) estimator (Landy & Szalay 1993), which is the most commonly used two-point correlation function estimator (Coil 2012). It is defined as

$$\omega(\theta) = \frac{1}{RR} \left[DD \left(\frac{n_R}{n_D} \right)^2 - 2DR \left(\frac{n_R}{n_D} \right) + RR \right], \quad (7)$$

where DD and DR are the pair count of galaxies in each separation bin in the data catalogue and between the data and random catalogues, and RR is the pair count for the random catalogue. n_D and n_R are

the number densities of galaxies in the data and random catalogues. In Kerscher, Szapudi & Szalay (2000), it has been shown that the LS estimator handles edge corrections better than other estimators on large scales and matches the performance on small scales; also the size of the random catalogue affects it less than other estimators (Coil 2012).

A ‘random catalogue’ was constructed to mirror the angular coverage of our observational data, maintaining the same sky coverage but populated with randomly distributed points. We used Pymangle, a PYTHON implementation derived from the Mangle C++ package (Swanson et al. 2008), to generate and distribute these random points and apply the star mask that was created based on *Gaia* stars, where the mask radius is a function of the magnitude of each star (Bellstedt et al. 2020) from *Gaia* DR2 (Gaia Collaboration 2018). This approach led to the production of random samples that contained ten times more points than our original catalogue, ensuring a comprehensive and statistically sound basis for our angular analysis. In this research, we used the TreeCorr PYTHON package (Jarvis, Bernstein & Jain 2004) to calculate two point correlation functions.

Our test involves estimating the angular correlation function for objects classified by our method and the baseline method into five categories for star-galaxy separation:

- (i) Objects identified as galaxies by both our method and the baseline method.
- (ii) All objects identified as galaxies by our method.
- (iii) Objects identified as galaxies by our method but ambiguous by the baseline method.
- (iv) Objects identified as stars by both methods.
- (v) Objects identified as stars by our method but ambiguous by the baseline method.

8.2 Results

The angular correlation functions for the above subsamples are plotted in Fig. 10 and the parameters of power-law fits to scales $\theta < 0.9^\circ$ are given in Table 7. The differences in clustering pattern for galaxies and stars for which both classification methods agree are clear. For sources classified as galaxies by both methods, for angular separation less than 0.9° , the fit parameters, a clustering amplitude A_ω of $(4.8 \pm 0.3) \times 10^{-3}$ and a slope δ of -0.76 ± 0.01 , suggest a robust clustering measurement, in line with the established large-scale structure of galaxies. Conversely, sources classified as stars by both methods exhibit a markedly shallower correlation function, with an amplitude of $(8.1 \pm 0.1) \times 10^{-2}$ and a slope δ of -0.05 ± 0.01 . The red and blue dashed lines in Fig. 10 represent the power-law fits for galaxies and stars, respectively, on which both methods agree in their classifications.

For sources identified as galaxies by our method, shown as grey points in the top left plot of Fig. 10, the clustering is similar to those sources agreed upon by both methods. This suggests that our classification of sources as galaxies remains uniform and consistent, irrespective of how the baseline method classifies them.

For objects identified as stars by our method and ambiguous by the baseline method, the bottom right plot in Fig. 10 reveals that these sources cluster similarly to stars. This suggests that our method effectively recognises star-like properties in these objects, which the baseline method categorizes as ambiguous, underscoring our method’s effectiveness in definitively identifying stars.

For objects identified as galaxies by our method, yet deemed ambiguous by the baseline method, the top right plot in Fig. 10 presents the angular correlation function. This function exhibits a

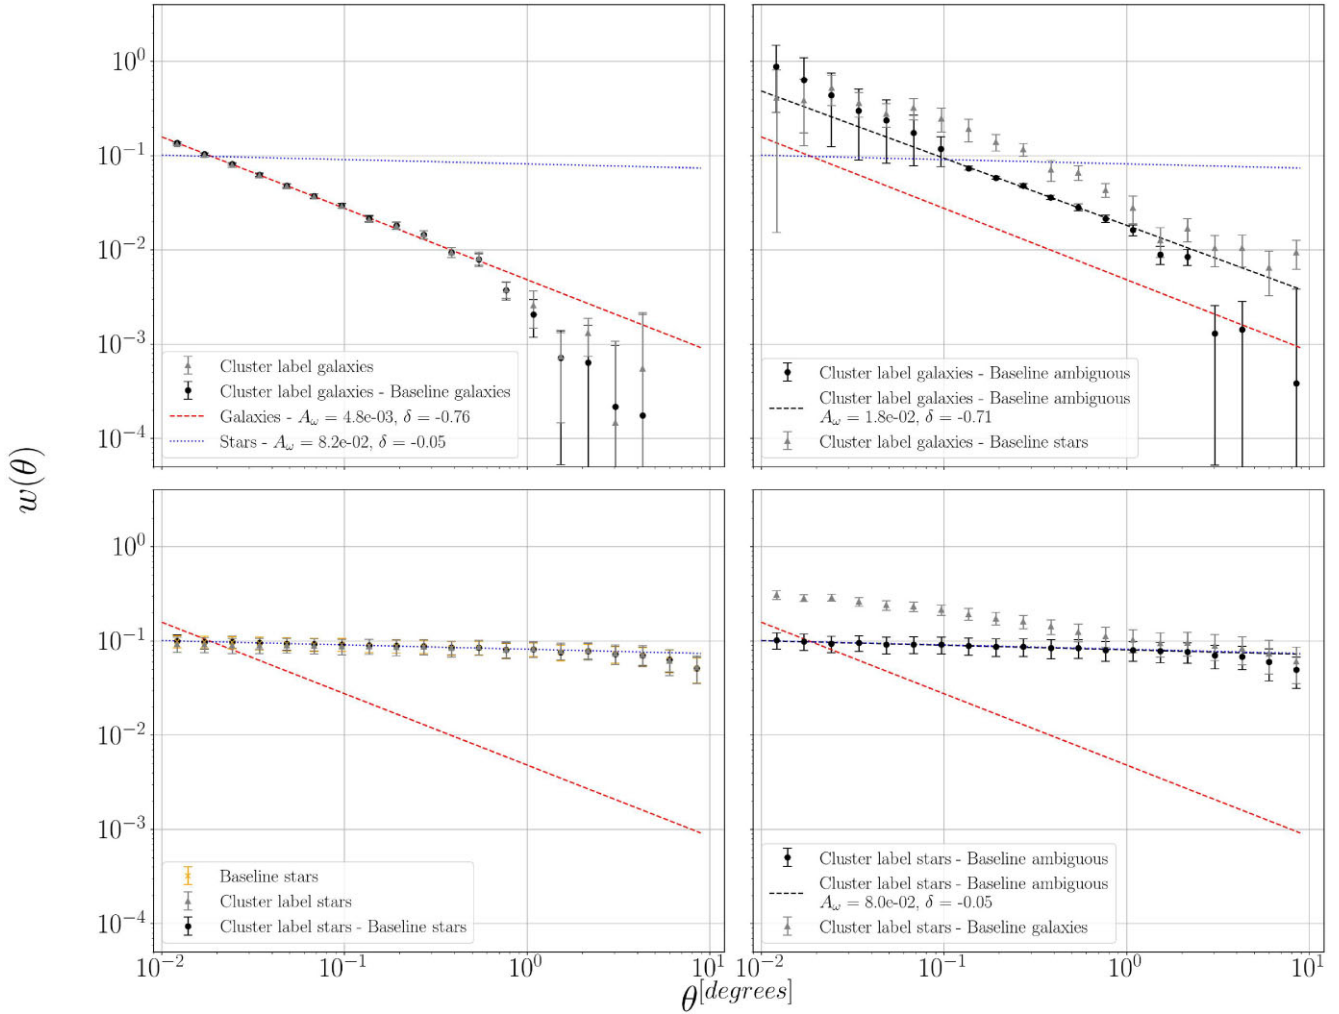


Figure 10. The angular two-point correlation function, $\omega(\theta)$, for various object classifications: galaxies, stars categorized by our method and the baseline method, and those objects which baseline method classified them as ambiguous. The red and blue dashed lines correspond to the best-fitting power-law models for galaxies and stars for which both methods agree in their classifications. The power law model lines for galaxies and stars have a steep and shallow negative slope, respectively.

Table 7. Power law fit (equation 6) for angular separation $\theta < 0.9^\circ$.

Power-law fit parameters				
Classification				
Cluster	Baseline	A_ω	δ	
i	Galaxy	Galaxy	$(4.8 \pm 0.2) \times 10^{-3}$	-0.76 ± 0.01
ii	Galaxy	All	$(4.7 \pm 0.3) \times 10^{-3}$	-0.76 ± 0.02
iii	Galaxy	Ambiguous	$(1.8 \pm 0.1) \times 10^{-2}$	-0.71 ± 0.02
iv	Star	Star	$(8.2 \pm 0.1) \times 10^{-2}$	-0.05 ± 0.02
v	Star	Ambiguous	$(8.0 \pm 0.1) \times 10^{-3}$	-0.05 ± 0.03

slope similar to that of galaxies, and it is parallel to the power-law fit, but there is an observed offset from the fit for sources classified as galaxies by both methods. This deviation is explored in Appendix D, where it is shown to arise from systematic variations in the seeing correction in the WAVES input catalogue adversely affecting the baseline classification.

Based on the results of this analysis, we can conclude that our method effectively identifies ambiguous sources that the baseline method was unable to classify definitively. This demonstrates the

robustness and precision of our approach in handling cases where the baseline classification was uncertain.

9 SUMMARY

We use unsupervised machine learning for star-galaxy separation of the WAVES-Wide input catalogue, and compare our results against a baseline method using a number of verification methods.

i) We construct a catalogue of ground truth data for verification, formed from Gaia stars with high signal-to-noise parallax measurements, and stars and galaxies from GAMA, SDSS, and DESI EDR spectroscopy. This gives us a sample of truth data even at faint magnitudes down to $Z < 21.2$. Fig. 2 demonstrates that the quantity of ground truth data at these faint magnitudes is orders of magnitudes smaller than the number of sources we wish to classify, giving us motivation to deviate from supervised machine learning in fear of bias from a training set.

ii) We utilize photometric features derived from the source-finding software PROFOUND, including KiDS and VIKING magnitudes, their colours, and sources' radii and axial ratios. A feature space is formed

and reduced using UMAP, a non-linear dimensionality reduction algorithm, and then clustered into stars and galaxies using HDBSCAN.

iii) Our method classifies 1672 758 fewer sources as galaxy or ambiguous compared to the baseline method, or 11.3 per cent, which is associated with an approximate reduction of 70 000 fewer 4MOST fibre hours after photometric redshift cuts, and fewer spurious stars. We compare the classification of the ground truth data to a baseline method, used in the star-galaxy separation for GAMA DR4 (Bellstedt et al. 2020), which uses a combination of morphological and colour classification. The full results can be seen in the confusion matrices of Fig. 7. They show a 5.86 per cent increase in the number of galaxies being correctly recovered, and a smaller amount being falsely classified as star or ambiguous sources. This can be quantified by an increase in F1 score from 0.9879 ± 0.0088 to 0.9971 ± 0.0018 . We also demonstrate that the F1 score is better for our method than the baseline method across all magnitudes in Fig. 8. This is mainly due to a major improvement in the purity of the sample.

iv) We assess the effectiveness of our method with ‘challenging’ galaxies, including quasars, compact galaxies and low surface brightness galaxies. We find that our method outperforms the baseline method for all three types. The baseline method only manages to classify 36.1 per cent of quasars as galaxies, with most being classified as ambiguous due to them being point sources. This compares to our method identifying 95.1 per cent of them correctly as galaxies. There is a minor increase in effectiveness in identifying compact galaxies, and a significant improvement in identifying low surface brightness galaxies due to their extraordinary colours.

v) Finally, we investigate the angular clustering of our sources and baseline sources. We find that the angular clustering of our stars and galaxies are consistent with what we expect. We also find that even for sources classified as ambiguous by the baseline method, if our method classifies these sources as stars or galaxies, then they exhibit the expected clustering.

Our method of unsupervised star-galaxy separation for the WAVES-Wide target catalogue shows promising results, improving on the baseline method in both purity and completeness, and saving valuable fibre hours on the 4MOST instrument. We hope the use of UMAP and other unsupervised machine learning techniques can be used in the future for other surveys, as we plan to incorporate all possible data into our target selection.

ACKNOWLEDGEMENTS

WAVES is a joint European-Australian project based around a spectroscopic campaign using the 4MOST. The WAVES input catalogue is based on data taken from the European Southern Observatory’s VST and VISTA telescopes. Complementary imaging of the WAVES regions is being obtained by a number of independent survey programmes including *GALEX* MIS, VST, *WISE*, *Herschel-ATLAS*, and ASKAP providing UV to radio coverage. WAVES is funded by the ARC (Australia) and the participating institutions. The WAVES website is <https://wavesurvey.org>. Based on observations made with ESO Telescopes at the La Silla Paranal Observatory under programme ID 179.A-2004 and ID 177.A-3016. TLC is supported by the Scientific and Technology Facilities Council (STFC), grant numbers [ST/T506461/1] and [ST/V507131/1]. SP is also supported by the STFC, grant number [ST/Y509620/1]. Some parts of this work used the DiRAC Data Intensive service (CSD3) at the University of Cambridge, managed by the University of Cambridge University Information Services on behalf of the STFC DiRAC HPC Facility (<https://www.dirac.ac.uk>). The DiRAC component of CSD3 at Cam-

bridge was funded by BEIS, UKRI, and STFC capital funding and STFC operations grants. DiRAC is part of the UKRI Digital Research Infrastructure. ET acknowledges the ETAg grant number PRG1006 and the CoE project TK202 funded by the HTM. RD gratefully acknowledges support by the ANID BASAL project FB210003. MB is supported by the Polish National Science Center through grant numbers 2020/38/E/ST9/00395, 2018/30/E/ST9/00698, 2018/31/G/ST9/03388, and 2020/39/B/ST9/03494, and by the Polish Ministry of Science and Higher Education through grant DIR/WK/2018/12. DG is supported by the Deutsche Forschungsgemeinschaft (German Research Foundation) under Germany’s Excellence Strategy – EXC-2094-390783311. AI and ML acknowledge financial support from INAF-Minigrant ‘4MOST-StePS: a Stellar Population Survey using 4MOST@VISTA’ (2022). GK is supported by the Polish National Science Center through grant number 2020/38/E/ST9/00395. We would also like to thank John F. Wu, Boudewijn F. Roukema, Letizia P. Cassarà, Guillaume Guiglion, Tayyaba Zafar, Alexander Fritz, and Andrew Humphrey for their proofreading, comments and suggestions.

DATA AVAILABILITY

The WAVES photometric input catalogue used in this work can be found under the ‘Input Cat Downloads’ tab of <https://wavesegview.icrar.org>. The data used to replicate the work in this paper can be found in that public website. A code tutorial of this method running on a subset of the WAVES photometric input catalogue can be found in the following GitHub repository: <https://github.com/toddlcook/UMAP-star-gal-separation>.

REFERENCES

- Abdurro’uf et al., 2022, *ApJS*, 259, 35
 Alexander D. M. et al., 2023, *AJ*, 165, 124
 Anders F., Chiappini C., Santiago B. X., Matijević G., Queiroz A. B., Steinmetz M., Guiglion G., 2018, *A&A*, 619, A125
 Assef R. J., Stern D., Noirot G., Jun H. D., Cutri R. M., Eisenhardt P. R. M., 2018, *ApJS*, 234, 23
 Bailer-Jones C. A. L., Fouesneau M., Andrae R., 2019, *MNRAS*, 490, 5615
 Baldry I. K. et al., 2010, *MNRAS*, 404, 86
 Baldry I. K., Sullivan T., Rani R., Turner S., 2021, *MNRAS*, 500, 1557
 Baqui P. O. et al., 2021, *A&A*, 645, A87
 Barro G. et al., 2013, *ApJ*, 765, 104
 Becht E., McInnes L., Healy J., Dutertre C.-A., Kwok I. W. H., Ng L. G., Ginhoux F., Newell E. W., 2019, *Nature Biotechnology*, 37, 38
 Bellstedt S. et al., 2020, *MNRAS*, 496, 3235
 Bertin E., Arnouts S., 1996, *A&AS*, 117, 393
 Brown A. G. A. et al., 2021, *A&A*, 649, A1
 Chari T., Pachter L., 2023, *PLOS Comput. Biol.*, 19, e1011288
 Chaussidon E. et al., 2023, *ApJ*, 944, 107
 Clarke A. O., Scaife A. M. M., Greenhalgh R., Griguta V., 2020, *A&A*, 639, A84
 Coil A. L., 2012, *Planets Stars Stellar Syst.*, 6, 387
 Cooper A. P. et al., 2023, *ApJ*, 947, 37
 de Jong J. T. A., Verdoes Kleijn G. A., Kuijken K. H., Valentijn E. A., 2013, *Exp. Astron.*, 35, 25
 de Jong R. S. et al., 2019, *The Messenger*, 175, 3
 DESI Collaboration, 2022, *AJ*, 164, 207
 DESI Collaboration, 2023, The Early Data Release of the Dark Energy Spectroscopic Instrument. Princeton Univ. Press, Princeton, NJ
 Davies L. J. M. et al., 2018, *MNRAS*, 480, 768
 Dawson K. S. et al., 2013, *AJ*, 145, 10
 Driver S. P. et al., 2011, *MNRAS*, 413, 971
 Driver S. P. et al., 2016, *MNRAS*, 455, 3911
 Driver S. P. et al., 2019, *The Messenger*, 175, 46

- Driver S. P. et al., 2022, *MNRAS*, 513, 439
- Edge A., Sutherland W., Kuijken K., Driver S., McMahon R., Eales S., Emerson J. P., 2013, *Messenger*, 154, 32
- Ester M., Kriegel H.-P., Xu X., 1996, *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, AAAI Press, Portland, Oregon, p. 226
- Fotopoulou S., 2024, *Astron. Comput.*, 48, 100851
- Gaia Collaboration, 2016, *A&A*, 595, A1
- Gaia Collaboration, 2018, *A&A*, 616, A10
- Gaia Collaboration, 2023, *A&A*, 674, A41
- Geach J. E., 2012, *MNRAS*, 419, 2633
- Greene J. E. et al., 2022, *ApJ*, 933, 150
- Groth E. J., Peebles P. J. E., 1977, *ApJ*, 217, 385
- Gruen D., Brimiouille F., 2017, *MNRAS*, 468, 769
- Guigliion G. et al., 2024, *A&A*, 682, A9
- Guy J. et al., 2023, *AJ*, 165, 144
- Hahn C. et al., 2023, *AJ*, 165, 253
- Holwerda B. W. et al., 2022, *MNRAS*, 513, 1972
- Holwerda B. W. et al., 2024, *MNRAS*, 259, 1067
- Hopkins A. M. et al., 2013, *MNRAS*, 430, 2047
- Ivezić Z. et al., 2002, *AJ*, 124, 2364
- Jarvis M., Bernstein G., Jain B., 2004, *MNRAS*, 352, 338
- Jolliffe I. T., Cadima J., 2016, *Phil. Trans. R. Soc. A: Math. Phys. Eng. Sci.*, 374, 20150202
- Kerscher M., Szapudi I., Szalay A. S., 2000, *ApJ*, 535, L13
- Kovács A., Szapudi I., 2015, *MNRAS*, 448, 1305
- Krakowski T., Małek K., Bilicki M., Pollo A., Kurcz A., Krupa M., 2016, *A&A*, 596, A39
- Kuijken K. et al., 2019, *A&A*, 625, A2
- Kurcz A., Bilicki M., Solarz A., Krupa M., Pollo A., Małek K., 2016, *A&A*, 592, A25
- Lan T.-W. et al., 2023, *ApJ*, 943, 68
- Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
- Lochner M., Bassett B. A., 2021, *Astron. Comput.*, 36, 100481
- Logan C. H. A., Fotopoulou S., 2020, *A&A*, 633, A154
- Maccacaro T., Gioia I. M., Stocke J. T., 1984, *ApJ*, 283, 486
- McInnes L., Healy J., Astels S., 2017, *J. Open Source Softw.*, 2, 205
- McInnes L., Healy J., Melville J., 2020, *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction*. Princeton Univ. Press, Princeton, NJ
- MacLeod C. L. et al., 2011, *ApJ*, 728, 26
- Malkan M., 1983, *ApJ*, 268, 582
- Masters D. et al., 2015, *ApJ*, 813, 53
- Miller A. A., Kulkarni M. K., Cao Y., Laher R. R., Masci F. J., Surace J. A., 2017, *AJ*, 153, 73
- Nakoneczny S. J. et al., 2021, *A&A*, 649, A81
- Odehahn S. C., Humphreys R. M., Aldering G., Thurmes P., 1993, *PASP*, 105, 1354
- Pedregosa F. et al., 2011, *J. Mach. Learn. Res.*, 12, 2825
- Planck Collaboration IX, 2013, *A&A*, 554, A139
- Prieto C. A. et al., 2020, *Res. Notes AAS*, 4, 188
- Queiroz A. B. A. et al., 2023, *A&A*, 673, A155
- Raichoor A. et al., 2023, *AJ*, 165, 126
- Reis I., Poznanski D., Baron D., Zasowski G., Shahaf S., 2018, *MNRAS*, 476, 2117
- Richards G. T. et al., 2002, *AJ*, 123, 2945
- Robotham A. S. G. et al., 2011, *MNRAS*, 416, 2640
- Robotham A. S. G., Davies L. J. M., Driver S. P., Koushan S., Taranu D. S., Casura S., Liske J., 2018, *MNRAS*, 476, 3137
- Ross N. P. et al., 2012, *ApJS*, 199, 3
- Siudek M. et al., 2018, *A&A*, 617, A70
- Slater C. T., Ivezić Z., Lupton R. H., 2020, *AJ*, 159, 65
- Soumagnac M. T. et al., 2015, *MNRAS*, 450, 666
- Stevens G., Fotopoulou S., Bremer M. N., Ray O., 2021, *J. Open Source Softw.*, 6, 3635
- Stoppa F. et al., 2023, *A&A*, 680, A109
- Strauss M. A. et al., 2002, *AJ*, 124, 1810
- Swanson M. E. C., Tegmark M., Hamilton A. J. S., Hill J. C., 2008, *MNRAS*, 387, 1391
- Taylor E. N. et al., 2011, *MNRAS*, 418, 1587
- Tempel E. et al., 2014, *A&A*, 566, A1
- Tempel E. et al., 2020a, *MNRAS*, 497, 4626
- Tempel E. et al., 2020b, *A&A*, 635, A101
- Traven G. et al., 2017, *ApJS*, 228, 24
- Tully R. B., 2015, *AJ*, 149, 171
- Turner S. et al., 2019, *MNRAS*, 482, 126
- Weir N., Fayyad U. M., Djorgovski S., 1995, *AJ*, 109, 2401
- Wright E. L. et al., 2010, *AJ*, 140, 1868
- Wright A. H. et al., 2016, *MNRAS*, 460, 765
- Wright A. H., Hildebrandt H., Busch J. L. v. d., Heymans C., 2020, *A&A*, 637, A100
- Yanny B. et al., 2009, *AJ*, 137, 4377
- York D. G. et al., 2000, *AJ*, 120, 1579
- Zhou R. et al., 2023, *AJ*, 165, 58

APPENDIX A: VISUALIZATION OF ALL UMAP PROJECTED DIMENSIONS

Fig. A1 shows a corner plot of the entire 10 projected dimensions of UMAP for the WAVES-Wide South region. This visualizes the entire feature space which HDBSCAN is clustering. UMAP 1 versus UMAP 2 appears to show the most variance, although higher dimensions also exhibit good clustering, such as UMAP 4 versus UMAP 7 and UMAP 6 versus UMAP 10.




Figure A1. The same as Fig. 4, but showing the entire 10-dimensional space of UMAP projections for the WAVES-Wide South region. Colours indicate the classification made by the baseline algorithm.

APPENDIX B: SOURCE NUMBER COUNTS

Fig. B1 shows the number counts of stars and galaxies according to our classifier as a function of brightness, radius and colour. The drop-off in the number of stars at magnitude $Z \sim 16$ is due to stars being masked out brighter than a *Gaia* magnitude $G < 16$.




Figure B1. The number counts of our classified stars and galaxies, as a function of Z -band magnitude in the upper panel, the log of half-light radius in the middle panel, and $J - K_s$ colour in the lower panel.

APPENDIX C: WAVES IMAGING OF EXEMPLAR ‘CHALLENGING’ GALAXIES

Fig. C1 shows some examples of the quasars, compact galaxies, and low surface brightness galaxies we classify in this work. All nine of these sources are correctly labelled by us as galaxies, but the baseline method (distinguished by the colour of the segments) is occasionally wrong. The imaging is taken from the WAVES segment viewer, <https://wavessegview.icrar.org/>, which can be used to view the sources identified by PROFOUND across the entire WAVES regions. It also showcases the imaging used for the input catalogue, a combination of KiDS, and VIKING photometry.




Figure C1. Nine exemplar ‘challenging’ galaxies from the ground truth set, which we successfully classify as galaxies, taken from <https://wavessegview.icrar.org/>, the WAVES segment viewer. The top row are quasars, the middle row are compact galaxies and the bottom row are low surface brightness galaxies. Galaxies are fainter going from left to right. The redshifts have been obtained from spectroscopy either from GAMA or DESI. The coloured segments are produced by PROFOUND, where purple, blue, and green segments indicate a baseline classification of galaxy, star, and ambiguous, respectively. The yellow segments indicate a source that has been masked. Each image is 1 arcmin across. Images are compiled using KiDS g and r bands, and VIKING Z band.

APPENDIX D: HEIGHTENED CLUSTERING AMPLITUDE FOR DISCREPANT CLASSIFICATIONS

We believe the heightened amplitude in the upper right plot of Fig. 10 (cluster galaxy, baseline ambiguous) to be due to a systematic issue within the WAVES photometric input catalogue. Plotting the sources which we label as star and the baseline method classes as galaxies (which exhibit the same heightened clustering) in the upper plot of Fig. D1, we see a systematic tiling pattern across square degrees. This can be traced to the process in which the photometric catalogue is built, in which square degree tiles of photometry are formed, and the seeing is averaged across the tile from different observational blocks. This potentially leads to a mis-estimation of the seeing across tiles. It can be seen in the lower panel of Fig. D1 sources with the 5th smallest percentile of seeing-subtracted radius, exhibiting an exaggerated tiling pattern across the square degrees. The baseline star-galaxy separation method is dependent on the seeing-subtracted radius of each source (see equation 2), which is why this tiling is evident in contradictory sources, in which our method and the baseline label disagree. The non-uniform distribution of these objects leads to an increase in the amplitude of the correlations at all scales.




Figure D1. Upper panel: RA and declination of sources in WAVES Wide North which have been classified by our method as star and by the baseline method as a galaxy. Lower panel: Sources with the smallest 5th percentile of seeing-subtracted radius, a metric that determines the baseline star-galaxy separation method. Both show evidence of systematic overdensities in square degree patterns.

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