



# Identifying Transient Hosts in LSST’s Deep Drilling Fields with Galaxy Catalogs

J. G. Weston<sup>1</sup> , D. R. Young<sup>1</sup> , S. J. Smartt<sup>1,2</sup> , M. Nicholl<sup>1</sup> , M. J. Jarvis<sup>2,3</sup> , and I. H. Whittam<sup>2,3</sup> 

<sup>1</sup>Astrophysics Research Centre, School of Mathematics and Physics, Queen’s University Belfast, BT7 1NN, UK; [jweston04@qub.ac.uk](mailto:jweston04@qub.ac.uk)

<sup>2</sup>Department of Physics, University of Oxford, Keble Road, Oxford, OX1 3RH, UK

<sup>3</sup>Department of Physics and Astronomy, University of the Western Cape, Robert Sobukwe Road, Bellville 7535, South Africa

Received 2025 November 27; revised 2026 February 23; accepted 2026 March 4; published 2026 March 30

## Abstract

The upcoming Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) will enable astronomers to discover rare and distant astrophysical transients. Host-galaxy association is crucial for selecting the most scientifically interesting transients for follow-up. LSST deep drilling field (DDF) observations will detect distant transients occurring in galaxies below the detection limits of most all-sky catalogs. Here, we investigate the use of preexisting, field-specific catalogs for host identification in the DDFs and a ranking of their usefulness. We have compiled a database of 70 deep catalogs that overlap with the Rubin DDFs and constructed thin catalogs to be homogenized and combined for transient-host matching. A systematic ranking of their utility is discussed and applied based on the inclusion of information such as spectroscopic redshifts and morphological information. Utilizing this data against a Dark Energy Survey sample of supernovae with pre-identified hosts in the XMM-Large Scale Structure and the Extended Chandra Deep Field-South fields, we evaluate different methods for transient-host association in terms of both accuracy and processing speed. We also apply light data-cleaning techniques to identify and remove contaminants within our associations, such as diffraction spikes and blended galaxies where the correct host cannot be determined with confidence. We use a lightweight machine learning approach in the form of extreme gradient boosting to generate confidence scores in our contaminant selections and associated metrics. Finally, we discuss the computational expense of implementation within the LSST transient alert brokers, which will require efficient, fast-paced processing to handle the large stream of survey data.

*Unified Astronomy Thesaurus concepts:* [Redshift surveys \(1378\)](#); [Sky surveys \(1464\)](#); [Transient detection \(1957\)](#); [Supernovae \(1668\)](#); [Catalogs \(205\)](#); [Transient sources \(1851\)](#); [Classification \(1907\)](#)

## 1. Introduction

The Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST) deep drilling field (DDF) mini-survey will observe five specific Rubin pointings with an enhanced cadence. These pointings have been selected due to the extensive amount of preexisting data in the regions from prior surveys carried out over the past 30 yr. With a higher cadence and greater depth than the Wide Fast Deep (WDF) survey, the DDF program will generate a greater number of transient alerts per square degree each night, making it a well-suited avenue for high-redshift transient science and cosmology (D. M. Scolnic et al. 2018; K. M. Hambleton et al. 2023; DES Collaboration et al. 2024; M. L. Graham et al. 2024; P. Gris et al. 2024). The rate of transient alert output (an estimated 10 million per night) will require extensive filtering, both to separate real candidates from fake detections and to distinguish more “interesting” transients from common ones (E. C. Bellm et al. 2020).

Efficient identification of real extragalactic transients and the subsequent selection of those meriting follow-up observations requires methods of “host matching,” wherein we determine the galactic host within which the transient occurred. With this additional contextual information, we can infer the reality of a transient, its redshift, and constraints on its likely spectroscopic or physical type. In the DDFs, we

expect to perform host matching thousands of times each night.

Multiple catalog-based methods of transient-host matching exist. A simple angular approach calculates the angular separation between a transient and its potential hosts to determine the “best” host as the one with the shortest distance (Y.-J. Qin et al. 2022). A redshift-based approach can be used to estimate physical distance from a given transient, provided redshift measurements are high quality (Y.-J. Qin et al. 2022). A morphology-based approach utilizes structure parameters such as galaxy axis lengths and position angle to determine a transient’s separation from the host in units of the galactic radius (M. Sullivan et al. 2006; R. R. Gupta et al. 2016). Image-based approaches utilize processed survey images in models such as convolutional neural networks to identify key associated structures (A. Gagliano et al. 2021; F. Förster et al. 2022). Implementations of transient-host-matching algorithms typically use a combination of the above approaches, depending on the amount of contextual data available near the transient location. In current alert brokers, such as Lasair (R. D. Williams et al. 2024), the crossmatching algorithm will incorporate wide-sky preexisting catalog data from prior surveys. Preexisting wide-field catalogs such as Pan-STARRS and the Sloan Digital Sky Survey (SDSS) are well matched to the depths reached by all-sky transient surveys and can therefore be used to identify the correct host for the majority of transients.

The current implementation of these pipelines may encounter issues with the depth of the LSST DDFs. It is unlikely that existing wide-field catalogs will contain the fainter and more distant galaxies that will be seen with Rubin.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Because the DDFs reach several magnitudes deeper than the Wide Fast Deep (WFD) survey, the surface density of detectable galaxies increases by more than an order of magnitude. As a result, any transient in a DDF will have far more potential hosts within a small angular radius than transients discovered in the WFD survey. This greatly increases the risk of ambiguous or incorrect associations unless deeper, field-specific catalogs are used. Until the first LSST data release, which is expected approximately 1.5 yr into the survey, we will not have wide-field catalogs that match the depth of the DDFs. However, the DDFs were selected by LSST for their preexisting deep coverage over a small area. In selecting the useful information from these existing narrow-field catalogs, we may be able to perform host matching at LSST depths within the DDFs from the beginning of the survey.

Seven independent Rubin alert brokers are currently designed to filter the LSST transient streams for a variety of science goals and provide contextual information, removing bogus transient candidates and promoting those most interesting for further science (E. C. Bellm et al. 2020; F. Förster et al. 2021; T. Matheson et al. 2021; A. Möller et al. 2021; R. D. Williams et al. 2024). Most involve some crossmatching with internal or external archival data. Real-time transient processing for large-scale surveys is a computationally expensive task, with LSST expected to detect thousands of extragalactic transient events each night in the WFD survey. Any changes to this crossmatching—whether to refine the methods of host association in the DDFs or to incorporate additional DDF data—must consider the impact on the broker pipeline’s efficiency.

This paper outlines the selection of preexisting catalogs within the Rubin DDFs for their use in transient-host matching and contextual classification in Rubin-LSST. We present a repository of the catalogs and the requisite Python code to produce summary analytics for ranking each set of data, before selecting the best catalogs for testing with samples of transients from the Dark Energy Survey (DES) in three of the DDFs. We compare a method that determines transient-host separation in units of the host’s semimajor axis against the computationally more expensive directional light radius (DLR) method (M. Sullivan et al. 2006; R. R. Gupta et al. 2016). We discuss the performance and cost of implementation within transient alert pipelines for LSST. Finally, we develop machine learning models to generate confidence scores for our associations and key features to indicate the reliability of a host match.

Following an introduction to LSST in Section 2, we give an outline of the transient alert brokers and contemporary methods for host matching in Section 3. Section 4 outlines the use and scope of the DDF catalogs, as well as our selection and ranking of the datasets. In Section 5, we apply host-matching methods to these datasets using samples of transient data from DES, Hyper Suprime-Cam (HSC), and the Asteroid Terrestrial Impact Last Alert System (ATLAS). Finally, in Section 6, we discuss the implementation of these catalogs and the attempted association techniques within the alert brokers.

## 2. LSST and the DDFs

The Vera C. Rubin Observatory’s LSST is a 10 yr all-sky survey aiming to cover the 19,600 deg<sup>2</sup> southern sky approximately 800 times with rapid revisit timescale

requirements (LSST Science Collaboration et al. 2009; Ž. Ivezić et al. 2019; SCOC 2025). The current observing strategy is set by the Survey Cadence Optimization Committee (SCOC) Phase 2 recommendations. The 3.2 gigapixel camera obtains observations in the six *ugrizy* filters, providing coverage between 320 and 1035 nm. A single LSST visit will reach a limiting magnitude of  $\sim 24$ – $25$  in the *g*, *r*, or *i* band (or somewhat shallower in the other bands); for co-added images, this depth approaches  $r \sim 28$ . It is expected that 30 Terabytes of nightly data will be generated.

The Rubin survey strategy divides its time between several types of observations. The WFD survey utilizes 80% of the total survey time to observe the majority (19,600 deg<sup>2</sup>) of the southern sky in areas with low-dust extinction or high stellar density. Mini-surveys target the north ecliptic spur, the dustier portions of the Galactic plane, and the south celestial pole. The smaller Near-Sun Twilight Microsurvey focuses observations around the ecliptic plane during twilight. Recommendations suggest that up to 3% of Rubin’s survey time be available for target of opportunity observations; the details of implementation are still under discussion.

### 2.1. The DDFs

The DDF program consists of five Rubin pointings observed at a high cadence with 30–40 visits per night over a total area of 57.6 deg<sup>2</sup> (D. M. Scolnic et al. 2018; P. Gris et al. 2024). The program provides several advantages for Rubin transient science. For a broader view of transients, the higher cadence will produce excellent light curves in six bands (all six being observed each night) that will help us better understand the physics of low-to-high-redshift supernovae (SNe) nuclear transients (NTs), and especially fast transients. The Dark Energy Science Collaboration aims to leverage the Type Ia supernovae (SNe Ia) data to extend the SN Hubble diagram by a factor of 2 past the WFD. Observations of SNe Ia in the DDFs will provide calibration for the high-quality light-curve templates required to fit the Rubin WFD data. The DDFs are chosen both for their large amounts of archival data from previous surveys across the spectrum and for their absence of bright foreground sources that cause ghosting, detector saturation, bleeds, and diffraction spikes. They consist of:

1. *COSMOS*. The Cosmic Evolution Survey (COSMOS) field (N. Scoville et al. 2007). COSMOS includes X-Ray to radio imaging from a range of space (Hubble, Spitzer, XMM, Chandra, among others) and ground-based telescopes (Keck, Subaru, MeerKAT, VISTA, etc.) (A. Finoguenov et al. 2007; D. B. Sanders et al. 2007; E. Schinnerer et al. 2007; Y. Taniguchi et al. 2007; H. J. McCracken et al. 2012; M. J. Jarvis et al. 2013; F. Civano et al. 2016; I. Heywood et al. 2022; J. R. Weaver et al. 2022). It is the most extensively covered of the DDFs and continues to be observed by these other instruments.
2. *ELAIS-S1*. The European Large-Area ISO Survey-S1 (ELAIS-S1; S. Oliver et al. 2000). ELAIS-S1 was the largest targeted survey by ESA’s Infrared Space Observatory using the ISOCAM and ISOPHOT instruments. Extensive archival data exist at radio, optical, near-infrared (NIR), and X-ray wavelengths

**Table 1**  
DDF Locations

Field	ELAIS-S1	XMM-LSS	ECDF-S	COSMOS	EDF-Sa	EDF-Sb
R.A. [deg]	9.45	35.57	52.98	150.11	58.9	63.6
Decl. [deg]	-44.02	-4.82	-28.12	2.23	-49.32	-47.6
Gal $l$ [deg]	311.29	171.1	224.07	236.78	257.9	254.48
Gal $b$ [deg]	-72.88	-58.91	-54.6	42.13	-48.46	-45.77
Eclip $l$ [deg]	346.66	31.59	40.81	151.39	32.0	40.97
Eclip $b$ [deg]	-43.2	-17.92	-45.44	-9.34	-66.61	-66.6


(F. La Franca et al. 2004; S. Berta et al. 2006; S. Puccetti et al. 2006).

3. *XMM-LSS*. The XMM-Large Scale Structure (LSS; M. Pierre et al. 2004). The XMM-LSS field has been observed by XMM-Newton in X-ray wavelengths. Subsequent follow-up observations of the field yield coverage up to radio, including the Canada-France-Hawaii Telescope (CFHT), Spitzer, the Low-Frequency Array, and Subaru (O. Le Fèvre et al. 2004).
4. *ECDF-S*. The Extended Chandra Deep Field-South (ECDF-S) or Wide Chandra Deep Field-South (W-CDF-S; B. D. Lehmer et al. 2005; Y. Q. Xue et al. 2016). The Chandra Deep Field-South (CDF-S), originally observed by the Chandra X-Ray Observatory, subsequently underwent extended-field follow-up observations by Chandra, the Hubble Space Telescope, the Very Large Telescope, and Spitzer, among others (C. N. Cardamone et al. 2010; J. D. Silverman et al. 2010).
5. *EDF-S*. The Euclid Deep Field-South (EDF-S; C. Scarlata et al. 2019; SCOC 2025). EDF-S has less multiwavelength data compared to the four DDFs above, and so was finalized as a collaborative observation between Rubin and the Euclid survey. The aim is to combine LSST observations with Euclid’s high-resolution optical/NIR imaging to yield better-calibrated photometry and more accurate photometric redshifts for cosmology (J. Rhodes et al. 2017; Euclid Collaboration et al. 2022).

The Phase 3 recommendations of the LSST SCOC are that 6%–7% of the total survey time be dedicated to the DDFs, with each receiving approximately 20,000 nightly visits over 10 yr (D. M. Scolnic et al. 2018; SCOC 2025). The exception to this is the COSMOS field, which will receive 20,000 visits by the end of Y3. There will then be a “complete” DDF with which to address the program’s science goals. EDF-S consists of two Rubin pointings but will be observed over its entire area in 20,000 visits, and as such, will only achieve half the depth of the other DDFs.

The DDF pointings are summarized in Table 1. For EDF-S, observations are divided between the two pointings, EDF-Sa and EDF-Sb. Enlarged pointings are displayed for reference in Figure 1.

The program provides several advantages for Rubin transient science. For a broader view of transients, the high cadence will produce excellent light curves in six bands for understanding the physics of low-to-high-redshift SNe, NTs, and especially fast transients. The Dark Energy Science Collaboration aims to leverage this SNe Ia data to extend the SN Hubble diagram by a factor of 2 past the WFD. Observations of SNe Ia in the DDFs will provide calibration for the high-quality light-curve templates required to fit the Rubin WFD data.



**Figure 1.** Pointings for the DDF mini-survey. Although the scale saturates at 1200 visits, the DDF pointings, clearly visible in yellow, will each receive 10,000–20,000 visits (P. Yoachim 2025).

### 3. Transient-host Matching

Crossmatching refers to the process of determining nearby “associated” records in preexisting data using observational data for a given object, such that historical information can inform our present understanding of the object, and vice versa. For example, a detected transient’s coordinates may be matched against a known cataclysmic variable (CV) or an active galactic nucleus (AGN). Host matching refers to crossmatching of an extragalactic transient (such as an SN) with an associated host galaxy. Contextual classification, in this case, refers to using the crossmatched data for an object to infer its reality and constrain its type. The detail of this classification can range from a simple real/bogus separation to distinguishing stellar from NTs, and at its most precise, to differentiating between specific SN types.

Different SN types are defined by their spectra, which can only be obtained through follow-up observations. While photometric analysis of the SN light curve provides estimates of an object’s redshift and type, these estimates are much less reliable, and the inclusion of photometrically classified SNe can contaminate spectroscopic sample data. In both cases, additional contextual information can assist in the ultimate classification; the secure identification of a host with a spectroscopic redshift further assists in the classification of the transient. Global host galaxy properties such as stellar mass, age, metallicity, and star formation rates can provide context for a transient’s evolution and (in the case of SNe) explosion mechanisms (J. P. Anderson et al. 2010; P. Wiseman et al. 2020).

As LSST returns a high number of potential transients each night, the allocation of follow-up observations will become

more competitive. The number of LSST transients that can be followed up spectroscopically will be orders of magnitude smaller than the number of transients discovered. It is therefore not only important that all available data is utilized, but that the method of crossmatching strikes the right balance between efficiency and accuracy to appropriately process the data.

Several approaches to crossmatching currently exist. The simplest is a spatial crossmatch, which returns the source with the smallest angular separation from the transient location. The object with the lowest angular separation is then selected as the most likely crossmatch. This is reasonable when the angular separation is so low that the transient is confidently identified as being spatially coincident with the cataloged object, as with AGN or variable stars (VSs).

For extragalactic transients outside of galaxy centers, a morphology-based approach may be better motivated. If we have information on the size and shape of a crossmatched galaxy (e.g., axis lengths and position angle in the direction of the transient), we can evaluate the likelihood of association based on the projected distance along the galaxy’s morphology (M. Sullivan et al. 2006; R. R. Gupta et al. 2016). If redshift information is available, we can also convert the angular separation to a physical distance and select only galaxies with a physically plausible distance for association. For truly accurate physical distance measurements, we must use either high-accuracy photometric redshifts or spectroscopic redshifts.

In an ideal scenario, we are able to combine these approaches, dependent on need, to reach a given accuracy threshold for our host selection. A host identified via a physical approach with a low-accuracy redshift may be verified via a morphology-based approach. In low-density areas, it may even be best to utilize a spatial approach as a lighter-weight method.

### 3.1. Sherlock

Sherlock (D. Young 2023a) is a Python package and MariaDB database of catalogs utilized by the Lasair transient alert broker (R. D. Williams et al. 2024), the ATLAS and Pan-STARRS survey pipelines (K. W. Smith et al. 2020; M. D. Fulton et al. 2025; H. F. Stevance et al. 2025), and the Public ESO Spectroscopic Survey for Transient Marshall (S. J. Smartt et al. 2015). Using a set of all-sky catalogs provides spatial crossmatching for a given astrophysical object to aid in contextual classification.

Sherlock uses the catalogs to identify transients that are synonymous with a cataloged source. These include VSs, known CVs, known AGN, and other transients found in the center of galaxies (NTs). These transients fall within the “synonym radius” of a cataloged source of the same type. Sherlock also attempts to associate nonsynonymous transients with their cataloged counterparts. A detection falling within an “association radius” of a cataloged star above a certain brightness will be classified as an artifact related to a bright star. In contrast, a detection found in association with a cataloged galaxy will be classified as an SN. If Sherlock fails to match the transient against any cataloged source, it is classified as an orphan.

The synonym radius used in classifications of VS, CV, AGN, and NT is mainly determined by the resolution of a given survey’s images. The association radius is more dependent on metrics such as the brightness of a cataloged source or (in the case of galaxies) its angular size. For potential hosts with morphological data, associations are rejected if a

transient is separated from a galaxy by more than 2.4 times the galaxy’s semimajor axis radius. In the case where there is a distance measurement or redshift for a galaxy, allowing us to calculate physical separations between the host-transient pair, Sherlock rejects matches where the distance is greater than 50 kpc (S. Srivastav et al. 2026).

Sherlock performs contextual classification of sources to return a list of possible synonyms/associations for each transient. A ranking algorithm is then used to select the most likely classification.

Sherlock’s database currently contains datasets from a number of all-sky surveys and source-specific catalogs. All-sky surveys include data from DESI Legacy Imaging Survey DR10 (A. Dey et al. 2019), Gaia DR3 (Gaia Collaboration et al. 2018), Pan-STARRS (Y. Tachibana & A. A. Miller 2018; K. C. Chambers et al. 2019), SDSS DR12 (S. Alam et al. 2015), 2MASS (M. F. Skrutskie et al. 2006) and the Guide Star Catalog (B. M. Lasker et al. 2008). For the source-specific data, Sherlock utilizes catalogs such as the Million Quasars Catalog (E. W. Flesch 2023), the Downes CV Catalog (R. A. Downes et al. 2001), and the Ritter Cataclysmic Binaries Catalogue (H. Ritter & U. Kolb 2003). In addition to the source catalogs above, Sherlock also contains spectroscopic and distance measurements from the NED-D Galaxy Catalog (I. Steer 2017), the NASA/IPAC Extragalactic Database via D. Young (2023b), and the LASr-GPS volume-limited galaxy catalog (D. Asmus et al. 2020). More catalogs are added to the library as they become publicly available; currently, the database consumes approximately 4.5 TB.

Currently, the above catalogs are well matched to wide-field, relatively shallow surveys. In the DDFs, we require adjustments—first, to account for the greater depth, field-specific catalogs must be implemented for associations to be achieved for higher-redshift sources. A large amount of data exists for the DDFs to achieve this. Second, to account for the more crowded nature of the catalog data in the DDFs, the synonym and association radii, as well as crossmatching methods, will need to be reviewed to ensure the software maintains a satisfactory level of accuracy in the new environment.

### 3.2. Host Matching in LSST

Rubin Alert packets are created via the detection of all sources in the corresponding LSST difference image with a signal-to-noise ratio greater than 5; it is predicted that 10 million total alerts per night will be produced. Not all of these will be real objects—with LSST expecting to discover 2 million quasars, 10 million SNe, and 50 million VSs across the 10 yr survey, the ratio of alerts to real objects may be high (M. L. Graham et al. 2024).

The Rubin alert brokers are a series of seven independent projects designed to filter the LSST transient streams for different science goals and provide additional contextual information (E. C. Bellm et al. 2020). Each broker has been developed independently, but typical broker tasks might include filtering alerts based on their content and history, and performing crossmatching with archival data. This filtered/enhanced data stream can then be used to identify and prioritize objects for follow-up observation, with or without a classification step.

The Lasair science platform (R. D. Williams et al. 2024) is the UK Community Broker for transient alerts. For host

matching and contextual classification, Lasair uses Sherlock to crossmatch against the preexisting data as detailed in the previous section.

### 3.3. Morphology-based Association

A morphology-based association determines the most likely host using galaxy size, shape, and position angle on the sky (measured east through north). By calculating which galaxy is “most” extended toward the transient, we can identify it as the most likely host association.

Calculation of the axis lengths and position angles of galaxies is typically done through software such as SExtractor (E. Bertin & S. Arnouts 1996). The program builds a catalog of astronomical sources from isophotal measurements of survey images. In calculating these measurements for an elliptical shape, SExtractor calculates the semimajor and semiminor axes A and B with the position angle THETA. A and B are the maximum and minimum spatial dispersion of the object profile along any direction. This dispersion and the “edge” of a galaxy are dependent on the detection threshold, often expressed in units of rms noise. A lower threshold accounts for fainter regions and can lead to larger radii.

Morphological crossmatching approaches can incorporate a range of details for higher accuracy at the cost of computational speed. The standard and straightforward method is to treat the elliptical shape of the galaxy as being strictly circular; i.e., the galaxy is equally extended in all directions (either at radius A (semimajor axis length) or B (semiminor axis length), an average of the two, or a mean half-light radius). We calculate the angular separation between the transient and the Galactic center and express it in units of this effective radius to give a dimensionless distance. In subsequent sections of this paper, we calculate the distance using the semimajor axis A and denote this calculated distance as the “A-value.”

A more sophisticated approach is to use the DLR method (M. Sullivan et al. 2006; R. R. Gupta et al. 2016). Here, we calculate the directional radius of the galaxy along the line connecting the transient position to the host center (Figure 2). Then the DLR along this direction is given as

$$\phi = \arctan 2(\Delta y, \Delta x) - \theta$$


$$\text{DLR} = \frac{ab}{\sqrt{(a \sin \phi)^2 + (b \cos \phi)^2}}$$

where  $\theta$  is the position angle measured east through north, and  $\phi$  is the angle from the galaxy’s major axis to the line connecting the galaxy center to the transient. We can calculate the dimensionless ratio  $d_{\text{DLR}}$  as

$$d_{\text{DLR}} = \frac{\text{transient–host angular separation (arcsec)}}{\text{DLR}}$$

In both approaches—DLR and A-value—we identify and select the lowest value as our “best candidate” host.

As we have seen in the case of Sherlock, these approaches are not strictly limited to identifying the host galaxy of a transient. While objects such as stars typically lack morphological parameters, the position of a transient relative to its galactic host can inform our understanding of its type. For example, a low DLR can be used to infer whether the transient is nuclear, an AGN, or a tidal disruption event. Outside the galactic nucleus, we can determine the likelihood of a supernova occurring at a given distance from its selected host



**Figure 2.** Comparison of the DLR and A-value methods. In the A-value method, a less detailed morphology creates “larger” galaxies from which to perform crossmatching.

center; outside of a given  $d_{\text{DLR}}$ , we can infer that the transient is hostless.

The DLR method requires significantly more steps in its calculation than the A-value method. As such, when discussing implementation, it is important to consider the benefits to host-matching accuracy gained from this approach. For instance, an immediate benefit is that the DLR method is less biased toward large, high-ellipticity galaxies, which the A-value method would treat as circular.

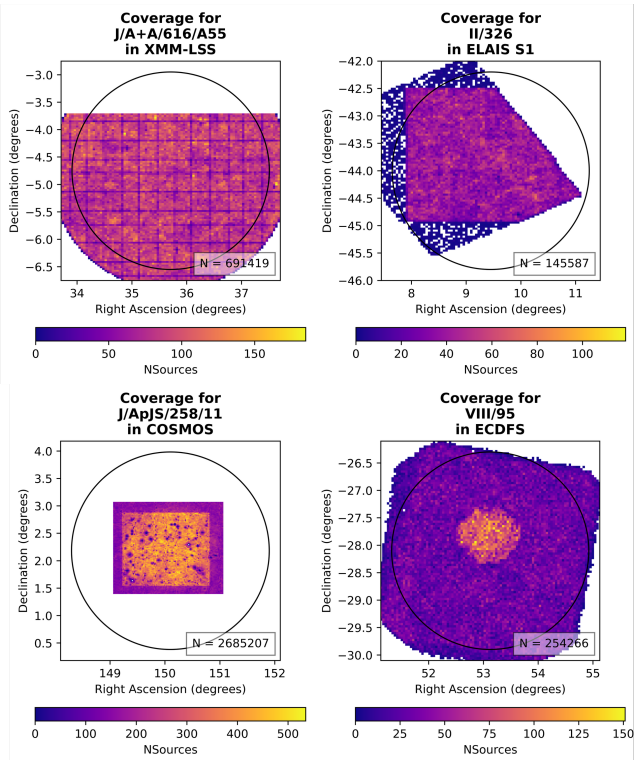
It is also important to consider, particularly when implementing this approach across multiple catalogs, that morphological approaches using isophotally derived metrics (i.e., radii and position angles) produce survey-dependent quantities. Images of galaxies and the subsequently derived axes and position angles depend on the observing filters, point-spread functions (PSFs), survey depth, or sensitivity, and the assigned detection threshold or chosen isophote. While we see little impact of these differences in our testing, it remains essential to consider and record the catalog from which the best associations for a given transient were obtained.

R. R. Gupta et al. (2016) take quantities used in calculating  $d_{\text{DLR}}$  (e.g., ratio between the lowest values, minimum magnitude, ellipticity, galaxy magnitude) and train a machine learning model to generate confidence scores for given associations. It is important that an association not only be the best available, but also reliable. While generating these scores is expensive in a live pipeline, offline machine learning analysis would enable assessment of association accuracy. We would expect that, in crowded fields with multiple nearby galaxies, the association would be assigned a lower confidence score. In the case of blended hosts, for example, it is difficult to determine without follow-up observation which galaxy the transient belongs to, even if the calculated  $d_{\text{DLR}}$  is low. In Lasair, there is no current metric for assessing host match quality, as the contextual match is provided for manual inspection by users who can verify the association.

## 4. The DDF Catalogs

We have previously discussed the extensive preexisting coverage of the DDFs. These catalogs typically summarize small, field-specific studies that go deeper than all-sky surveys. These catalogs allow us to examine fainter, higher-redshift galaxies and other objects, and often provide multiwavelength data. In this section, we collect available data from deep catalogs overlapping the LSST DDFs and assess the utility of each catalog for crossmatching in LSST alert brokers.

Catalogs were obtained from a number of sources, including VizieR via the Strasbourg Astronomical Data Centre (CDS)

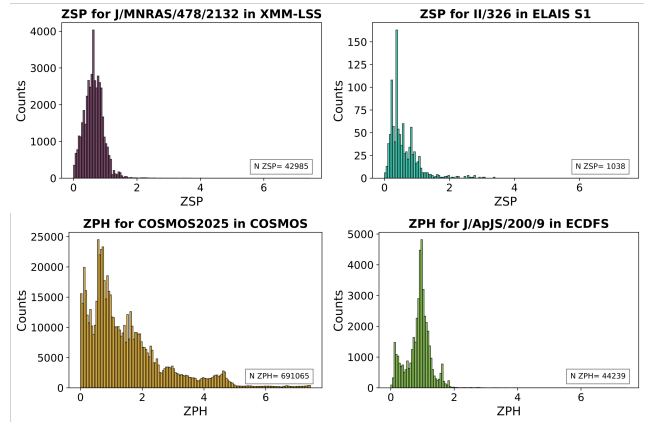


**Figure 3.** Example field coverage plots for a DDF catalog in four of the fields. Lestrade performs a cone search in the region surrounding the central pointing for each catalog to identify subsets of data for each field. Top left: the CFHQSR survey (S. Pipien et al. 2018) in XMM-LSS. Top right: revised SWIRE Photometric Redshift Catalogue in ELAIS-S1 (M. Rowan-Robinson et al. 2013). Bottom left: the COSMOS2020 catalog in COSMOS (J. R. Weaver et al. 2022). Bottom right: the Herschel Multi-tiered Extragalactic Survey catalog (Hermes Team et al. 2017) in ECDF-S.

service (F. Genova 2013), SIMBAD (M. Wenger et al. 2000), CosmoHub (J. Carretero et al. 2017; P. Tallada et al. 2020), and The Spitzer Spectroscopic Data Fusion (M. Vaccari 2025). As of this writing, we have collected a total of 67 catalogs from various surveys across all fields, except Euclid-South. Of these, 56 contain more than 1000 entries, and 44 contain more than 10,000 entries. Of these 44, 28 contain redshift measurements (spectroscopic or photometric). For the morphology-based approach, we identify 15 catalogs across COSMOS, ECDF-S, and XMM-LSS that provide detailed spatial extent information (Semimajor axis length, semiminor axis length, and position angle). More catalogs contain other morphology data, such as the effective radius. Finally, for contextual classification, we identify 29 catalogs with classification flags for different astronomical objects, such as star/galaxy, AGN, and QSO flags. A complete summary of the catalog collection and coverage plots can be found in Figures 3 and 4, and in Tables 2, 3, 4, and 5 in the Appendix.

We have neglected to include a table for EDF-S, which so far is covered by the Euclid Quick Data Release 1 (Euclid Collaboration et al. 2025). This release contains a photometric redshift and spectroscopic catalog for the three Euclid deep fields, including EDF-S.

We have developed a software package, Lestrade, to process these catalogs (J. Weston 2026). Lestrade provides two options for catalog analysis. CDS bibcode-IDs may be input for automatic analysis of the corresponding README file for authors, summaries, and schema. The software generates



**Figure 4.** Example redshift coverage plots for a DDF catalog in four of the fields. Top left: spectroscopic redshifts in the XMM-SERVS survey (C. T. J. Chen et al. 2018) in XMM-LSS. Top right: spectroscopic redshifts in the Revised SWIRE Photometric Redshift Catalogue in ELAIS-S1 (M. Rowan-Robinson et al. 2013). Bottom left: photometric redshifts in the COSMOS2025 Catalog in COSMOS (M. Shuntov et al. 2025). Bottom right: photometric redshifts in the ACS-GC catalog (R. L. Griffith et al. 2012) in ECDF-S.

Markdown files, which are published to the repository wiki. From the schema, the software determines key spectroscopic and photometric redshifts, the existence of morphological or class-related metrics, and coverage within each DDF. These catalog summaries can then be adjusted and re-entered by users back into Lestrade. Alternatively, the key metrics and information can be input manually by the user in the first instance to generate a more detailed wiki page with plots for the relevant redshift metrics, redshift quality metrics, and type flags.

Within this collection of catalogs, we can sort and rank each dataset by the quality and depth of its information. In ranking them, we prioritize catalogs for implementation within the crossmatching Sherlock service of Lasair.

1. *A+*. Spectroscopic redshift catalogs with additional contextual information. Spectroscopic redshifts provide the most reliable estimation of host-transient separation and are a priority for implementation within crossmatching pipelines. There is high value in the contextual information regarding redshift and the type of galaxy (or any other object) the transient has been crossmatched with. Redshift and luminosity distance are extremely valuable for the rapid calculation of absolute magnitudes of transients.
2. *A*. Spectroscopic redshift catalogs with no additional contextual information. While not as well-suited to our morphology-based approach, the spectroscopic redshifts still provide significant value for crossmatching studies.
3. *B+*. Photometric redshift catalogs with additional contextual information. While not as reliable for crossmatching, photometric redshifts still provide useful context, both individually and when used in tandem with non-luminosity-distance approaches.
4. *B*. Photometric redshift catalogs with no additional contextual information. In these cases, a crossmatch using photometric redshifts may be less reliable and would be improved by utilizing the additional contextual data from another catalog.

5. *C+/C*. Photometric catalogs. Such catalogs are less valuable than those containing redshift information and are therefore the lowest priority for implementation.

This ranking is not exhaustive and exists only to provide a broad overview of considerations regarding catalog implementation, which increase in complexity when discussing the implementation of multiple catalogs. A significant omission is that of catalog size, depth, and spatial coverage.

The full ranking of the catalogs above can be found on the Lestrade wiki,<sup>4</sup> with each dataset separated by its usefulness in crossmatching. New catalogs are continually being added to the repository as they are published.

## 5. Crossmatching in the DDFs

### 5.1. Crossmatching in XMM-LSS and ECDF-S

We test our catalog crossmatching in the DDFs using a sample of SNe from DES, observed over a 5 yr period within the XMM-LSS, ECDF-S, and ELAIS-S1 DDFs (B. O. Sánchez et al. 2024). The sample contains accurate coordinates for each transient, along with hosts identified using the DLR method (P. Wiseman et al. 2020). Spectroscopic redshifts for these host galaxies were then obtained primarily from OzDES, along with other sources (F. Yuan et al. 2015; M. J. Childress et al. 2017; C. Lidman et al. 2020). The sample contains 8274 transients divided between the three DDFs, with a redshift range of  $0.0001 < z < 6.84$  and a median redshift of 0.56. The DES SN program was carried out to a limiting magnitude of  $\sim 24$  in all bands; the expected median  $5\sigma$  point-source depth for LSST Y1 will be  $\sim 26$ . We process each DES transient through our own host-matching pipeline and compare the identified galaxies to the host galaxy provided in the DES sample to provide a performance measure for both the DLR and *A*-value methods.

Lestrade constructs “primary catalogs” that combine data within the specified region from the available published datasets. For morphological host matching, we only require the host coordinates, axis lengths, and position angle. Thus, six catalogs with a wide range of data spanning multiple fields can be combined and condensed into a single five-column catalog covering a single DDF. The pipeline takes a list of the catalogs to be imported alongside any “calibrations” required to make the data homogeneous. If the morphological data are provided in units of pixels, for instance, an instrument axis scale can be provided to allow for the necessary conversions between units.

Of the three DDFs covered by the DES sample, we lacked ELAIS-S1 catalogs with morphological information at the time of testing. We identify XMM-LSS as the best case study, with extensive preexisting catalog coverage and 2664 transients to test crossmatching approaches. From our catalog coverage, we select two datasets within the region as suitable catalogs for testing. The SPLASH-SXDF Multiwavelength catalog was obtained via the Spitzer Large Area Survey with Hyper Suprime-Cam (SPLASH; V. Mehta et al. 2018). The million-row dataset contains corresponding photometric redshifts and over 10,000 spectroscopic redshifts, with relevant morphological information for each object. The CFHQSIR survey is a *Y*-band extension of the CFHT Legacy Survey containing data from MegaCam and WIRCam (S. Pipien et al. 2018). This catalog contains 690,000 rows with no redshift coverage but

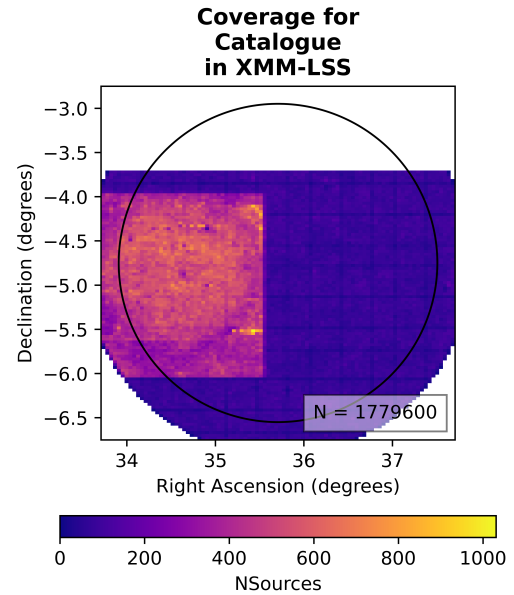


Figure 5. Primary catalog coverage for XMM-LSS.

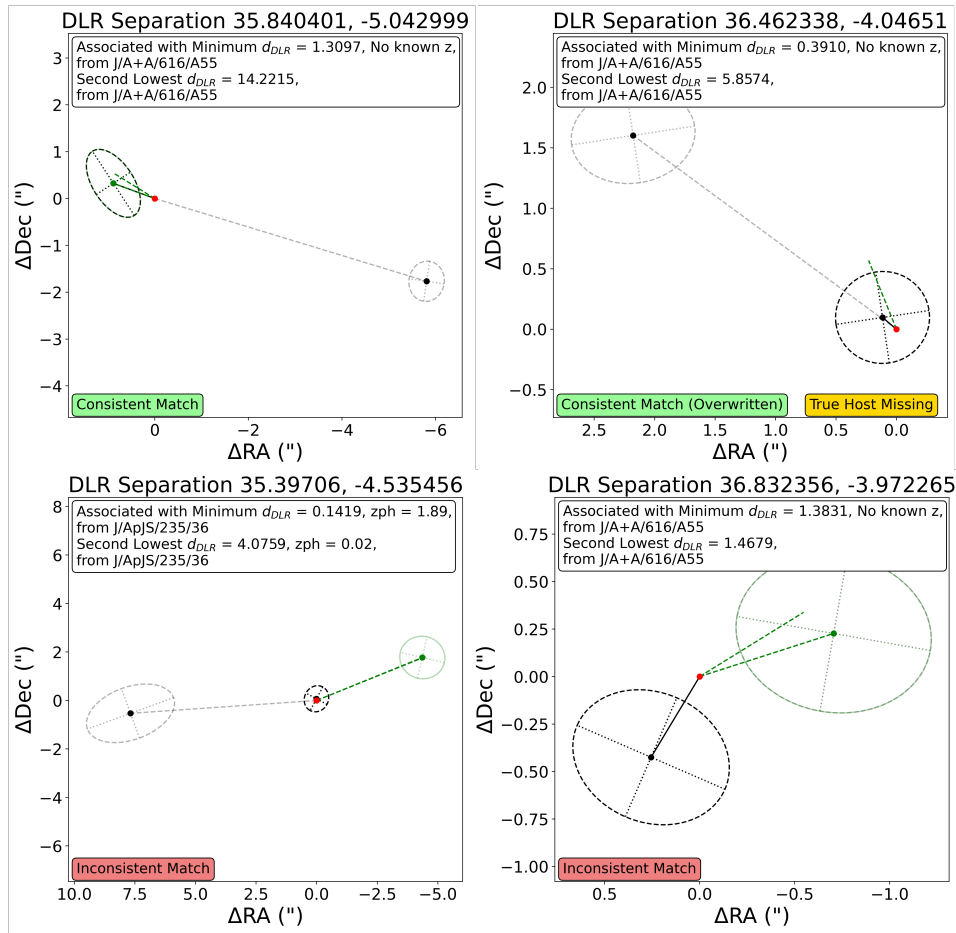
provides morphological information for faint ( $Y \leq 23$  mag) galaxies. Together, these catalogs summarize 1.5 million galaxies with extensive morphological coverage and some redshift information across the XMM-LSS DDF (Figure 5).

Lestrade imports the DES sample data within XMM-LSS and constructs the master catalog described above. For each transient, every potential host within a given radius (which we will tune) is identified with a corresponding  $d_{DLR}$  (or *A*-value) calculated. These values are then sorted, with the best and second-best hosts identified as those with the lowest and second-lowest  $d_{DLR}$ . Here, the pipeline checks the proximity of these two hosts—if their shapes are found to have significant overlap (such that the two morphologies may describe the same or a sufficiently blended host), the second-best host is discarded, and the next-best host is selected to remove duplicate catalog entries. A real and unique second-best host provides valuable context in determining the accuracy of our best choice.

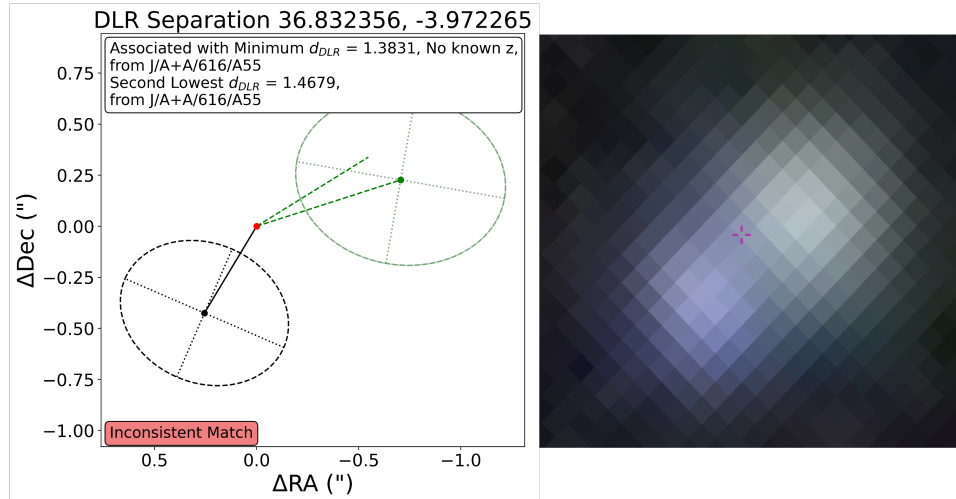
In cases such as our DES sample, where the true (DES-identified) host is known, the code then inspects its best choice against the coordinates of the “true” host. If the coordinates fall within the ellipse of our selected host (or within a given distance of the ellipse edge), we designate it a successful match (or at least consistent with the prior selection by DES). If the matches are inconsistent, we designate the attempt as unsuccessful. When the true host coordinates do not fall within any galaxy in our catalog, we classify it as an “Incomplete” case and exclude it from the study. The user may then examine each result and choose to override any case as they see necessary. In the case of XMM-LSS, we identify 22 “inconsistent” matches for inspection. These were a mixture of cases where (a) the “true” host coordinates lay just outside the best identified host (i.e., the coordinates are less than half an arcsecond outside the ellipse), and (b) the “true” host is so far away from the transient (multiple arcminutes away) that we state the closer host identified by our code is a better match.

Upon selecting the two best hosts for each transient and verifying the match, we output the results in both tabular form and as individual plots (see Figure 6). These plots provide a

<sup>4</sup> <https://github.com/joshgithubbin/Sherlock-DDF/wiki>



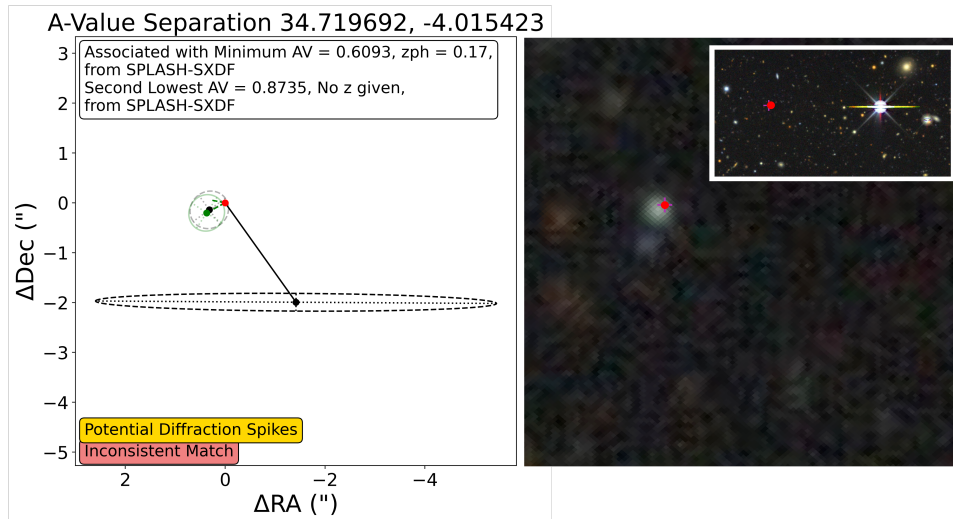
**Figure 6.** Lestrade visual aid plots for transient-host matching with the DES sample in XMM-LSS. The red marker denotes the transient location, with the DLR method’s identified best host (black) and second-best host (gray). A green dashed line links to the true host coordinates in the DES sample, and the morphology of the host exists within the catalog. First from top: a consistent match by the DLR method to the host identified by DES. Second from top: a consistent match by the DLR method where the match has been manually labeled as correct following visual inspection. Third from top: an inconsistent match by the DLR method where the true host was neither identified as the best nor the second-best candidate host. Fourth from top: an inconsistent match by the DLR method where the true host was identified as the second-best host.



**Figure 7.** Blended hosts in a DLR match.

simplified visual of the relationship between the transient and the host, with additional contextual data such as redshift, redshift type, and the host source catalog. Via inspection, we select certain cases for further annotation. Blended hosts, for

example, are unsuitable for testing crossmatching methods; without follow-up observations, it is nearly impossible to determine the true transient host, and so we cannot use these cases to compare the DLR and A-value methods. We use the



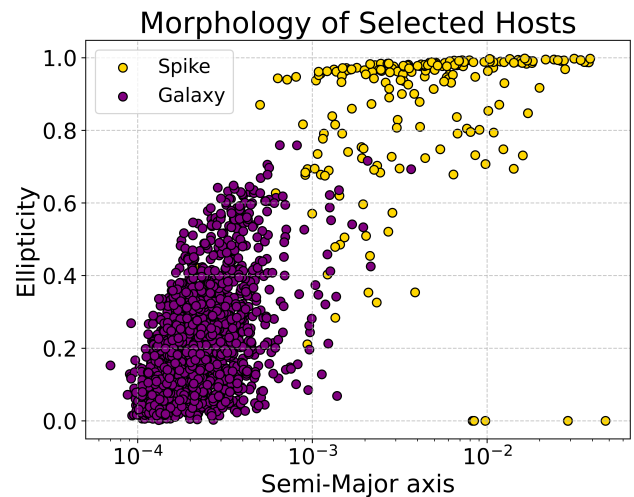
**Figure 8.** Diffraction spike source identified in an A-value match. Note the large bright source along the axis of the diffraction spike (inset); a rotating variable (BD-04 376) with magnitude  $V = 10.05$ .

visual aids to identify them in our results and remove them from subsequent analysis (Figure 7).

Another larger subgroup in our results has identified hosts in the catalogs with very large semimajor axes and high ellipticity. On examination of survey image data, these “hosts” have been identified as diffraction spikes, see Figures 8 and 9. These spikes need to be removed to make the best use of the catalogs. The large semimajor axes of these spikes provide very low  $d_{DLR}$  values not directed toward the transient, “beating” closer, more realistic hosts. For the A-value method, where ellipticity is ignored, this problem is amplified as a large, arcminute-scale “galaxy” dominates over all galaxies in its region. We identify these again via the visual aids in tandem with survey image data. We perform a rudimentary culling of diffraction spikes from our XMM-LSS master catalogs (SPLASH-SXDF and CFHQSIR) by deleting rows with ellipticity greater than 0.8 and angular size greater than  $10^{-3}$ , and rerun the same pipeline. The 18 real galaxies removed by this cut are large enough and bright enough ( $K \sim 12$ ) to be found in the larger catalogs already within the Sherlock database.

### 5.2. Performance Comparison

Confusion matrices are produced for the DES XMM-LSS Sample as seen in Figure 10. We provide the matrices for both the sample with incomplete cases and blended hosts (Figure 7) removed, and for the sample with no cases removed. In an ideal case with an uncontaminated (e.g., diffraction spike-free) catalog and exhaustive coverage, these matrices would converge to near or complete equivalence; we provide both to showcase the impact of “indeterminable cases” (or blended hosts and incomplete cases). Across all cases, we see both methods fail to identify approximately 10% of hosts. The DLR method fails to identify four hosts where the A-value method is successful, while the A-value method fails to identify two hosts where the DLR method is successful. Looking at the filtered confusion matrix, we see that two of these A-value specific successes are, in fact, diffraction spikes spanning the area of the true host. For the unique successes of the A-value method, one case involves the true host being a distant but extended host that would have benefited from being treated as “larger.”

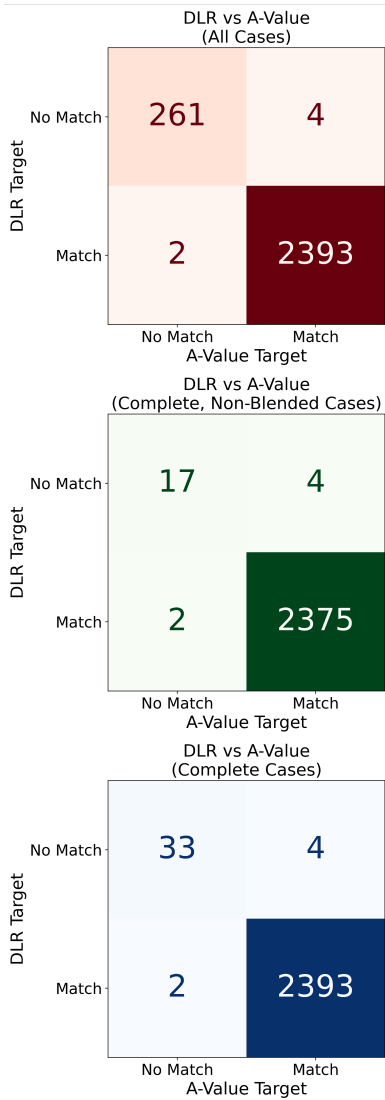


**Figure 9.** Diffraction spike sources identified in the initial application of the DLR and A-value methods.

In the second case, the A-values for the best and second-best candidates are 2.9057 and 2.9455, with DLR values of 3.4030 and 3.0322, respectively. For the unique DLR successes, both galaxies have a similar angular distance and benefit from the attention to detail provided by the approach.

We also have sufficient coverage in the ECDF-S field to compare the DLR and A-value methods. This test is carried out on a smaller subsample of 188 DES transients. While we have a large number of morphological catalogs for this DDF, the majority of the coverage is limited to within half a degree of the central pointing. Subsequently, we only use the CDF-S multicolor dataset (C. Wolf et al. 2004). The set contains the requisite morphological information (semimajor and semiminor axes and position angle) and photometric redshifts for each object. As before, we produce confusion matrices for the field using the pipeline outlined above (Figure 11). The DLR method slightly outperforms the A-value method with two unique successful matches by the former compared to 0 by the latter.


In the case of the DDFs, where the number of large hosts is limited, we can state that while the DLR method may examine



**Figure 10.** Confusion matrices comparing DLR and A-value approaches for DES transients in XMM-LSS on catalogs where more extreme diffraction spikes have been culled.

morphologies to greater detail than the A-value method, there is insufficient evidence to suggest a statistically significant improvement. What is apparent is that the DLR method, being more sensitive to the ellipticity and angle of a potential host, provides greater robustness against contaminants such as diffraction spikes than the A-value approach. As with all morphology-focused approaches, both methods struggle with blended hosts; the visual aid plot will help users identify these cases, but the pipeline is limited in its ability to classify them automatically. The DLR approach may also be well-suited to scenarios where galaxy semimajor axes are similar and at similar angular separations from the transient. For XMM-LSS, we find that 90% of transients that could not be successfully matched did not contain the correct host in our catalogs, and 92% for ECDF-S. As a greater number of surveys focus on the DDFs—LSST included—the number of incomplete cases should decrease, and may already be decreasing by implementing more catalogs or other host-matching approaches as outlined in Section 3.

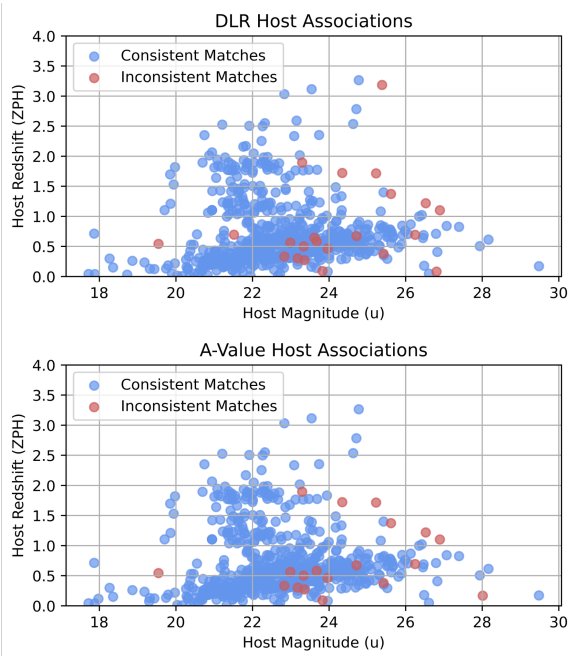
Both approaches correctly identify hosts between 18 and 28 mag (Figure 12). While the majority of incorrect matches



**Figure 11.** Confusion matrices comparing DLR and A-value approaches for DES transients in ECDF-S on catalogs where more extreme diffraction spikes have been culled.

appear at  $m \geq 23$ , where many of the LSST DDF transient hosts will be, most hosts are still correctly identified (251 of 269, or 93%, for the DLR method). This selection will improve as more galaxies are incorporated within the pipeline as potential hosts following further observations of the DDFs. At  $m \leq 23$ , only 1% of the hosts are incorrectly identified.

The majority of hosts are within  $z \sim 1$ , consistent with transient redshift distribution within the DES sample. Below this, the host magnitudes and photometric redshifts follow a sensible correlation with higher-redshift galaxies appearing fainter. A more diffuse group at  $z \geq 1$  contains galaxies that ignore this relationship and appear bright for their given redshift. Higher host photometric redshifts (between 1.5 and 3.5) are low in frequency and can be attributed to uncertainties in the photometric redshift estimation. Despite this, the correct association is made in the majority of cases, and there is no correlation between the success rate and photometric redshift. We do not see a large number of faint, high-redshift hosts, which we would expect to associate with the higher-redshift transients in the DDFs. This selection bias will be improved as



**Figure 12.** Successful and unsuccessful matches by the DLR method (top) and the A-value method (bottom).

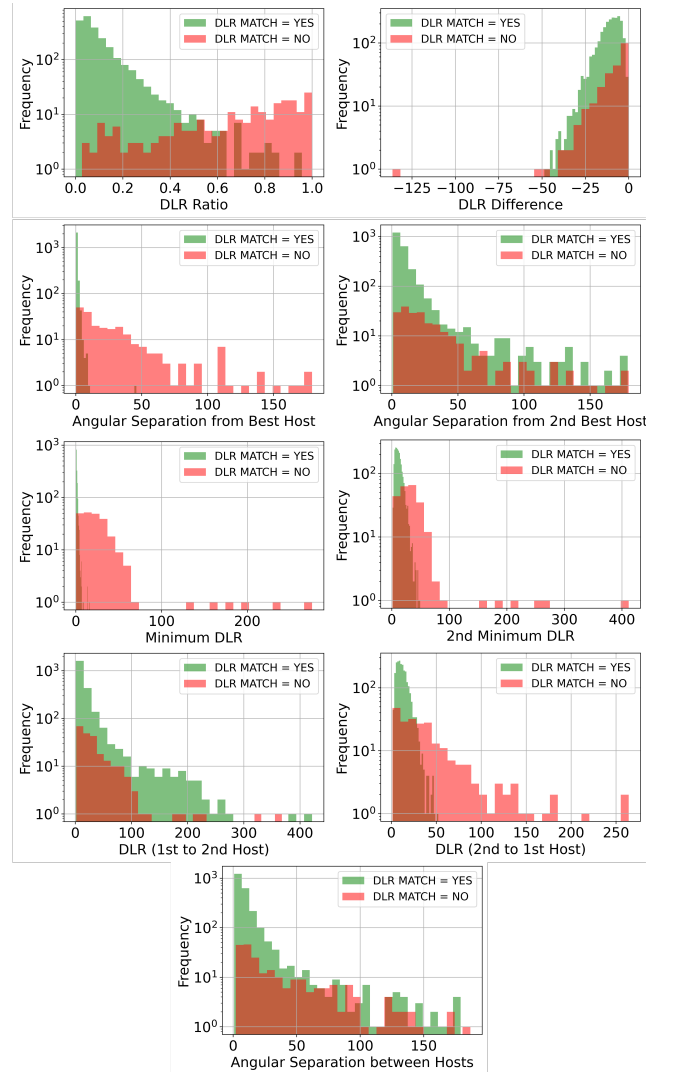
we continue to identify high-redshift phenomena in the LSST data releases.

### 5.3. Crossmatching Confidence Scores

Given comparable performance from both the DLR and A-value approaches, it is worth examining the success of each method in greater detail. We have manually inspected each case to highlight types of failures, such as diffraction spikes and “incomplete” cases. We have also overwritten successful cases that were initially automatically classed as incorrect. This annotated set of host matches and their outcomes—successful or unsuccessful—can be used to train a machine learning binary classification model that will assign a “confidence” score for each match made. The better the two classes are separated by each approach (DLR and A-value) in the score distribution of our samples, the greater the confidence we can place in each method. Provided a well-performing model is trained for each model, we can reduce the time spent manually inspecting host matches by automatically classifying unlabeled host candidates as “successful” or “unsuccessful,” either in real-time or offline.

For our analysis, we continue to use the DES sample in the XMM-LSS DDF. We keep the overwritten cases within the sample as successes and do not remove incomplete cases; there is a benefit in all types of unsuccessful matches being included in the training set for the model, even if we try to avoid such cases in the pipeline. For the same reason, we retain blended hosts, though we expect their impact on our training to be minimal, as the group’s volume remains small compared to other successful/unsuccessful cases.


We select a subset of features obtained during cross-matching for analysis, plotting histograms of their distributions by class (see Figure 13). For the DLR method, we identify several useful features: above a  $d_{\text{DLR}}$  of 20, all matches were unsuccessful, with most successful cases occurring at  $d_{\text{DLR}} < 10$ . For the second minimum  $d_{\text{DLR}}$ , we see that all



**Figure 13.** Feature distributions of successful/unsuccessful matches made by the DLR method.

cases at  $d_{\text{DLR}2} > 50$  are unsuccessful. For the angular separation between the transient and the selected best host, we see no successful matches above  $15''$ . Finally, the  $d_{\text{DLR}}$  of the second host in the direction of the best host also provides a clear difference in the distributions between successful and unsuccessful matches, with no value of  $d_{\text{DLR}2 \rightarrow 1}$  above 50 corresponding to a success. We see similar cases in the A-value approach for the minimum A, the angular separation between the transient and the best host, and the A-value of the second-best host in relation to the center of the first  $A_{2 \rightarrow 1}$  (Figure 14).

For our machine learning confidence model, we employ the XGBoost (extreme gradient-boosting) model found in the scikit-learn Python package. An XGBoost model is a regularizing gradient-boosting algorithm that sequentially builds decision trees to arrive at an optimized process. We divide our data into a training and testing set, reserving 80% of our sample for the testing process, and train the model on the remaining 20%, using our preselected features. We use hold-out validation to characterize classifier performance against a fixed score threshold. A full model would incorporate cross-validation once additional, non-morphology-based metrics are included. The XGBoost approach successively trains decision




**Figure 14.** Feature distributions of successful/unsuccessful matches made by the A-value method.

trees to correct mistakes by the preceding tree over a fixed number of boosting rounds by minimizing the logistic loss:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)],$$

where  $y_i$  is the true label (0 or 1),  $\hat{y}_i$  is the predicted probability for class 1, and  $N$  is the number of samples. We end training after the training and testing set logistic losses begin to diverge (in this case, 10 epochs (Figure 15)). As the divergence is small, the model’s performance on the two sets remains comparable, indicating that significant overfitting has not yet occurred. The final model assigns a score to each given set of host-matching parameters as a prediction of confidence: 1 indicating a successful match and 0 indicating an incorrect match. We run the trained model on our testing set, plotting model score against frequency and labeling by true population label.

Following a successful model training, feature importance metrics can be computed as the number of times each feature is used in any split across all trees within the model. Initially, we utilize the method’s best and second-best

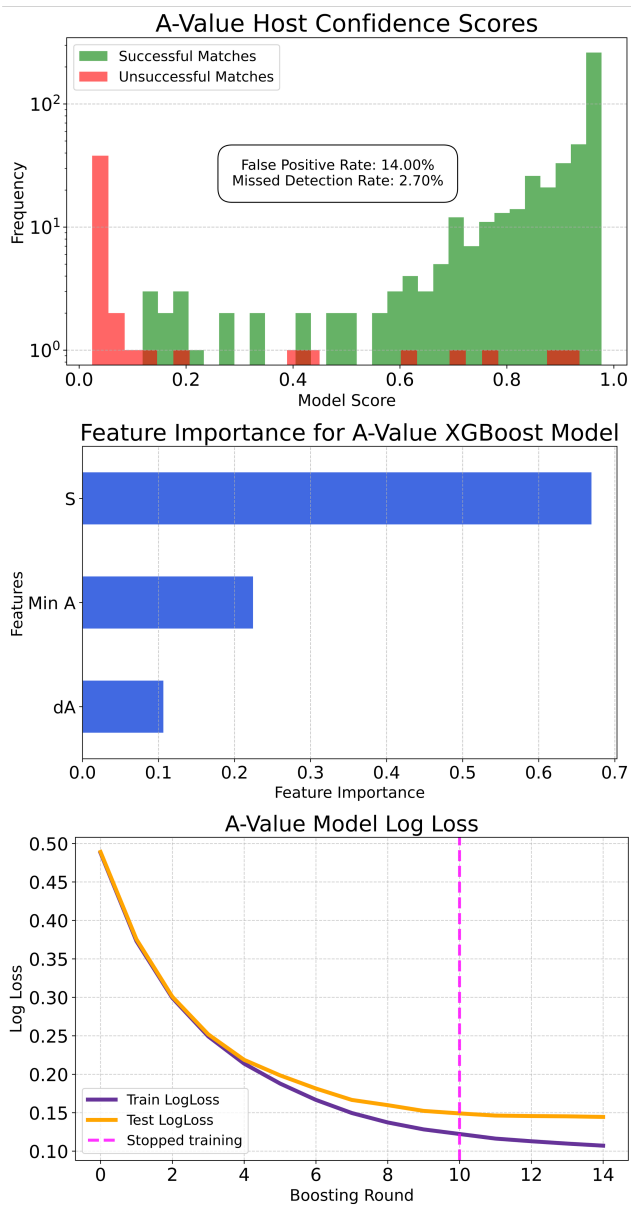


**Figure 15.** Score distribution of successful/unsuccessful matches by the DLR Confidence model.

distance value ( $d_{\text{DLR}}$  or  $A$ ), the difference ( $d\text{DLR}$  and  $dA$ ), and ratio ( $r\text{DLR}$  and  $rA$ ) between these values, and the angular separation between the transient and selected host (S) metric. Given the need for efficient real-time host association of transients in LSST, we want to reduce the size and complexity of our models and therefore reduce redundant metrics. Ranking the inputs by feature importance for our “alpha” models, we see a clear emphasis on value in the angular separation and minimum distance value, with the difference between the two lowest distance values in third. For our final inputs, we select these values alone, which, in both the DLR and A-value cases, appear to improve model performance.

We can define our model performance by the false positive rate (FPR) and missed detection rate (MDR):

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \times 100\%.$$



**Figure 16.** Score distribution of successful/unsuccessful matches by the A-value Confidence model.

$$\text{MDR} = \frac{\text{FN}}{\text{FN} + \text{TP}} \times 100\%.$$

The former refers to the percentage of incorrect matches (false positives (FPs) and true negatives (TNs)) that are incorrectly selected as successful, while the latter refers to the percentage of correct matches (false negatives (FNs) and true positives (TPs)) that are erroneously selected as unsuccessful. Examining the DLR model with a score threshold of 0.4, we have a missed detection rate of 2.69%, with only four FPs in total corresponding to a 8.16% FPR. The A-value metrics are less successful in separating successful and unsuccessful cases (Figure 16). For the same score threshold of 0.4, we now have seven FPs corresponding to an FPR of 14.00%. We also have 13 missed detections corresponding to an MDR of 2.70%.

Of the 11 FPs flagged across both models, four are duplicates under both approaches. Two of these identify a likelier host than the one listed in the sample and should be

overwritten. The remaining cases have distant first and second hosts identified by both methods (DLR,  $A\text{-value} \geq 5$ ). The three FPs unique to the A-value methods have very similar first and second host A-values, where greater detail in the morphology may have been beneficial.

Of 26 total missed detections by both the DLR and A-value methods, seven are duplicate misses by both approaches. Of these, one is a case that was overwritten when a likelier host was identified by the approaches. Of the remaining six, two have high DLR or A-values, indicating either a large transient-host separation or an incorrect match. The remaining four have low DLR or A-value differences.

Of the remaining six missed detections by the DLR model, three have high minimum DLRs, with the other three having similar angular separations between the two hosts and low DLR differences. Of the remaining six missed detections by the A-value model, one has a high A-value, with the other five having low A-value differences. While the DLR model is more successful in discerning successful matches from unsuccessful matches, it should be noted that this is not an indicator of the performance by either approach: it is simply an indicator of which method will be easier to manually inspect. As the DLR model more reliably scores successful matches highly and rarely provides a low score to a successful match, we can spend less time manually inspecting each crossmatch. Should the morphology-based approach be combined with other approaches (such as the redshift-based approach), the model score can be used to determine how well supported a potential host is by its morphological information, used in tandem with scores or metrics from other data.

#### 5.4. Crossmatching in the Wider Sky

Galaxies in the DDF catalogs are mostly at greater distances than those in wide-field catalogs, and therefore have a smaller angular size. This makes it difficult to resolve their shape, as they may appear comparable to or smaller than the PSF of the instrument used to image them. For nearby galaxies with a larger angular extent and more detailed morphology, the difference between the DLR and A-value approaches might be even more apparent.

To test this, we use the ATLAS100 sample of transients discovered by ATLAS (Srivastav et al. 2026, in preparation) (K. W. Smith et al. 2020) with known host galaxies, and match them against the SDSS Data Release 16 cataloged galaxies (R. Ahumada et al. 2020). We find that of 723 transients in the sample for which the true host is in SDSS, 572 are successfully matched by both methods. 37 are matched by only the DLR, and 20 are matched only by the A-value method (Figure 17). Inspection of each case shows a number of SDSS contaminants, including knots in the spiral arms of large galaxies, being identified as individual galaxies. Although at times both methods select these knots rather than the true galaxy, we see the DLR method as being more robust against these contaminants. While an attempted removal of some of these galaxy components yielded a more similar performance between the two methods, the perfect implementation of these morphological approaches for wider-sky crossmatching is beyond the scope of this study. When galaxies are better resolved, a combination of methods is required to identify the correct host galaxy. Among the morphological techniques discussed, the DLR method remains the better-performing

DLR vs A-Value  
(Complete Cases)

DLR Target	No Match	94	20
	Match	37	572
		No Match	Match
		A-Value Target	

**Figure 17.** Confusion matrices comparing DLR and A-value approaches for ATLAS transients on the SDSS.

approach in both DDFs and wide-sky applications, even with contaminated data.

## 6. Implementation

The initial implementation priority is to install the A+ and A-ranked catalogs with more than 10,000 rows into the Sherlock database. Owing to the small size of these catalogs in relation to the currently stored data ( $\sim 0.1\%$ ), the true challenge to implementation is the human labor required to translate the original catalog schema for use by Sherlock. For COSMOS, the selected catalogs—the COSMOS/UVI and the FourStar Galaxy Evolution survey—contain a combined 296,079 rows. XMM-LSS, FourStar, and VIPERS contain 351,771 rows. The MUSYC Photometric Redshift Catalog and FourStar data in ECDF-S provide 136,431 rows for this field. As no A or A+ ranked catalogs have been found for ELAIS-S1, we are left with approximately 784,000 rows across four catalogs to implement for three of the DDFs. In the future, we will expand this implementation to lower-tier catalogs.

It is ultimately difficult to select an optimal method between the DLR and A-value approaches. Both maintain a level performance in the DDFs, with a low number of unique successes and failures (and a low number of mutual failures overall). From a solely morphology-based perspective, either approach is optimal.

When implementing real-time transient-host matching, the trade-off between computational expense and performance must be considered. The DLR method involves more pipeline operations and requires more cataloged information, such as the galaxy position angle and semiminor axis. While not a significant constraint for offline matching, this will compound as more transients are processed in the Rubin pipeline. Running the Sherlock software 21 times on a set of  $\sim 24,000$  transients, we found that the classification rate with the DLR method implemented was 418.5 transients/sec, with a DLR calculation overhead of  $18 \mu\text{s}/\text{transient}$ , increasing the Sherlock runtime by only 0.8%. Given the minimal contribution of the DLR calculation to the total runtime, the method has now been integrated into Sherlock.

We must also consider how each method handles “contaminated” catalogs. The DLR method is more robust against contaminants and less likely to misidentify diffraction spikes as real objects, whereas the A-value approach is more susceptible to artifacts. These contaminants should be removed prior to catalog implementation, though some may survive

culling and only be identified once selected as a potential host. If contaminants are expected in some or all of the cataloged data, the DLR method will be more successful at identifying true hosts.

Preliminary testing of batch inference in Section 5.3 indicates a score calculation time of  $5 \mu\text{s}/\text{transient}$ . This does not account for feature creation or any subsequent filtering based on the generated score. Any implementation of a classifier to assess Sherlock association quality would also look to include magnitude and redshift data where available, further increasing computational time. As the success rate of Sherlock and the DLR method is high, and most incorrect associations can be identified by Lasair broker users on visual inspection, we do not see the development and implementation of a classifier as a priority. Nonetheless, the model shows promising results and would be well applied to offline filtering following data extraction with the Lasair API.

## 7. Conclusion

In this paper, we have explored a method for analyzing and utilizing existing high-quality galaxy catalogs for transient-host matching in LSST’s DDFs. We identify catalogs from various surveys across four DDFs, with a large number of photometric and spectroscopic redshifts, morphology measurements, and object type flags. These analyses and the requisite code have been uploaded to a repository wiki as a resource for transient scientists to identify useful catalogs.

We compare the DLR method with the A-value method of transient-host matching and test their performance on a pre-assessed sample of DES transients in the XMM-LSS and ECDF-S fields. In carrying out our own host association, we identify a similar performance between each approach, but highlight that the DLR approach is more robust against large-scale contaminants such as diffraction spikes. We discuss other benefits of each method, such as the A-value method’s advantage over poorly fitted morphologies and the DLR method’s advantage for evenly separated galaxies.

We examine machine learning-based avenues for assessing host selection confidence, again focusing on the DLR and A-value methods. We find that the DLR confidence model correctly identifies successful matches with a lower missed detection rate than the A-value model. While this should not be used to promote one approach’s inherent superiority, it highlights a way of identifying successful matches for unlabeled data. We note the difficulty caused by catalog contamination from artifacts (such as diffraction spikes) when assessing host-matching methods, and the problem in identifying the true host between a blended pair.

Finally, we look to integrate the A+ and A-ranked catalogs into Sherlock for host crossmatching in the DDFs by Lasair, which contain 784,000 sources in total across COSMOS, XMM-LSS, and ECDF-S, with additional catalogs to be added later. Implementing the DLR method within Sherlock contributes very little to the total runtime (0.8%) at a rate of  $18 \mu\text{s}$  per transient. The DLR method is now implemented within Sherlock.

An avenue for future research would be to formally assess the performance of a redshift-based host-matching approach compared to our morphology-based approach, or to a simple angular-separation-based approach. Combining these approaches and incorporating a host confidence model, as we have begun exploring above, may be a useful exercise.

We also note a lack of publicly available and accurate transient-host samples for testing transient science in the LSST DDFs. The construction of such a sample would prove useful in late-stage preparations for real-time transient science with Rubin.

### Acknowledgments

We thank two anonymous referees for providing comments on the manuscript. J.W. acknowledges a studentship funded by the Leverhulme Trust through the Leverhulme Interdisciplinary Network on Algorithmic Solutions (LINAS) at Queen's University Belfast. M.J.J. acknowledges the support of a UKRI Frontiers Research grant [EP/X026639/1], which was selected by the European Research Council. M.J.J., I.H.W., S.J.S., and K.W.S. acknowledge support from the Oxford Hintze Centre for Astrophysical Surveys, which is funded through generous support from the Hintze Family Charitable Foundation. M.N. is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 948381). S.J.S., D.R.Y., and K.S. acknowledge funding from STFC grants ST/X001253/1, ST/Y001605/1, and a Royal Society Research Professorship.

This work has made use of CosmoHub, developed by PIC (maintained by IFAE and CIEMAT) in collaboration with ICE-CSIC. It received funding from the Spanish government (grant EQC2021-007479-P funded by MCIN/AEI/10.13039/501100011033), the EU NextGeneration/PRTR (PRTR-C17. I1), and the Generalitat de Catalunya.

### Data Availability

The catalog ingest, host matching, and classification in this study were carried out by the Lestrade software (J. Weston 2026). The data underlying this article are drawn from publicly available astronomical catalogs, primarily accessed via the CDS (VizieR and SIMBAD) and survey-specific data releases cited throughout the paper. A full list of catalogs by field and associated references can be found in the appendices. The DES SN sample used for validation is publicly available as described in B. O. Sánchez et al. (2024).

The compiled DDF catalog summaries and the software developed for catalog processing and host-matching analysis (including Lestrade and associated scripts) are available at the project GitHub repository: <https://github.com/joshgithubbin/Sherlock-DDF/wiki>. Derived data products generated during this study are available from the corresponding author upon reasonable request.

### Appendix Additional Figures

We outline the analysed catalogs for each DDF in Tables 2, 3, 4, and 5. Each table lists the number of objects per catalog, spectroscopic and photometric redshift information where available, and flags for morphology/class information. The catalogs are ranked by the hierarchy outlined in Section 4.

**Table 2**  
DDF Catalogs in COSMOS

Reference	Catalog	No. Objects	zspec No.	zspec 90	zspec 50	zphot No.	zphot 90	zphot 50	Morphology	Type_Flag	Rank
R. L. Griffith et al. (2012)	The ACS-GC catalog	304,688	10,013	1.42	0.61	251,592	1.99	0.83	Yes	Yes	A+
DESI Collaboration et al. (2025)	DESI DR1	204,977	199,523	1.41	0.72	...	...	...	...	Yes	A
A. Muzzin et al. (2013)	COSMOS/UltraVISTA Ks-selected catalogs	262,615	5537	0.84	0.5	0	...	...	...	Yes	A
C. M. S. Straatman et al. (2016)	FourStar galaxy evolution survey	33,464	465	3.07	0.73	0	...	...	Yes	...	A
S. J. Lilly et al. (2007)	zCOSMOS-bright catalog, DR3	20,689	19,148	0.93	0.53	...	...	...	...	...	A-
G. Hasinger et al. (2018)	DEIMOS 10K spectroscopic survey in the COSMOS field	10,873	8786	3.1	0.89	...	...	...	...	...	A-
J. R. Weaver et al. (2022)	The COSMOS2020 Farmer catalog	2,685,207	...	...	...	1,439,880	3.35	1.18	Yes	Yes	B+
M. Shuntov et al. (2025)	The COSMOS2025 Catalog	784,016	...	...	...	681,218	3.45	1.14	Yes	...	B
B. Darvish et al. (2017)	Cosmic web of galaxies in the COSMOS field	45,421	...	...	...	45,421	1.13	0.84	...	Yes	B
A. Molino et al. (2016)	ALHAMBRA Survey	38,427	...	...	...	38,427	1.75	0.9	Yes	...	B
P. W. Hatfield et al. (2022)	Photometric redshifts in COSMOS	995,049	...	...	...	995,049	3.98	2.16	...	...	B-
P. Capak et al. (2007)	COSMOS Multiwavelength Photometry Catalog	438,226	...	...	...	438,226	1.98	0.74	...	...	B-
O. Razim et al. (2021)	COSMOS2015 dataset machine learning photo-z	214,398	...	...	...	214,398	1.09	0.75	...	...	B-
H. Nayyeri et al. (2017)	Multiwavelength data in the CANDELS COSMOS field	38,671	0	0	...	...	...	...	Yes	...	C
C. L. Hale et al. (2025)	MIGHTEE DR1	20,886	...	...	...	...	...	...	Yes	...	C
H. J. McCracken et al. (2012)	The fourth UltraVISTA data release	451,587	...	...	...	...	...	...	...	...	C-
Hermes Team et al. (2017)	Herschel Multi-tiered Extragalactic Survey	124,343	...	...	...	...	...	...	...	...	C-
Y.-Y. Chang et al. (2021)	Machine learning predicted AGNs in the HSC-Wide region	112,609	...	...	...	...	...	...	...	...	C-
L. Cabayol et al. (2019)	Star-galaxy multi-narrowband classification	38,427	...	...	...	...	...	...	...	...	C-
V. Smolčič et al. (2017)	VLA-COSMOS 3 GHz Large Project	10,830	...	...	...	...	...	...	...	...	C-

**Table 3**  
DDF Catalogs in XMM-LSS

Reference	Catalog	No. Objects	zspec No.	zspec 90	zspec 50	zphot No.	zphot 90	zphot 50	Morphology	Type_Flag	Rank
C. T. J. Chen et al. (2018)	XMM-LSS SERVS. New XMM-Newton point-source cat.	396,287	429,85	0.99	0.63	390,900	1.65	0.85	...	Yes	A
B. Garilli et al. (2014)	VIMOS Public Extragalactic Survey (VIPERS) DR1	316,717	13,525	0.98	0.71	...	...	...	...	Yes	A
DESI Collaboration et al. (2025)	DESI DR1	41,119	39,353	1.49	0.80	...	...	...	...	Yes	A
L. Cabayol et al. (2019)	The PAU Survey	218,280	29,124	0.92	0.62	218,280	1.01	0.58	...	Yes	A
C. M. S. Straatman et al. (2016)	FourStar Galaxy Evolution Survey	35,054	152	1.49	0.75	0	...	...	Yes	...	A
M. Rowan-Robinson et al. (2013)	Revised SWIRE photometric redshifts	311,643	3628	1.15	0.63	...	...	...	...	...	A-
V. Mehta et al. (2018)	SPLASH-SXDF multiwavelength photometric catalog	1,088,181	...	...	...	1,084,522	2.97	1.02	Yes	Yes	B+
T. Moutard et al. (2016)	VIPERS Multi-Lambda Survey	740,712	...	...	...	740,160	1.65	0.75	...	Yes	B
P. W. Hatfield et al. (2022)	Photometric Redshifts in XMM-LSS	1,674,689	...	...	...	1,674,689	4.07	1.98	...	...	B-
M. Lacy et al. (2021)	Spitzer Survey of DDFs	694,991	...	...	...	69,2977	1.03	2.33	...	...	B-
H. J. McCracken et al. (2008)	VIRMOS deep imaging survey. VVDS-F02 catalog	2,028,704	...	...	...	...	...	...	Yes	Yes	C+
M. J. Jarvis et al. (2013)	VISTA Deep Extragalactic Observations Survey (VIDEO)	127,9857	...	...	...	...	...	...	Yes	Yes	C+
S. Pipien et al. (2018)	CFHQSIR survey	69,1419	...	...	...	...	...	...	Yes	Yes	C+
C. L. Hale et al. (2025)	MIGHTEE DR1	53,393	...	...	...	...	...	...	Yes	...	C
Hermes Team et al. (2017)	Herschel Multi-tiered Extragalactic Survey	258,612	...	...	...	...	...	...	...	...	C-

**Table 4**  
DDF catalogs in ECDF-S

Reference	Catalog	No. Objects	zspec No.	zspec 90	zspec 50	zphot No.	zphot 90	zphot 50	Morphology	Type_Flag	Rank
C. N. Cardamone et al. (2010)	MUYSC Photometric Redshifts	84,402	3988	2.7	1.04	...	...	...	Yes	Yes	A+
R. L. Griffith et al. (2012)	The ACS-GC catalog	70,446	6955	2.02	0.69	44,239	1.27	0.93	Yes	Yes	A+
C. M. S. Straatman et al. (2016)	FourStar galaxy evolution survey	52,029	1226	2.67	0.97	0	...	...	Yes	...	A
A. Grazian et al. (2006)	GOODS-MUSIC sample multicolor catalog	18,296	1275	1.34	0.73	14,646	3.14	1.19	...	Yes	A
E. N. Taylor et al. (2009)	MUSYC Survey	16,910	2914	1.23	0.68	13,547	135.94	24.96	...	Yes	A
M. Rowan-Robinson et al. (2013)	Revised SWIRE photometric redshifts	149,766	728	1.1	0.52	...	...	...	...	...	A-
D. A. Rafferty et al. (2011)	SMBH Growth in Starburst Galaxies with Chandra	100,318	2587	2.31	0.74	100,318	3.44	1.62	...	...	A-
M. Lacy et al. (2021)	Spitzer Survey of DDFs	799,606	...	...	...	787,398	1.10	2.41	...	...	B-
M. J. Jarvis et al. (2013)	VISTA Deep Extragalactic Observations Survey (VIDEO)	1,138,485	...	...	...	...	...	...	Yes	Yes	C+
C. Wolf et al. (2004)	CDF-S multicolor data	63,501	...	...	...	...	...	...	Yes	Yes	C+
Y. Guo et al. (2013)	GOODS-S CANDELS multiwavelength catalog	34,930	...	...	...	...	...	...	...	Yes	C
C. L. Hale et al. (2025)	MIGHTEE DR1	17,866	...	...	...	...	...	...	Yes	...	C
Hermes Team et al. (2017)	Herschel Multi-tiered Extragalactic Survey	254,266	...	...	...	...	...	...	...	...	C-

**Table 5**  
DDF catalogs in ELAIS-S1

Reference	Catalog	No. Objects	zspec No.	zspec 90	zspec 50	zphot No.	zphot 90	zphot 50	Morphology	Type_Flag	Rank
M. Rowan-Robinson et al. (2013)	Revised SWIRE photometric redshifts	145,587	145,587	−99	−99	...	...	...	...	...	A−
M. J. Jarvis et al. (2013)	VISTA Deep Extragalactic Observations Survey (VIDEO)	842,337	...	...	...	...	...	...	Yes	Yes	A−
Hermes Team et al. (2017)	Herschel Multi-tiered Extragalactic Survey	86,779	...	...	...	...	...	...	...	...	C−

## ORCID iDs

J. G. Weston  <https://orcid.org/0009-0002-9460-9900>  
 D. R. Young  <https://orcid.org/0000-0002-1229-2499>  
 S. J. Smartt  <https://orcid.org/0000-0002-8229-1731>  
 M. Nicholl  <https://orcid.org/0000-0002-2555-3192>  
 M. J. Jarvis  <https://orcid.org/0000-0001-7039-9078>  
 I. H. Whittam  <https://orcid.org/0000-0003-2265-5983>

## References

- Ahumada, R., Prieto, C. A., Almeida, A., et al. 2020, *ApJS*, **249**, 3  
 Alam, S., Albareti, F. D., Allende Prieto, C., et al. 2015, *ApJS*, **219**, 12  
 Anderson, J. P., Covarrubias, R. A., James, P. A., Hamuy, M., & Haberman, S. M. 2010, *MNRAS*, **407**, 2660  
 Asmus, D., Greenwell, C. L., Gandhi, P., et al. 2020, *MNRAS*, **494**, 1784  
 Bellm, E. C., Blum, R., Graham, M. L., et al. 2020, Plans and Policies for LSST Alert Distribution Tech. Rep. LDM-612, The LSST Project <https://LDM-612.lsst.io/>  
 Berta, S., Rubele, S., Franceschini, A., et al. 2006, *A&A*, **451**, 881  
 Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393  
 Cabayol, L., Sevilla-Noarbe, I., Fernández, E., et al. 2019, *MNRAS*, **483**, 529  
 Capak, P., Aussel, H., Ajiki, M., et al. 2007, *ApJS*, **172**, 99  
 Cardamone, C. N., van Dokkum, P. G., Urry, C. M., et al. 2010, *ApJS*, **189**, 270  
 Carretero, J., Tallada, P., Casals, J., et al. 2017, in Proc. European Physical Society Conf. on High Energy Physics, ed. P. Collins & K. Kruger (SISSA), 488  
 Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2019, arXiv:1612.05560  
 Chang, Y.-Y., Hsieh, B.-C., Wang, W.-H., et al. 2021, *ApJ*, **920**, 68  
 Chen, C. T. J., Brandt, W. N., Luo, B., et al. 2018, *MNRAS*, **478**, 2132  
 Childress, M. J., Lidman, C., Davis, T. M., et al. 2017, *MNRAS*, **472**, 273  
 Civano, F., Marchesi, S., Comastri, A., et al. 2016, *ApJ*, **819**, 62  
 Darvish, B., Mobasher, B., Martin, D. C., et al. 2017, *ApJ*, **837**, 16  
 DES Collaboration, Abbott, T. M. C., Acevedo, M., et al. 2024, *ApJL*, **973**, L14  
 DESI Collaboration, Abdul-Karim, M., Adame, A. G., et al. 2025, arXiv:2503.14745  
 Dey, A., Schlegel, D. J., Lang, D., et al. 2019, *AJ*, **157**, 168  
 Downes, R. A., Webbink, R. F., Shara, M. M., et al. 2001, *PASP*, **113**, 764  
 Euclid Collaboration, Aussel, H., Tereno, I., et al. 2025, arXiv:2503.15302  
 Euclid Collaboration, Scaramella, R., Amiaux, J., et al. 2022, *A&A*, **662**, A112  
 Finoguenov, A., Guzzo, L., Hasinger, G., et al. 2007, *ApJS*, **172**, 182  
 Flesch, E. W. 2023, *OJAp*, **6**, 49  
 Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., et al. 2021, *AJ*, **161**, 242  
 Förster, F., Muñoz Arancibia, A. M., Reyes-Jainaga, I., et al. 2022, *AJ*, **164**, 195  
 Fulton, M. D., Smartt, S. J., Huber, M. E., et al. 2025, *MNRAS*, **542**, 541  
 Gagliano, A., Narayan, G., Engel, A., Carrasco Kind, M., & LSST Dark Energy Science Collaboration 2021, *ApJ*, **908**, 170  
 Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1  
 Garilli, B., Guzzo, L., Scodreggio, M., et al. 2014, *A&A*, **562**, A23  
 Genova, F. 2013, *DatSJ*, **12**, WDS56  
 Graham, M. L., Bellm, E. C., Guy, L., et al. 2024, LSST Alerts: Key Numbers, Tech. Rep. DMTN-102, LSST Data Management <https://dmtn-102.lsst.io/>  
 Grazian, A., Fontana, A., de Santis, C., et al. 2006, *A&A*, **449**, 951  
 Griffith, R. L., Cooper, M. C., Newman, J. A., et al. 2012, *ApJS*, **200**, 9  
 Gris, P., Awan, H., Becker, M. R., et al. 2024, *ApJS*, **275**, 21  
 Guo, Y., Ferguson, H. C., Gialvalisco, M., et al. 2013, *ApJS*, **207**, 24  
 Gupta, R. R., Kuhlmann, S., Kovacs, E., et al. 2016, *AJ*, **152**, 154  
 Hale, C. L., Heywood, I., Jarvis, M. J., et al. 2025, *MNRAS*, **536**, 2187  
 Hambleton, K. M., Bianco, F. B., Street, R., et al. 2023, *PASP*, **135**, 105002  
 Hasinger, G., Capak, P., Salvato, M., et al. 2018, *ApJ*, **858**, 77  
 Hatfield, P. W., Jarvis, M. J., Adams, N., et al. 2022, *yCat*, *J/MNRAS/513/3719*  
 Hermes Team, Oliver, S. J., Bock, J., et al. 2017, *yCat*, **VIII**, 103  
 Heywood, I., Jarvis, M. J., Hale, C. L., et al. 2022, *MNRAS*, **509**, 2150  
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111  
 Jarvis, M. J., Bonfield, D. G., Bruce, V. A., et al. 2013, *MNRAS*, **428**, 1281  
 La Franca, F., Gruppioni, C., Matute, I., et al. 2004, *AJ*, **127**, 3075  
 Lacy, M., Surace, J. A., Farrah, D., et al. 2021, *MNRAS*, **501**, 892  
 Lasker, B. M., Lattanzi, M. G., McLean, B. J., et al. 2008, *AJ*, **136**, 735  
 Le Fèvre, O., Mellier, Y., McCracken, H. J., et al. 2004, *A&A*, **417**, 839  
 Lehmer, B. D., Brandt, W. N., Alexander, D. M., et al. 2005, *ApJS*, **161**, 21  
 Lidman, C., Tucker, B. E., Davis, T. M., et al. 2020, *MNRAS*, **496**, 19  
 Lilly, S. J., Le Fèvre, O., Renzini, A., et al. 2007, *ApJS*, **172**, 70  
 LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201  
 Matheson, T., Stubens, C., Wolf, N., et al. 2021, *AJ*, **161**, 107  
 McCracken, H. J., Radovich, M., Iovino, A., et al. 2008, *yCat*, **II**, 286  
 McCracken, H. J., Milvang-Jensen, B., Dunlop, J., et al. 2012, *A&A*, **544**, A156  
 Mehta, V., Scarlata, C., Capak, P., et al. 2018, *ApJS*, **235**, 36  
 Molino, A., Benitez, N., Moles, M., et al. 2016, *yCat*, *J/MNRAS/441/2891*  
 Möller, A., Peloton, J., Ishida, E. E. O., et al. 2021, *MNRAS*, **501**, 3272  
 Moutard, T., Arnouts, S., Ilbert, O., et al. 2016, *yCat*, *J/A+A/590/A102*  
 Muzzin, A., Marchesini, D., Stefanon, M., et al. 2013, *ApJS*, **206**, 8  
 Nayyeri, H., Hemmati, S., Mobasher, B., et al. 2017, *ApJS*, **228**, 7  
 Oliver, S., Rowan-Robinson, M., Alexander, D. M., et al. 2000, *MNRAS*, **316**, 749  
 Pierre, M., Valtchanov, I., Altieri, B., et al. 2004, *JCAP*, **2004**, 011  
 Papien, S., Basa, S., Cuby, J. G., et al. 2018, *yCat*, *J/A+A/616/A55*  
 Pucetti, S., Fiore, F., D'Elia, V., et al. 2006, *A&A*, **457**, 501  
 Qin, Y.-J., Zabludoff, A., Kisley, M., et al. 2022, *ApJS*, **259**, 13  
 Rafferty, D. A., Brandt, W. N., Alexander, D. M., et al. 2011, *ApJ*, **742**, 3  
 Razim, O., Cavuoti, S., Brescia, M., et al. 2021, *MNRAS*, **507**, 5034  
 Rhodes, J., Nichol, R. C., Aubourg, É., et al. 2017, *ApJS*, **233**, 21  
 Ritter, H., & Kolb, U. 2003, *A&A*, **404**, 301  
 Rowan-Robinson, M., Gonzalez-Solares, E., Vaccari, M., & Marchetti, L. 2013, *MNRAS*, **428**, 1958  
 Sánchez, B. O., Brout, D., Vincenzi, M., et al. 2024, *ApJ*, **975**, 5  
 Sanders, D. B., Salvato, M., Aussel, H., et al. 2007, *ApJS*, **172**, 86  
 Scarlata, C., Capak, P., Finkelstein, S., et al. 2019, The Euclid Deep Field South, Spitzer Proposal ID #14235,  
 Schinnerer, E., Smolčić, V., Carilli, C. L., et al. 2007, *ApJS*, **172**, 46  
 SCOC 2025, Survey Cadence Optimization Committee's Phase 3 Recommendations, Tech. Rep. PTSN-056, The LSST Project, <https://pstin-056.lsst.io/>  
 Scolnic, D. M., Lochner, M., Gris, P., et al. 2018, arXiv:1812.00516  
 Scoville, N., Aussel, H., Brusa, M., et al. 2007, *ApJS*, **172**, 1  
 Shuntov, M., Akins, H. B., Paquereau, L., et al. 2025, *A&A*, **704**, A339  
 Silverman, J. D., Mainieri, V., Salvato, M., et al. 2010, *ApJS*, **191**, 124  
 Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, **131**, 1163  
 Smartt, S. J., Valenti, S., Fraser, M., et al. 2015, *A&A*, **579**, A40  
 Smith, K. W., Smartt, S. J., Young, D. R., et al. 2020, *PASP*, **132**, 085002  
 Smolčić, V., Novak, M., Bondi, M., et al. 2017, *A&A*, **602**, A1  
 Srivastava, S., Smartt, S. J., Moore, T., et al. 2026  
 Steer, I. 2017, AAS Meeting, **229**, 224.03  
 Stevance, H. F., Smith, K. W., Smartt, S. J., et al. 2025, *ApJ*, **990**, 201  
 Straatman, C. M. S., Spitler, L. R., Quadri, R. F., et al. 2016, *ApJ*, **830**, 51  
 Sullivan, M., Le Borgne, D., Pritchet, C. J., et al. 2006, *ApJ*, **648**, 868  
 Tachibana, Y., & Miller, A. A. 2018, *PASP*, **130**, 128001  
 Tallada, P., Carretero, J., Casals, J., et al. 2020, *A&C*, **32**, 100391  
 Taniguchi, Y., Scoville, N., Murayama, T., et al. 2007, *ApJS*, **172**, 9  
 Taylor, E. N., Franx, M., van Dokkum, P. G., et al. 2009, *ApJS*, **183**, 295  
 Vaccari, M. 2025, The Spitzer Spectroscopic Data Fusion - Merged Spectroscopic Redshift Catalogs in Spitzer Fields, v20 March 2025 , Zenodo, doi:10.5281/zenodo.15176245  
 Weaver, J. R., Kauffmann, O. B., Ilbert, O., et al. 2022, *yCat*, *J/ApJS/258/11*  
 Wenger, M., Oksenbein, F., Egret, D., et al. 2000, *A&AS*, **143**, 9  
 Weston, J. 2026, joshgithubbin/Sherlock-DDF: Lestrade, v0.10.0 , Zenodo, doi:10.5281/zenodo.18378596  
 Williams, R. D., Francis, G. P., Lawrence, A., et al. 2024, *RASTI*, **3**, 362  
 Wiseman, P., Smith, M., Childress, M., et al. 2020, *MNRAS*, **495**, 4040  
 Wolf, C., Meisenheimer, K., Kleinheinrich, M., et al. 2004, *yCat*, **II**, 253A  
 Xue, Y. Q., Luo, B., Brandt, W. N., et al. 2016, *ApJS*, **224**, 15  
 Yoachim, P. 2025, lsst-sims/sims\_featureScheduler\_runs5.0: Initial Release, Zenodo, doi:10.5281/zenodo.15832326  
 Young, D. 2023a, Sherlock. Contextual classification of astronomical transient sources, v2.2.0, Zenodo, doi:10.5281/zenodo.8289325  
 Young, D. 2023b, neddy, v0.3.0 , Zenodo, doi:10.5281/zenodo.8037632  
 Yuan, F., Lidman, C., Davis, T. M., et al. 2015, *MNRAS*, **452**, 3047