Astronomy & Astrophysics

Identifying astrophysical anomalies in 99.6 million source cutouts from the *Hubble* legacy archive using `AnomalyMatch`

David O'Ryan* and Pablo Gómez

European Space Agency (ESA), European Space Astronomy Centre (ESAC), Camino Bajo del Castillo s/n, 28692 Villaneuva de la Cañada, Madrid

ABSTRACT

Aims. Astronomical archives contain vast quantities of unexplored data that potentially harbour rare and scientifically valuable cosmic phenomena. We leverage new semi-supervised methods to extract such objects from the *Hubble* Legacy Archive.
Methods. We have systematically searched approximately 100 million image cutouts from the entire *Hubble* Legacy Archive using the recently developed `AnomalyMatch` method, which combines semi-supervised and active learning techniques for the efficient detection of astrophysical anomalies. This comprehensive search rapidly uncovered a multitude of astrophysical anomalies presented here that significantly expand the inventory of known rare objects.
Results. Among our discoveries are 86 new candidate gravitational lenses, 18 jellyfish galaxies, and 417 mergers or interacting galaxies. The efficiency and accuracy of our iterative detection strategy allows us to trawl the complete archive within just 2–3 days, highlighting its potential for large-scale astronomical surveys.
Conclusions. We present a detailed overview of these newly identified objects, discuss their astrophysical significance, and demonstrate the considerable potential of `AnomalyMatch` to efficiently explore extensive astronomical datasets, including, for example, the upcoming *Euclid* data releases.

Key words. methods: data analysis – catalogs – galaxies: general – galaxies: interactions – galaxies: peculiar

1. Introduction

With the launch of *Euclid* (Euclid Collaboration 2025b), the beginning of operations of the Vera C. Rubin observatory (Ivezić et al. 2019), and the construction of the *Nancy Grace Roman Space Telescope* (Mosby et al. 2020), the size of astronomical datasets is growing rapidly. These observatories will survey a large fraction of the sky in an agnostic way, leading to the potential discovery of a large number of new objects of astrophysical interest allowing us to expand our catalogues of rare galaxy morphology types, including cosmological effects like gravitational lenses or galaxies undergoing the effects of dense environments.

Rare objects, often termed astrophysical anomalies, are particularly informative for improving our understanding of galaxy evolution and cosmology. For example, strong lensing – a gravitational effect of chance alignment of galaxies – allows precise testing of the gravitational potential of the foreground galaxy, as well as the in depth study of the background galaxy from magnification effects (Shajib et al. 2024). When observed around a galaxy cluster, such lensing is an excellent test of cosmology, allowing us to probe their dark matter halo. Other examples include jellyfish galaxies (Poggianti et al. 2016, 2017; Durret et al. 2021), galaxy mergers (Pearson et al. 2019; Ackermann et al. 2018; Wei et al. 2024), edge-on protoplanetary disks (Berghea et al. 2024), and Voorwerps (Lintott et al. 2009).

What many of these rare systems have in common are their methods of discovery. Often they are found by experts 'manually' exploring their data, and detecting odd morphologies or serendipitously finding objects of interest (Pearce-Casey et al. 2025; Acevedo Barroso et al. 2025). Another common approach leverages the contribution of citizen scientists. Using platforms such as Zooniverse from the Galaxy Zoo collaboration (Lintott et al. 2008), volunteers are able to mark any galaxies with anomalous morphologies as 'odd'.

However, these two searching strategies are time consuming and difficult to scale to large datasets. Additionally, they suffer particularly from the problem of subjectivity and the difficulty in defining what an 'odd' morphology is. To address this, the Galaxy Zoo: Weird and Wonderful project (Mantha et al. 2024) formalised the distinction between 'interesting' and 'non-interesting' anomalous galaxies in its classification scheme. Yet, volunteers may not be able to identify an anomaly as reliably as an expert in the field of extragalactic astrophysics.

By their nature, anomalies are rare and difficult to find. Hence, applying commonly used machine learning techniques – such as supervised convolutional neural networks – is difficult due to limited training data. Samples containing the wanted anomaly are often small, and it is difficult to train models on limited training data for a highly imbalanced search space with mostly "uninteresting" data.

Approaches such as `Astronomaly` (Lochner & Bassett 2021) aim to resolve this problem. Rather than training to find individual types of anomalies, they utilise a combination of isolation forests (Liu et al. 2008) or local outlier factor (Breunig et al. 2000) with active learning to identify anomalies specifically sought by a user. In further study, this can be combined with feature spaces produced by machine learning models (e.g `Zoobot`; Walmsley et al. 2022, 2023) to identify further examples of

---

an anomaly. This is effectively done with the ulisse tool (Doorenbos et al. 2022), where these feature spaces can be analysed such that objects of similar morphologies appear close together and can be extracted from areas of interest.

Related works in the area often use unsupervised learning to search for astrophysical anomalies. Unsupervised methods in this context leverage unlabelled data and its data structure to create groups of similar objects expecting anomalies to cluster or fall outside clusters. Such an approach was successfully applied in Stein et al. (2022). In this work, a self-supervised machine learning method was used to mine the entire Dark Energy Spectroscopic Instrument survey Data Release 9 for strong gravitational lenses. Their self-supervised method created an output representation space which was linked to the morphology of the objects trained upon. Once this feature space was created, the representation vector of a single image of a gravitational lens (output by the same model) was then used to conduct a similarity search through 76 million images, identifying 1192 candidate gravitational lenses with excellent computational efficiency.

In this work, we utilise a different approach. We frame anomaly detection as an imbalanced binary classification problem. This is not a fundamentally different task as anomaly detection inherently involves distinguishing rare instances from common ones, but it allows us to leverage semi-supervised learning techniques specifically designed for extreme class imbalance, enabling effective learning from minimal labelled anomaly examples. Our method utilises semi-supervised learning (SSL) techniques such as pseudo-labelling (Lee 2013) and consistency regularisation (Bachman et al. 2014) to make binary classifications on an anomalous or nominal source. Additionally, we use active learning, where experts iteratively validate and label predictions of the top ranking scored images. This hybrid approach offers key advantages: it requires very few initial labels (even fewer than ten anomalies), enables expert validation throughout the process, and effectively utilises both unlabelled data and expert knowledge to progressively improve detection performance.

Our approach, named AnomalyMatch, is fully described in the companion paper (Gómez et al. 2025). There, we conduct thorough benchmarking and validation of the approach on established datasets such as mini-ImageNet (Vinyals et al. 2016) and GalaxyMNIST (Walmsley et al. 2022). In parallel to the development, we actively applied this approach to searching for anomalies in the entire *Hubble Space Telescope* Legacy Archive[1] (HLA).

We used the European Space Agency's (ESA) science data platform ESA Datalabs[2] (Navarro et al. 2024) to conduct this search. In initial testing, we attempted to identify edge-on protoplanetary disks. These systems form very distinct 'hamburger' or 'butterfly' shapes which are easily detected by image classification algorithms. However, only very few of them are known in the literature. As we developed and updated the algorithm, we expanded our search via active learning and marking other objects of astrophysical interest. This was due to multiple factors. First, was a limitation of the data we are using. We search primarily in a previously created dataset of *F*814*W* filter *HST* sources observed with the Advanced Camera for Surveys/Wide Field Channel (ACS/WFC; hereafter only ACS). In this dataset, the number of edge-on protoplanetary disks is limited as they are often observed at other wavelengths, in the Wide-Field Planetary Camera-2 (WFPC2) or are very difficult to resolve in *HST* data

(e.g. Burrows et al. 1996; Krist et al. 1998; Stapelfeldt et al. 2003; Ricci et al. 2008; Angelo et al. 2023). The second was that our algorithm began to detect numerous other objects of interest that we wished to pursue. These included lensed quasars, mergers, and gravitational lenses.

In this work, we discuss and display the wealth of new objects found in the development and application of AnomalyMatch and discuss the implications for our model. Overall, we report 1176 newly found anomalies which span 19 different classes. These include galaxy mergers, gravitational lenses and arcs, edge-on protoplanetary disks, and a host of rare galaxy morphologies.

This paper is laid out as follows. Section 2 describes the imaging data we use from the HLA, its creation, its potential limitations and the resultant training set we use. We briefly describe AnomalyMatch in Section 3, but primarily focus on our methods of anomaly extraction from the HLA and of searching the literature. We also briefly describe ESA Datalabs and the efficiency of our model searching through our dataset. Sections 4 and 5 show the anomalies we find through the HLA, we discuss their representation in the literature and the new objects that we have discovered. Finally, Section 6 concludes this paper and we describe our plan for future implementation of AnomalyMatch.

For an excellent description of an anomaly, see the introduction of Ruff et al. (2020). In this work, we define an astrophysical anomaly (or simply anomaly) to be an astrophysical object that shows morphological characteristics which deviate notably from the general population.

Our contributions are as follows:

– We have conducted the first comprehensive systematic anomaly search of the entire HLA, comprising approximately 100 million image cutouts, using the recently developed semi-supervised and active learning method AnomalyMatch.

– We present a substantial catalogue of newly identified astrophysical anomalies, significantly expanding the known populations of rare cosmic phenomena: 417 previously unknown galaxy mergers, 138 candidate gravitational lenses, 18 jellyfish galaxies, and 2 collisional ring galaxies.

– We have demonstrated the exceptional efficiency and accuracy of our approach, processing the entire HLA dataset within just 2–3 days, highlighting its strong potential for rapid anomaly detection in upcoming large-scale astronomical surveys such as *Euclid*.

2. Data and training set

In this work, we used the source cutouts created in O'Ryan et al. (2023). In that work, we searched the HLA for interacting and merging galaxies, but created cutouts of every extended source within the archive. In this section, we briefly describe the creation process and comment on its limitations in finding other anomalous objects in the HLA.

2.1. Source cutouts of the HLA

The HLA contains all *Hubble Space Telescope* (*HST*) observations. These observations were obtained across many different dates, instruments, and filters. Therefore, to create a consistent dataset, O'Ryan et al. (2023) elected to only use observations from the Advanced Camera for Surveys (ACS) Wide Field Channel with the *F*814*W* filter. They also only used observations at Calibration Level 3 – these are science-ready mosaics from the

---

[1] HLA: https://hla.stsci.edu/
[2] ESA Datalabs: https://datalabs.esa.int/

general reduction pipeline. The source of these was the *Hubble* Advanced Product dataset. This subsection of the HLA specifically contains curated *HST* science-ready mosaics from across the sky.

With these selections, ~10k observations were available. While these could be queried from well known Table Access Protocol services via a host like the Mikulski Archive for Space Telescopes this would lead to downloading terabytes of data. Therefore, the science platform ESA Datalabs was employed. This platform provides a direct link to various science archives and a Jupyter Notebook environment to give users direct access to *HST* observations. The *HST* data are stored in specific data volumes dependent on the instrument of observation. The level 3 *F814W* mosaics can be easily accessed with no downloading required.

Rather than conducting computationally expensive source extraction across such a large set of observations, O'Ryan et al. (2023) employed the *Hubble* Source Catalogue (HSC; Whitmore et al. 2016). This is a large, publicly available catalogue. Whitmore et al. (2016) applied `SourceExtractor` (Bertin & Arnouts 1996) across each observation in the HLA. Selecting only extended sources provided 99.6 million source cutouts in the image dataset.

For each cutout, a fixed size of 150×150 pixels was used ($7.5'' \times 7.5''$), with a `LinearStretch` and `ZScaleInterval` from Astropy (Astropy Collaboration 2013; Astropy Collaboration 2018, 2022). Each image was stored as a one-channel greyscale JPEG image. While this does lead to a loss of quality of the image, it allows for more efficient storage. The image creation process was prioritised such that the morphological features of each galaxy were visible. This approach, however, was not optimised to find low surface brightness features, potentially impacting our detection of anomalies such as jellyfish galaxies.

Upon the creation of the dataset in O'Ryan et al. (2023), it was found that many extended sources were shredded. This could result in an extended source appearing in the dataset multiple times, but from different centring points. The number of duplicates of extended sources depended on the source size and, therefore, the redshift of the source. The largest sources, we found, could be duplicated up to five times. These large sources would often be large enough to fill the 150×150 pixel source cutouts, and therefore, would show little detail. These were then given a low anomaly score by `AnomalyMatch`.

Smaller sources, from $z > 0.2$, were less affected by shredding and duplication. In a minority of cases, these could be shredded up to three times. Whitmore et al. (2016) did provide a `MatchID`, which attempted de-duplication of the HSC, but O'Ryan et al. (2023) opted to conduct their own de-duplication after classification. This was motivated by preventing the removal of multiple galaxies that were involved in interactions and preserving the identification of merging galaxies with multiple cores. We conduct de-duplication of the data after the identification of anomalies. When we create our training set, we ensure there are no duplicate sources, or multiple shredded cutouts, of the source.

While only extended sources were selected in the HSC, many dense star fields – such as those in deep ACS observations of Andromeda and the Magellanic Clouds or globular clusters – were also flagged as extended. This occurred because the individual point sources within these regions were so densely packed that they were blended into a single extended source. This is a specific example of an image artifact which would make other methods of anomaly detection difficult. We identify this in
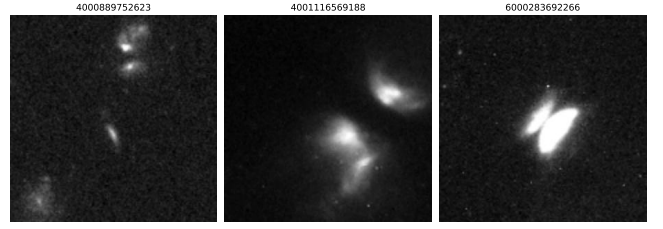

Fig. 1. Initial three images labelled anomaly used to train `AnomalyMatch`. These three images containing edge-on protoplanetary disks which we initially aimed to search additional instances of in the HLA. During active learning, this set was expanded to include sources with odd morphologies like mergers, lenses, and jellyfish galaxies which were serendipitously discovered. Titles are the Source IDs of the objects found in O'Ryan et al. (2023).

the active learning step of training our model, and ensure that `AnomalyMatch` gives these a low anomaly score.

For ease of storage, the 99.6 million cutouts were are stored in ~1000 HDF5 files, containing ~100 000 images each. By using this storage format, the images can be accessed and read to the memory efficiently.

2.2. Training set

Our initial goal was to find more examples of edge-on protoplanetary disks. Therefore, we started our search with three examples of these anomalies, a set 128 labelled nominal data and ~99.6 million unlabelled images. The nominal imaging was selected by visual inspection of randomly selected images from the full dataset. We selected images of individual galaxies, star fields, and any imaging artifacts that appeared in the sample. Figure 1 shows the three anomalies we initially used to train `AnomalyMatch`. However, upon conducting active learning, we found that the odd morphology of these systems led to the discovery of other objects of interest also being given high anomaly scores. With these new objects, we began to pivot our anomaly search to other interesting objects which were also more likely to appear in *F814W* data.

From this attempt to search for edge-on protoplanetary disks, we had constructed a second catalogue of interesting objects through active learning. The anomalies at this stage were classified based on their morphology alone, rather than by searching the literature for specific kinds of objects. In total, the training set consisted of 1400 images: 375 anomalies and 1025 nominal images. We continued building the nominal labelled set during active learning, adding to it substantially to represent further un-interesting morphologies and unaccounted artifacts that were scored highly. The anomalies were primarily mergers (178); galaxies which are undergoing close interactions and coalescing. However, we also uncovered many gravitational lenses that were incorporated into the training set (63). Figure 2 shows a subset of the training set that was built during the development of `AnomalyMatch`.

Increasing the size of the training set and generalising it led to the identification of many different sub-classes of anomalies. However, we were unable to identify new edge-on protoplanetary disks that were not already present in the literature. First, this was due to the known ones becoming a subset of our training set as we add other types of anomalies to it. Secondly, was a limitation of the data we were searching, as observations of edge-on protoplanetary disks in the *F814W* are rare in the literature.
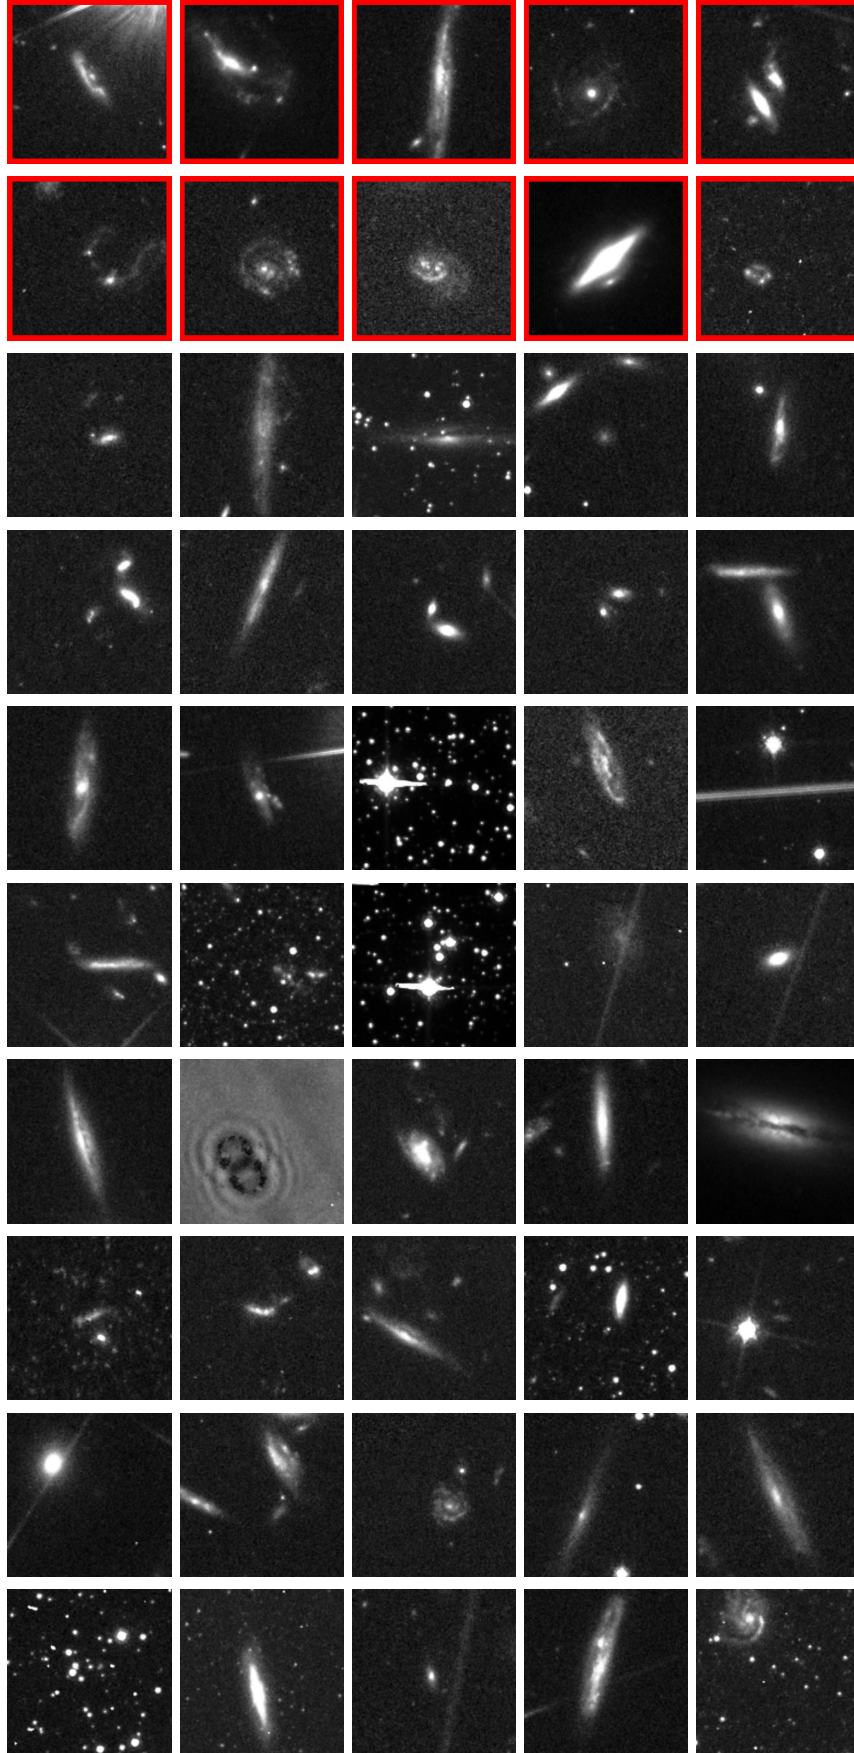
Fig. 2. Fifty examples of the final training set used in applying `AnomalyMatch` to the HLA. The top two rows, highlighted in red, show ten examples of the anomaly class. These are made up of mergers, lenses, edge-on proto-planetary disks as well as some galaxies showing odd morphology. The remaining 40 images are then examples of 'nominal' data. This is primarily isolated galaxies, star fields, and artifacts.

3. Methods

3.1. AnomalyMatch

`AnomalyMatch` is an approach combining active learning and SSL that we have developed and optimised for the purpose of identifying anomalies in large datasets. It is fully described, benchmarked, and validated in the companion paper Gómez et al. (2025). We provide a brief description of it here.

`AnomalyMatch` is based on the `FixMatch` approach (Sohn et al. 2020) as iterated on in `MSMatch` by Gómez & Meoni (2021). The latter adopted an EfficientNet backbone, support for multispectral data and has demonstrated high accuracy using limited labelled samples through consistency regularisation and pseudo-labelling, effectively handling challenging scenarios such as remote sensing with multispectral data.

Building upon these principles, `AnomalyMatch` explicitly frames anomaly detection as a binary classification problem, distinguishing rare anomalies from abundant nominal data, efficiently leveraging a minimal number of labelled anomalies alongside extensive unlabelled data. To effectively address severe class imbalance, the supervised component employs binary cross-entropy loss combined with oversampling of the minority anomaly class, ensuring robust learning from sparse labelled samples.

The unsupervised component applies weak augmentations to generate high-confidence pseudo-labels for unlabelled images, enforcing prediction consistency with strongly augmented versions of these images to exploit intrinsic dataset structure. Crucially, `AnomalyMatch` integrates an active learning approach via an interactive user interface (UI), enabling iterative expert validation and labelling of high-confidence anomaly candidates, progressively refining the model and significantly enhancing anomaly detection performance on large datasets with few known anomalies.

This divide between large quantities of unlabelled data and small samples of labelled data is often the case in astronomy. This is especially true for anomalies, where the number of known examples of an object in certain datasets can be less than ten. Therefore, `AnomalyMatch` provides a method to expand the catalogues where examples of specific objects are very few.

`AnomalyMatch` improved on `MSMatch` by incorporating an active learning loop. The user provides `AnomalyMatch` with a set of labelled and unlabelled data and then conducts an initial round of training. Predictions are then made on the unlabelled data and an 'anomaly score' is given to each image based on the model's classification. The images are ranked by this score, and then shown to the user in an easy to use UI. An example of this is shown in Gómez et al. (2025). The user can then provide additional labels for consecutive training on a dataset that is progressively expanded and refined. This process of active learning is where the the bulk of our initial training set has been collected. Figure 3 shows the full `AnomalyMatch` workflow. The `AnomalyMatch` code will be published on GitHub pending completion of an ESA open-source licensing process.[3]

3.2. Anomaly identification

To identify anomalies, the only initial information we have is their morphology. During active learning and inference over the HLA, there was no direct access to further ancillary data. The underlying distribution of anomaly score is highly skewed to nominal images – meaning the model is very confident that an

object is not anomalous. Given the strong class imbalance and training on only a limited subset of data, model scores were not calibrated and they are not probabilities. We do not apply a cut to this score but investigate their ranked order.

After training on all labelled data and a subset of the unlabelled data, we visually inspect the 5000 sources with the highest predicted anomaly scores of the entire HLA. Based on either the morphology of the galaxy or the object in the image, we are able to make an initial estimation of the kind of anomaly we find. However, to confirm or reject this, we must turn to the literature. We made use primarily of SIMBAD[4] and the ESASky platform[5]. SIMBAD is an excellent tool to programmatically check if the anomalies we found have any associated papers. We conducted a cone search about the coordinates of each anomaly with a radius of 3″. We then checked if any source within this radius had was associated to any work at the time of writing. This could then be followed up quickly and efficiently using ESASky.

In many cases, no literature was associated with the anomalies we discovered. For these objects, we determined a classification based solely on the morphology of the system. This classification would be informed by comparing to others we had made where literature was available. Below, we break down our classification scheme for defining different categories of anomalies. Example images for each class are shown in Figure 7 (Section 4).

3.2.1. Galaxy mergers

Galaxy interaction and merging is a well studied phenomenon, with many different algorithms developed for detecting them. Mergers are identifiable as two or more galaxies lying at approximately the same redshift and exhibit signs of a gravitational effect upon one another. Samples are often plagued by incompleteness or contamination from galaxy pairs which are only close together by projection (e.g. discussions in Ackermann et al. 2018; Pearson et al. 2019; Margalef-Bentabol et al. 2024). Distinguishing between physically interacting galaxy pairs and close pairs is challenging without ancillary data such as velocity or redshift information. Samples of these objects typically include a few thousand systems (Darg et al. 2010; Pearson et al. 2022).

While there are many stages to a galaxy merger, going from interacting pairs to coalescence, we define any kind of interacting galaxy system as a merger. To identify mergers, we rely on the existence of tidal features between the two systems. These features form through the distortion of the galactic disks as the interaction progresses. This makes our approach primarily sensitive to late-stage mergers, where the systems have already passed each other, resulting in a highly distorted galactic disk with one or more cores visible within it.

3.2.2. Overlapping galaxies

Overlapping, or 'backlit', galaxies are systems that appear merging in the 2D projection but are at large 3D separations in the plane of the sky. The galaxies show little morphological distortion, with their disks overlapping. This distinguishes them from 'close pair' contamination in our merging sample, where the two galaxies are close together in the projection of the sky but not overlapping. Samples of such systems are often small, with the largest being 2000 from Galaxy Zoo 2 (Keel et al. 2013). They

---

[3] `AnomalyMatch`: https://github.com/esa/AnomalyMatch

[4] https://simbad.u-strasbg.fr/simbad/
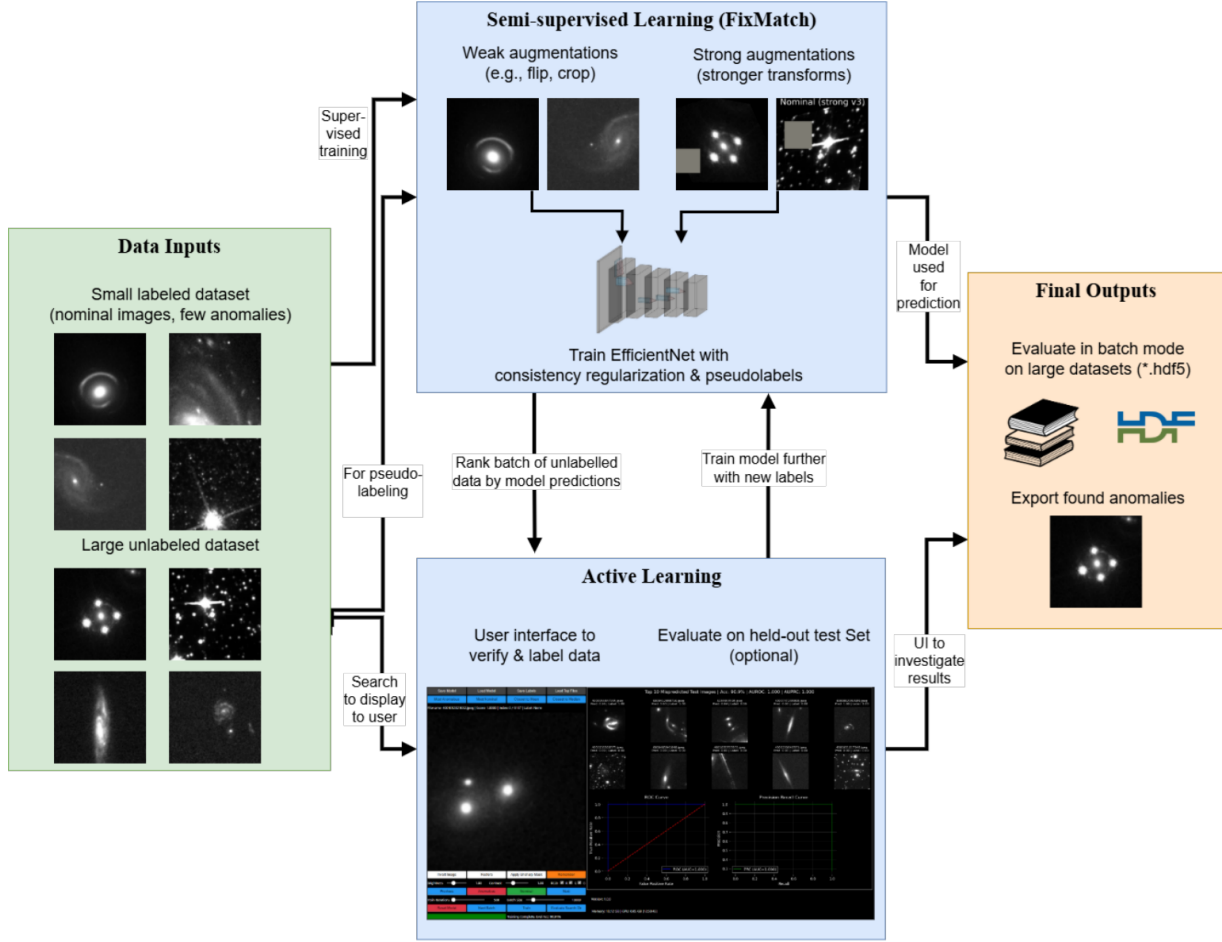[5] https://sky.esa.int/esasky/

Fig. 3. Workflow when using `AnomalyMatch`. We leverage both labelled and unlabelled data from a user dataset to train an EfficientNet architecture and include an active learning loop. Here, the unlabelled data is ranked by anomaly score, the user can extract more examples of the object they are searching for and add them to their training data. Once the desired model metrics are achieved, the model can be saved and then run across all images in their dataset.

are primarily used in the study of dust attenuation independent of dust temperature and the substructure of galaxies.

3.2.3. Gravitational lenses

Gravitational lenses are the prime example of anomalies in the literature, with many recent works searching for them. They are the product of chance alignment between two galactic systems. The background galaxy's light path is altered in the foreground galaxy's gravitational potential and smeared into an arc around the foreground galaxy. We can observe these gravitational lenses (or, rings if the alignment is complete) around the foreground galaxy, and reconstruct the image of the background galaxy in higher detail than observing it normally. This is due to the lensing having a magnification effect on the background galaxys' light.

The most recent works searching for gravitational lenses used *Euclid* data and have employed both expert labelling (Pearce-Casey et al. 2025) and machine learning techniques (Euclid Collaboration 2025a). We classify a gravitational lens about an object when we observe these structures around the galaxies in our sample.

3.2.4. Gravitational arcs

Gravitational arcs and lenses are results of the same light-bending phenomenon. However, in this work, we distinguish between a strong lens about a galaxy and a strong arc about a galaxy cluster. In our dataset, gravitational arcs are found in source cutouts themselves, rather than closely wrapped around a galaxy within the image. Famous examples of gravitational arcs include the Cosmic Snake (Cava et al. 2018) and those used in dedicated galaxy cluster lensing searches (Postman et al. 2012; Coe et al. 2019).

3.2.5. Lensed quasars

Strongly lensed quasars (often in an Einstein Cross configuration) are the final type of gravitational lensing effect that we classify among our anomalies. These occur when there is the chance alignment between a foreground galaxy and a background highly luminescent quasar, forming distinctive and bright point sources around the foreground galaxy. Very few of these objects are known (<200), with the discovery of a new system often leading to a publication (e.g. Agnello et al. 2018; Bettoni et al. 2019; Tubín-Arenas et al. 2023; Lemon et al. 2023).

3.2.6. Jellyfish galaxies

Jellyfish galaxies are a unique galaxy type found in dense galaxy cluster environments. They are galaxies with high internal gas fractions which are undergoing intense ram pressure stripping (RPS). The dense intracluster medium acts on the internal gas of

the galaxy and strips it into long tendrils in the opposite direction of travel. A distinctive bow-shock forms in the direction of travel of the galaxy, as the gas within is compressed and star formation is likely enhanced (Rohr et al. 2023; Zhu et al. 2024). Across the literature, samples of these objects rarely exceed 200 systems, although in recent works they are beginning to increase (Poggianti et al. 2017; Roberts et al. 2021; George et al. 2024; Foster et al. 2025).

As we classify our anomalies, we find many galaxies which show signs of RPS, for instance spiral galaxies with large gas reservoirs around them, or long tails with no obvious companion to have created it by interaction or merging. We, therefore, only identify jellyfish galaxies if they show the distinctive bow shock and are located in a cluster environment.

3.2.7. Clumpy galaxies

Clumpy galaxies, owing their name to the presence of large and luminous star-forming clumps within them, are gas rich systems which often reside at $1 < z < 3$ (Bournaud et al. 2008; Shibuya et al. 2016). Samples typically range from 100 to 1000 systems (Guo et al. 2015; Adams et al. 2022). The clumps can easily be mistaken for a second core within the galactic disk leading to possible misclassification as a merger. We ensure that when making either the merger or clumpy classification, we consider for any morphological disturbance to the galactic disk. Also, the existence of more than one secondary core in the galaxy is likely an indication that it is a clumpy galaxy rather than a merger.

3.2.8. Active galactic nuclei

Active galactic nuclei (AGNs) are the central supermassive black holes of galaxies which are undergoing accretion and growth. They are usually classified from spectral information using diagnostic measurements to identify them based on their emission (e.g. Kauffmann et al. 2003). Samples of AGNs obtained this way are usually large – a few thousand systems. We, however, do not have access to this information and only use the morphology of the system. In rare cases, an AGN can be so luminous that the centre of the galaxy detected as a bright point source. Upon visual inspection, systems hosting a point source at their galactic centre are classified as hosting an AGN.

3.2.9. Ring and collisional ring galaxies

We specify two different kinds of ringed galaxies in our anomaly classification scheme: 'normal' and collisional ring galaxies. Normal ring galaxies are systems that host a large, empty gap in the galactic disks while a large ring is formed on the outskirts. The formation of such rings is an active area of research. Samples of such rings typically contain several hundred objects (Buta 1995; Comerón et al. 2014; Timmis & Shamir 2017).

Collisional rings are formed in galaxy mergers with specific orientations of the impact, when the two galaxies go directly through each other meaning sample sizes are limited. The interaction causes a shockwave to move through the disk incurring a burst of star formation (Appleton & Struck-Marcell 1996; Mapelli et al. 2008). We distinguish collisional ring galaxies from ring galaxies as they are far more luminous. They also show disruption to their disks, and the ring itself can be bent or host features. The existence of a secondary galaxy nearby is also accounted for in making this morphology classification.

3.2.10. Edge-on protoplanetary disks

As stated in Section 3, we originally deployed `AnomalyMatch` to search for edge-on protoplanetary disks. These systems are exceptionally rare, with samples not exceeding 25 objects across the literature (Villenave et al. 2020). They exhibit a single, dark dust lane across their centres with a distinctive butterfly shape extending in perpendicular directions. They often host a jet, which is also perpendicular to the dust lane. These objects, when known, are very well studied. We, therefore, use this distinctive morphology to identify them in our sample of anomalies. We also closely crossmatch potential candidates with the literature.

3.2.11. Galaxies hosting a jet

Galaxies hosting a jet are not uncommon in the literature. Often, in the process of supermassive blackhole growth, long jets will form on either side of the galaxy from the galactic core. While common in the radio or X-ray bands, finding these at optical wavelengths is rare. The most famous example of this is the jet of M87 (Biretta et al. 1999), which has been extensively studied with *HST*. Optical jets have been identified at various redshifts, typically in samples on the order of tens, and are often used to study properties of AGNs (e.g. Blandford et al. 2019; Kravchenko et al. 2025).

3.2.12. High redshift galaxies

During the creation of the images in O'Ryan et al. (2023), the redshift distribution of the underlying sources was unknown. Therefore, there are many examples in our dataset of sources detected with very low signal-to-noise. In these systems, the morphology of the galaxy can be difficult to distinguish against the background.

When these low signal-to-noise systems are marked as anomalous by `AnomalyMatch`, we broadly label them as high-redshift galaxies, without further classification into anomaly subtypes. Some of these are simply described as 'high-redshift galaxies' in the literature as well.

3.2.13. Odd galaxies and unknown galaxies

Finally, we introduce two general classifications that we apply to anomalies that do not easily fit above categories. We label as 'Odd' galaxies those which have been scored highly because they host morphological abnormalities or are highly irregular, so they cannot be classified as nominal galaxies, but do not fit our other criteria to be classified as a specific anomaly. Many of the galaxies that meet this definition are located in clusters. In such environments, effects like RPS, harassment and merging lead to major changes in their morphology. However, they do not host a bow shock to be a jellyfish galaxy and they are not in the process of interacting or merging to be classified as such. These systems also have no available literature to aid in their classification.

The final class is 'Unknown' galaxies. These galaxies have a morphology which completely defies classification based on morphology. In many cases, they may not be galaxies. They have no associated literature to aid in their classification. Therefore, we leave these as unknown objects that could potentially serve as new anomalies for further sub-classification, but we are unable to make such a classification in this work.
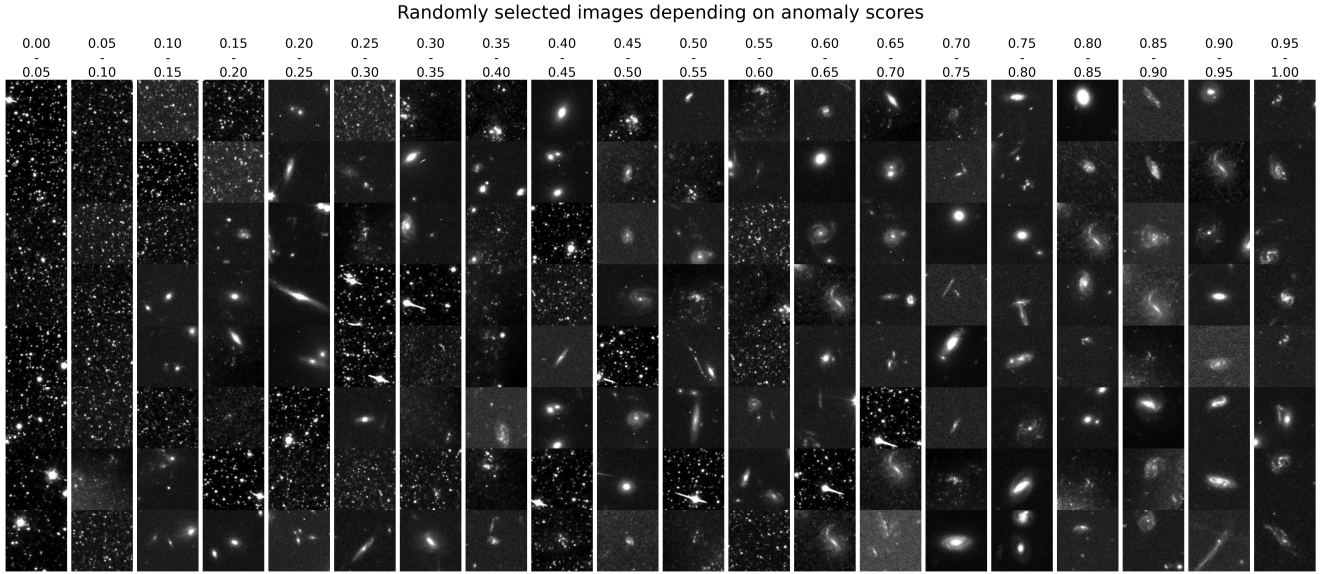
Randomly selected images depending on anomaly scores

Fig. 4. Exemplary anomaly scores on a random subsample of the data. Notably, artefacts are clearly isolated and increasing scores correspond well with increasingly interesting data. The model shows robustness against varying brightness, image noise or differing sizes of objects in the images.

3.3. Anomaly score and SHAP distributions

Upon training `AnomalyMatch`, we are able to test the model across subsets of our unlabelled data. To demonstrate the performance and output of `AnomalyMatch`, we apply it to five random of the stored 1000 HDF5 files. `AnomalyMatch` gives each image a score between 0 and 1 – the anomaly score. Figure 4 shows representative examples of different anomaly scores from this distribution. This clearly shows that many of the images in the HLA sample are of star fields, which are ranked very low in anomaly score. As we increase the score, more and more galaxies with ordinary morphologies are revealed.

Figure 5 shows the resultant distribution of anomaly scores across all five HDF5 files (~500 000 sources). The majority of our sources are given scores ~0. While we do not know the underlying distribution or occurrence rates of anomalies in the dataset, we do not expect many anomalies to appear in a random selection of 0.5% of the data. At an anomaly score of ~1, we observe anomalies which could be of interest.

To investigate how our model makes its classifications and verify that it behaves as expected, we take examples of the top and bottom scored anomalies and overlay saliency plots on each image. These plots highlight which areas of an image are informing the model to make a classification. We map out changes in the SHapely Additive exPlanations (SHAP Lundberg & Lee 2017) values across each image. This, essentially, shows a weighting of which pixels were relatively more important to the model's classification.

Figure 6 shows the results of our saliency mapping. We show three correctly identified anomalies, one false positive, and one example of a low scored image. In the correctly identified anomalies, SHAP values are highest at the pixels where the anomaly is within the image. For the lens, this covers the arc itself. For the merger, this is over the tidal features. For the edge-on protoplanetary disk, this is over the butterfly shape of the object.

The SHAP values for the false positive have an interesting distribution. `AnomalyMatch` has highly weighted the core of the galaxy – which could be an AGN in this case – and the extended, loosely wound spiral arm. This has the appearance of a tidal
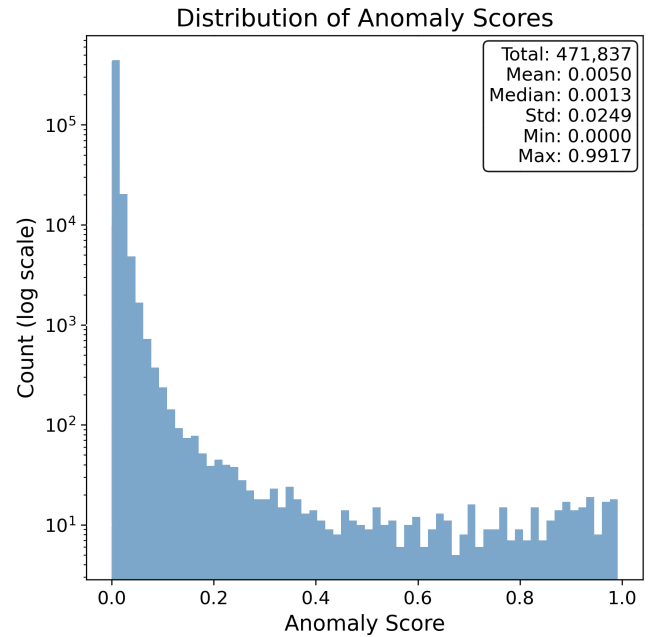

Fig. 5. Anomaly score distribution obtained by applying our final trained `AnomalyMatch` model on a random subset of the HLA of ~500 000 cutouts. We find that the distribution is highly weighted to zero, as expected. The majority of our sources are not anomalous.

feature in a merger beyond the disk and towards a second small system which may be a satellite.

Finally, the last image is of a star field where we can see the SHAP values either weight up or down different areas of the image. However, it has mainly down-weighted the pixels with many stars with empty space between them.

Figure 6 demonstrates that `AnomalyMatch` has successfully been trained to make anomaly scores based on source(s) within the image, rather than on another property of the image. Even in images containing only one source, it weights up pixels based on dividual sections of the galaxy rather than the entire object
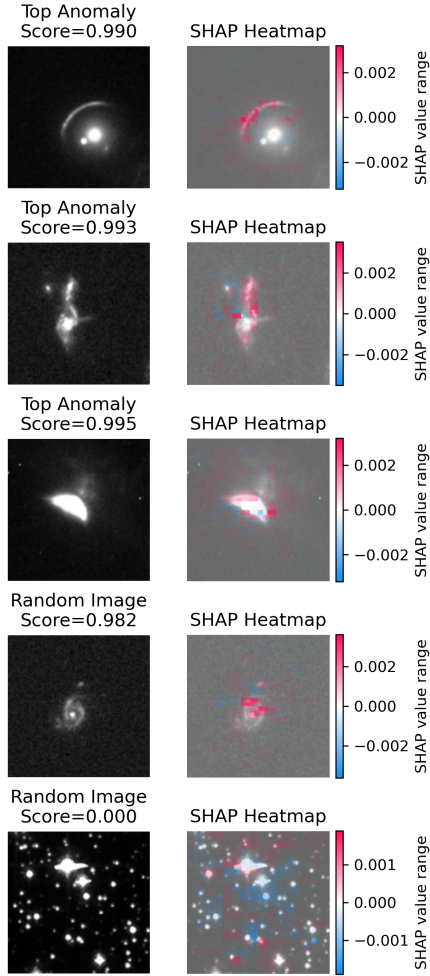
Top Anomaly
Score=0.990

SHAP Heatmap

Top Anomaly
Score=0.993

SHAP Heatmap

Top Anomaly
Score=0.995

SHAP Heatmap

Random Image
Score=0.982

SHAP Heatmap

Random Image
Score=0.000

SHAP Heatmap

Fig. 6. Saliency plots for three highly scored anomalies, one high scoring false positive, and one low-scoring image, presented as heat maps of the SHAP values assigned to each pixel. These visualisations highlight the regions most relevant to the models predictions, often corresponding to features a human might also consider significant (or lack thereof, in the case of the final image).

in the image. This is shown with the lens image (image 1), where `AnomalyMatch` primarily uses the lens itself to make the classification instead of the galaxy within the image.

3.4. ESA Datalabs

To store our data, conduct training, and complete inference over the images we utilised the platform ESA Datalabs. This platform allowed for direct access to data from many of ESA's observatories and missions where it is a partner with other agencies. This removed the requirement for downloading observations or the movement of data. We took advantage of the newly integrated GPU cluster which was available via the platform, and allowed us to both train our models and conduct inference efficiently, as proved by the timescale of this project.

We fully trained the model over our training set (1400 labelled and ~99 000 unlabelled images) and took less than four hours on a GPU via ESA Datalabs. To make classifications over all ~100 million images took 2.5 days running the code on ESA Datalabs with no user input, or intervention. The main bottleneck in efficiency was loading the images into memory to then

be inferred over. This was done via HDF5 files in batches of 100 k.

This high level of efficiency shows that ESA Datalabs is an ideal platform for large-scale data exploration. With the addition of GPUs to the platform, applying machine learning algorithms to large quantities of observational data is straightforward. This will facilitate further exploration of a wide range of observational datasets, including JWST, *Euclid*, and *HST*.

4. Results

To perform the anomaly sub-classification, we initially selected the top 5000 scored anomalies from the HLA. First, we apply de-duplication to these 5000 samples. The output from `AnomalyMatch` is a CSV file containing the filename of the source and the score given to the source. In our case, the source filename is the source ID found in the HSC. We cross match each of our sources with the HSC based on this, and extract their coordinates. We then apply an aggressive radial cross match between each of the sources within 10″.

We applied such a large cut as the likelihood that two of our anomalies are within this separation was low. It allowed us to completely de-duplicate the final sample, and ensure we were not using expert time re-inspecting duplicate entries. This reduced the number of anomalies from 5000 to 1339 unique images. This also highlights the high level of shredding and duplication in the HSC. Upon de-duplication, DOR acted as the expert classifier and classified each anomaly according to the system defined in Section 3.

4.1. Detected anomalies

The unique 1339 anomalies taken from the ranking of highest anomaly score were classified into their sub-classifications. This was done by a combination of accounting for morphology and searching the literature for works related to each anomaly. Table 1 shows a breakdown of these classification of the anomalies. Figure 7 shows a representative sample of all sub-classifications where we found at least five anomalies, except for our 'odd', 'nominal', and 'unknown' classifications, which will be presented later in the section. The 'nominal' classification is our false positive rate, where `AnomalyMatch` returned a normal source.

The largest population of anomalies we find are merging or interacting galaxies (~50%). These are likely the most common type of anomaly that we are searching for, as well as the most distinctive. With our aggressive de-duplication, these represent 629 merging systems rather than individual systems. The cutouts used in `AnomalyMatch` contain a field of view of $7.5'' \times 7.5''$, therefore galaxy mergers contained in the same source cutout will be de-duplicated under one ID. Following our definition, not all of our classified merger anomalies have a secondary. Many are highly disturbed merging or post-merger systems, with only one system in the image.

It is also important to note that for some merger classifications, the full extent of the system is not visible in the field of view of the cutout. However, the tidal features of the system are and, therefore, the anomaly score was ranked highly. Looking at these systems within ESASky reveals the full system and associated literature when the field of view is adjusted.

The second most numerous anomaly we detect is gravitational lenses (and arcs, if classified under the same category). These are currently the main anomaly being searched for in the
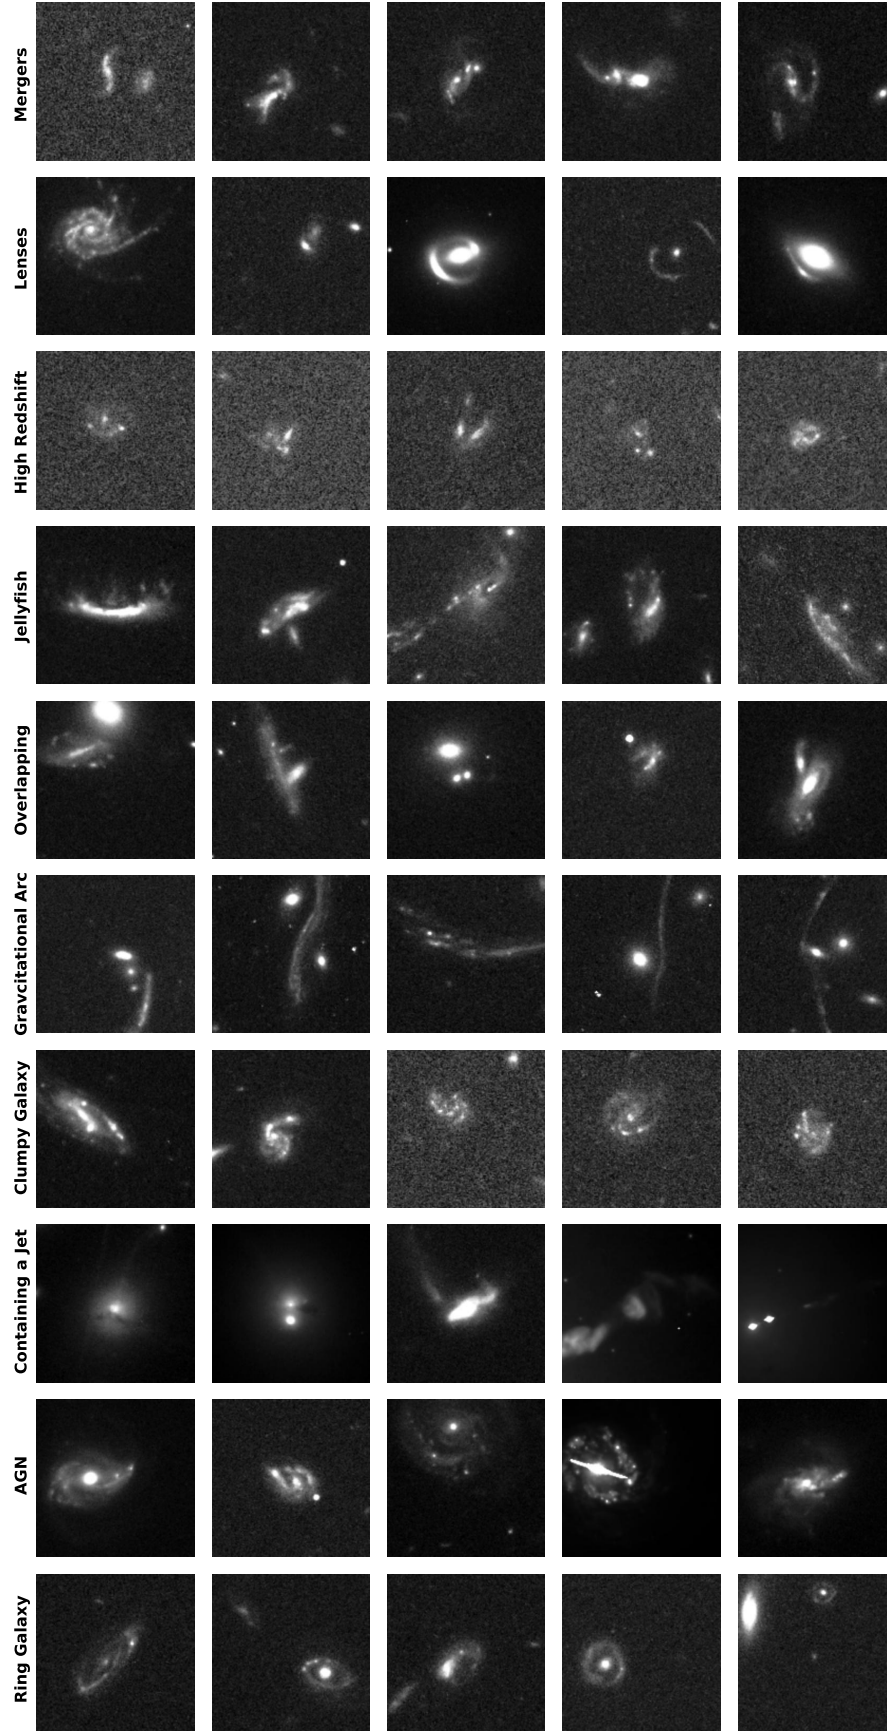
Fig. 7. Five examples of every anomaly sub-class for which we found at least five objects (excluding lensed quasars for later discussion), selected as representative of that sub-class.

Table 1. Breakdown of anomalies found in the development of the `AnomalyMatch` algorithm.

Classification	N Found	N Referenced
Merging Galaxy	629	212
Gravitational Lenses	140	54
Odd Galaxy	164	60
Nominal Galaxy	163	60
Unknown Morphology	43	0
High Redshift Galaxy	28	7
Jellyfish Galaxy	35	17
Overlapping Galaxy	39	26
Gravitational Arc	39	19
Clumpy Galaxy	11	5
Galaxy Hosting a Jet	13	5
Galaxy Hosting an AGN	8	4
Ring Galaxy	12	5
Lensed Quasar	5	5
Collisional Ring Galaxy	2	0
Edge-On Protoplanetary Disk	2	2
Galaxy Hosting Supernova	2	2
Submillimetre Galaxy	1	1

field, and therefore, made up a large portion of the training set. Due to the shape and position of lenses, the main source of contamination in our results will be spiral arms, or other features which appear like arcs around galaxies. The different luminosity of the lens aids in the classification, however this is not always the case. We identify many gravitational lenses that are already identified in the literature – but many candidate new lenses. As we did not attempt to model and confirm them we point out that these are only candidate lenses which require either follow-up observations or modelling.

We do not include our sample of 39 gravitational arcs from our lens classification here. Often, these arcs – formed by clusters rather than galaxies – are so large, that they extend out of the field of view of our source cutouts. We see the true extent of these systems in ESASky, although only parts of them are flagged as anomalies here. Other cutouts containing the arc are removed in our aggressive cross matching and de-duplication of 10″. We also ensure that each of our 39 arcs is unique using ESASky.

We find a population of what we have termed 'high redshift' galaxies ($z > 1$). These are systems that appear highly disturbed, or affected in the image, but are at the threshold of being detected. Looking at those which are referenced, we find that these systems are representative of high redshift galaxies. We also find them to be small and clumpy.

The clump classification we make provides a sample of 11 systems. These galaxies could also be mistaken for merging galaxies, as they often look irregular compared to the general galaxy population. However, the disks often host many more than one other core – which would be expected in mergers – and often these are contained in spiral arms, which would be destroyed in the merging process.

We find 35 different jellyfish galaxies in our 1339 unique systems. These have been classified as such as they clearly have stripping occurring, and they are residing in a dense cluster environment. On each of the examples shown in Figure 7, the bow shock is present in the direction of travel.

Next, we identify a similar number of overlapping galaxy systems, with the whole system included in the field of view.

A very small sample (13) of our found anomalies are galaxies which host relativistic jets. These can rarely be detected in the $F814W$ filter of *HST* (the most famous being M87, which is shredded and appears multiple times in the HSC). Our systems are much smaller and often show a small jet or object moving away from the galactic core. This classification was made by comparing to the literature, where galaxies with active jets made up just two of the samples we found. However, upon recognising this morphology, we were able to make more classifications.

Our final two classifications are AGN-host galaxies (8) and ring galaxies (12). Our classification of an AGN is primarily from the literature, where all but one of these systems has an associated reference. Otherwise, the requirement for ancillary data would make this classification difficult with `AnomalyMatch` alone. Finally, we identify ring galaxies by their morphology as the structure is easily recognisable compared to others.

`AnomalyMatch` successfully detected several types of anomalies for which we had provided no explicit training examples, such as lensed quasars. While these unseen anomaly types likely share morphological features with our trained classes (e.g. lensed quasars may exhibit arc-like features similar to gravitational lenses or multiple bright regions like mergers), their detection demonstrates the method's ability to generalise to morphologically similar but distinct anomaly types. However, we acknowledge that unseen anomalies with fundamentally different morphologies from our training set would likely be missed. Figure 8 shows the four systems that we found in HLA. We initially provided `AnomalyMatch` with no examples of these, however, even in the first training iterations, it identified these anomalies and scored them highly. Therefore, we added one to the training set and identified four additional lensed quasars.

Figure 9 shows the identified collisional rings, edge-on protoplanetary disks, and galaxies hosting a supernova. While ancillary data are required for the classifications of supernova hosts and edge-on protoplanetary disks, we confirm these with the literature here. This is also true for our classification of a submillimetre galaxy, where the object we identify has a large repository of literature associated with it. However, our collisional rings are not represented in the literature and are, therefore, classifications we have made by checking examples of known collisional rings and confirm their morphology is similar.
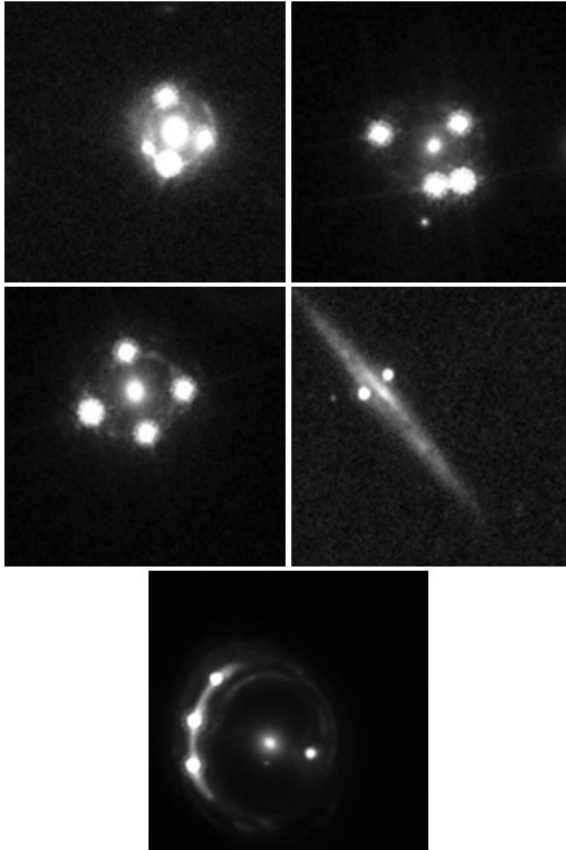
Finally, we include three general classifications in our released catalogues: odd, nominal, and unknown galaxy morphologies. As stated in our classification scheme, we define odd galaxies as those which show an odd morphology which could be identified as due to interaction or RPS. However, we find that they do not fit this criteria. Often, they are galaxies within a dense environment showing some distortion. Figure 10 shows ten representative examples of this distortion possibly due to harassment by other cluster members. Some of these odd galaxies are not in a cluster, and simply are irregular galaxies that `AnomalyMatch` has scored highly. It is not unreasonable for these to have been detected with our extended training set of different anomalies, particularly if including mergers.

Galaxies we classify as nominal by visual inspection, but that received a high score by `AnomalyMatch`, represent the contamination in our final anomaly sample. We find a contamination rate of ∼10%. Figure 11 shows a representative example of ten sources. Primarily, these sources are very small sources which exhibit some potential merging behaviour, but it is difficult to make a specific classification. Other examples include sources with artifacts from the observation itself, aligned with a point source and star fields. Star fields with this morphology are often either from the Panchromatic *Hubble* Andromeda Treasury

Table 2. First five rows of the MRT released with this work.

SourceID	RA	Dec	Classification	Referenced
4000862191779	8.461384	−7.833325	Lens	False
4000705851114	186.877395	33.544637	Lens	False
6000532112741	342.211252	−44.513375	Merger	True
4000905661818	34.404996	−5.224907	Lens	True
6000265754287	53.175990	−27.773749	Merger	True

Notes. The table contains the (1) Source ID from the HSC, (2) Right Ascension in degrees, (3) Declination in Degrees, (4) classification in this work, and (5) whether the classification was made using the literature.

Fig. 8. Five lensed quasars (Einstein crosses) that were found in this work. All of which are represented in the literature. However, using *HST*, we would not expect to find unreferenced systems.

(Dalcanton et al. 2012) or studies of large globular clusters with the ACS instrument. Therefore, many of the nominal images we find here are anomalous when compared to the rest of the data, but they are of limited value to the astrophysical community.

Finally, we find 43 objects with morphologies defying classification. Some of these objects may not be galaxies but rather other objects that we have limited expertise in classifying. Figure 12 shows ten example sources we have been unable to classify. We release these to the community for further discussion, or use, but do not attempt to make morphology classifications here. None of these objects have definitions in the literature.

We release each of these catalogues as a machine-readable table (MRT). The table contains the source identifier from the HSC, its right ascension and declination, our classification, and whether the source has been classified by using the literature. An example of the first five rows of the table are shown in Table 2.

Fig. 9. Anomalies where less than five objects were found in this work. We find two collisional ring systems, two further edge-on protoplanetary disks, and two galaxies containing supernovae.

For the ease of use of our catalogue, we also release all of the sources and table on Zenodo (DOI: `10.5281/zenodo.15298641`). Here, the images are contained in separate .zip files based on their anomaly classification. The filenames are the Source IDs of the sources.

5. Discussion

Conducting robust and complete anomaly detection is nontrivial. As stated previously, utilising supervised machine learning methods suffer from a lack of training data, and the complex and variable nature of anomalies to be found. Using this test run of the neural network `AnomalyMatch` during development, we have been able to extract a range of anomalous objects from a limited training set. With this SSL approach, `AnomalyMatch` is
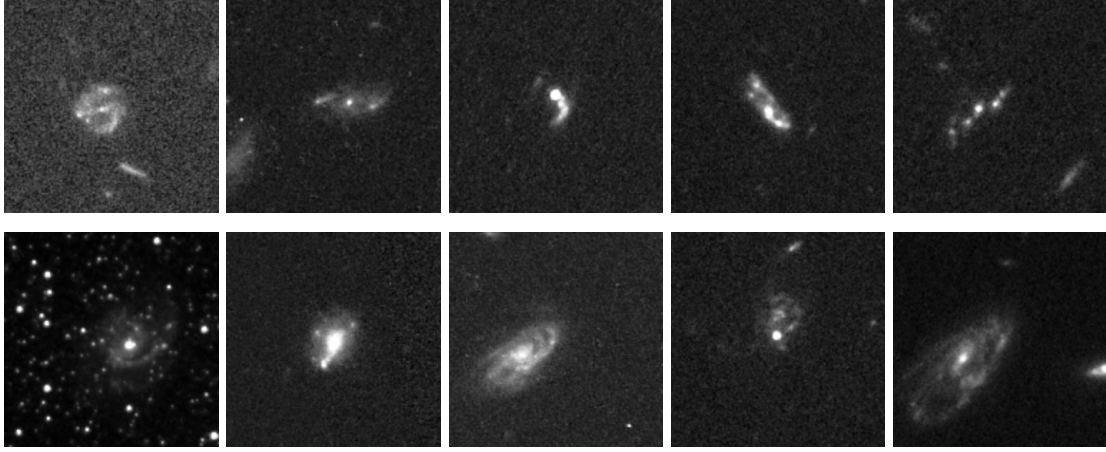
Fig. 10. Ten examples of sources to which `AnomalyMatch` gave a high anomaly score and we have classified as odd morphologies. These galaxies are close to being classified as nominal galaxies – for instance, of little interest to the community – but have odd morphological features. These are often due to these galaxies residing in dense environments.
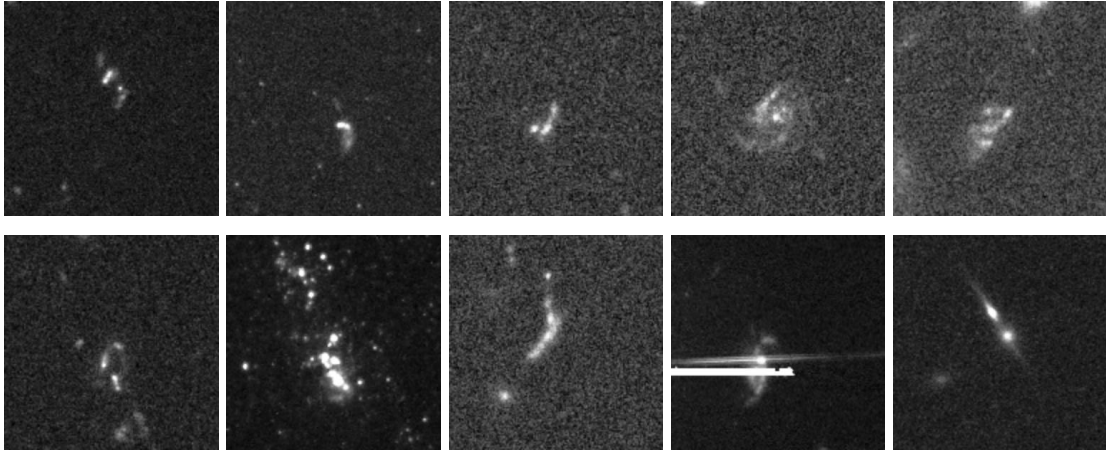

Fig. 11. Ten examples of sources to which `AnomalyMatch` gave a high anomaly score, but we have visually classified as nominal images. These include sources which exhibit some features, but are too small in the field of view to give definitive classification. It also includes some star fields, as well as artifacts in unique alignments.

Fig. 12. Ten examples of sources to which `AnomalyMatch` gave a high anomaly score, but we have been unable to morphologically classify. Many of these could be jellyfish galaxies, as they appear to show RPS signatures. However, they do not reside in a dense cluster environment. Others show curved morphologies, which are do not fall into our classification scheme here.

able to distinguish unique morphologies from general ones using an incredibly limited set of labels.

Other works focused on anomaly detection often must begin from using similarity searches, to citizen scientists to identify large samples for further searches (Lintott et al. 2008; Lochner & Bassett 2021; Mantha et al. 2024; Euclid Collaboration 2025c). Conversely, `AnomalyMatch` can be used as a baseline approach to grow existing samples of anomalies from the literature. It can also be employed in a flexible manner, where other anomalies not included in the initial labelled dataset can be detected and extracted during the hunt for other objects.

This is the main difference between the envisioned implementation of `AnomalyMatch` and what we find here. We started by searching for edge-on protoplanetary disks – as they are very rare objects – but expanded our training set to include other objects like mergers, lenses, and jellyfish galaxies upon their serendipitous discovery. This approach results in a varied output which will contain objects a user is not interested in, but the inclusion of the active learning loop in the training process allows them to filter out unwanted anomalies and retain those which are desired. For instance, in this work, the majority of the anomalies we find are mergers. Mergers are reasonably common – although difficult to build large samples of. In the future, we would use the active learning loop to remove these objects, increasing purity and accuracy in finding a desired anomaly.

In this work, we find 811 unreferenced objects in our sample of 1339 objects. However, by the nature of using the HLA and, specifically, *HST* to search for anomalies, we are limiting our capability for identifying new objects. The *HST* is not a survey telescope: observations must be applied for, and PIs specifically request the coordinates and objects they wish to observe. Therefore, any anomalies we do find are either in the background of observations of another astrophysical source of interest, or are the target of the observation in question. The latter is particularly true for edge-on proto-planetary disks, where these targets of exceptional interest are rarely observed 'by accident' in the background of targeted observations. The fact that ~65% of our anomalies are not represented in the literature thus shows the exceptional potential of `AnomalyMatch` for application to survey telescopes' data.

With survey telescopes, much of the data will be completely unexplored and rarely looked at besides with machine learning algorithms. This role is where `AnomalyMatch` will excel and, as shown in this work, be able to identify objects of various morphologies that can then be classified into different categories. Missions such as *Euclid* and Vera C. Rubin observatories are perfect for this role, where Terabytes of data will be collected and processed per night.

In the future, `AnomalyMatch` will be used to search for specific types of anomalies that the user can hone in on. By ranking objects in this way, we will be able to rapidly find and increase the size of our samples of different anomalies while using visual inspection. As found in Gómez et al. (2025), anomalies identified with the 1% highest scores have a high precision, or a high chance to be only or almost only anomalies – assuming such objects are present in the data of them in the first place. However, extracting them still requires visual inspection by the user.

6. Conclusion

We have developed and deployed an innovative approach based on semi-supervised and active learning for the purpose of identifying sources with anomalous morphologies. The method – `AnomalyMatch` – has been created as an out-of-the-box tool,

which has been integrated into the ESA Datalabs platform. It has an intuitive and interactive UI, which seamlessly combines semi-supervised learning methods with active learning so a user can identify objects of interest efficiently and with few labels. The full benchmarking and testing of this algorithm is described in Gómez et al. (2025).

In this work, we release the anomalies that we discovered in the HLA while developing this approach. Using a large test set of 99.6 million image cutouts in the *F*814*W* filter of *HST*, we ranked them all by anomaly score. To start, we aimed to identify more edge-on protoplanetary disks. Using a training set of just three labelled objects, a large sample of nominal and unlabelled data, we searched the archive for objects of interest. During active learning, we identified many other objects of astrophysical interest and grew the training set based on what we found.

Upon completion of development, training this model, and applying it across the HLA, we selected the top 1339 de-duplicated sources. We then classified these into different sub-classes of anomalies based on morphology. These included 629 mergers, 140 gravitational lenses, 35 jellyfish galaxies, and 2 edge-on protoplanetary disks. Using SIMBAD and ESASky, we conducted a literature search based on the coordinates of each object, and to investigate if they had been included in other catalogues and samples. We find that ∼ 65%, 811, of the objects do not appear in the literature.

We release a machine readable table containing source IDs, positions and classifications of all sources we have discussed in this work. We also release both the tables and images of each object on Zenodo (DOI: `10.5281/zenodo.15298642`), where these samples are available for download.

The large number of unreferenced sources is excellent and shows `AnomalyMatch`'s capability for identifying new systems. This is especially notable given that our data comes from *HST*, where each observation is targeted – meaning that either a PI requested to observe the anomaly we have detected, or it is a previously unnoticed background system.

This shows the potential of `AnomalyMatch` for use in large surveys, such as *Euclid* or Vera C. Rubin Observatory, which will observe large areas of the sky with no knowledge of the objects contained within. The large volume of data that will be produced per night will be impossible to manually search and visually inspect. Therefore, using algorithms where small amounts of data are required for training will be paramount for anomaly detection in these large surveys.

`AnomalyMatch` thus fills a particularly fruitful niche. We will be able to use this method to gradually grow our small samples of different anomalies, while also uncovering anomalies with completely new morphologies.

`AnomalyMatch` is freely available on the ESA Datalabs platform with an open-source release pending a successful licensing process currently underway at ESA. The software will be available on GitHub[6]. It has been developed with GPU capabilities to be able to seamlessly explore ESA data archives and search them for anomalies of interest to other users. This powerful addition and implementation of SSL will enhance our ability to detect anomalies across the field.

Data availability

A copy of the catalog is available at the CDS via `https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/704/A227`.

---

[6] `https://github.com/esa/AnomalyMatch`

References

Acevedo Barroso, J. A., O'Riordan, C. M., Clément, B., et al. 2025, A&A, 697, A14

Ackermann, S., Schawinski, K., Zhang, C., Weigel, A. K., & Turp, M. D. 2018, MNRAS, 479, 415

Adams, D., Mehta, V., Dickinson, H., et al. 2022, ApJ, 931, 16

Agnello, A., Schechter, P. L., Morgan, N. D., et al. 2018, MNRAS, 475, 2086

Angelo, I., Duchene, G., Stapelfeldt, K., et al. 2023, ApJ, 945, 130

Appleton, P. N., & Struck-Marcell, C. 1996, Fund. Cosmic Phys., 16, 111

Astropy Collaboration (Robitaille, T. P., et al.) 2013, A&A, 558, A33

Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, AJ, 156, 123

Astropy Collaboration (Price-Whelan, A. M., et al.) 2022, ApJ, 935, 167

Bachman, P., Alsharif, Q., & Precup, D. 2014, Adv. Neural Inform. Process. Syst., 27

Baines, D., Giordano, F., Racero, E., et al. 2017, PASP, 129, 028001

Berghea, C. T., Bayyari, A., Sitko, M. L., et al. 2024, ApJ, 967, L3

Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393

Bettoni, D., Falomo, R., Scarpa, R., et al. 2019, ApJ, 873, L14

Biretta, J. A., Sparks, W. B., & Macchetto, F. 1999, ApJ, 520, 621

Blandford, R., Meier, D., & Readhead, A. 2019, ARA&A, 57, 467

Bournaud, F., Daddi, E., Elmegreen, B. G., et al. 2008, A&A, 486, 741

Breunig, M. M., Kriegel, H.-P., Ng, R. T., & Sander, J. 2000, SIGMOD Rec., 29, 93

Burrows, C. J., Stapelfeldt, K. R., Watson, A. M., et al. 1996, ApJ, 473, 437

Buta, R. 1995, ApJS, 96, 39

Cava, A., Schaerer, D., Richard, J., et al. 2018, Nat. Astron., 2, 76

Coe, D., Salmon, B., Bradač, M., et al. 2019, ApJ, 884, 85

Comerón, S., Salo, H., Laurikainen, E., et al. 2014, A&A, 562, A121

Dalcanton, J. J., Williams, B. F., Lang, D., et al. 2012, ApJS, 200, 18

Darg, D. W., Kaviraj, S., Lintott, C. J., et al. 2010, MNRAS, 401, 1043

Doorenbos, L., Torbaniuk, O., Cavuoti, S., et al. 2022, A&A, 666, A171

Durret, F., Chiche, S., Lobo, C., & Jauzac, M. 2021, A&A, 648, A63

Euclid Collaboration (Bergamini, P., et al.) 2025a, https://doi.org/10.1051/0004-6361/202554577

Euclid Collaboration (Mellier, Y., et al.) 2025b, A&A, 697, A1

Euclid Collaboration (Walmsley, M., et al.) 2025c, A&A, accepted [arXiv:2503.15324]

Foster, L. M., Parker, L. C., Gwyn, S., et al. 2025, ApJ, 982, 120

George, K., Poggianti, B. M., Omizzolo, A., et al. 2024, A&A, 690, A337

Giordano, F., Racero, E., Norman, H., et al. 2018, Astron. Comput., 24, 97

Gómez, P., & Meoni, G. 2021, arXiv e-prints [arXiv:2103.10368]

Gómez, P., et al. 2025, arXiv e-prints [arXiv:2505.03509]

Guo, Y., Ferguson, H. C., Bell, E. F., et al. 2015, ApJ, 800, 39

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357

Hunter, J. D. 2007, Comput. Sci. Eng., 9, 90

Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111

Kauffmann, G., Heckman, T. M., Tremonti, C., et al. 2003, MNRAS, 346, 1055

Keel, W. C., Manning, A. M., Holwerda, B. W., et al. 2013, PASP, 125, 2

Kravchenko, E. V., Pashchenko, I. N., Homan, D. C., et al. 2025, MNRAS, 538, 2008

Krist, J. E., Stapelfeldt, K. R., Burrows, C. J., et al. 1998, ApJ, 501, 841

Lee, D.-H. 2013, ICML 2013 Workshop : Challenges in Representation Learning (WREPL)

Lemon, C., Anguita, T., Auger-Williams, M. W., et al. 2023, MNRAS, 520, 3305

Lintott, C. J., Schawinski, K., Slosar, A., et al. 2008, MNRAS, 389, 1179

Lintott, C. J., Schawinski, K., Keel, W., et al. 2009, MNRAS, 399, 129

Liu, F. T., Ting, K. M., & Zhou, Z.-H. 2008, IEEE, 413

Lochner, M., & Bassett, B. A. 2021, Astron. Comput., 36, 100481

Lundberg, S. M., & Lee, S.-I. 2017, in Advances in Neural Information Processing Systems, 30, eds. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Curran Associates, Inc.), 4765

Mantha, K., Krawczyk, C., Roberts, H., et al. 2024, Citizen Sci.: Theory Pract., 9, 1

Mapelli, M., Moore, B., Ripamonti, E., et al. 2008, MNRAS, 383, 1223

Margalef-Bentabol, B., Wang, L., La Marca, A., et al. 2024, A&A, 687, A24

McKinney, W. 2010, https://conference.scipy.org/proceedings/scipy2010/pdfs/mckinney.pdf

Mosby, G., Rauscher, B. J., Bennett, C., et al. 2020, J. Astron. Telesc. Instrum. Syst., 6, 046001

Navarro, V., del Rio, S., Angel Diego, M., et al. 2024, in Space Data Management. Studies in Big Data, 141 (Springer Nature), 1

O'Ryan, D., Merín, B., Simmons, B. D., et al. 2023, ApJ, 948, 40

Pearce-Casey, R., Nagam, B. C., Wilde, J., et al. 2025, A&A, 696, A214

Pearson, W. J., Wang, L., Trayford, J. W., Petrillo, C. E., & van der Tak, F. F. S. 2019, A&A, 626, A49

Pearson, W. J., Suelves, L. E., Ho, S. C. C., et al. 2022, A&A, 661, A52

Poggianti, B. M., Fasano, G., Omizzolo, A., et al. 2016, AJ, 151, 78

Poggianti, B. M., Moretti, A., Gullieuszik, M., et al. 2017, ApJ, 844, 48

Postman, M., Coe, D., Benítez, N., et al. 2012, ApJS, 199, 25

Ricci, L., Robberto, M., & Soderblom, D. R. 2008, AJ, 136, 2136

Roberts, I. D., van Weeren, R. J., McGee, S. L., et al. 2021, A&A, 650, A111

Rohr, E., Pillepich, A., Nelson, D., et al. 2023, MNRAS, 524, 3502

Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., et al. 2021, Proc. IEEE, 109, 756

Shajib, A. J., Vernardos, G., Collett, T. E., et al. 2024, Space Sci. Rev., 220, 87

Shibuya, T., Ouchi, M., Kubo, M., & Harikane, Y. 2016, ApJ, 821, 72

Sohn, K., Berthelot, D., Li, C.-L., et al. 2020, arXiv e-prints [arXiv:2001.07685]

Stapelfeldt, K. R., Ménard, F., Watson, A. M., et al. 2003, ApJ, 589, 410

Stein, G., Blaum, J., Harrington, P., Medan, T., & Lukić, Z. 2022, ApJ, 932, 107

Timmis, I., & Shamir, L. 2017, ApJS, 231, 2

Tubín-Arenas, D., Lamer, G., Krumpe, M., et al. 2023, A&A, 672, L9

Villenave, M., Ménard, F., Dent, W. R. F., et al. 2020, A&A, 642, A164

Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., & Wierstra, D. 2016, in Advances in Neural Information Processing Systems. 29 (NeurIPS), 3630

Walmsley, M., Lintott, C., Géron, T., et al. 2022, MNRAS, 509, 3966

Walmsley, M., Scaife, A. M. M., Lintott, C., et al. 2022, MNRAS, 513, 1581

Walmsley, M., Allen, C., Aussel, B., et al. 2023, J. Open Source Softw., 8, 5312

Wei, S., Song, X., Zhang, Z., et al. 2024, ApJS, 274, 23

Wenger, M., Ochsenbein, F., Egret, D., et al. 2000, A&AS, 143, 9

Whitmore, B. C., Allam, S. S., Budavári, T., et al. 2016, AJ, 151, 134

Zhu, J., Tonnesen, S., & Bryan, G. L. 2024, ApJ, 960, 54