



The LSST AGN Data Challenge: Selection Methods

Đorđe V. Savić^{1,2}, Isidora Jankov³, Weixiang Yu⁴, Vincenzo Petrecca^{5,6}, Matthew J. Temple^{7,16}, Qingling Ni⁸, Raphael Shirley^{9,10}, Andjelka B. Kovačević^{3,11,17}, Mladen Nikolić³, Dragana Ilić^{3,12,18}, Luka Č. Popović^{2,3}, Maurizio Paolillo^{5,6}, Swayamtrupta Panda^{13,14,19}, Aleksandra Čiprijanović¹⁵, and Gordon T. Richards⁴

¹Institut d'Astrophysique et de Géophysique, Université de Liège, Allée du 6 Août 19c, B-4000 Liège, Belgium; djsavic@aob.rs

²Astronomical Observatory, Volgina 7, 11000 Belgrade, Serbia

³University of Belgrade - Faculty of Mathematics, Department of Astronomy, Studentski trg 16, Belgrade, Serbia

⁴Drexel University, Department of Physics, 32 S. 32nd Street, Philadelphia, PA 19104, USA

⁵Department of Physics, University of Napoli "Federico II," via Cinthia 9, I-80126 Napoli, Italy

⁶INAF—Osservatorio Astronomico di Capodimonte, via Moiariello 16, I-80131 Napoli, Italy

⁷Instituto de Estudios Astrofísicos, Universidad Diego Portales, Av. Ejército Libertador 441, Santiago, Chile

⁸Max-Planck-Institut für extraterrestrische Physik (MPE), Gießenbachstraße 1, D-85748 Garching bei München, Germany

⁹Astronomy Centre, Department of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, UK

¹⁰Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

¹¹Key Laboratory for Particle Astrophysics, Institute of High Energy Physics, Chinese Academy of Sciences, 19B Yuquan Road, 100049 Beijing, People's Republic of China

¹²Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, D-21029 Hamburg, Germany

¹³Laboratório Nacional de Astrofísica—MCTI, R. dos Estados Unidos, 154—Nações, Itajubá—MG, 37504-364, Brazil

¹⁴Center for Theoretical Physics, Polish Academy of Sciences, Al. Lotników 32/46, 02-668 Warsaw, Poland

¹⁵Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510, USA

Received 2022 November 21; revised 2023 June 24; accepted 2023 June 27; published 2023 August 11

Abstract

Development of the Rubin Observatory Legacy Survey of Space and Time (LSST) includes a series of Data Challenges (DCs) arranged by various LSST Scientific Collaborations that are taking place during the project's preoperational phase. The AGN Science Collaboration Data Challenge (AGNSC-DC) is a partial prototype of the expected LSST data on active galactic nuclei (AGNs), aimed at validating machine learning approaches for AGN selection and characterization in large surveys like LSST. The AGNSC-DC took place in 2021, focusing on accuracy, robustness, and scalability. The training and the blinded data sets were constructed to mimic the future LSST release catalogs using the data from the Sloan Digital Sky Survey Stripe 82 region and the XMM-Newton Large Scale Structure Survey region. Data features were divided into astrometry, photometry, color, morphology, redshift, and class label with the addition of variability features and images. We present the results of four submitted solutions to DCs using both classical and machine learning methods. We systematically test the performance of supervised models (support vector machine, random forest, extreme gradient boosting, artificial neural network, convolutional neural network) and unsupervised ones (deep embedding clustering) when applied to the problem of classifying/clustering sources as stars, galaxies, or AGNs. We obtained classification accuracy of 97.5% for supervised models and clustering accuracy of 96.0% for unsupervised ones and 95.0% with a classic approach for a blinded data set. We find that variability features significantly improve the accuracy of the trained models, and correlation analysis among different bands enables a fast and inexpensive first-order selection of quasar candidates.

Unified Astronomy Thesaurus concepts: Active galactic nuclei (16); Quasars (1319); Active galaxies (17); Neural networks (1933); Classification (1907); Random Forests (1935); Support vector machine (1936); Catalogs (205); Convolutional neural networks (1938); Clustering (1908); Sky surveys (1464)

1. Introduction

A few percent of galaxies show enhanced emission from the nucleus that typically surpasses the stellar emission from the rest of the galaxy (e.g., Macuga et al. 2019); such sources are known as active galactic nuclei (AGNs). Emission from AGNs is produced by an accretion disk and ionized clouds

surrounding a central supermassive black hole (Salpeter 1964; Zel'dovich & Novikov 1964; Antonucci 1993; Netzer 2015). AGNs emit across the whole electromagnetic spectrum (Padovani et al. 2017) and are readily observed at large distances due to their high luminosity, with potential to be used as probes of cosmology (Panda et al. 2019; Risaliti & Lusso 2019; Czerny et al. 2022). AGNs have profound effects on the life and evolution of their entire host galaxy (Ferrarese & Merritt 2000; Gebhardt et al. 2000; Kormendy & Ho 2013). Outflows and jets interact with the local environment and release a large amount of energy capable of driving away the nearby gas, hence terminating star formation (Fabian 2012). Moreover, AGNs also have an impact on the surrounding hot intergalactic medium and play an active role in the evolution of the host galaxy clusters (Eckert et al. 2021). Therefore, each successful detection and observation of AGNs and

¹⁶ Fondecyt Fellow.

¹⁷ PIFI Research Fellow.

¹⁸ Humboldt Research Fellow

¹⁹ CNPq Fellow.



measurement of their physical properties is crucial for many areas of modern astrophysics and cosmology.

Ongoing and forthcoming large-scale photometric surveys (e.g., Zwicky Transient Facility—Bellm 2014; Pan-STARRS—Chambers et al. 2016; Gaia—Gaia Collaboration et al. 2016) will produce catalogs for a vast number of sources, which brings the astronomy into the new era of “big data.” The Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) is designed to address the main challenges for probing dark energy and dark matter, exploring the solar system, exploring the transient optical sky, and mapping the Milky Way (LSST Science Collaboration et al. 2017; Ivezić et al. 2019). With a state-of-the-art 3.2 gigapixel flat-focal array camera mounted on an 8.4 m telescope, LSST will cover the whole observable sky approximately every four nights in the optical/near-infrared *ugrizy* bands. The expected data volume of LSST is ~ 300 PB of raw data and $\sim 4 \times 10^{10}$ objects after 10 years of planned survey (Ivezić et al. 2019). Every night, LSST will monitor tens of millions of AGNs over an area of $\sim 18,000$ deg² (Luo et al. 2017; De Cicco et al. 2021). Although the actual number of AGNs that will be detected depends on the optimal observing strategy (Bianco et al. 2022), LSST will produce an AGN sample that supersedes the largest current AGN samples by more than an order of magnitude. This present work is a preparatory step toward producing a high-purity AGN sample with LSST.

To identify AGNs within LSST, the main challenge is to separate them from normal galaxies and stars. Construction of LSST’s AGN census will build upon a considerable volume of past work, making use of colors, proper motion, variability, and image morphologies. The idea of performing a multifaceted quasar selection (i.e., combining information from multiple observables) has long been proposed (Koo et al. 1986). However, the quality, quantity, and type of data from LSST will allow for a more complete AGN selection and thus these approaches should be considered from scratch.

Color selection has been widely used as the gold standard for identification of unobscured AGNs since their discovery (Koo & Kron 1982; Warren et al. 1991; Richards et al. 2002), but we expect AGN colors to change as a function of luminosity as we probe fainter objects toward LSST-like depths (e.g., Temple et al. 2021). While an application of modern statistical techniques to color data could be used to select AGNs, we expect the addition of multiparameter data to result in a purer and more complete AGN selection function.

As some AGNs and stars have similar colors, the fact that AGNs lack proper motions (unlike Galactic stars) has long been used as a discriminant (Sandage & Luyten 1967; Kron & Chiu 1981). LSST’s use of astrometric data will be no different in that regard. One way that LSST will be unique is in its ability to take advantage of differential chromatic refraction (DCR) of AGNs (Kaczmarczik et al. 2009; Yu et al. 2020), which is by making use of the astrometric offset of an emission-line object from that expected (in the astrometric solution) for a power-law source to break degeneracies in photometric redshifts of luminous AGNs (henceforth quasars or QSOs).

Selection of AGNs via time series due to their variability is another well known method (Bonoli et al. 1979; Trevese et al. 1989; Butler & Bloom 2011; De Cicco et al. 2019; Poulain et al. 2020). As quasars display higher fractional variability in their brightness than the average star and with different

characteristics than the typical variable star, variability will be a cornerstone of AGN classification for LSST (Suberlak et al. 2021). For luminous quasars, it has been shown that variability combined with colors works better for selection than variability alone (Peters et al. 2015). Lower-luminosity AGNs are expected to have the most variable nuclei, but increased contamination from the host galaxy could compromise variability-selection methods if insufficient care is taken. The implementation of difference image analysis (DIA) in the LSST reduction pipelines will completely revolutionize the detection of AGNs through variability by removing the contribution from the host galaxy (Zebrun et al. 2001; Bramich 2008; Kozłowski et al. 2010a, 2016). Finally, the addition of high-resolution images is expected to considerably increase the performance of the selection methods (Doorenbos et al. 2022).

Several data challenges (DCs) have been created in the past to facilitate the preparation of LSST, by other groups and for science use cases other than the study of AGNs (e.g., Hložek et al. 2020; Sánchez et al. 2020). In 2021, the LSST AGN Science Collaboration (AGNSC²⁰) organized a DC to get more people involved in the work needed for the AGN science with the upcoming LSST data. The main goal of the LSST AGN DC was to address the problem of AGN selection.

Unlike the previous data challenges, which relied on simulated data sets, this AGN DC utilizes real observational data. Major tasks also include establishing public training/test sets that will be used as a benchmark to test different machine learning (ML) algorithms. Five proposed solutions were submitted to the DC: one using a classical approach, and four applying ML-based AGN selection. We will present the solution using the classical approach and three ML-based solutions, while the one remaining ML-based solution is addressed by Doorenbos et al. (2022).

The paper is organized as follows: in Section 2 we address the data retrieval and the construction of training and blinded data sets. We elaborate on applied ML methods in Section 3. The results obtained from the various methodologies employed in this DC are presented in Section 4. We discuss further issues relevant to our work and summarize our findings in Section 5.

2. Construction of Data Sets

LSST will deliver three levels of data products and services: prompt data products that are computed and released within 24 hr of observation, data release data products that are computed during annual processing campaigns, and user-generated data products. The Data Products Definitions Document (DPDD²¹) is the ultimate reference for descriptions of the planned LSST data products and pipelines (see also Ivezić et al. 2019).

The input data for the AGNSC-DC were modified such that the column names and units used for different measurements (e.g., flux) comply with the DPDD standards for data release catalogs (DPDD, Section 4.3). We refer to distinct astrophysical bodies that emit light detected as “Objects” and individual instances (detection) of those objects as “Sources.” Observations from a specific point in time will appear in the `Source` tables in the data releases, while “coadded” (averaged/summed over time) information will appear in the `Object` tables. So-called “light curves” (brightness as a function of time) will appear in the `ForcedSource` tables

²⁰ <https://agn.science.lsst.org/>

²¹ <https://docushare.lsst.org/docushare/dsweb/Get/LSE-163>

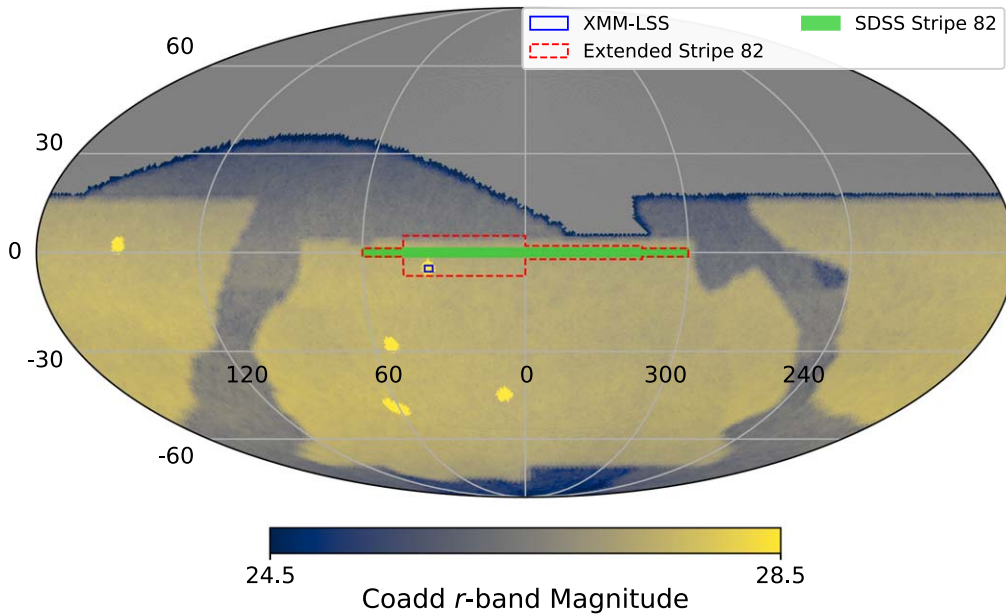


Figure 1. SDSS Stripe 82 (filled green), extended Stripe 82 (dashed red), and XMM-LSS (solid blue) areas projected on the LSST observable sky. The color map indicates the final depth of the coadds in the r band. The XMM-LSS region coincides completely with one of the LSST deep drilling fields (bright yellow regions). The map was generated using the LSST simulation project `OpSim` (Delgado et al. 2014).

(with summary statistics in the `Object` tables). “Simulated” training data will attempt to follow this data structure as closely as possible.

The data sets released in this data challenge are pulled from different sources (public archives) and put together to mimic the architecture of future LSST data release catalogs as much as possible, but without taking into account the expected number of objects that will appear in certain regions of LSST sky. Details on how the tables are constructed can be found under the “docs” folder in the main github repository (Yu & Richards 2021). Here, we provide a brief overview of the data sets.

The astronomical objects included in the release data set are drawn from three main sources: spectroscopically identified objects in an extended Sloan Digital Sky Survey (SDSS;²² York et al. 2000) Stripe 82 area with the spectroscopy collected from the 16th data release of SDSS (DR16; Ahumada et al. 2020), X-ray-detected and -classified objects in the XMM-LSS²³ (Pierre et al. 2007) region, and unidentified variable objects in the original SDSS Stripe 82 area²⁴ (Ivezic et al. 2007). Figure 1 illustrates the footprint of the source survey region on the LSST sky for the baseline observing strategy. The XMM-LSS area is encompassed by one of the LSST deep drilling fields.

The total number of objects (both combined) in the `Object` table is $\sim 440,000$ after removing ~ 5000 duplicates found in more than one of the sources above. The total number of epochs in the `ForcedSource` table is ~ 5 million. The total number of features (parameters) in the object table is 374. Features are divided into main categories, with the number of features in each category indicated in the parentheses:

1. `Astrometry(5)`: R.A., decl., proper motion, and parallax.

2. `Photometry(48)`: point and extended source photometry, in both AB magnitudes and fluxes (nJy).
3. `Color(10)`: derived from the flux ratio between different photometric bands.
4. `Morphology(6)`: a real-valued quantity between 0 and 1. Values closer to 1 are for extended sources while values closer to 0 indicate pointlike sources.
5. `Light Curve Features(302)`: extracted from the SDSS light curves if available.
6. `Redshift(2)`: both spectroscopic and photometric, wherever available.
7. `Class Labels(1)`: Star/Gal/Qso (Agn, highZQso), wherever available.

Distributions of a subset of 30 features are shown in Figure 2. It is notable that stars, galaxies, and quasars overlap in the feature space, but they may be separated by combining a selection of features containing the most information about interclass difference. For example, galaxies occupy different parts for morphology features when compared to stars and quasars, since the latter two are observed as point sources (Figure 2, fourth row, middle and right panels); and similarly for quasars when compared to stars and galaxies for variability features (Figure 2, bottom row, from left to right).

Astrometry measurements were obtained by matching the main catalogs (SDSS Stripe 82 and XMM-LSS) with Gaia EDR3²⁵ (Early Data Release 3, Gaia Collaboration et al. 2016, 2021) and the NOIRLab Source Catalog (NSC) data release 2 (DR2;²⁶ Nidever et al. 2021). Sources in NSC are extracted from reprocessed public images drawn from the NOIRLab Astro Data Archive.²⁷ The astrometry for NSC DR2 is calibrated using Gaia DR2 (Gaia Collaboration et al. 2018); its proper motion measurement achieves an accuracy of 0.2 mas yr^{-1} and a precision of 2.5 mas yr^{-1} relative to Gaia

²² <https://www.sdss.org/>

²³ <https://www.cosmos.esa.int/web/xmm-newton>

²⁴ <http://faculty.washington.edu/ivezic/sdss/catalogs/S82variables.html>

²⁵ <https://www.cosmos.esa.int/web/gaia/early-data-release-3>

²⁶ <https://datalab.noirlab.edu/nscdr2/index.php>

²⁷ <https://astroarchive.noirlab.edu>

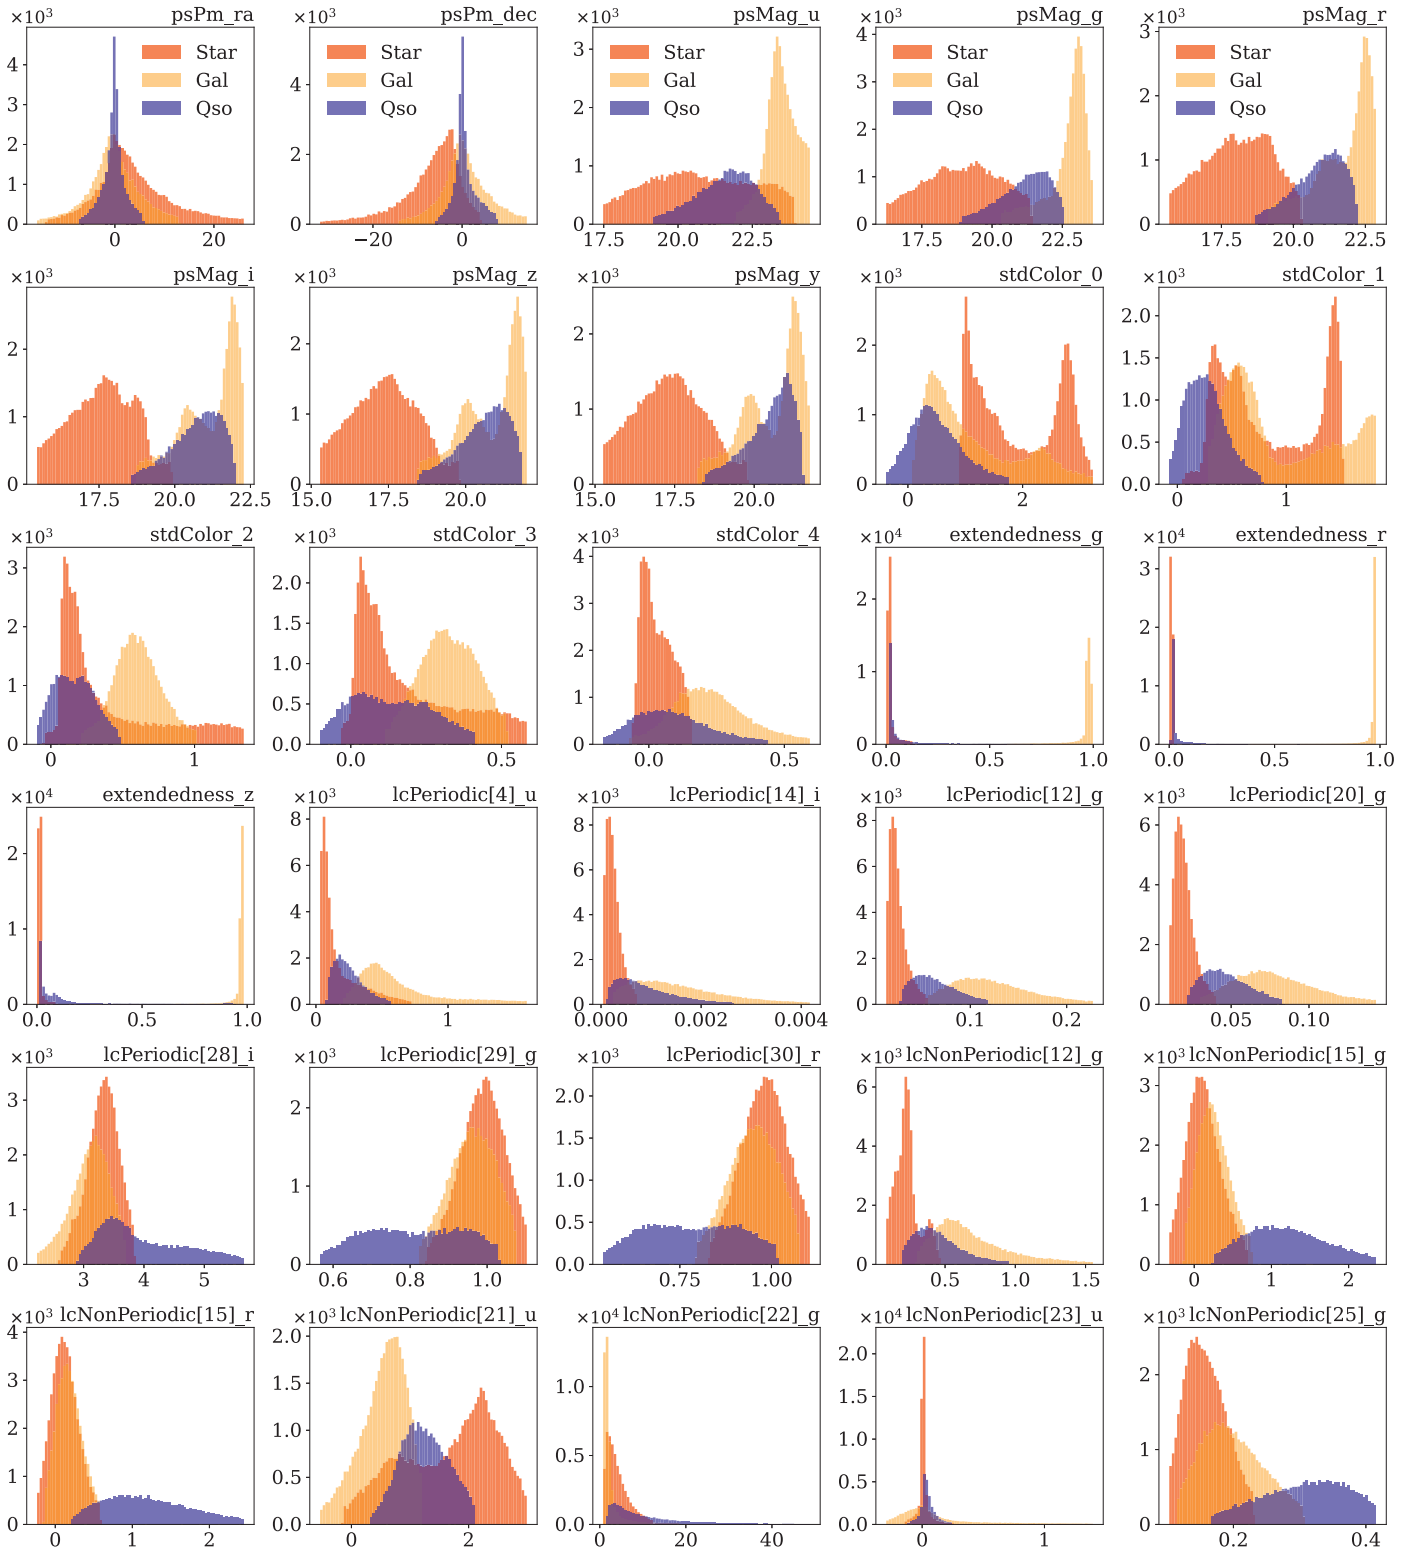


Figure 2. Distributions of 30 features drawn from astrometry (proper motion `psPm_ra` and `psPm_dec`); photometry (point-source magnitudes `psMag_u/g/r/i/z/y` and colors `stdColor_0/1/2/3/4`); morphology (`extendedness` `extendedness_g/r/z`); LC features (`lcPeriodic[12/20/29]_g`, `lcPeriodic[14/28]_i`, `lcPeriodic[30]_r`, `lcNonPeriodic[12/15/22/25]_g`, `lcNonPeriodic[4/21/23]_u`, `lcNonPeriodic[15]_r`). Color-coded per class: star (red), galaxy (orange), quasar (blue).

DR2 (Nidever et al. 2021). For objects with astrometry measurements found in both catalogs, we used the values from Gaia.

The photometry was assembled following a mix-and-match approach. In the extended Stripe 82 region, we cross-matched

our sources against the Dark Energy Survey (DES; Dark Energy Survey Collaboration et al. 2016) Data Release 2 (DR2; DES Collaboration et al. 2021) photometry catalog, the SDSS Stripe 82 coadded photometry catalog (Annis et al. 2014), and the SDSS DR16 single-epoch photometry catalog




Figure 3. Image cutouts for a star (top), a galaxy (middle), and a quasar (bottom), all randomly selected. From left to right are colors $u - g$, $g - r$, and $r - i$. Images are normalized to unity for each color.

(Ahumada et al. 2020). DES provides photometry in *grizY* bands and SDSS provides photometry in *ugriz* bands. When a source is matched with photometry in the same band from catalogs, we choose the photometry following the precedence of: DES DR2 > SDSS Stripe 82 coadd > SDSS DR16. In the XMM-LSS region, the *griz* photometry was collected from the HSC-VISTA joint catalog (see Section 2.1.2) and the *u*-band photometry comes from the Canada–France–Hawaii Telescope Legacy Survey catalog (Gwyn 2012).

The broadband colors were first derived using the flux ratios of two adjacent bands and then converted into magnitudes. The error on the color was computed following the standard method for uncertainty propagation. For the photometry taken from SDSS, the morphology (i.e., extendedness in LSST’s nomenclature) is defined as $1 - \text{probPSF}$. For DES photometry, we define extendedness as $1 - \text{class_star}$. Both probPSF and class_star describe how close the photometry is to a true point source. For the photometry taken from the HSC-VISTA catalog, the LSST pipeline directly outputs the extendedness column.

All objects, except for ~ 130 of them, in the `Object` table have pre-generated image thumbnails/cutouts of size 64×64 pixels from SDSS DR16. The objects without image cutouts were not removed from the main sample. Examples of image cutouts for a star, a galaxy, and a quasar, all randomly selected, are shown in Figure 3.

We list the percentage per class label and the assigning method:

1. `Star`(24.5%): spectroscopically confirmed stars (both variable and nonvariable)
2. `Gal`(55.9%): spectroscopically confirmed galaxies
3. `Qso`(19.0%): spectroscopically confirmed AGNs or quasars
4. `highZQso`(0.3%): a separate catalog of high-redshift ($z > 4.5$) quasars

5. `Agn`(0.3%): X-ray-classified AGNs in the XMM-LSS region and spectroscopically classified galaxies having emission properties consistent with being a Seyfert or LINER in the extended Stripe 82 region.

In addition to the described data sets that are publicly available, a distinct blinded data set, which is $\sim 10\%$ of the training data set, was constructed and set aside in order to evaluate the performance of the ML and non-ML methods submitted by the participants of the DC. The total object count for a blinded data set is $\sim 45,100$, among which ~ 1000 come from the XMM-LSS region, $\sim 44,000$ come from the Stripe 82 region, and 100 come from the separate high- z quasar catalog. Class labels and spectroscopic redshift were removed for those objects in the blinded data set. Around 21,000 objects have precomputed variability features.

We stress important caveats:

1. About $\sim 13,000$ objects have no labels, and they come from the SDSS Stripe 82 unidentified variables source catalog (Ivezić et al. 2007). About ~ 2500 objects from the XMM-LSS are classified using X-ray data and infrared photometry (Chen et al. 2018), and also have no spectroscopic redshift available.
2. Approximately $\sim 1\%$ of the objects do not have an optical counterpart, i.e., they are bright in X-rays, but are too faint to be detected in the optical band.
3. CARMA(1,0) and CARMA(2,1) fits are presented as is. Potential bad fits (e.g., perhaps due to limited temporal sampling and/or poor signal-to-noise ratio (S/N) of the photometry) were not removed.²⁸

In the following two subsections, we describe multiwavelength data sets, the high-redshift quasar catalog, and how the light-curve features are computed.

2.1. Multiwavelength Data

Multiwavelength tables are provided along with the `Object` table. The multiwavelength data are obtained by performing a positional cross-match between our source positions and other catalogs. A summary of class label distribution per observing mission is listed in Tables 1 and 2.

2.1.1. XMM-LSS

In the X-ray band of the XMM-LSS field, a number of XMM-Newton surveys of different sensitivities have been collected (e.g., Figure 3 of Brandt & Alexander 2015, also Table 2 of Chen et al. 2018). The X-ray source catalog from Chen et al. (2018) is adopted for our data set, which makes use of XMM-Newton observations taken from 2000 to 2017 in the XMM-LSS²⁹ field, including the 1.3 Ms new endeavor from the XMM-SERVS survey (Chen et al. 2018; Ni et al. 2021). With the XMM-SERVS survey, a flux limit of $6.5 \times 10^{-15} \text{ erg cm}^{-2} \text{ s}^{-1}$ over 90% of the XMM-LSS area is achieved in the 0.5–10 keV band. We only include X-ray sources that are classified as AGNs in the data set (Section 6 of Chen et al. 2018) for the source classification details. The X-ray AGNs are matched to other optical/IR catalogs with

²⁸ A robust goodness-of-fit metric for CARMA models is not available. The definition for a bad fit can also change given the problem at hand.

²⁹ <https://personal.psu.edu/wnb3/xmmservs/xmmservs.html>

Table 1
The Distribution of Class Labels per Catalog

Class Label	Agn/Qso	highZQso	Gal	Star
Stripe 82	73,000	90	213,000	106,000
GALEX	36,000	10	45,000	38,000
UKIDSS	36,000	30	87,000	92,000
Spitzer	12,000	20	27,000	30,000
Herschel	2500	10	39,000	12,000
FIRST	2000	50	43,000	250,000

Table 2
The Distribution of Class Labels per Catalog for Objects Found in the XMM-LSS Data Set

Class Label	Agn/Qso	Gal	Star
XMM-Newton	4000	N/A	N/A
GALEX	400	300	100
VISTA/VIDEO	3500	3500	370
Spitzer	3500	3500	350
Herschel	1500	1500	50

likelihood-ratio matching methods as described in Section 4 of Chen et al. (2018).

2.1.2. HSC and VISTA Joint Catalogue

We have provided a jointly processed optical and near-infrared data set from the HSC³⁰ (Hyper Suprime-Cam) Public Data Release 2 (Aihara et al. 2019) deep and ultra-deep regions, and the VISTA³¹ (Visible and Infrared Survey Telescope for Astronomy) VIDEO (Jarvis et al. 2012) surveys. The data set was produced using the LSST Science Pipelines as described in Bosch et al. (2018). An object detected in any one of the ten bands across these two surveys is measured in every band, ensuring that each object will have a measurement in each band. This data set is a prototype developed in preparation for the upcoming LSST data.

2.1.3. UKIDSS

In the near-infrared *YJHK* bands, we include 2'' diameter aperture magnitudes (A_{perMag3}) where available from the UKIDSS DR11plus (Hewett et al. 2006; Lawrence et al. 2007; Hambly et al. 2008; Hodgkin et al. 2009). SDSS Stripe 82 is partially covered by the UKIDSS Large Area Survey (LAS) in the *YJHK* bands to approximate depths of 20.2, 19.6, 18.8, and 18.2 respectively. The UKIDSS-LAS covers the original Stripe 82 footprint (shown in green in Figure 1) but not the full extended area used in this challenge. We remove duplicate detections from overlapping tiles using `PriOrSec`, and remove noise and saturated detections using `mergedClass`.

2.1.4. Herschel Forced Photometry

Far-infrared measurements from the Herschel³² Space Observatory come from HELP: The Herschel Extragalactic

Legacy Project (Shirley et al. 2021). This data set was produced by taking a prior list of Spitzer IRAC detections and providing full Bayesian probability posterior distributions on the object fluxes to account for blending in the low-resolution far-infrared maps at 250, 350, and 500 μm .

2.2. High-redshift Quasars

The catalog of high-redshift known quasars is constructed by collecting all quasars at $z \geq 4.5$ known before 2020 October. These quasars are mainly selected using the optical/near-infrared colors, based on the wide-field optical and infrared photometric surveys, e.g., SDSS (York et al. 2000), the Pan-STARRS1 survey (PS1, Chambers et al. 2016), the DESI Legacy Imaging Surveys (Dey et al. 2019), the Hyper Suprime-Cam Subaru Strategic Program survey (Aihara et al. 2018), the UKIRT Hemisphere Survey (Dye et al. 2018), the UKIDSS-LAS (Lawrence et al. 2007), the VISTA Hemisphere Survey (McMahon et al. 2013), and the Wide-field Infrared Survey Explorer (WISE, Wright et al. 2010).

About half of the $z < 6$ quasars are from SDSS quasar catalogs and the rest are from several major quasar surveys (e.g., McGreer et al. 2013; Bañados et al. 2016; Wang et al. 2016; Yang et al. 2019a). The $z > 6$ quasars were mostly collected from quasar surveys such as the SDSS high-redshift quasar survey (e.g., Fan et al. 2006; Jiang et al. 2016), the Canada–France High- z Quasar Survey (e.g., Willott et al. 2010), the PS1 distant quasar survey (e.g., Venemans et al. 2015; Bañados et al. 2016; Mazzucchelli et al. 2017), the Subaru High- z Exploration of Low-Luminosity Quasars project (e.g., Matsuoka et al. 2018), the DES quasar survey (e.g., Reed et al. 2017), and the reionization-era quasar survey (e.g., Yang et al. 2019b; Wang et al. 2019). Quasars included in this catalog are all identified through spectroscopic observations. Their redshifts are mainly from the quasar broad emission lines in the rest-frame UV (e.g., Ly α , Si IV, C IV, and Mg II), which result in a redshift uncertainty up to ~ 0.05 . A small number of these quasars have [C II]-based redshifts, with a typical uncertainty of < 0.001 .

2.3. Light-curve Features

The light curve (LC) features are computed for the sources that have corresponding time-domain data from SDSS. These features populate the `lcPeriodic` and `lcNonPeriodic` columns. The majority of the features computed are described in Richards et al. (2011, hereafter R11). Some additional features computed include those introduced by the Feature Analysis for Time Series (FATS) project³³ and best-fit CARMA (1,0) (continuous-time autoregressive moving average model, otherwise known as a damped random walk; DRW) and CARMA(2,1) (otherwise known as a damped harmonic oscillator; DHO) parameters (Kelly et al. 2009; Kasliwal et al. 2017; Moreno et al. 2019; Yu et al. 2022) obtained using EZTAO³⁴ (Yu & Richards 2022). Both the R11 and FATS features are computed using the CESIUM³⁵ software package. Examples of LCs for a random star, galaxy, and QSO for the SDSS *ugriz* bands are shown in Figure 4.

We note that a 5σ clipping (in magnitude) was applied before the variability metrics were computed. This decision is

³⁰ <https://www.naoj.org/Projects/HSC/>

³¹ <https://www.eso.org/public/teles-instr/paranal-observatory/surveytelescopes/vista/>

³² <https://www.herschel.caltech.edu/>

³³ <https://isadoranun.github.io/tsfeat/FeaturesDocumentation.html>

³⁴ <https://github.com/ywx649999311/EzTao>

³⁵ <https://cesium-ml.org/>

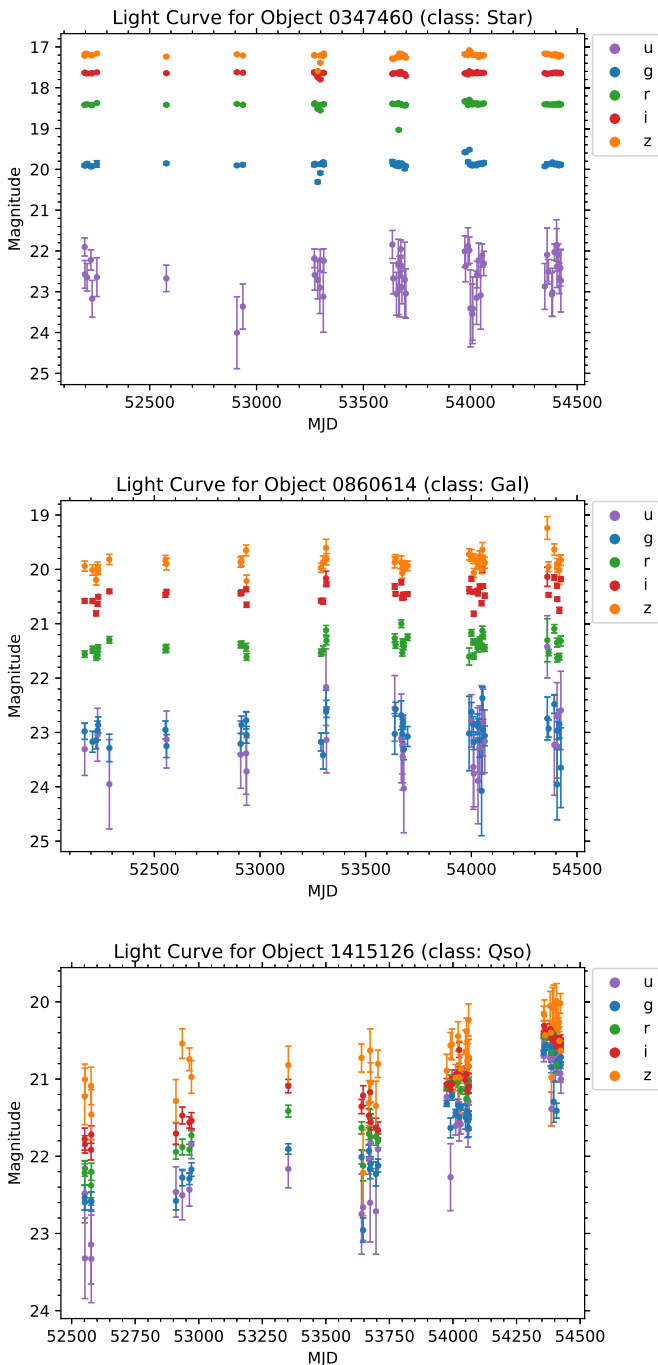


Figure 4. Example LCs for the SDSS *ugriz* bands: *u* band (purple), *g* band (blue), *r* band (green), *i* band (red), and *z* band (orange) for a star (top panel), galaxy (middle panel), and QSO (bottom panel). Time units are given in modified Julian date (MJD).

mostly driven by the “spurious dimming” of SDSS light curves as discussed by Schmidt et al. (2010). The total count of time-domain objects in Stripe 82 is $\sim 210,000$. Each SDSS filter, for which LC features have been computed, has the same number of visits and the same cadence. These numbers vary from object to object. The distribution of the number of visits per object is shown in Figure 5. The bulk of objects have between 30 and 70 visits. For comparison, LSST sources will have a larger number of observations (Bianco et al. 2022; Kovacevic et al. 2022; Raiteri et al. 2022; Czerny et al. 2023; Pozo Nuñez et al. 2023).

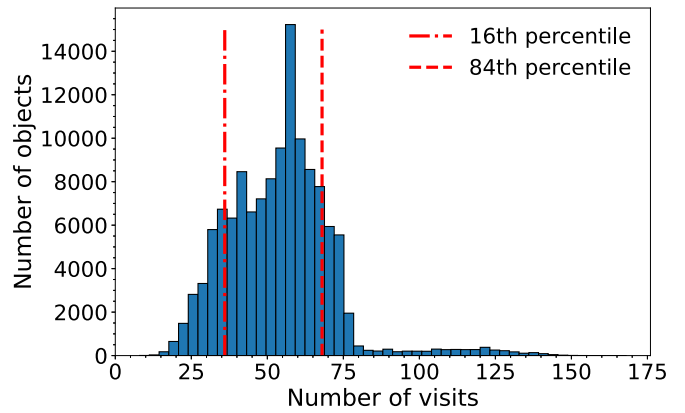


Figure 5. Distribution of the number of visits in the `ForcedSource` table for objects with known class label. Vertical red lines mark the 16th and 84th percentile values.

Table 3 summarizes the LC features and notation. A full description of the LC features is also publicly available.³⁶

3. Machine Learning Methods

The classification of sources from wide-field surveys is one of the most fundamental problems in astronomy. Efficient classification will be difficult if using classical techniques for manual inspection, and it is expected that ML applications will be helpful in automating the process. When trained on big astronomical data, ML methods tend to outperform traditional methods based on explicit programming (e.g., Banerji et al. 2010; Lochner et al. 2016; Baron 2019).

Broadly speaking, ML methods can be divided into supervised and unsupervised methods (Berry et al. 2019). Supervised learning is an ML paradigm that relies on labeled data for acquiring the input–output relationship information for classification and regression problems. The supervised ML methods we used are: support vector machine (SVM, Cortes & Vapnik 1995), random forest (RF, Ho 1995), extreme gradient boosting (XGB, Chen & Guestrin 2016), and artificial neural networks (ANNs, Cybenko 1989). These methods have been widely applied to numerous classification problems in astronomy (e.g., Baron 2019) and for AGN classification (Carballo et al. 2008; Cavuoti et al. 2014; Doert & Errando 2014; Chang et al. 2021; Chen et al. 2021; De Cicco et al. 2021; Poliszczuk et al. 2021). For more details on each of supervised ML methods we used, we refer to Berry et al. (2019).

Unsupervised methods discover hidden patterns in the data without the need for human intervention and are mostly used for clustering and dimensionality reduction. Commonly used are: *k*-means (Lloyd 1982) and Gaussian mixture of models (Reynolds 2009) for clustering; and principal component analysis (PCA, Jolliffe 1986), autoencoders (Kramer 1991), and t-distributed stochastic neighbor embedding (t-SNE, van der Maaten & Hinton 2008) for dimensionality reduction.

Deep learning (DL) is another broader family of ML methods that relies exclusively on the use of ANNs with a large number of hidden layers, hence deep. The abundance of imaging data in astronomy has naturally led to the application of deep convolutional neural networks (CNNs, Lecun et al. 2015).

³⁶ https://github.com/RichardsGroup/AGN_DataChallenge/blob/main/docs/04_LC_features.ipynb

Table 3
LC Features Divided into Two Groups Consisting of 33 Periodic and 29 Nonperiodic Features

LC feature	Description
lcPeriodic 0-3	Four best-fit CARMA (2, 1) /DHO parameters, fitted in flux using a three-point median filter with a 5σ clipping for removing outliers. Fitted data are for g -, r -, and i -band light curves only. Fitted light curves have more than 30 epochs.
lcPeriodic 4-32	Generalized Lomb–Scargle fit/parameters as described in the R11 paper. The exact matching from index to features is given below. Note that first relative phase (<code>rel phase 0</code>) is not included since it is negligible.
lcNonPeriodic 0-20	Nonperiodic features introduced in the R11 paper.
lcNonPeriodic 21	Variance divided by the median.
lcNonPeriodic 22	Reduced χ^2 for a constant model with given degrees of freedom:
	$\chi^2/\text{d.o.f.} = \frac{1}{N-1} \sum_{i=1}^N \left(\frac{m_i - \bar{m}}{\sigma_i} \right)^2,$
lcNonPeriodic 23	where \bar{m} is the inverse variance weighted average. Excess variance defined by Bernstein et al. (2018):
	$\sigma_{\text{sys}}^2 \equiv \langle \Delta m_i^2 \sigma_{\text{stat},i}^2 \rangle$ $\Delta m_i \equiv \frac{m_i - \bar{m}}{\sqrt{1 - w_i / \sum w_j}},$ $w_i \equiv \sigma_{\text{stat},i}^{-2},$ $\bar{m} \equiv \frac{\sum w_j m_j}{\sum w_j},$
lcNonPeriodic 24	where σ_{sys} is the excess variance and $\sigma_{\text{stat},i}$ is the photometric uncertainty. Normalized excess variance (Allevato et al. 2013):
	$\sigma_{\text{sys, norm}}^2 \equiv \frac{\sigma_{\text{sys}}^2}{N \bar{m}^2}.$
lcNonPeriodic 25	Range of a cumulative sum (Kim et al. 2011):
	$R_{\text{CS}} = \max(S) - \min(S),$ $S_l = \frac{1}{N\sigma} \sum_{i=1}^l (m_i - \bar{m}).$
lcNonPeriodic 26	The von Neumann ratio:
	$\eta = \frac{1}{(N-1)\sigma^2} \sum_{i=1}^{N-1} (m_{i+1} - m_i)^2.$
lcNonPeriodic 27-28	The best-fit CARMA(1,0)/DRW parameters. Light curves with fewer than 10 epochs were not fitted.
lcNonPeriodic 27	The driving amplitude σ , or β_0 in the CARMA notation (Kelly et al. 2009).
lcNonPeriodic 28	The characteristic timescale described by the DRW model, or $1/\alpha_1$ in the CARMA notation.

Note. Feature numeration starts with 0.

One advantage of DL is that it allows for simultaneous training of models with multiple inputs (e.g., catalog features + images) and for multiple outputs (e.g., regression and classification, Chollet et al. 2015). An important property of DL is the application of transfer learning (Tan et al. 2018), i.e., a model trained for one task is repurposed on a second, related task (Doorenbos et al. 2022). For the application of DL in astronomy, we refer to a review by Smith & Geach (2022).

Many ML methods are developed as a combination of the above-mentioned methods. Therefore, it is common to train a dimensionality reduction model and then train supervised or unsupervised models using the latent features. We used one unsupervised model built in this manner: deep embedding clustering (DEC: Xie et al. 2015; Guo et al. 2017), which consists of an ANN based on an autoencoder for dimensionality reduction and the preservation of the latent space, followed by clustering using latent features. Additionally, latent features are used for visualizing complex multidimensional data space in two or three dimensions (Clarke et al. 2020; Jankov et al. 2021), which is often the very beginning of the setup of the ML experiment.

4. Results

In this section, we summarize four out of the five “solutions” submitted to the DC (since one using CNNs and transfer learning has been already presented by Doorenbos et al. 2022). One submission used a non-ML method that extended the traditional approach based on color–color diagrams by adding magnitude standard deviation and the coefficients of correlation between light curves in two wave bands. A second submission trained supervised models—XGB, RF, SVM, ANN—and one unsupervised model (DEC) for the Star/Gal/QSO classification and clustering respectively by adding LC features. The third submission trained CNNs that use time-series data projected onto 2D images for Star/Gal/QSO classification. A final submission trained an RF on a subsample consisting of stars and quasars only, when all data features, except for flags and redshift, are taken into account. A summary of the contributions, ML methods, sizes of training and test samples, dimensionality (number of features), and model performance is listed in Table 4, indicating which of the coauthors submitted the solution. For a measure of performance, we report accuracy and completeness.

Table 4
Summary of the Contributions, Methods, Sample Sizes, Dimensionality, and Performance

Contribution	ML Method	Data Type	Sample Size		Dim.	Performance			
			Training	Blinded		Accuracy		Completeness	
						Training	Blinded	Training	Blinded
V.P. and M.P.	WEDGE+EXT+COL	tabular	10,000	6000	10	0.932(1)	0.950(2)	0.818(3)	0.831(4)
Đ.S., I.J., and SER-SAG	XGB	tabular	380,000	44,000	64	0.980(1)	0.970(1)	0.894(2)	0.834(1)
		tabular	128,000	15,000	64	0.983(4)	0.978(3)	0.929(3)	0.882(3)
		RF	128,000	15,000	64	0.982(4)	0.976(2)	0.920(4)	0.866(2)
		SVM	128,000	15,000	64	0.982(5)	0.976(3)	0.919(5)	0.870(2)
		ANN	128,000	15,000	64	0.982(4)	0.975(4)	0.914(5)	0.858(6)
		ANN	128,000	15,000	64	0.982(5)	0.974(4)	0.913(5)	0.852(6)
W.Y.	CNN	tabular+im.	152,000	17,000	1070	0.975(3)	0.975(3)	0.900(5)	0.860(4)
		tabular	61,000	3000	380	0.995(1)	0.994(1)	0.946(5)	0.924(5)
G.T.R.	RF	tabular	61,000	3000	380	0.995(1)	0.994(1)	0.946(5)	0.924(5)
L.D.	CNN	images	350,000	25,000	50,176	N/A	N/A	N/A	N/A

Notes. Columns from left to right: contributions, listed by participants; ML methods used; data type used for model training (tabular, images, or tabular+images); sizes of training and test samples; dimensionality or the number of features used; and the performance (purity and completeness) for the training and blinded data sets. The sample sizes are rounded down to the nearest thousand. The performance of the CNN model by L.D. (bottom row) is described by Doorenbos et al. (2022), and the model is not applicable (N/A) to images in our data sets due to low resolution.

In the following subsections, we elaborate in detail each of the submitted solutions while following the main workflow: feature selection, preprocessing (if needed), model training, and evaluation on test and blinded data sets. The uncertainties on the performance were estimated with k -fold cross-validation (Stone 1974). We split the training data set into $k = 10$ subsamples. One of the groups is used for testing while the rest are used for training. This process is repeated k times, with each group being used once for testing. The evaluation results are then averaged to give an overall training performance. During the cross-validation process, after every training iteration, we also evaluate the performance on the blinded data set that is always kept aside.

4.1. A Preliminary Classical Approach—V.P. and M.P.

As already mentioned, variability is an intrinsic property of AGNs and a promising selection tool, because both the timescales of the variations and the overall trends are different from those of other, mainly stellar, sources. Power spectra of AGNs are characterized by a typical red noise behavior (Kelly et al. 2009; Kozłowski et al. 2010b; Zu et al. 2013), with most of the variation arising on long timescales (Uttley et al. 2002). This timescale behavior means that the longer the observed baseline, the better the selection through variability (see, e.g., De Cicco et al. 2019). The LSST, with its dense and long coverage, will be the best survey to exploit this selection technique. Moreover, the implementation of DIA on the entire data set will also enable the selection of low-luminosity AGNs dominated by their host galaxy.

Unfortunately, the AGN DC does not contain difference images, and the population of confirmed AGNs in the region where optical light curves are available (SDSS Stripe 82) is biased toward bright quasars. In spite of the limitation of the archival data, which prevented us from testing the capabilities of DIA on AGN science, we examined some of the light-curve features not included in the original data set, before using ML with all the available LSST-like data products.

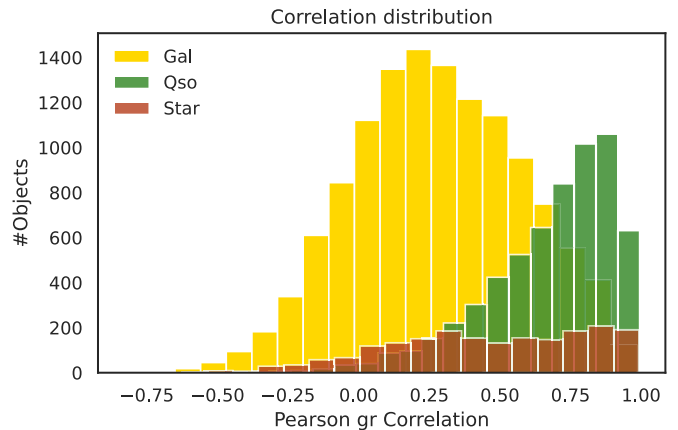


Figure 6. Distribution of the Pearson correlation coefficients for the gr bands for galaxies (yellow), QSOs (green), and stars (red).

Intensive X-ray, UV, and optical monitoring campaigns of AGNs show that, whatever the intrinsic physical mechanism (thermal propagating fluctuations, hydrodynamical instabilities, reprocessing of high-energy coronal photons by the accretion disk; McHardy et al. 2018) we can expect correlation between adjacent regions of the electromagnetic spectrum, such as the LSST bluer and redder bands. As SDSS images in the $ugriz$ bands were taken close in time (at 71.7 s intervals; Gunn et al. 1998), we calculated the Pearson correlation coefficient between pairs of light curves in different bands, along with the average magnitude and the standard deviation for each band. We restricted our analysis to the gri bands because they have the highest signal-to-noise ratio. A preliminary analysis that includes the uz bands has shown that they bring no improvement.

The distributions of the correlation coefficients for sources labeled as QSO, Star, or Galaxy clearly show that they belong to three different populations (Figure 6). We selected random samples of 10,000 objects and looked at how they were distributed on a plot of the standard deviation of g -band

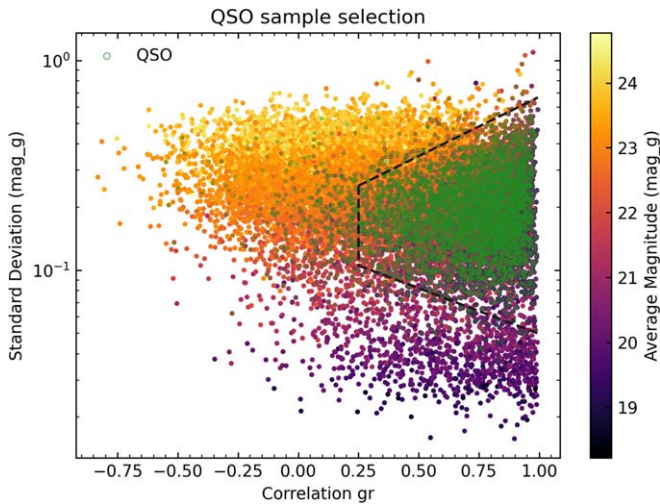


Figure 7. Standard deviation of the light curves vs. $g - r$ band correlation for a random sample of sources. Points are color-coded according to the g -band average magnitude. The black dashed lines define the *wedge* where QSOs (green points) tend to group.

magnitude versus the correlation coefficient of gr bands. The majority of quasars (more than 90%) tend to group in a well defined wedge of this space,³⁷ as shown in Figure 7. Thus, we identified the region of interest for our sources and tested the wedge selection on the blinded data set that was provided at the end of the AGN DC. We find that the light-curve variance and correlation among bands alone allow us to produce samples of AGNs with a completeness³⁸ of 90.9%, albeit with a low purity³⁹ of 52.0%. The low purity is expected since these two features alone tend to identify intrinsic correlated variability above a certain variance threshold, but do not include a full characterization of AGN properties to allow one to disentangle them from, e.g., stars.

It is possible to remove most of the contaminants and reach a purity of 95% (and decrease the completeness by less than 10%) by adding the extendedness and color information, which is particularly useful in the selection of candidate AGNs (Richards et al. 2002). As our AGNs were mainly bright quasars, we made a cut requiring the LSST extendedness parameter to be greater than 0.95. Then, we used an $r - i$ versus $g - r$ color-color diagram to select candidate quasars among the sources deriving from the wedge+extendedness criteria by defining a box in which they tend to group⁴⁰ (see Figure 8).

In order to evaluate the contribution of these three selection criteria, we tested them both singularly and combined together (see Table 5). The best result overall is obtained by the combination of all the selection criteria, with a purity of 95.0% (Tables 4 and 5). All the methods alone suffer strongly from contamination, and the extendedness seems to show the best performance. However, this is mainly due to the bias toward pointlike quasars in the AGN DC sample and does not reflect the true diversity of the AGN population. For this reason, it is also worth noting that the pair wedge+color returns completeness and purity of 83.0% and 85.7% respectively, without making any assumption on the morphology of the source

³⁷ Comprised in $x > 0.25$, $y < e^{1.3x - 1.7}$, $y > e^x - 2$.

³⁸ The ratio of QSOs selected to the total number of QSOs in the sample.

³⁹ The ratio of the selected QSOs to the total number of selected sources.

⁴⁰ Comprised in $-0.2 < g - r < 0.8$ and $-0.2 < r - i < 0.6$.




Figure 8. Color-color diagram showing the distribution of galaxies, QSOs, and stars selected by the wedge, with a cut in LSST extendedness > 0.95 . Gray dots represent the total sample, while the black dashed lines highlight the box where QSOs tend to group.

(Table 5). This rate of success is extremely promising from the LSST perspective of applying the selection directly to sources detected on difference images, where sources will be pointlike and contamination will only be due to transients, variable stars, or bogus events.

We point out that the lower performance with respect to ML approaches presented in the following subsections has to be expected, since we did not use any advanced light-curve feature. However, in spite of the lower performance, we demonstrate that correlation analysis among different bands enables a very fast and cheap first-order selection of candidate QSOs (and possibly less luminous AGNs). Furthermore, in the case of LSST where the low-luminosity AGN population will be detectable through DIA, correlation will help to probe the low-S/N regime, disentangling intrinsic variability from spurious uncorrelated noise.

4.2. AGN, Galaxy, Star Classification—D.S., I.J., and SER-SAG

Initial feature selection was done through trial and error. Bearing in mind that XGB supports missing (null) values by default, we first train a few XGB classifiers on a total set of $\sim 380,000$ containing all objects with known labels using a smaller subset of features. We further examined feature histograms with good visual separation between at least two classes (especially star-QSO and galaxy-QSO). After numerous tests, we find 64 optimal features for which we report accuracy of $98.0\% \pm 0.1\%$ and $97.0\% \pm 0.1\%$ on test and blinded data sets. A large fraction of the data set contains missing values for many of the features. We have performed a few methods of data imputation: median, hot deck, and k th nearest neighbor imputation, but the preliminary results are poor. Therefore, we keep the same 64 features and filter out objects with missing values, for which we obtain a subset of $\sim 128,000$ objects divided into $\sim 55,000$ stars, $\sim 45,000$ galaxies, and $\sim 28,000$ quasars. With such a setup, we proceed with ML steps before training separate XGB, RF, SVM, and NN models. The 25 most important features in the data set from the initial XGB analysis are given in Figure 9. More than half

Table 5
Results of the Selection of QSOs on the Blinded Data Set, Using a Classical Approach and a Combinations of Different Selection Criteria

	Galaxy	QSO	Star	Total	Completeness (QSO)	Purity (QSO)
Sample	13,066	6177	1955	21,238		
Wedge	4767	5614	404	10,789	90.9%	52.0%
Extendedness	1385	6029	1796	9211	97.6%	65.4%
Color	2763	5538	956	9257	90.3%	59.8%
Wedge+Ext	325	5492	331	6148	89.5%	89.3%
Col+Ext	612	5492	950	7054	89.5%	77.8%
Wedge+Col	716	5093	130	5989	83.0%	85.7%
Wedge+Ext+Col	138	5055	125	5318	82.4%	95.0%




Figure 9. Top 25 features ranked by importance for XGB classifiers run on a full data set.

of the top 25 features are LC features, both periodic and nonperiodic.

We divide the data subset(s) into training, validation, and test sets, which account for 70%, 20%, and 10% of the subset object count respectively. Data preprocessing is required before applying any of the ML methods. We standardize the training data to zero mean and a standard deviation of unity. Using mean and standard deviations obtained for each training feature, we normalize validation and test sets. We used both supervised and unsupervised ML methods with the goal of comparing the performance of the methods in order to establish a solid foundation for building more advanced models. For supervised learning, we were able to achieve a high accuracy for each method (e.g., XGB accuracy of 98.3(0.4)% and 97.8(0.3)% for the training and blinded data sets respectively) when the light-curve features are taken into account. However, the data set is dominated by bright quasars. The XGB and RF perform the best overall on this subset of data, which is very often the case for ML applications on tabular data. The performance rating measured by classification accuracy is XGB > RF > SVM > NN (Table 4). Confusion matrices are illustrated in Figure 10.

Feature importance for the latter four methods follows a similar trend to for the initial XGB analysis. Applying the same methods on the same set when the LC features are ignored reduces the classification accuracy to 95%.

In addition to tabular data, we investigated whether the addition of pixel information from image cutouts could contribute to achieving accuracy >99%. We investigated deep ANNs with joint image and tabular data as inputs, as well as

the use of an autoencoder with a simple bottleneck architecture for reduction of image dimensionality to ~ 10 latent features that are concatenated to the tabular data. No improvement has been observed in comparison with the models trained using tabular data only. The pixel resolution is simply too low for any morphology traits to be learned by the model.

Unsupervised methods are usually statistically weaker and are in general outperformed by supervised methods on similar problems. Essentially, clustering is grouping the instances based on their similarity without the help of class labels. In this ML paradigm, the model learns the similarities (i.e., the distance between instances in multidimensional parameter space), not the mapping function between input features and class label as is the case with supervised methods. The class naturally emerges in the representation of the data clusters. In this way, we can expect that latent features of stars will differ at a measurable level from the latent features of galaxies and QSOs. However, the number of clusters is in principle unknown. We set the number of clusters to be three, which is the number of ground-truth categories (star/galaxy/QSO). When the objects are passed to the model, it outputs soft assignments (the probabilities that given objects belong to given clusters). Each object is assigned a label of the most probable cluster. The performance of the DEC model is evaluated by unsupervised clustering accuracy (Xie et al. 2015, Equation (10)) computed in the following way. The optimal one-to-one matching between the set of cluster labels and the set of class labels is found, so as to maximize the agreement between the mapped cluster labels and true class labels on the given set of objects. Such mapped cluster labels constitute the predictions of the clustering model. Then the accuracy of the predictions given the true classes is computed.

We obtained 97.3(0.8)% and 95.9(0.6)% (Table 4) clustering accuracy on the same test and blinded data sets respectively as for the supervised methods. When applying k -means or a Gaussian mixture of models on the latent space, a clustering accuracy of $\sim 94\%$ is obtained; however, these numbers are heavily influenced by the initial weights of the ML model. Confusion matrices for DEC are shown in Figure 11.

For the data sets used in our analysis, we expect a similar ranking of methods with respect to measured performance within a possible framework developed by Shy et al. (2022), which attempts to incorporate measurement error in astronomical classification problems. We expect a drop in performance if such experiments are repeated.

4.3. Density Maps—W.Y.

We also explored the approach of first projecting time-series data onto 2D images (i.e., density maps) and then performing

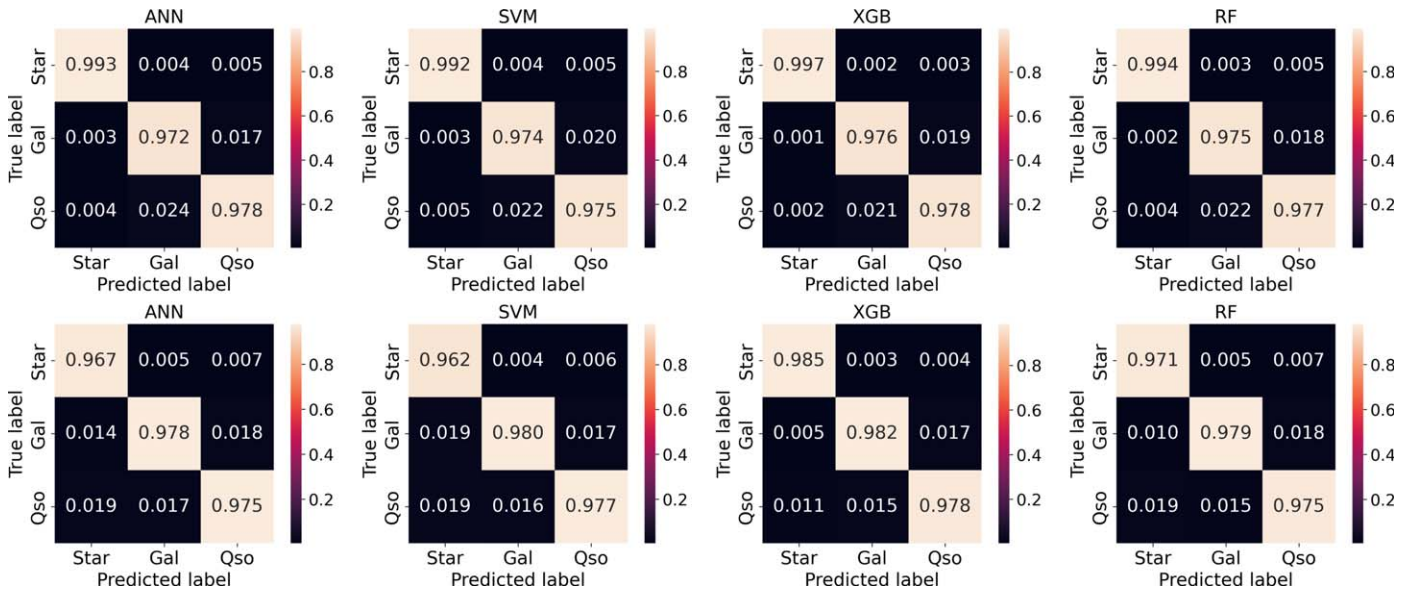


Figure 10. Confusion matrices normalized by purity for tabular data. From left to right: ANN, SVM, XGB, RF. Upper panels were computed on a test set, while lower panels were computed for a blinded set. True labels are placed on the vertical axis, while the predicted labels are on the horizontal axis.




Figure 11. Same as Figure 10, but for the DEC model. The performance is shown for the test set (left) and the blinded set (right).

classification using the constructed density maps (e.g., Mahabal et al. 2017). Density maps are essentially a 2D distribution of variability power in the timescale and magnitude space; example density maps and the distinguishing power brought by this method are shown in Figure 12. The actual algorithm used to generate those density maps is a modified version of the one used in Mahabal et al. (2017). Projecting time-series data onto 2D images of fixed size allows us to exploit the power of CNNs (Lecun et al. 2015)—the gold standard for ML-based image classification.

The total size of the subsample in the density map is $\sim 152,000$, divided into $\sim 32,000$ stars, $\sim 52,400$ galaxies, and $\sim 67,600$ QSOs. We removed objects with fewer than five epochs in their r -band light curves and dropped objects having missing values (e.g., NaN) for the features utilized to train the model (see next paragraph). Objects with missing values typically fall into one of two categories: (1) too faint to have been detected in all “ $ugrizY$ ” bands; (2) too extended (and/or faint) to have a reliably determined proper motion from Gaia EDR3 (Gaia Collaboration et al. 2021) or NSC DR2 (Nidever et al. 2021). Two versions of the above data set are created to test our ML model: one containing only objects with low temporal sampling (fewer than 30 epochs in the r band) and one containing 20% of all objects randomly selected from the parent sample (without any additional cut regarding temporal sampling).




Figure 12. Computed density map in the r band for one randomly selected quasar (left), galaxy (middle), and star (right).

Our final classification model starts with extracting latent features from the constructed density maps using a convolutional neural network and then concatenating the latent features with a selection of features/columns from the `Object` table to form the final input for the ANN that is used for classification. The columns picked from the `Object` table are the five optical colors (i.e., `stdColor[0-4]`), proper motion, and eight time-series features (i.e., `lcNonPeriodic[9, 13, 15, 19, 20, 22, 25, 26]`). The time-series features were cherry-picked from the feature importance rank of a random-forest classifier trained/tested on the full catalog of 374 features. A very high accuracy of 97.5% (Table 4) was achieved for both low-cadence and high-cadence subsamples, which demonstrates the robustness of the classification trained on density maps that are computationally inexpensive. Confusion matrices are shown in Figure 13.

4.4. Separating Quasars from Stars—G.T.R.

In this experiment, we use all data features except for flags and the redshift, which results in a subsample of $\sim 10,000$ quasars and $\sim 50,000$ stars after excluding all objects with any missing values. The excluded objects either fall into the same two categories of objects containing missing values listed in Section 4.3 or have fewer than 30 epochs in their r -band light curves. The strict cut on temporal sampling is given by the fact that only light curves with more than 30 epochs are fitted with `CARMA(2,1)`, whose parameters are `lcPeriodic[0-3]`. The main difference between this experiment and the one in




Figure 13. Confusion matrices for a CNN model with density maps included. The performance is shown for the test set (left) and the blinded set (right).

Section 4.2 is that the training data set contains more data features while at the same time fewer objects due to filtering out those with missing and/or NaN values. We used RF for separating quasars from stars with an incredibly high accuracy of 99.6% (Table 4) and a similar performance on the blinded data set. The white dwarfs and M-type stars can be seen to not be confused with quasars, which is a common source of error. Confusion matrices are shown in Figure 14.

5. Summary and Discussion

The LSST AGNSC-DC was designed to fulfill the growing need for efficient ML-based AGN selection methods. While we focus on simulated LSST data, these tools should benefit a much wider AGN community. We addressed the problem of star, galaxy, and quasar classification from both a classical perspective and an ML one while mimicking the future LSST catalog data. We used both supervised and unsupervised ML approaches. We followed the standard procedure of dividing the data set into training, validation, and test data sets. The blinded data set was revealed much later, after the submission of proposed solutions was finished. The performance of each method, the sample size, and the dimensionality are listed in Table 4. We obtain a high performance for supervised models and a slightly lower one for the unsupervised models for training and test data sets. We obtain a slightly lower performance for the blinded data set. The addition of LC features engineered from the *gri* bands significantly improves the classification accuracy by a few percent, but the data set is skewed toward bright quasars. Fainter AGNs where the host galaxy contamination is strong, as well as sources in the regions where colors overlap, will benefit strongly from variability (De Cicco et al. 2019). Further improvements of the ML methods that rely on LC features will be the inclusion of LC features computed using the missing *uzy* bands. We are limited by the number of visits per object, which is between 30 and 70 (Figure 5), and uneven temporal sampling (Figure 4), which are the two main obstacles to direct application of DL but, at the same time, the main reason for adopting LC feature engineering. The number of visits, higher by at least an order of magnitude, will allow us to train DL models directly on the light curves.

In addition, we find that 64×64 pixel image cutouts are not sufficient for extracting pixel-level information. However, in the parallel effort done on low-redshift/low-luminosity AGNs (Doorenbos et al. 2022), it was found that images of resolution 224×224 pixels can be used to directly extract meaningful information that can help in disentangling AGNs from non-AGN hosts. Moreover, the quality of images in all six bands that will be provided by LSST will allow us to develop DL




Figure 14. Confusion matrices for a star/quasar RF classification model. The performance is shown for the test set (left) and the blinded set (right).

models for separating AGNs from galaxies in a highly efficient manner at the cost of additional computing power.

In addition to LSST data, we can also make use of the data from other surveys—not just in the optical, but also in the X-ray, ultraviolet, infrared, and radio wave bands. Such data allow us to make use of differences in the spectral energy distributions of objects across a more extended wavelength range. For example, Myers et al. (2015) have shown that adding just a single infrared data point to SDSS *ugriz* optical bands significantly improves the classification probability for AGNs versus stars. However, these other surveys do not cover the same area of sky, and adding multiwavelength data will have the effect of breaking our largely monolithic LSST survey area into considerably smaller regions. Surveys at different wavelengths most often have different spatial resolutions, leading to situations where a single source in one survey is associated with multiple sources in another wave band. Forced photometry and probabilistic cross-matching are techniques to mitigate this problem (Budavári & Szalay 2008; Lang et al. 2016; Nyland et al. 2017; Buchner et al. 2021; Salvato et al. 2022). Simple recognition of the problem may enable it to be avoided in many cases; however, the depth of LSST means that there will be few sources that are truly isolated.

We conclude that LSST-like data enable the development of highly accurate star/galaxy/quasar classifiers, mainly using only *gri* bands and without spectral information. Promising results have been achieved in spite of limitations in resolution, survey area, cadence, and baseline of the SDSS data (compared to what LSST will be capable of). Thus the new survey will allow us not only to improve the performances of the algorithms successfully tested within the DC, but also to develop more sophisticated techniques, e.g., exploitation of pixel-level information on source cutouts, analysis of the population of weak AGNs dominated by their host galaxy through DIA, and identification of new types of AGNs. The AGN DC collection of multiwavelength data is an important legacy for future research beyond AGN science and will be included in the follow-up work.

Acknowledgments

We thank the anonymous referees for valuable comments that improved the quality of the manuscript. Prizes for participating in data challenges were funded by the LSST Corporation’s Enabling Science Program. Đ.S. acknowledges the support by the F.R.S. FNRS under grant PDR T.0116.21. Đ.S. and L.C.P. acknowledge support by the Astronomical Observatory (the contract N^o 451-03-68/2022-14/200002), through the grants by the Ministry of Education, Science, and Technological Development of the Republic of Serbia. Đ.S.

acknowledges support by the Science Fund of the Republic of Serbia, PROMIS N^o 6060916, BOWIE. D.I., A.B.K., and L.C.P. acknowledge funding provided by the University of Belgrade—Faculty of Mathematics (the contract N^o 451-03-68/2022-14/200104) through the grants by the Ministry of Education, Science, and Technological Development of the Republic of Serbia. D.I. acknowledges the support of the Alexander von Humboldt Foundation. A.B.K. and L.C.P. thank the support of the Chinese Academy of Sciences President’s International Fellowship Initiative (PIFI) for visiting scientists. M.J.T. acknowledges support from ANID (Fondecyt Proyecto 3220516). S.P. acknowledges financial support from the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) Fellowship (N^o 164753/2020-6) and the Polish Funding Agency National Science Centre, project N^o 2017/26/A/ST9/00756 (MAESTRO 9). A.C. acknowledges support from the Fermi Research Alliance, LLC under Contract N^o DE-AC02-07CH11359 with the U.S. Department of Energy (DOE), Office of Science, Office of High Energy Physics.

The authors thank Feige Wang and Jinyi Yang for constructing and providing the `highZQso` catalog.

This research makes use of the SciServer science platform (sciserver.org). SciServer is a collaborative research environment for large-scale data-driven science. It is being developed at, and administered by, the Institute for Data Intensive Engineering and Science at Johns Hopkins University. SciServer is funded by the National Science Foundation through the Data Infrastructure Building Blocks (DIBBs) program and others, as well as by the Alfred P. Sloan Foundation and the Gordon and Betty Moore Foundation.

Software: PYTHON (Van Rossum & Drake 1995), JUPYTER (Kluyver et al. 2016).

ML packages: NUMPY and SCIPY (van der Walt et al. 2011), PANDAS (McKinney et al. 2010), SCIKIT-LEARN (Pedregosa et al. 2011), KERAS (Chollet et al. 2015), TENSORFLOW (Abadi et al. 2016).

Data visualization: MATPLOTLIB (Hunter 2007), SEABORN (Waskom et al. 2017).

Data Statement





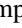


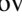


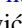



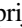
The DC was hosted on the SciServer⁴¹ platform. Each participant of the DC had a verified account on the platform. The data sets released in the DC are publicly available on Zenodo⁴² (Yu et al. 2022) under a Creative Commons Attribution 4.0 International Public License. The notebooks that reproduce the results of this work are designed as end-to-end and follow the flowchart of this work. The notebooks are publicly available on GitHub.⁴³ The notebooks have been successfully executed on different servers with different operating systems that can also support GPU acceleration, including SciServer.

Author Contribution Statement

Đ.S. and I.J. submitted the winning solution to the data challenge and drafted the manuscript; W.Y. and G.T.R. curated the data challenge; M.J.T., Q.N., and R.S. supplied the multiwavelength data and associated paper text; Đ.S., I.J., W.Y., V.P., A.K., M.N., D.I., L.P., M.P., A.C. and G.T.R.

contributed to solutions to the data challenge. All authors contributed to the writing and editing of the draft.

ORCID iDs

Đorđe V. Savić  <https://orcid.org/0000-0003-0880-8963>
 Isidora Jankov  <https://orcid.org/0000-0002-0431-2357>
 Weixiang Yu  <https://orcid.org/0000-0003-1262-2897>
 Vincenzo Petrecca  <https://orcid.org/0000-0002-3078-856X>
 Matthew J. Temple  <https://orcid.org/0000-0001-8433-550X>
 Qingling Ni  <https://orcid.org/0000-0002-8577-2717>
 Raphael Shirley  <https://orcid.org/0000-0002-1114-0135>
 Andjelka B. Kovačević  <https://orcid.org/0000-0001-5139-1978>
 Mladen Nikolić  <https://orcid.org/0000-0003-3105-7037>
 Dragana Ilić  <https://orcid.org/0000-0002-1134-4015>
 Luka Č. Popović  <https://orcid.org/0000-0003-2398-7664>
 Maurizio Paolillo  <https://orcid.org/0000-0003-4210-7693>
 Swayamtrupta Panda  <https://orcid.org/0000-0002-5854-7426>
 Aleksandra Čiprijanović  <https://orcid.org/0000-0003-1281-7192>
 Gordon T. Richards  <https://orcid.org/0000-0002-1061-1804>

References

- Abadi, M., Barham, P., Chen, J., et al. 2016, arXiv:1605.08695
 Ahumada, R., Prieto, C. A., Almeida, A., et al. 2020, *ApJS*, **249**, 3
 Aihara, H., AlSayyad, Y., Ando, M., et al. 2019, arXiv:1905.12221
 Aihara, H., Arimoto, N., Armstrong, R., et al. 2018, *PASJ*, **70**, S4
 Allevalo, V., Paolillo, M., Papadakis, I., & Pinto, C. 2013, *ApJ*, **771**, 9
 Annis, J., Soares-Santos, M., Strauss, M. A., et al. 2014, *ApJ*, **794**, 120
 Antonucci, R. 1993, *ARA&A*, **31**, 473
 Bañados, E., Venemans, B. P., Decarli, R., et al. 2016, *ApJS*, **227**, 11
 Banerji, M., Lahav, O., Lintott, C. J., et al. 2010, *MNRAS*, **406**, 342
 Baron, D. 2019, arXiv:1904.07248
 Bellm, E. 2014, in The Third Hot-wiring the Transient Universe Workshop, ed. P. R. Wozniak et al., **27**
 Bernstein, G. M., Abbott, T. M. C., Armstrong, R., et al. 2018, *PASP*, **130**, 054501
 Berry, M. W., Mohamed, A., & Yap, B. W. 2019, *Supervised and Unsupervised Learning for Data Science* (1st ed.; Berlin: Springer)
 Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, *ApJS*, **258**, 1
 Bonoli, F., Braccisi, A., Federici, L., Zitelli, V., & Formiggini, L. 1979, *A&AS*, **35**, 391
 Bosch, J., Armstrong, R., Bickerton, S., et al. 2018, *PASJ*, **70**, S5
 Bramich, D. M. 2008, *MNRAS Lett.*, **386**, L77
 Brandt, W. N., & Alexander, D. M. 2015, *A&ARv*, **23**, 1
 Buchner, J., Salvato, M., Budavári, T., & Fotopoulou, S. 2021, nway: Bayesian Cross-matching of Astronomical Catalogs, Astrophysics Source Code Library, ascl:2102.014
 Budavári, T., & Szalay, A. S. 2008, *ApJ*, **679**, 301
 Butler, N. R., & Bloom, J. S. 2011, *AJ*, **141**, 93
 Carballo, R., González-Serrano, J. I., Benn, C. R., & Jiménez-Luján, F. 2008, *MNRAS*, **391**, 369
 Cavauoti, S., Brescia, M., D’Abrusco, R., Longo, G., & Paolillo, M. 2014, *MNRAS*, **437**, 968
 Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
 Chang, Y.-Y., Hsieh, B.-C., Wang, W.-H., et al. 2021, *ApJ*, **920**, 68
 Chen, B. H., Goto, T., Kim, S. J., et al. 2021, *MNRAS*, **501**, 3951
 Chen, C. T. J., Brandt, W. N., Luo, B., et al. 2018, *MNRAS*, **478**, 2132
 Chen, T., & Guestrin, C. 2016, arXiv:1603.02754
 Chollet, A. O., et al. 2015, Keras, GitHub, <https://github.com/fchollet/keras>
 Clarke, A. O., Scaife, A. M. M., Greenhalgh, R., & Griguta, V. 2020, *A&A*, **639**, A84
 Cortes, C., & Vapnik, V. 1995, *Machine Learning*, **20**, 273
 Cybenko, G. V. 1989, *Math. Control Signal Systems*, **2**, 303
 Czerny, B., Cao, S., Jaiswal, V. K., et al. 2022, arXiv:2209.06563
 Czerny, B., Panda, S., Prince, R., et al. 2023, arXiv:2301.08975
 De Cicco, D., Bauer, F. E., Paolillo, M., et al. 2021, *A&A*, **645**, A103
 De Cicco, D., Paolillo, M., Falocco, S., et al. 2019, *A&A*, **627**, A33

⁴¹ <https://www.sciserver.org/>

⁴² <https://doi.org/10.5281/zenodo.6878414>

⁴³ https://github.com/RichardsGroup/AGN_DataChallenge

- Delgado, F., Saha, A., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, **9150**, 915015
- Dey, A., Schlegel, D. J., Lang, D., et al. 2019, *AJ*, **157**, 168
- Doert, M., & Errando, M. 2014, *ApJ*, **782**, 41
- Doorenbos, L., Torbaniuk, O., Cavuoti, S., et al. 2022, *A&A*, **666**, A171
- Dye, S., Lawrence, A., Read, M. A., et al. 2018, *MNRAS*, **473**, 5113
- Dark Energy Survey Collaboration, Abbott, T., Abdalla, F. B., et al. 2016, *MNRAS*, **460**, 1270
- DES Collaboration, Abbott, T. M. C., Adamow, M., et al. 2021, arXiv:2101.05765
- Eckert, D., Gaspari, M., Gastaldello, F., Le Brun, A. M. C., & O'Sullivan, E. 2021, *Univ*, **7**, 142
- Fabian, A. C. 2012, *ARA&A*, **50**, 455
- Fan, X., Strauss, M. A., Richards, G. T., et al. 2006, *AJ*, **131**, 1203
- Ferrarese, L., & Merritt, D. 2000, *ApJL*, **539**, L9
- Gebhardt, K., Kormendy, J., Ho, L. C., et al. 2000, *ApJL*, **543**, L5
- Gunn, J. E., Carr, M., Rockosi, C., et al. 1998, *AJ*, **116**, 3040
- Guo, X., Gao, L., Liu, X., & Yin, J. 2017, in Proc. Twenty-Sixth International Joint Conf. on Artificial Intelligence, 1753
- Gwyn, S. D. J. 2012, *AJ*, **143**, 38
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021, *A&A*, **649**, A1
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, **595**, A1
- Hambly, N. C., Collins, R. S., Cross, N. J. G., et al. 2008, *MNRAS*, **384**, 637
- Hewett, P. C., Warren, S. J., Leggett, S. K., & Hodgkin, S. T. 2006, *MNRAS*, **367**, 454
- Hložek, R., Ponder, K. A., Malz, A. I., et al. 2020, arXiv:2012.12392
- Ho, T. K. 1995, Proc. 3rd Int. Conf. on Document Analysis and Recognition, Vol. 1 (Piscataway, NJ: IEEE), 278
- Hodgkin, S. T., Irwin, M. J., Hewett, P. C., & Warren, S. J. 2009, *MNRAS*, **394**, 675
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ivezić, Ž., Smith, J. A., Miknaitis, G., et al. 2007, *AJ*, **134**, 973
- Jankov, I., Ilić, D., & Kovačević, A. 2021, in Proc. XIX Serbian Astronomical Conf. (Belgrade: Publications of the Astronomical Observatory of Belgrade), 241
- Jarvis, M. J., Bonfield, D. G., Bruce, V. A., et al. 2012, *MNRAS*, **428**, 1281
- Jiang, L., McGreer, I. D., Fan, X., et al. 2016, *ApJ*, **833**, 222
- Jolliffe, I. T. 1986, *Principal Component Analysis* (Berlin: Springer)
- Kaczmarski, M. C., Richards, G. T., Mehta, S. S., & Schlegel, D. J. 2009, *AJ*, **138**, 19
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- Kasliwal, V. P., Vogeley, M. S., & Richards, G. T. 2017, *MNRAS*, **470**, 3027
- Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, *ApJ*, **698**, 895
- Kim, D.-W., Protopapas, P., Byun, Y.-I., et al. 2011, *ApJ*, **735**, 68
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87, <https://eprints.soton.ac.uk/403913/>
- Koo, D. C., & Kron, R. G. 1982, *A&A*, **105**, 107
- Koo, D. C., Kron, R. G., & Cudworth, K. M. 1986, *PASP*, **98**, 285
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, **51**, 511
- Kovačević, A. B., Radović, V., Ilić, D., et al. 2022, arXiv:2208.06203
- Kozłowski, S., Kochanek, C. S., Ashby, M. L. N., et al. 2016, *ApJ*, **817**, 119
- Kozłowski, S., Kochanek, C. S., Stern, D., et al. 2010a, *ApJ*, **716**, 530
- Kozłowski, S., Kochanek, C. S., Udalski, A., et al. 2010b, *ApJ*, **708**, 927
- Kramer, M. A. 1991, *AICHe*, **37**, 233
- Kron, R. G., & Chiu, L. T. G. 1981, *PASP*, **93**, 397
- Lang, D., Hogg, D. W., & Mykytyn, D. 2016, The Tractor: Probabilistic astronomical source detection and measurement, Astrophysics Source Code Library, ascl:1604.008
- Lawrence, A., Warren, S. J., Almaini, O., et al. 2007, *MNRAS*, **379**, 1599
- Lecun, Y., Bengio, Y., & Hinton, G. 2015, *Natur*, **521**, 436
- Lloyd, S. 1982, *ITIT*, **28**, 129
- Lochner, M., McEwen, J. D., Peiris, H. V., Lahav, O., & Winter, M. K. 2016, *ApJS*, **225**, 31
- Luo, B., Brandt, W. N., Xue, Y. Q., et al. 2017, *ApJS*, **228**, 2
- LSST Science Collaboration, Marshall, P., Anguita, T., et al. 2017, arXiv:1708.04058
- Macuga, M., Martini, P., Miller, E. D., et al. 2019, *ApJ*, **874**, 54
- Mahabal, A., Sheth, K., Gieseke, F., et al. 2017, arXiv:1709.06257
- Matsuoka, Y., Iwasawa, K., Onoue, M., et al. 2018, *ApJS*, **237**, 5
- Mazzucchelli, C., Bañados, E., Venemans, B. P., et al. 2017, *ApJ*, **849**, 91
- McGreer, I. D., Jiang, L., Fan, X., et al. 2013, *ApJ*, **768**, 105
- McHardy, I. M., Connolly, S. D., Horne, K., et al. 2018, *MNRAS*, **480**, 2881
- McKinney, W., et al. 2010, in Proc. 9th Python in Science Conf., ed. S. van der Walt & J. Millman, 51
- McMahon, R. G., Banerji, M., Gonzalez, E., et al. 2013, *Msngr*, **154**, 35
- Moreno, J., Vogeley, M. S., Richards, G. T., & Yu, W. 2019, *PASP*, **131**, 063001
- Myers, A. D., Palanque-Delabrouille, N., Prakash, A., et al. 2015, *ApJS*, **221**, 27
- Netzer, H. 2015, *ARA&A*, **53**, 365
- Ni, Q., Brandt, W. N., Chen, C.-T., et al. 2021, *ApJS*, **256**, 21
- Nidever, D. L., Dey, A., Fasbender, K., et al. 2021, *AJ*, **161**, 192
- Nyland, K., Lacy, M., Sajina, A., et al. 2017, *ApJS*, **230**, 9
- Padovani, P., Alexander, D. M., Assef, R. J., et al. 2017, *A&ARv*, **25**, 2
- Panda, S., Martínez-Aldama, M. L., & Zajaček, M. 2019, *FrASS*, **6**, 75
- 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
- Pierre, M., Chiappetti, L., Pacaud, F., et al. 2007, *MNRAS*, **382**, 279
- Poliszczuk, A., Pollo, A., MaPlek, K., et al. 2021, *A&A*, **651**, A108
- Poulain, M., Paolillo, M., De Cicco, D., et al. 2020, *A&A*, **634**, A50
- Pozo Nuñez, F., Bruckmann, C., Deesamutara, S., et al. 2023, *MNRAS*, **522**, 2002
- Raiteri, C. M., Carnerero, M. I., Balmaverde, B., et al. 2022, *ApJS*, **258**, 3
- Reed, S. L., McMahon, R. G., Martini, P., et al. 2017, *MNRAS*, **468**, 4702
- Reynolds, D. 2009, *Gaussian Mixture Models* (Boston, MA: Springer), 659
- Richards, G. T., Fan, X., Newberg, H. J., et al. 2002, *AJ*, **123**, 2945
- Richards, J. W., Starr, D. L., Butler, N. R., et al. 2011, *ApJ*, **733**, 10
- Risaliti, G., & Lusso, E. 2019, *NatAs*, **3**, 272
- Salpeter, E. E. 1964, *ApJ*, **140**, 796
- Salvato, M., Wolf, J., Dwelly, T., et al. 2022, *A&A*, **661**, A3
- Sánchez, J., Walter, C. W., Awan, H., et al. 2020, *MNRAS*, **497**, 210
- Sandage, A., & Luyten, W. J. 1967, *ApJ*, **148**, 767
- Schmidt, K. B., Marshall, P. J., Rix, H.-W., et al. 2010, *ApJ*, **714**, 1194
- Shirley, R., Duncan, K., Campos Varillas, M. C., et al. 2021, *MNRAS*, **507**, 129
- Shy, S., Tak, H., Feigelson, E. D., Timlin, J. D., & Babu, G. J. 2022, *AJ*, **164**, 6
- Smith, M. J., & Geach, J. E. 2022, arXiv:2211.03796
- Stone, M. 1974, *J. Royal Stat. Soc. B*, **36**, 111
- Suberlak, K. L., Ivezić, Ž., & MacLeod, C. 2021, *ApJ*, **907**, 96
- Tan, C., Sun, F., Kong, T., et al. 2018, arXiv:1808.01974
- Temple, M. J., Hewett, P. C., & Banerji, M. 2021, *MNRAS*, **508**, 737
- Trevese, D., Pittella, G., Kron, R. G., Koo, D. C., & Bershadsky, M. 1989, *AJ*, **98**, 108
- Uttley, P., McHardy, I. M., & Papadakis, I. E. 2002, *MNRAS*, **332**, 231
- van der Maaten, L., & Hinton, G. 2008, *J. Mach. Learn. Res.*, **9**, 2579
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Van Rossum, G., & Drake, F. L., Jr 1995, *Python Reference Manual* (Amsterdam: Centrum voor Wiskunde en Informatica Amsterdam)
- Venemans, B. P., Bañados, E., Decarli, R., et al. 2015, *ApJL*, **801**, L11
- Wang, F., Wu, X.-B., Fan, X., et al. 2016, *ApJ*, **819**, 24
- Wang, F., Yang, J., Fan, X., et al. 2019, *ApJ*, **884**, 30
- Warren, S. J., Hewett, P. C., Irwin, M. J., & Osmer, P. S. 1991, *ApJS*, **76**, 1
- Waskom, M., Botvinnik, O., O'Kane, D., et al. 2017, mwaskom/seaborn: v0.8.1 (September 2017), v0.8.1, Zenodo, doi:10.5281/zenodo.883859
- Willott, C. J., Delorme, P., Reylé, C., et al. 2010, *AJ*, **139**, 906
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, **140**, 1868
- Xie, J., Girshick, R., & Farhadi, A. 2015, arXiv:1511.06335
- Yang, J., Wang, F., Fan, X., et al. 2019a, *ApJ*, **871**, 199
- Yang, J., Wang, F., Fan, X., et al. 2019b, *AJ*, **157**, 236
- York, D. G., Adelman, J., Anderson, J. E. J., et al. 2000, *AJ*, **120**, 1579
- Yu, W., Richards, G., Buat, V., et al. 2022, LSSTC AGN Data Challenge 2021, v1.1, Zenodo, doi:10.5281/zenodo.6878414
- Yu, W., & Richards, G. T. 2021, AGN_DataChallenge, GitHub, https://github.com/RichardsGroup/AGN_DataChallenge
- Yu, W., & Richards, G. T. 2022, EzTao: Easier CARMA Modeling Astrophysics Source Code Library, ascl:2201.001
- Yu, W., Richards, G. T., Vogeley, M. S., Moreno, J., & Graham, M. J. 2022, *ApJ*, **936**, 132
- Yu, W., Richards, G. T., Yoachim, P., & Peters, C. 2020, *RNAAS*, **4**, 252
- Zebrun, K., Soszynski, I., Wozniak, P. R., et al. 2001, *AcA*, **51**, 317
- Zel'dovich, Y. B., & Novikov, I. D. 1964, *SPhD*, **9**, 246
- Zu, Y., Kochanek, C. S., Kozłowski, S., & Udalski, A. 2013, *ApJ*, **765**, 106