

# Harnessing the Hubble Space Telescope Archives: A Catalog of 21,926 Interacting Galaxies

David O’Ryan<sup>1,2</sup> , Bruno Merín<sup>2</sup> , Brooke D. Simmons<sup>1</sup> , Antónia Vojteková<sup>2</sup>, Anna Anku<sup>2</sup>, Mike Walmsley<sup>3</sup> , Izzy L. Garland<sup>1</sup> , Tobias Géron<sup>4</sup> , William Keel<sup>5</sup> , Sandor Kruk<sup>6</sup> , Chris J. Lintott<sup>4</sup> , Kameswara Bharadwaj Mantha<sup>7</sup> , Karen L. Masters<sup>8</sup> , Jan Reerink<sup>2</sup>, Rebecca J. Smethurst<sup>4</sup> , and Matthew R. Thorne<sup>1</sup> 

<sup>1</sup> Department of Physics, Lancaster University, Bailrigg, Lancaster, LA1 4YB, UK; [d.ryan@lancaster.ac.uk](mailto:d.ryan@lancaster.ac.uk)

<sup>2</sup> European Space Agency (ESA), European Space Astronomy Centre (ESAC), Camino Bajo del Castillo s/n, E-28692, Villanueva de la Cañada, Madrid, Spain

<sup>3</sup> Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Oxford Road, Manchester, M13 9PL, UK

<sup>4</sup> Oxford Astrophysics, Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford, OX1 3RH, UK

<sup>5</sup> Department of Physics and Astronomy, University of Alabama, Box 870324, Tuscaloosa, AL 35487, USA

<sup>6</sup> Max-Planck-Institut für Extraterrestrische Physik (MPE), Giessenbachstrasse 1, D-85748 Garching bei München, Germany

<sup>7</sup> School of Physics and Astronomy, University of Minnesota, 116 Church Street SE, Minneapolis, MN 55455, USA

<sup>8</sup> Departments of Physics and Astronomy, Haverford College, 370 Lancaster Avenue, Haverford, Pennsylvania 19041, USA

Received 2022 November 25; revised 2023 February 10; accepted 2023 February 28; published 2023 May 4

## Abstract

Mergers play a complex role in galaxy formation and evolution. Continuing to improve our understanding of these systems requires ever larger samples, which can be difficult (even impossible) to select from individual surveys. We use the new platform ESA Datalabs to assemble a catalog of interacting galaxies from the Hubble Space Telescope science archives; this catalog is larger than previously published catalogs by nearly an order of magnitude. In particular, we apply the `Zoobot` convolutional neural network directly to the entire public archive of HST F814W images and make probabilistic interaction predictions for 126 million sources from the Hubble Source Catalog. We employ a combination of automated visual representation and visual analysis to identify a clean sample of 21,926 interacting galaxy systems, mostly with  $z < 1$ . Sixty-five percent of these systems have no previous references in either the NASA Extragalactic Database or Simbad. In the process of removing contamination, we also discover many other objects of interest, such as gravitational lenses, edge-on protoplanetary disks, and “backlit” overlapping galaxies. We briefly investigate the basic properties of this sample, and we make our catalog publicly available for use by the community. In addition to providing a new catalog of scientifically interesting objects imaged by HST, this work also demonstrates the power of the ESA Datalabs tool to facilitate substantial archival analysis without placing a high computational or storage burden on the end user.

*Unified Astronomy Thesaurus concepts:* [Interacting galaxies \(802\)](#); [Computational methods \(1965\)](#); [Catalogs \(205\)](#)

*Supporting material:* machine-readable table

## 1. Introduction

Interacting and merging galaxies are important to our current theory of  $\Lambda$  cold dark matter ( $\Lambda$ CDM) cosmology, in which structure typically assembles hierarchically (Abadi et al. 2003; Springel et al. 2005; De Lucia & Blaizot 2007; Guo & White 2008). Galaxy interaction leads to highly disturbed morphologies (Toomre & Toomre 1972; Hernández-Toledo et al. 2005; Wallin et al. 2016), intense starbursts (Mihos & Hernquist 1996; Springel 2000; Saitoh et al. 2009; Moreno et al. 2021), and, potentially, quenching of some systems (Hopkins et al. 2013; Smethurst et al. 2018; Hani et al. 2020; Das et al. 2022). In general, galaxies undergoing interaction are observed to have higher star formation rates than those that exist in the field (Ellison et al. 2008; Scudder et al. 2012; Pearson et al. 2019). Interaction also has a direct impact on the gas angular momentum within each galaxy, causing it to decrease. This, potentially, leads to funnelling of gas into their nuclear regions and igniting activity. This could be a connection with active galactic nuclei (AGNs; Ellison et al. 2008, 2011; Li et al. 2008; Comerford et al. 2015). However,

such a connection remains debated (Alonso et al. 2007; McKernan et al. 2010; Marian et al. 2020). Thus, understanding galaxy interaction is crucial to testing theories of galaxy evolution itself.

Interacting galaxies have long been explored with different samples of galaxies. Examples include constraining merger rates as a function of redshift (Lotz et al. 2008), inferring the contribution of minor mergers to the cosmic star formation budget (Kaviraj 2014a, 2014b), and examining interactions as a function of their local environments, internal properties, and AGN activity (Darg et al. 2010a). These studies (and many others; for further examples, see Barton et al. 2000; Alonso et al. 2004; Ellison et al. 2013; Holincheck et al. 2016; Silva et al. 2021) illustrate the complex parameter space involved in understanding the role of interaction in galaxy evolution. Thus, to effectively study interacting galaxies, we need observed data sets of such a size that they can sample a wide range of various parameters of interest.

The first large-scale catalogs of interacting galaxies are from the mid 20th century (Vorontsov-Velyaminov 1959, 1977; Arp 1966). These catalogs primarily used visual inspection to identify mergers (e.g., de Mello et al. 1997; Nair & Abraham 2010) and generally found from hundreds to thousands of systems. The largest set of interacting galaxies identified by a single expert classifier contains 2565 relatively



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.

nearby systems (Arp & Madore 1987). Citizen science techniques can extend this number, as was presented by Darg et al. (2010b) who used them to find a catalog of 3003 interacting galaxies.

The inclusion of automated classification shows promise to continue this expansion. The use of machine learning in classifying galaxy morphology is well established (Ardizzone et al. 1996; Abd El Aziz et al. 2017; Barchi et al. 2020; Ghosh et al. 2020; Cheng et al. 2021). The workhorse algorithm is the convolutional neural network (CNN; for an introduction, see O’Shea & Nash 2015), most often used in image recognition and feature extraction. CNNs can be used for general classification (e.g., early- versus late-type galaxies) or to extract specific morphological features of galaxies, such as bars and spiral arms; many works have demonstrated their effectiveness at this (e.g., Ackermann et al. 2018; Jacobs et al. 2019; Bickley et al. 2021; Buck & Wolf 2021; Walmsley et al. 2022a). Pearson et al. (2022) demonstrated the power of CNNs for finding interacting and merging galaxies specifically, finding 2109 in 5.4 deg<sup>2</sup> of Hyper Suprime-Cam imagery—a large sample for the small area covered.

However, issues with using CNNs in classifying interacting galaxies have been found on numerous occasions. The primary concern, is that—without due care—classifying interacting galaxies by morphology alone can be highly contaminated. For example, CNNs often confuse chance alignments of galaxy pairs on the sky for interacting systems. This leads to many predicted interacting systems being thrown away after visual inspection (in some cases up to 60%; Bottrell et al. 2019; Pearson et al. 2022).

In this work, we aim to use machine learning to create a large, high-confidence catalog of interacting systems, drawn entirely from existing astronomical imagery. We search through the European Space Agency’s Hubble Space Telescope (HST) Science Archive<sup>9</sup> using a CNN to predict whether an image contains an interacting system, from among the 126 million extended objects in the Hubble Source Catalog (Whitmore et al. 2016, hereafter **HSC**). The feature extraction we implement is focused on finding tidal features or morphological disturbance caused by the interaction. The tidal features prioritized include tidal tails, tidal bridges, or tidal debris. As stated previously, this runs the risk of introducing high levels of contamination by close pairs. We thus implement further automated and manual methods, which significantly reduce this. The systems we find are often in the background of previous deep surveys (such as the Cosmic Evolution Survey, COSMOS, Scoville et al. 2007; the Great Observatories Origins Deep Survey, GOODS, Giavalisco et al. 2004; and the Pancromatic Hubble Andromeda Treasury Survey, PHAT, Dalcanton et al. 2012), where spectroscopic coverage varies. Therefore, while our final catalog reduces contamination to  $\sim$ 3%, definitively removing all contamination by close pairs remains a challenge following this work.

This paper is laid out as follows: Section 2 describes the **HSC** and all of the criteria we applied to create the images we predict over. It also introduces ESA Datalabs<sup>10</sup>, a new platform that allows the user to directly access the Hubble Science Archive. Section 3 gives an in-depth description of the Zoobot CNN we utilize for our predictions, and how it differs

from a commonly used CNN. Section 4 explains the process of creating the training set for our CNN to find interacting galaxies, and Section 5 shows how well it performed and provides the diagnostics of the CNN. We also use Section 5 to investigate the contamination in our catalog. Section 6 describes our results and discusses the final catalog as well as defines interesting systems or objects that we have found. We also explore some basic properties of the catalog here. Finally, Section 7 summarizes our results and conclusions.

Where necessary, we use a flat  $\Lambda$ CDM cosmology with  $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$  and  $\Omega_M = 0.3$ . Hereafter in this paper, when referring to an interacting galaxy, we are referring to a galaxy that has undergone one or multiple flybys by a secondary galaxy and caused tidal disturbance. A merging galaxy is the final state of these flybys, where two or more systems have coalesced to form a highly morphologically irregular system.

## 2. Data

### 2.1. The Hubble Archives and ESA Datalabs

The observational data is directly from the Hubble Science Archive and is accessed from the new ESA Datalabs platform. The repository contains approximately 100 TB of data from HST. This repository spans all HST instruments and filters. ESA Datalabs provides a direct interface between users and the data. On this platform, every observations’ FITS file can be accessed. To streamline our pipeline, we applied criteria to the observations as not all filters have the same number of observations, some instruments are not as sensitive to the low surface brightness regime as others, or the field of view of certain instruments would not be ideal for measuring galaxy morphology. Finally, we do not conduct source extraction from each FITS file ourselves but use the **HSC** to define the center of each source cutout.

The criteria we apply are: the observational data must be from the Advanced Camera for Surveys (ACS), it must be final product data of HST (i.e., within a .drc file, where the data has been drizzle combined and had charge-transfer-efficiency corrections applied; Avila et al. 2015), observed within the F814W filter, and must be flagged as an extended source in the **HSC**. This offloads sky subtraction, cosmic-ray rejection, and charge efficiency calculations to the original HST pipeline and removes costly steps from our cutout creation process. We utilize all final product data of the F814W filter from HST as this was the filter that contained the most FITS files, and therefore observations. The F814W filter contained 9527 final product FITS files, which could be used for source extraction, whereas the closest second (the F606W filter) contained  $\approx$ 6000. By using the filter with the most files, we are confident that we cover a majority of the **HSC**. Applying this criteria gives 126 million sources to predict over.

We must create 126 million source cutouts from 9507 different FITS files. Creating a data set of cutouts at this magnitude in conventional methods (such as AstroQuery or Table Access Protocol, TAP, services) would be impractical due to making many network calls and long FITS file download times. Instead, we use the ESA Datalabs platform, which is due to be released in Q3 of 2023. This platform has been developed to allow us to “mount” the Hubble Science Archive onto it. In practice, this provides access to the entire Hubble Science Archives as local files for the user to

<sup>9</sup> See <http://hst.esac.esa.int/ehst/>.

<sup>10</sup> <https://datalabs.esa.int/>

manipulate while on the platform. This bypasses network calls to servers to download our required FITS files, a process that could have taken minutes per download. Having direct access to the files, and quickly matching source coordinates to FITS files (described in Section 2.2) allows us to open an FITS file and create all source cutouts from it without having to close or reopen it. Therefore, we were able to create on the order of 10,000 cutouts in the same order of time taken to download a single file.

The source cutouts were created as F814W grayscaled  $150 \times 150$  ( $7.5'' \times 7.5''$ ) pixel images using the [HSC](#) source coordinates as the center. The image size was set and standardized to streamline the pipeline. The majority of cutouts are centered on the source but, in a minority, misalignment between source and image center occurs. This is a result of the drizzling process, with incorrect alignment sometimes being significant. However, the target source was always present in the cutout and we, therefore, did not attempt to rectify this. To scale the images, we used a ZScaleInterval with a hard set contrast of 0.05 and a linear stretch following the default parameters in the [Astropy](#) (Astropy Collaboration et al. 2013, 2018) package. These were binned to  $300 \times 300$  pixels (pixel resolution is  $3.25'' \times 3.25''$ ) with a linear interpolation from the [CV2](#) python package. The images were created at  $150 \times 150$  to minimize storage required on the early version of [ESA Datalabs](#) being used. Creating the images at half the size allowed us to scale up to  $300 \times 300$  pixels without any effects of the interpolation.

## 2.2. The Shapely Python Package

A large computational expense in our pipeline was matching FITS files to sources. Conventionally, the [Astropy](#) `CONTAINS` function would be used to match source coordinates to the FITS file WCS. We instead use the [Shapely](#)<sup>11</sup> Python package. [Shapely](#) is a geometry oriented package primarily focused on geospatial data. We found converting the FITS image footprints into [Shapely](#) Polygons and the source coordinates to [Shapely](#) Points and then checking if they overlapped had significant speed up. Per iteration, [Astropy](#)’s `CONTAINED_BY` function matches a source to an FITS file on the order of 500 ms. Using [Shapely](#)’s `CONTAINS` function, the same process is on the order of 6  $\mu$ s.

## 3. Utilizing a Convolutional Neural Network

We must choose a CNN that would best suit our needs to classify them into interacting galaxies or not. We select the newly developed CNN [Zoobot](#) (Walmsley et al. 2022a, 2022b). [Zoobot](#) is a CNN specifically trained to classify galaxies based on morphology into many different types (spiral, disk, elliptical, barred, nonbarred, etc). We retrain it to only classify galaxies into interacting or noninteracting. Instead of training [Zoobot](#) from scratch and creating a new model, we use transfer learning to fine-tune existing [Zoobot](#) models to classify our data for our particular question. This allows us to retain information from [Zoobot](#)’s previous training. More importantly, it requires a significantly smaller training set to achieve high accuracy.

<sup>11</sup> [Shapely](#) docs: <https://shapely.readthedocs.io/en/stable/manual.html>.

### 3.1. Zoobot

The version of [Zoobot](#) we use is a deep CNN that was trained and tested on Galaxy Zoo volunteer classifications over three different Galaxy Zoo: Dark Energy Camera Legacy Survey (DECaLS, GZD; described in Dey et al. 2019) campaigns. These were GZD-1, GZD-2, and GZD-5—each number corresponding to the DECaLS data release. For training [Zoobot](#), DECaLS imaging was selected using the NASA-Sloan Atlas (NSA), which was itself constructed with Sloan Digital Sky Survey (SDSS) Data Release 8 (DR8) images. This also introduced implicit cuts to the training data, as SDSS cannot get to the depths of DECaLS. This introduces implicit magnitude and redshift cuts on the training data. Specifically, SDSS DR8 and the NSA cover galaxies brighter than  $m_r > 17.77$  and closer than  $z < 0.15$ . In Section 3.2 we describe using transfer learning to use [Zoobot](#) effectively outside of this magnitude and redshift range.

Walmsley et al. (2022a) used the 249,581 volunteer classifications from the GZD-5 campaign to train [Zoobot](#) to answer all 34 questions (example shown in Figure 4 of Walmsley et al. 2022a) in the remaining campaigns. GZD-5 was used as it had a slightly different volunteer decision tree, having an expanded question on potential different galaxy merger stages. Each galaxy image had been shown to volunteers as a three-color ( $g, r, z$ ) of  $424 \times 424$  cutout. Each image pixel scale was an interpolation between the measured Petrosian 50%- and 90%-light radius. The measured full Petrosian radius had to be at least  $3''$  to be shown to the volunteers. When inputting into [Zoobot](#), these cutouts were scaled and grayscaled to  $300 \times 300 \times 1$  images, averaging over the three-color channels to remove color information and avoid biasing the morphology predictions. [Zoobot](#) utilized the Adam (Kingma & Ba 2014) optimizer to train.

By training [Zoobot](#) in this way, combining the approach of answering many questions at once with Bayesian representation learning, it learns a generalizable summary of many types of galaxies. These generalized summaries are lower-dimensional descriptions of galaxy types and are referred to as representations. These representations change depending on the galaxy type, morphology, or environment in an image and lead to similar images being closer together in a representation space than dissimilar ones. This representation approach on a very broad classification problem is found to increase accuracy and generality of [Zoobot](#), giving it an edge over conventional CNNs. A more detailed breakdown of this approach, as well as further details about [Zoobots](#)’ architecture, can also be found in Walmsley et al. (2022a).

[Zoobot](#) was trained to give a prediction score to an image of a galaxy based on the question it is answering. The type of prediction score is set by the user’s choice of the model final layer in [Zoobot](#). We elect to use a SOFTMAX output, which returns an output score as a float between 0 and 1. This prediction score is not a probability score, although it may seem analogous. A well-behaved prediction score will map to probability, though not necessarily linearly. The mapping between prediction score and probability is not considered in this work, and we use the prediction score as an indicator of [Zoobot](#)’s confidence that a source is an interacting galaxy.

We are only interested in the “Is the galaxy merging or disturbed?” question from the Galaxy Zoo: DECaLS workflow, where the answer can be “merging,” “major disturbance,” “minor disturbance” or “none,” and we only want our version

of Zoobot to return the answer to this. Our version of Zoobot is also not trained to predict over HST data, which differs from DECaLS data (different resolutions, filter bandwidths, etc). If we were to use our version of Zoobot as downloaded, we would likely lose accuracy. We utilize transfer learning to optimize accuracy of just our question as well as to classify HST data. Since this work, Zoobot has been trained on HST data so the transfer learning step would not be needed in future with the new models. How we apply transfer learning is discussed in Section 3.2, but an excellent review and discussion of applying transfer learning for detecting galaxy mergers can be found in Ackermann et al. (2018).

### 3.2. Transfer Learning

Transfer learning (or fine-tuning) is a method of applying the same machine-learning model to a similar problem that it was originally trained on. Rather than having to completely retrain all parameters in a model and essentially create a new one, we can use the original model architecture and the parameters it has learned from its previous training. In the case of Zoobot, we keep the parameters it has learned from training on the DECaLS data set and freeze all sections of the model responsible for feature extraction and recognition.

We construct a classification section that maximizes accuracy and only allow the weights of this section to change. As the classification section has fewer parameters than the feature extraction section (the classification section contains 86,209 parameters compared to the feature extraction sections’ 4,048,989 parameters), we need significantly less data to completely retrain it (in our case, a factor of 15 less). Once this retraining is complete, the weights of the feature extraction sections of the model can be unfrozen and tweaked using our smaller data set with a very low learning rate to further boost overall model accuracy.

An example of taking an existing model and applying it to a new problem with transfer learning is shown in Walmsley et al. (2022c). Here, they take the trained model and fine-tune it to finding ring galaxies. They retain an accuracy of 89% while only needing to train the model on  $10^3$  ring galaxies. This significantly reduces computational expense and training time of the model, while keeping the required training set very small. Interacting galaxies are rare, and interacting galaxy catalogs not expansive. So retraining the full network on hundreds of thousands of interacting galaxies is not feasible. Using transfer learning, and following the example from Walmsley et al. (2022c), we only need to create a training set of  $10^3$ – $10^4$  interacting galaxies to achieve an accuracy of  $\approx 90\%$ .

## 4. Creating the Training Set

We create a large training set of interacting galaxies following the criteria described in Section 2 to train our model. Therefore, we need a large, labeled set of interacting and noninteracting galaxies. We elect to follow the methodology of fine-tuning as described in Walmsley et al. (2022c), and aim to create a balanced training set. This has the advantage that it significantly improves the performance and accuracy of machine-learning classifiers, but the disadvantage that it can bias our final model if few interacting galaxies exist compared to the general population. However, such a bias will be mitigated by using a high prediction cutoff to define an interacting galaxy. This is discussed in Section 5.1. To create

this large training set, we use the Galaxy Zoo collaboration (initial data release described in Lintott et al. 2008).

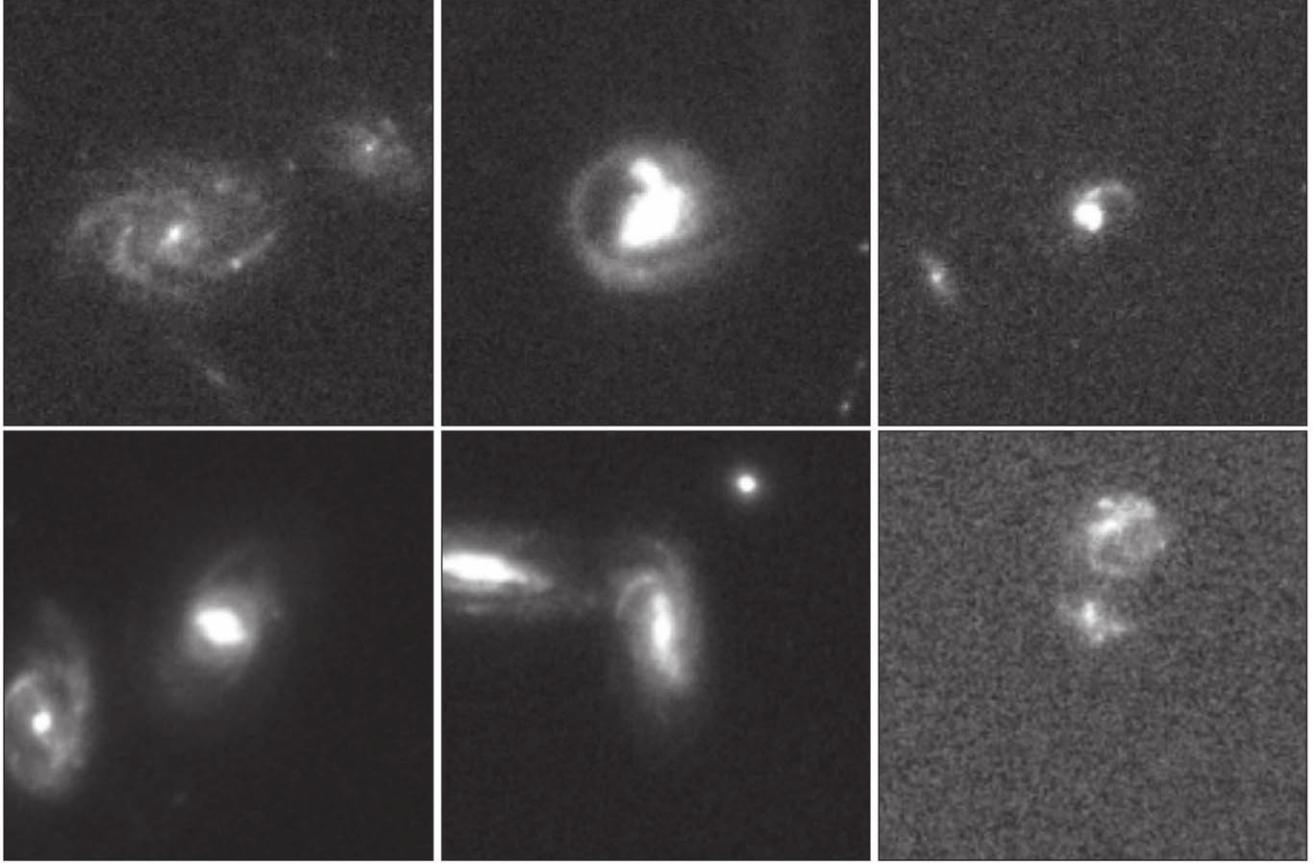
### 4.1. Interacting Galaxies and Galaxy Zoo

The data in Galaxy Zoo is volunteer classifications on galaxy images spanning multiple projects. We incorporate classifications from all major Galaxy Zoo projects; Galaxy Zoo 1 (Lintott et al. 2008), Galaxy Zoo 2 (Willett et al. 2013), Galaxy Zoo: Hubble (Willett et al. 2017), Galaxy Zoo: CANDELS (Simmons et al. 2017), and GZD (Walmsley et al. 2022a). These projects contain a total of 1,367,760 labeled galaxy images that we must extract the interacting galaxies from. We only use labels that are from citizen scientists, and no labels generated by previous versions of Zoobot. We apply three criteria to each interacting or noninteracting label. First, it must have greater than 20 volunteer votes on it. Applying this allows us to use a statistically robust weighted vote from a crowd answer rather than trusting any volunteers individually. Second, the calculated weighted vote (i.e., the combination of the 20 or greater votes) must then be greater than 75% in favor of being an interacting galaxy or less than or equal to 25% for it not to be; this ensured purity in our training set. If the question given to volunteers was more specific (such as “Is this a minor disturbance?” and “Is this a major disturbance?”), then if either answer was the majority vote we classified it as an interacting galaxy. Third, the object must exist in the Hubble footprint so that we could make a cutout of it.

Checking if each training source existed in the Hubble footprint was only possible in an efficient way because of ESA Datalabs. Rather than querying every coordinate and making network calls to TAP services, we extract every final product F814W observation footprint and check if each labeled galaxy exists in at least one file. We make this check by creating a Shapely Polygon for each observational footprint and a Shapely Point for each labeled galaxy central coordinate. Using the Shapely Polygon CONTAINS function, we check if a labeled galaxy’s Point overlaps with an observations’ footprint Polygon. This returns a list of files that contain the training source. If a training source was not found in any observational footprint, we discarded it. We make no attempt here to check if our sources have other photometry available to them, and only create one-color images with the F814W data. We provide the images to Zoobot as one-color grayscaled cutouts.

Upon applying these criteria, we found 3167 labeled interacting galaxies in Galaxy Zoo: Hubble project, the largest contribution to our training set. These were paired with 3167 labeled noninteracting systems (following the previous criteria) to balance the training set. From all other projects, we found 869 labeled interacting systems that fit the creation criteria. The primary limiting factor for Galaxy Zoos 1 and 2 was that many found interacting galaxies did not exist in the Hubble footprint. For Galaxy Zoo: CANDELS and GZD, the limiting factor was the required calculated weighted vote. These labeled interacting systems were then paired with 869 labeled noninteracting systems, ensuring that each labeled noninteracting system came from the same project as its labeled interacting system counterpart.

Each of these projects has a varied redshift range: Galaxy Zoo: Hubble is  $z < 1$ , Galaxy Zoo: CANDELS  $1 < z < 3$ , and Galaxy Zoos 1 and 2 and GZD are  $z < 0.15$ . This introduces a redshift bias into our model, where the morphology and



**Figure 1.** Example images of the labeled interacting galaxy systems used to train *Zoobot*. Each galaxy had a weighted vote fraction  $\geq 0.75$  in Galaxy Zoo. Top row: three examples from the Galaxy Zoo: Hubble project of the training set. Bottom row: three examples from the other Galaxy Zoo projects. These are, from left to right, Galaxy Zoo 2, Galaxy Zoo CANDELS, and GZD. The priority with this training set was that the interactors had clear tidal features and disruption so *Zoobot* would learn to highly weight them and not misclassify close pairs.

brightness of interacting sources change with  $z > 1$ . This is only partially rectified by including Galaxy Zoo: CANDELS, which provided 322 labeled interacting systems.

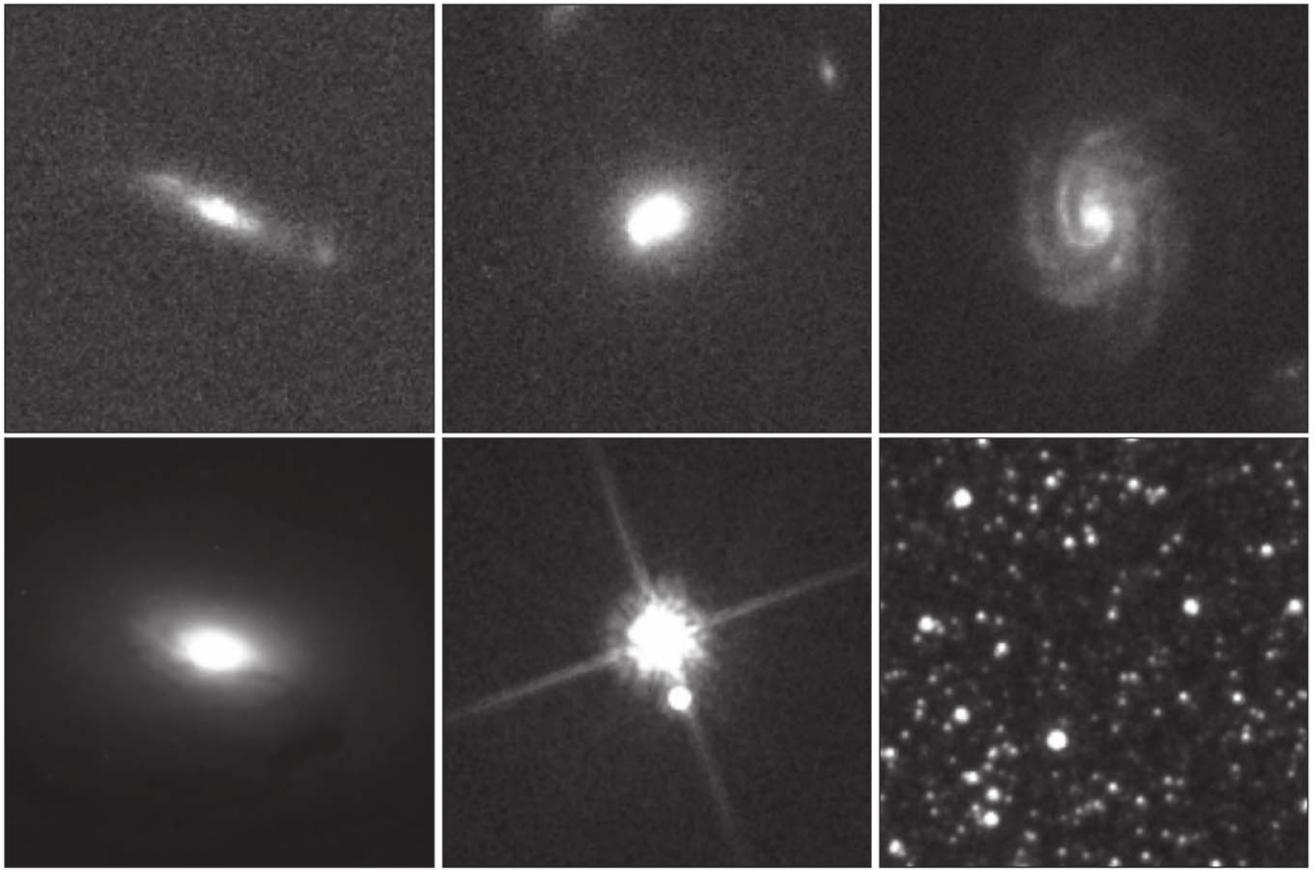
From all Galaxy Zoo projects, we find a training set of 4036 labeled interacting galaxies and combine them with their matched 4036 labeled noninteracting galaxies, giving a total training set size of 8072. Figures 1 and 2 show six examples of our labeled interacting and noninteracting galaxy training set. As we require *Zoobot* to learn to weight tidal features or disturbances highly, it is important that such structures dominate the training set. Previous works, such as Pearson et al. (2022), have found that final catalogs produced by CNNs are often heavily contaminated by sources that are simply close pairs by projection effects and chance alignment in the sky. By focusing our CNN on tidal features, we aim to minimize this contamination. We ran an initial test of the prediction pipeline on the first 500,000 sources that had been created from the [HSC](#) to initially test our *Zoobot* model. We investigate any source that was given a prediction score  $\geq 0.75$  and, to further increase the size of our training set, conduct one step of active learning.

#### 4.2. One Active Learning Cycle

To enlarge our training set further, we conduct one step of active learning to find interacting galaxies. An active learning cycle involves an “expert” checking the predictions made by the model, correcting any incorrect predictions and then

feeding it back into the model as additional labeled images to a training set. We complete fine-tuning of *Zoobot* on our initial training set of 8072 galaxies and make predictions on the first 500,000 sources from the [HSC](#) (created under the criteria previously discussed). We visually inspect the sources *Zoobot* gives a prediction score of  $\geq 0.75$  to and correct any wrong predictions. These corrected labeled sources and those *Zoobot* correctly labeled are then added to the training set. Not only does this step allow us to add more labeled interacting galaxies to the training set, but it also allows us to evaluate *Zoobot*’s behavior and check if it consistently predicts a type of source or galactic morphology incorrectly.

From the first 500,000 sources, a total of 6198 sources were given a prediction score of  $\geq 0.75$ . We correct the predictions *Zoobot* made and balance this set to 5698. During this cycle, a large number of globular clusters/star fields/open clusters were given a very high prediction score. Figure 2 shows an example of these contaminating star fields. We created sources of 1250 star fields and added these into the training set, labeling them as noninteracting. Adding the balanced 5698 sources plus the 1250 star fields to our training set gave us an unbalanced training set of 15,020 sources. To then balance the training set, we took 1250 labeled interacting galaxies from the Galaxy Zoo: Hubble project and made random image augmentations with the [TensorFlow](#) Python package. These augmentations were simple rotations, cropping, and resizing. With these extra sources, our training set contains 16,270 sources. Of these, 50% (8135) were labeled images of interacting galaxy systems.



**Figure 2.** Example images of the labeled noninteracting galaxy systems used to train Zoobot. Top row: three examples from the Galaxy Zoo: Hubble project of the training set. Bottom row: three examples from the other Galaxy Zoo projects. These are, from left to right, Galaxy Zoo 2, Galaxy Zoo CANDELS, and a star field from the active learning cycle. Star fields/globular clusters/open clusters existed throughout the [HSC](#) flagged as extended sources. One thousand two hundred and fifty images of star fields were added to the training set so Zoobot would give them a very low score.

## 5. Diagnostics

### 5.1. Model Performance

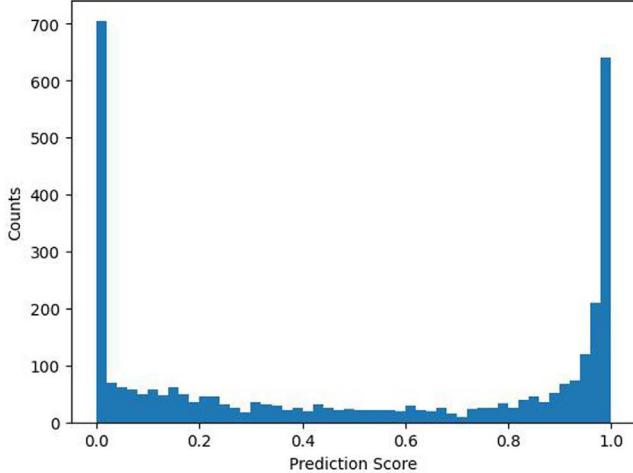
Upon fine-tuning Zoobot, we validate its performance. We reuse the validation set that Zoobot automatically creates when training. This set is created by putting aside a random set of 20% of the training set. Zoobot then uses it to validate its performance in training. We record which images Zoobot selected, and extract these from the training set for further diagnostics. This provides us with a validation set of 3270 images, containing 1648 noninteracting galaxies and 1622 interacting galaxies.

Zoobot gave a prediction score between 0 and 1 to each of the validation images, and Figure 3 shows the resulting distribution. This distribution shows that our model has high confidence in what is or is not an interacting system due to the high counts at very low and very high probability scores. It is likely the use of a balanced training set, and the very low volunteer score needed to define a source as noninteracting that leads to a strongly bimodal prediction score distribution. Using a balanced training set is an intrinsic trade-off between ease of training and potential biases introduced. Having a balanced data set does not reflect reality, and leads Zoobot to overpredict interacting galaxies. Using very stringent volunteer classification cutoffs also leaves few ambiguous systems in the validation set, further enhancing this bimodality.

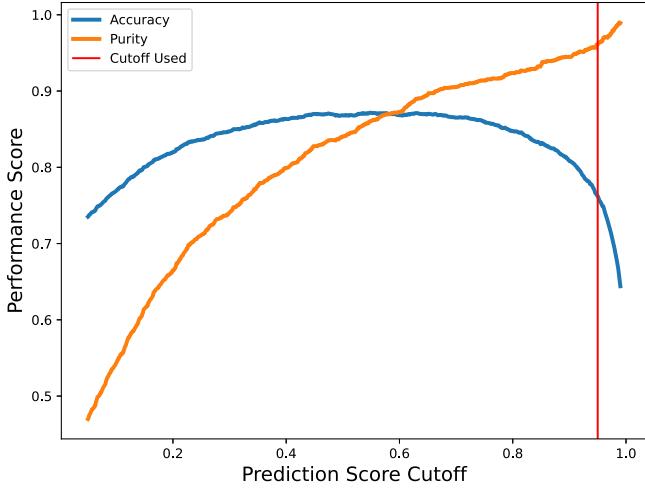
The prediction score must be reduced to a binary classification for our problem. We use Figure 3 to define a prediction

score above which a source is classified as an interacting galaxy. We measure the accuracy of Zoobot for different cutoffs, where the accuracy is the fraction of labels correctly predicted over the total number of labels predicted on. Figure 4 shows this change in accuracy. We find that our model is most accurate with a prediction score cutoff of 0.55 with an accuracy of 88.2%. Figure 4 also shows the change in the purity of our catalog with changing prediction cutoff. Here, purity is the ratio of the number of true interacting galaxies to total sources in the final catalog. These scores can be combined into the F1 score of our model, shown in Figure 15 in Appendix A.

Figure 5 also shows a measure of accuracy for our model at different cutoffs using confusion matrices. Importantly, it also shows how our model is getting labels wrong: either giving false positives (where a labeled noninteracting galaxy is predicted to be interacting) or false negatives (where a labeled interacting galaxy is predicted to be a noninteracting). The number of incorrect positive and negative predictions change based on the prediction cutoff, with a very low cutoff giving many false positives and a very high cutoff giving many false negatives. Figure 5 shows that with a cutoff of 0.50, we would return a high level contamination in our final catalog. Of the 1622 galaxies predicted to be interacting, 218 would be noninteracting systems—approximately 13%. Our main aim in this work is to present a highly pure, large interacting galaxy catalog that can be used for statistical exploration of interacting galaxy parameter space. Therefore, we use a very stringent cutoff of 0.95.



**Figure 3.** The distribution of prediction scores given to our validation set of 3270 labeled sources set aside by *Zoobot* in training. These were split into 1648 noninteracting sources and 1622 interacting sources. As can be seen from the distribution, our model is often confident when a source does or does not contain an interacting galaxy by the strong bimodality. This is likely due to the very stringent vote weightings used when selecting the training set. Using this distribution, we decide the prediction score to use as a cutoff to give us our final binary classification: interacting galaxy or not.



**Figure 4.** A measure of accuracy and purity against prediction score. The accuracy (in blue) is a direct measure of the number of sources *Zoobot* correctly predicted vs. the total number of predictions made. The measure of purity (in orange) is the number of predictions *Zoobot* correctly made vs. the total number of predictions for an interacting galaxy. The cutoff score (in red) shows the point above which we would define an interacting galaxy and below which we would not. At this point, the accuracy appears lower due to *Zoobot* making many false-negative predictions while successfully making true-negative predictions. This is confirmed by the maximization of purity. Due to the number of sources *Zoobot* is predicting over, the size of the catalog will exceed any previous catalogs. Therefore, we use this very conservative cutoff to maximize purity over the completeness of our catalog. These measures can also be shown with the F1 score. Figure 15 shows this change with prediction cutoff in Appendix A.

Using a cutoff of 0.95 reduces contamination significantly. Figure 5 shows the final contamination in our validation catalog would be  $\approx 2\%$ , where Figure 4 shows that we are maximizing the purity in our sample at the expense of accuracy. The aim of this work is not to create a general tool to be used by the community, but to find a large catalog of interacting galaxies. As we are investigating 126 million sources, despite removing  $\approx 50\%$  of interacting galaxies from

		Cutoff = 0.25		Cutoff = 0.50	
		Non-Interacting	Interacting	Non-Interacting	Interacting
True Labels	Non-Interacting	1200 72.82%	448 27.18%	1430 86.77%	218 13.23%
	Interacting	99 6.10%	1523 93.90%	199 12.27%	1423 87.73%
True Labels	Non-Interacting	1539 93.39%	109 6.61%	1617 98.12%	31 1.88%
	Interacting	356 21.95%	1266 78.05%	750 46.24%	872 53.76%
		Cutoff = 0.75		Cutoff = 0.95	
True Labels	Non-Interacting	1539 93.39%	109 6.61%	1617 98.12%	31 1.88%
	Interacting	356 21.95%	1266 78.05%	750 46.24%	872 53.76%
Zoobot Predicted Labels					
Non-Interacting					

**Figure 5.** Confusion matrices of four different cutoffs of prediction score defining a binary classification of interacting galaxy or not. Confusion matrices break down our accuracy measurement into how *Zoobot* is misclassifying sources. At a cutoff of 0.50, the accuracy is highest at 88.2%. However, at this cutoff,  $\approx 10\%$  of our final catalog would contain contamination. We elect to use the very stringent prediction cutoff of 0.95 for the rest of this work, as it will return the lowest contamination.

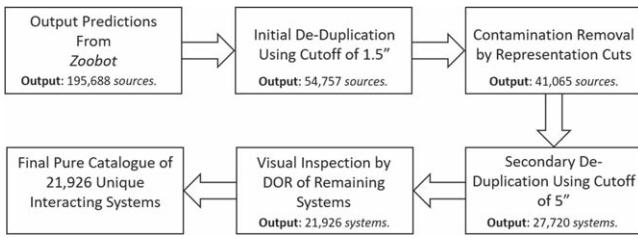
the final catalog, we are certain that we can find a catalog larger than previous works.

Using such a high cutoff also reduces any risk of any biases introduced by using a balanced training set. While using such a training set often increases the accuracy and speeds up training, it can bias the model toward one conclusion. In our case, the true rate of interacting galaxies will be much smaller than 50%. Therefore, our model will be biased to labeling a source as an interacting galaxy. This will be particularly true for edge cases, which could be ambiguous to even an expert classifier. By using such a high cutoff score, this bias will be mitigated by only labeling the most clearly interacting objects as interacting.

## 5.2. Duplication Removal

The fully trained *Zoobot* made predictions on  $\approx 126$  million extended sources from the *HSC* that had passed our creation criteria. Of these, 195,688 sources were given a score of 0.95 or greater,  $\approx 0.2\%$  of the total number of sources. Upon visually inspecting a subset of sources, it is clear that our *Zoobot* model had predicted for an interacting galaxy even if it was not the central (and, therefore, target) source in the image. This is due to the misalignment of sources from the center in the training set as described in Section 4. *Zoobot* learned to classify an image as an interacting galaxy if it contained one, and not just if it was the central source. Therefore, many interacting systems were duplicated in our final catalog, appearing in cutouts where the central source was not interacting.

Another source of further duplication was the *HSC* itself. In the *HSC*, many extended objects have multiple source IDs applied to them. This is due to bright clumps in extended sources being assigned a new ID, sources that had been found



**Figure 6.** Flow diagram of our contamination and duplication removal process. De-duplication used agglomerative clustering based on sky separation. The first step of de-duplication uses a cutoff of  $1.5''$ . This significantly reduced duplication in the catalog, as well as the size of the catalog to 54,757 interacting galaxies. We then applied contamination removal to this de-duplicated catalog. Upon visual inspection, a small number of duplicated systems still existed in the catalog. To ensure a pure catalog of unique systems, we applied an agglomerative clustering again with a cutoff of  $5''$ . This gave us a catalog of 27,720 unique interacting systems. The final step to ensure purity was visual inspection by coauthor D.O.R., removing any remaining contamination. This gave the final pure catalog of 21,926 unique interacting systems.

but did not exist in reality, or background sources that existed in extended systems. We find that of the 195,688 Source IDs given a prediction score of 0.95 or greater, approximately 3.6 Source IDs, were matched to a single real object. To refine the catalog and remove the duplication, we use spatial clustering of each source with agglomerative clustering (an introduction and description of hierarchical clustering, including agglomerative clustering, can be found in Nielsen 2016).

Agglomerative clustering is a method of hierarchical clustering based on a distance metric between the sources. We set the maximum distance between points to define a cluster, i.e., any sources within a defined distance on the sky from each other will be merged under one source ID. This approach means we do not need any knowledge of how many clusters of sources exist in the data set or the level of duplication within it, as would be the case in many other clustering approaches. We create distance matrices of the angular separation of every source using the Astropy Python package. These projected sky separations are then used as a Euclidean distance in the clustering algorithm with a EUCLIDEAN\_LINKAGE. The new ID of a cluster is the first source ID in the cluster.

Initially, we utilize a limiting sky separation of  $1.5''$  to remove the duplication. This reduced the size of our potential catalog to 54,757 interacting galaxy candidates. We then applied contamination removal as described in Section 5.3. Once contamination removal was completed, the catalog size was 41,065 interacting galaxies. Visual inspection found further duplication, indicating our initial de-duplication had not been aggressive enough. To ensure the catalog was of unique systems, we opted to use a final aggressive limiting sky separation of  $5''$  completely removing the duplication in our catalog. This aggressive de-duplication further reduced the size of our catalog to 27,720 candidate interacting systems. However, we could be certain that each of these candidate systems was unique. Figure 6 shows a full breakdown of the steps in our de-duplication and contamination removal process.

### 5.3. Bad Predictions and Removal

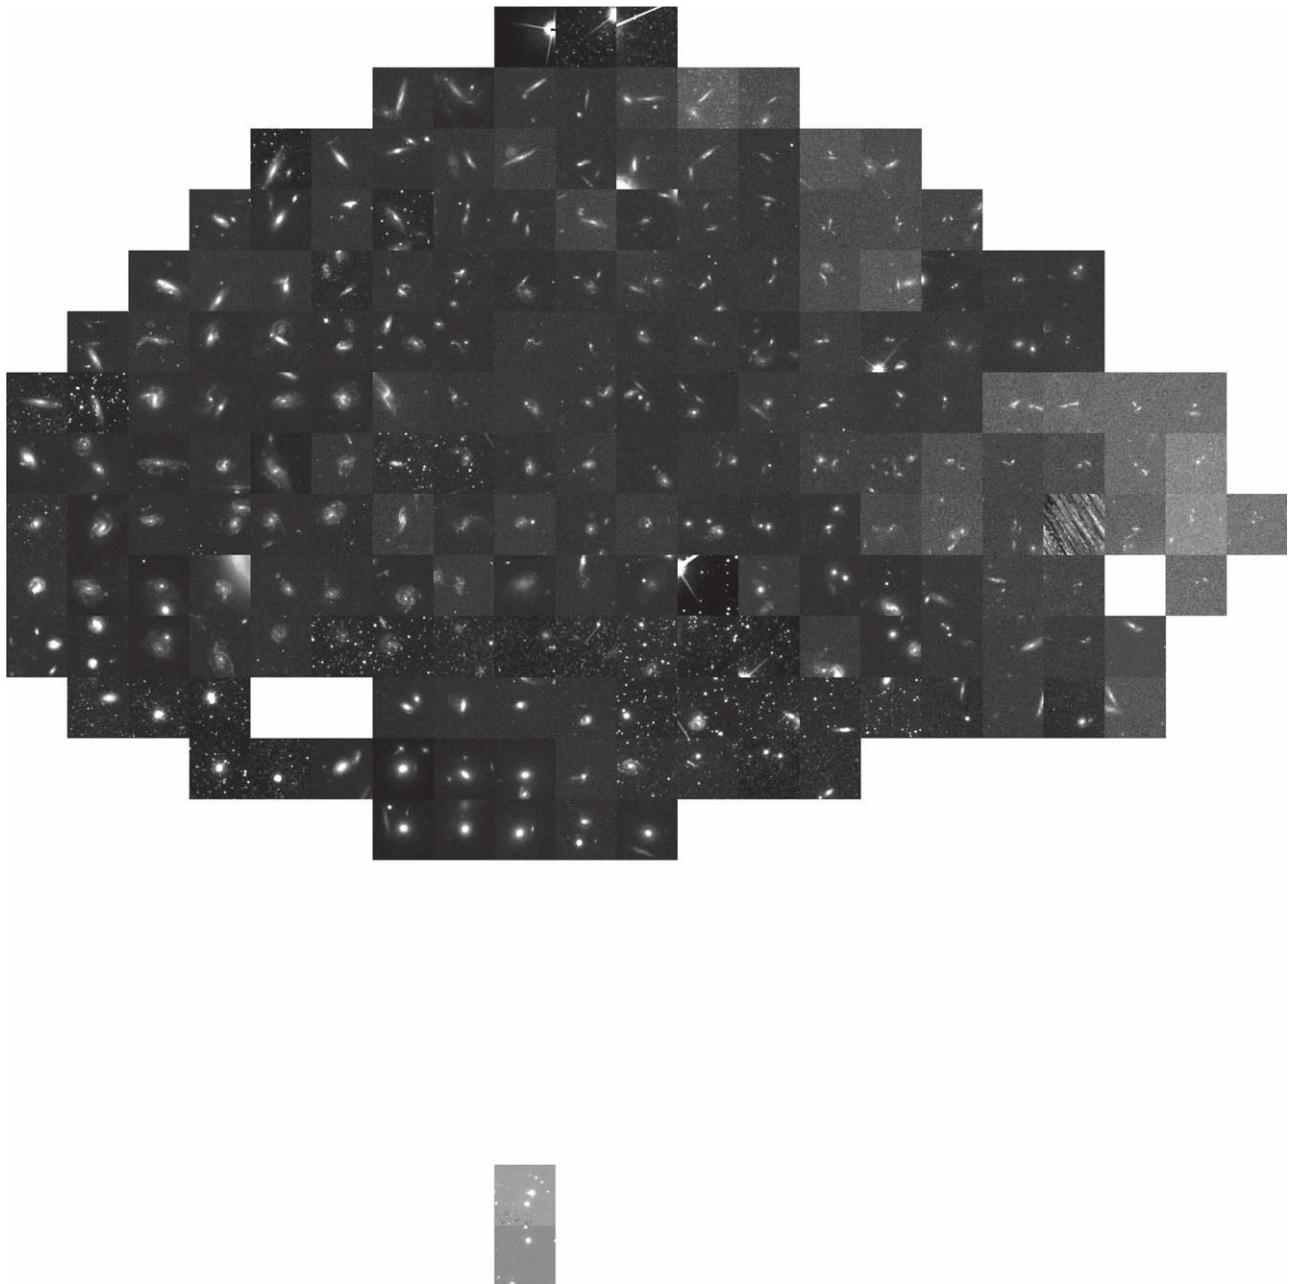
After the initial step of de-duplication, we begin removal of contamination from the catalog. A major, and expected, source of contamination is by close pairs of galaxies. These are systems where chance alignment in the sky appears that

galaxies are close together but are actually at different redshifts. Other sources of contamination include large central galaxies with satellite galaxies about them, star fields with extended sources in them, and objects with strange morphologies that Zoobot predicted were tidal features.

Upon applying the clustering by sky projection of  $1.5''$ , the catalog contained 54,757 candidate interacting galaxies. Our primary concern is contamination by close pairs. Creating catalogs of interacting galaxies with CNNs is notorious for suffering from this problem, where a significant number of candidates must be removed from otherwise large final catalogs (Bottrell et al. 2019; Pearson et al. 2022). The decisive way to remove this contamination is to compare redshift measurements of each galaxy in the candidate interacting system. However, this is impractical for our catalog where the majority of candidates have no redshift measurements. To find close pairs, and remove them effectively, we take advantage of the representations Zoobot learns of each image. As described previously, Zoobot was trained to answer every question in GZD simultaneously for every galaxy. It therefore learns a generalizable representation of many kinds of galaxies. In this representation space, morphologically similar galaxies will exist close together in clusters while those that are dissimilar will be farther apart. We extract the features Zoobot has learned of each candidate, and plot its representation.

We remove the classification head of Zoobot and directly output the final layer of the feature learning section of the model. This gives 1280 features (the representations) for each of our 27,720 candidate systems. However, there will be much redundant information in this very high-dimensional feature space. We compress this using incremental principal component analysis (Ross et al. 2008). An excellent demonstration of using this approach can be found in Walmsley et al. (2022c). We reduce the dimensionality from 1280 to 40 (as in Walmsley et al. 2022c), and input the resultant components into the auto-encoder UMAP (McInnes et al. 2018). UMAP projects the 40 dimensional components of each candidate system onto a 2D manifold. The position of each galaxy on this manifold is directly linked to its visual morphology. Close pairs have similar visual features, which will then appear as a cluster in our representation space.

Figure 7 shows the representation distribution of our 54,757 candidates after compression with UMAP. A random image in each bin has been selected to show the morphology of the objects within the bin. There are three clear gradients that exist in the representation distribution: one of source size, one of the source inclination, and one of image contrast between the source and the background. The gradient of source size is clear from left to right. This is also true of contrast between the source and background. The gradient of source inclination is from top to bottom. The top shows very inclined sources, and even the diffraction spikes of stars, while along the bottom we find face-on sources that take up a larger part of the cutout center. At the very bottom of the figure (away from the main body), a cluster of very poorly contrasted sources with the background that are face-on can be found. The gradients of inclination and source size are expected, while that of contrast is less so. This gradient is likely a result of how we created our images using a linear stretch with fixed contrast. The effect of this is that dimmer sources have brighter backgrounds, a particular issue at high redshift.

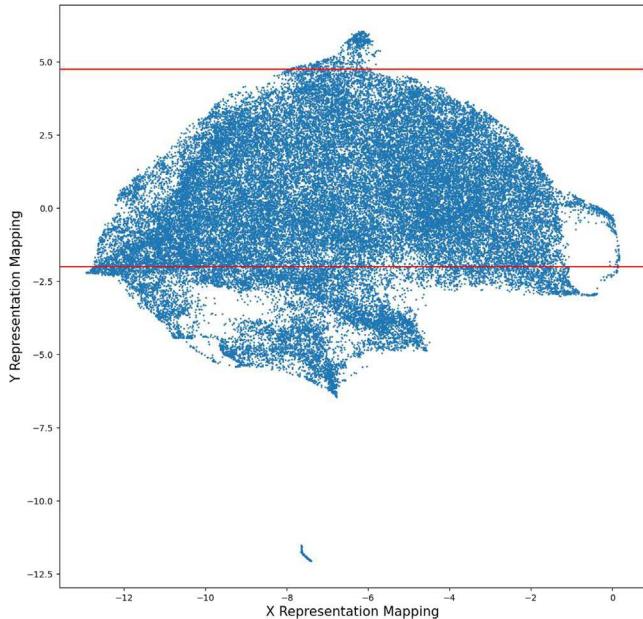


**Figure 7.** The representation distribution of 54,757 candidate interacting galaxies. This distribution is the compressed 2D representation of the 1280 dimensional representation that *Zoobot* has learned for each image. Each image is a randomly selected one from sources within each bin in the distribution. The  $X$ - and  $Y$ -axes on this plot are the 2D mapping on the manifold given by UMAP for the 40 dimensional principal components of each source, and not physical parameters. Three gradients are clear in this distribution: first; from left to right there is a distinct gradient in the contrast of the images. The images to the left are local galaxies with low redshift, while those on the right are dimmer sources at much higher redshift. This is an effect of how the images are created using a linear scaling function and a fixed contrast. The second feature, also from left to right, is a gradient of larger source size to smaller source size. This is a feature *Zoobot* has learned based on the redshift of the source as well. The third, from top to bottom, is a gradient of the inclination of the source. With the most inclined (and even diffraction spikes) of the sources appearing at the top, while at the bottom the sources are face-on. Along the bottom of the representation plot, there are close paired sources as well as many star fields. Along the very top, there is contamination in the form of isolated stars in star fields. Thus, we make aggressive cuts along the top and bottom of our representation space to remove as much contamination in a general way. The full representation plot, with all sources and the cuts, is shown in Figure 8.

Figure 7 has many areas of similar morphology. On the left, we have isolated objects: disturbed spirals or large galaxies with tidal disturbance to them. Along the bottom, we see isolated bright objects with satellites about them. On the bottom right, we see our area of representation space dominated by close pairs. In the center, we see the population of interacting galaxies that *Zoobot* was trained to find. The areas of representation space that are dominated by clear sources of contamination are cut. Figure 8 shows a scatter plot

of the representation distribution and the cuts we make. They are made such that any source with a  $Y$ -mapping of  $-2 \leq Y \leq 4.75$  will be kept in the catalog. The choice of these cuts has been made by eye, and then bootstrapping the remaining images to check contamination removed. After applying these cuts, we retain 41,065 systems in our catalog.

We estimate  $\approx 25\%$  of sources in the greater-than-0.95-prediction bin are close pairs. This may seem lower than previous works, but is due to our very conservative prediction



**Figure 8.** Scatter plot showing the precise distribution of each representation of sources in the remaining 54,757 sources. This is the unbinned version of Figure 7. The two red lines show the cutoffs utilized to remove the majority of close pairs by projection as well as the very obvious contamination of stars and stellar fields at the top of the representation distribution. The number of candidate interacting systems in the catalog was reduced to 41,065 systems.

cutoff. The general cuts to our population based on their position in representation space make it very likely that we retain some close pairs in the catalog, while also removing interacting galaxy systems.

As described in Section 5.2, we then apply a 5'' to the 41,065 remaining candidates, further reducing our catalog to 27,720 systems. With such an aggressive sky projection cut, many individual interacting galaxies are now identified under the same ID as the secondary galaxy in the system. To remove remaining contamination in the catalog, a final visual classification step was conducted. This visual inspection was conducted by coauthor D.O.R. Any systems removed at this stage were classified into three categories: interacting system, contamination, and gems. The gems subcategory became necessary as many sources of contamination that were being removed were objects of other astrophysical interest, and is described in Section 6.2.

## 6. Results and Discussion

### 6.1. An Interacting Galaxy Catalog

Upon de-duplication and contamination removal described in Sections 5.2 and 5.3, our final catalog contains 21,926 interacting systems. Figure 9 shows a random sample of 50 of the systems from our catalog. In these examples we can see highly distorted or currently interacting systems, precisely what we trained Zoobot to highly predict. Some cutouts are of the full interacting system, containing both the primary and secondary galaxies in the interaction. Some source cutouts only show one of the interacting galaxies, though these systems remain highly disturbed. Due to the constraints in our training set, highly weighting disturbance or tidal features in our predictions, we are sampling interaction from all epochs except the approach to the initial pass. At this initial stage, there will be no tidal features formed or disturbance in the disks as the

two galaxies approach each other. Separating them from close pairs would be difficult without kinematic or redshift information, not available for the majority of these sources.

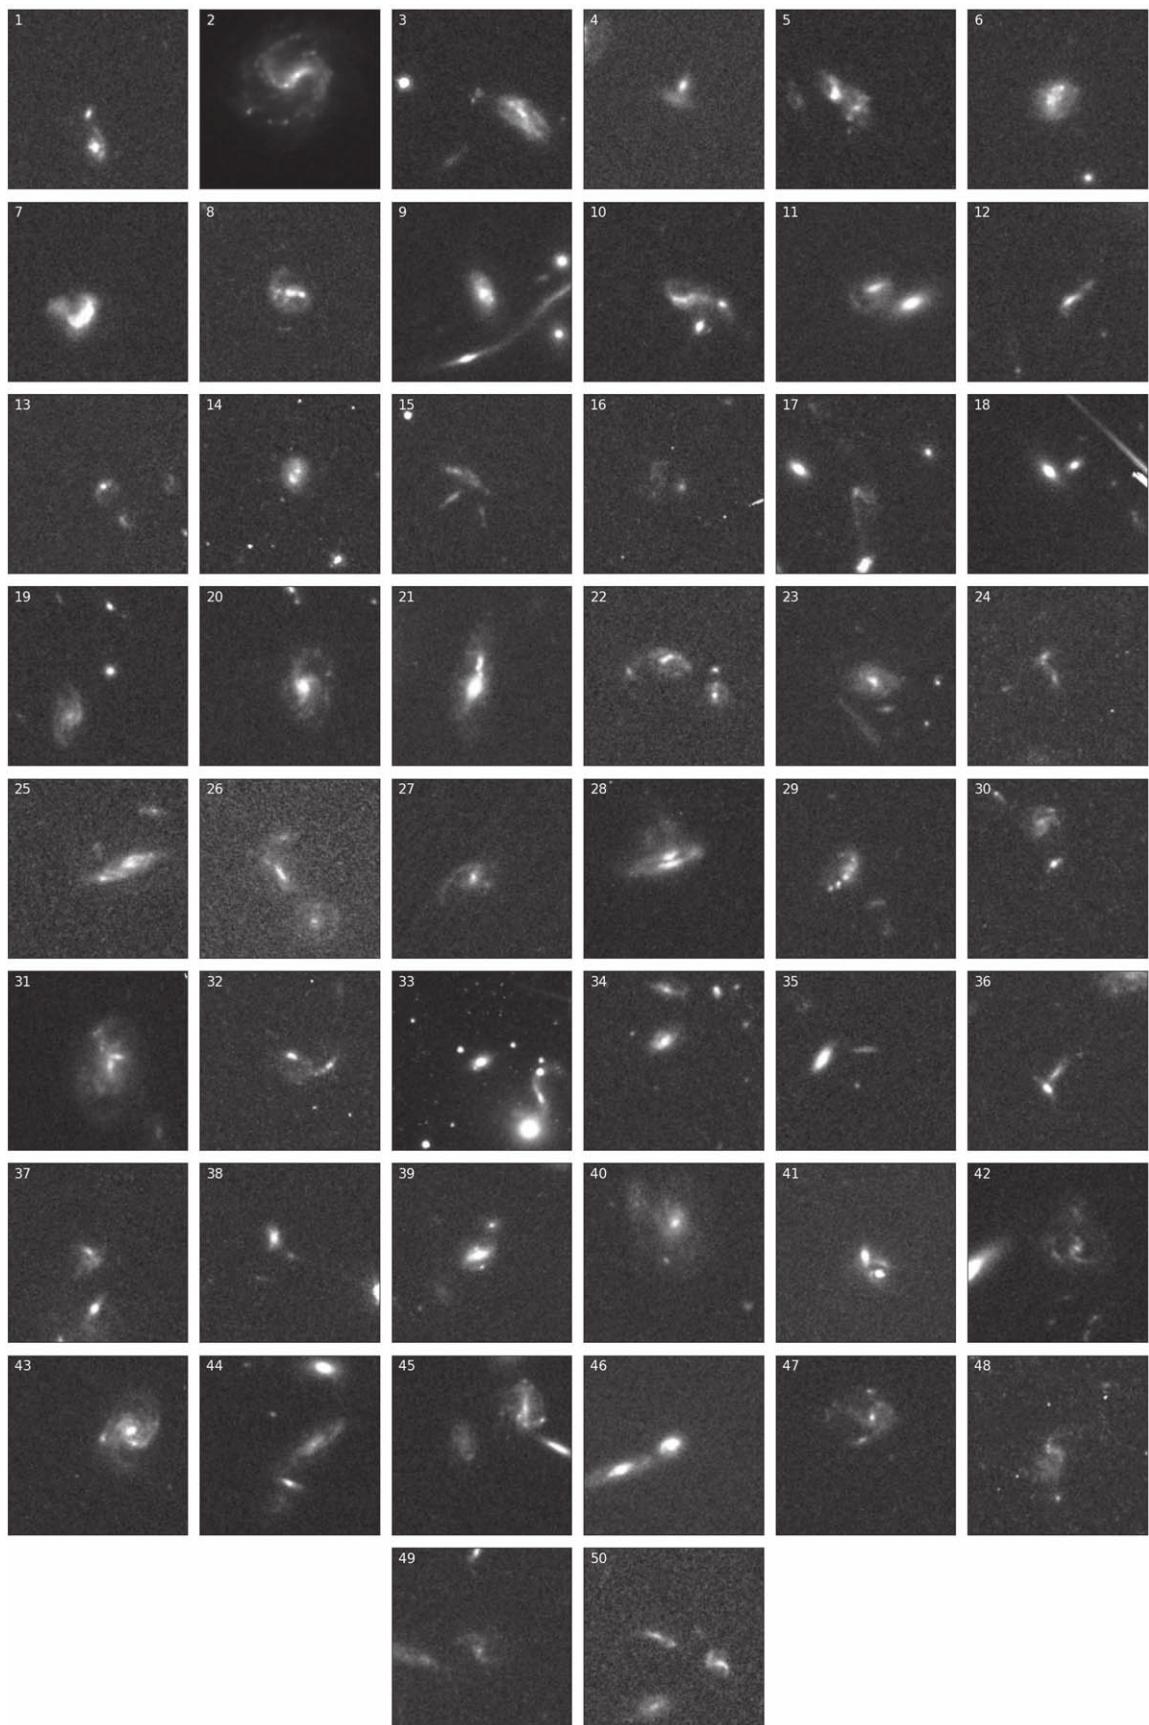
We investigate which of the systems in our catalog have previous references in astrophysical literature. To search the literature, we use the AstroQuery Python package with a coordinates-based search of cutoff radius 5''. We search the astronomical databases Simbad (Wenger et al. 2000), the NASA Extragalactic Database (NED; NASA/IPAC Extragalactic Database 2019), and VizieR (Ochsenbein et al. 2000) for references to our interacting systems. These return either a list of references, or an empty list showing no references associated with the system. We find that 7522 of our systems have at least one reference associated with them, while 14,404 do not. A flag exists in the catalog data release, which shows whether a system has references associated with it or it could be considered a “new” system. We, however, do not claim that we discovered these systems. These systems have always existed in the backgrounds of large surveys or observations and been discovered by others; it is only with ESA Datalabs that we can apply a methodology such as in this work to extract those systems from these observations. We also do not claim that these unreference systems are particularly interesting or phenomenal. It is most likely that these systems are the very faint background galaxies in surveys or observations whose main objective was something other than finding interacting galaxies. This will be further discussed in Section 6.3.

Figure 10 shows the distribution of our catalog in the sky. HST is able to observe the majority of the sky, so the catalog sources are scattered throughout it. We find that the sources cluster in different parts of the sky that correspond to major surveys conducted using the HST involving ACS/WFC and the F814W filter. We also mark the centers of the seven surveys that correspond to the major clustering of interacting systems in the sky. These were the COSMOS, the GOODS North, GOODS South, PHAT, CANDELS, AEGIS, and Spitzer Space Telescope FLSV Region (Morganti et al. 2004) surveys.

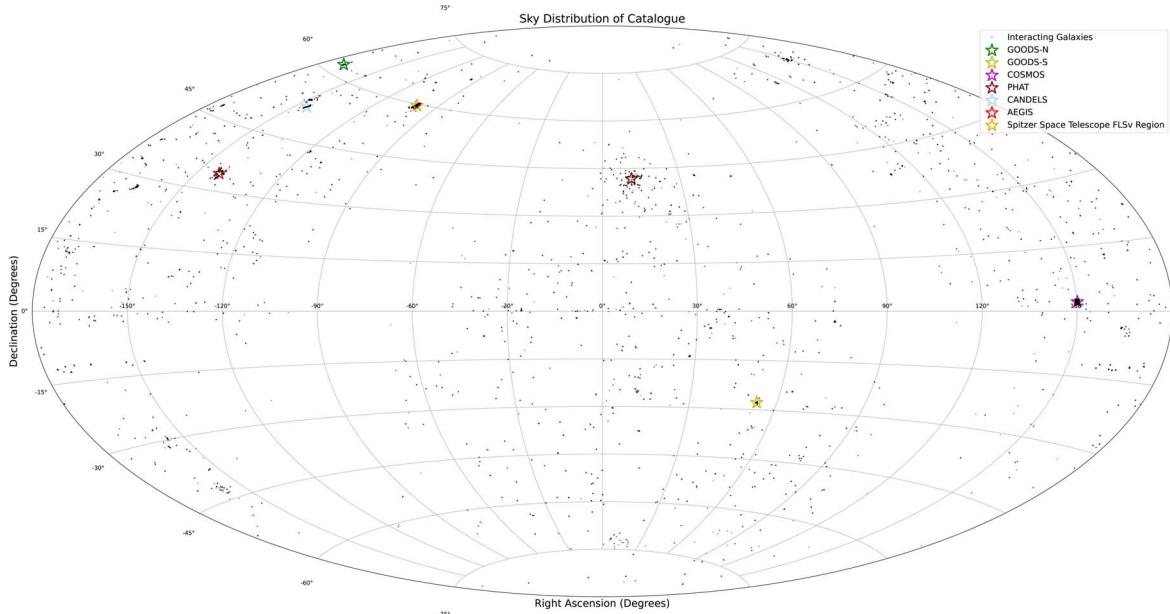
The full catalog and data product are found on Zenodo at the following DOI where it is freely accessible to the community: doi:10.5281/zenodo.7684876. Table 1 shows an example of the data and format of the 50 sources shown in Figure 9. We also bootstrap the final catalog as an estimate of contamination remaining. As described in Section 5.3, the final step of contamination removal was visual inspection by coauthor D.O.R. of the 27,720 candidate interacting systems to remove the remaining 5794 contaminants from the final catalog. Visual inspection by a single expert at this scale is not perfect. We extract random sources from the catalog in batches of 500 and manually reclassify them again. This bootstrapping reveals that  $\approx 3\%$  of our interacting system in the final catalog remains contaminated.

### 6.2. The Gems

By conducting a visual inspection of the 27,720 candidate systems, we were able to directly identify many other objects of astrophysical interest. As Zoobot was trained to highly predict objects with irregular morphologies, we also find many other astrophysical objects with strange morphologies that may be of interest to the community. We call these sources of contamination “gems.” We make 16 subcategories of these: AGN/quasars, submillimeter galaxies, galaxy groups, high-redshift galaxies, jellyfish galaxies, galaxy jets, gravitational



**Figure 9.** An example of 50 of the final interacting systems found with Zoobot. These were selected randomly from the de-duplicated and de-contaminated 21,926 sources. Each of these examples have extended tidal features and distortion. Not all of the final interacting systems have two galaxies within them (for example, image 2), but are clearly very disturbed by a tidal event. These were kept in as they would form a large part of the interacting galaxy population and would be flagged as disturbed or interacting in Galaxy Zoo. Each of these images is a one-color image using the F814W HST filter.



**Figure 10.** Sky distribution of our catalog, with marked positions of well-known deep surveys conducted by HST. HST is able to observe almost the entire sky, and therefore the interacting galaxies are scattered throughout. Large clusters of sources are found in the locations of surveys. This shows that often our sources are in the background of larger surveys and observations.

lenses/lensing galaxies, Ly $\alpha$  emitters, overlapping galaxies, edge-on protoplanetary disks, radio halos, ringed galaxies, supernova remnants, transitional young stellar objects, young stellar clusters, and unknown objects.

Each subcategory has been defined by checking Simbad and VizieR for references within a 5'' radius of each source and using astrophysical literature for a definition of the source. Coauthor D.O.R. classified any unreferenced objects by morphological similarity to other defined objects. The platforms ESASky<sup>12</sup> (Merín et al. 2017.), NED, and the SDSS were also used to investigate any unreferenced objects. ESASky was of paramount importance as we could investigate many objects across a range of wavelengths with many instruments.

The only objects that were classified by other means than visual morphology were AGN/quasars, submillimeter galaxies, and the six unknown objects. We attempt to confirm the unreferenced AGN/quasars as candidates by investigating the source in Chandra or XMM-Newton for hard or soft X-Ray emission. The submillimeter candidates were also investigated using Herschel or Planck measurements. If there was a positive signal in their positions, they were classified as such. Further work will be needed to confirm these classifications.

The final category that required further inspection was that of the unknown objects. These are objects that have unusual morphology, which marks them out from the rest of the sample; but there are no references associated with them in Simbad or VizieR. They also did not appear in NED, meaning they could not be confirmed to be galaxies. These objects are shown in Appendix C.

Table 2 shows a breakdown of the total number of objects found and the number that were referenced or unreferenced. We have released catalogs of each subcategory in the same format as that of the main catalog without the interaction

prediction column. Each of these catalogs can also be found at the same Zenodo link.

### 6.3. Source Redshifts and Photometry

We investigate the redshift distribution and photometric properties of sources in our catalog. We extract all sources with preexisting data, querying Simbad, VizieR, the HSC via the Milkulski Archive for Space Telescopes (MAST), and NED. Our queries use a 5'' search radius within the Python package AstroQuery. The existing data from each of these databases has undergone heterogeneous selection and analysis procedures by the various studies we extract them from; we do not try to reconcile these here. Rather than a detailed physical analysis of these sources, our priority in this subsection is to highlight how to explore and use this catalog, as well as any difficulties that may arise.

Of the 21,926 interacting systems in our high-confidence sample, 3037 of the 7522 referenced sources have a measured redshift. Figure 11 shows the redshift distribution of this subset of our catalog. We found that 42.5% of the sources have a redshift  $z \leq 0.5$ , 45.1% have a redshift  $0.5 < z < 1$ , and 12.4% have a redshift  $z > 1$ . In fact, a small fraction (15) of these sources are found to be at  $z \geq 5$ . Upon investigation of these sources, two of their redshifts have been measured photometrically, while the remaining 13 sources did not have the method of measurement recorded in the archive. Therefore, this finding of very-high-redshift interacting galaxies is uncertain at best.

It is important to note that the small sample with redshift information is affected by the selection biases of the combined studies publishing these values, and therefore the distribution may not be representative of the full sample. In addition, above redshift  $z = 1$ , the F814W filter begins to only capture rest-frame UV flux, and therefore  $z > 1$  galaxies with low star formation rates are more likely to fall below the flux limits of our detection images. Sampling only the rest-frame UV also changes a galaxy’s observed brightness and morphology

<sup>12</sup> ESASky: <https://sky.esa.int/>

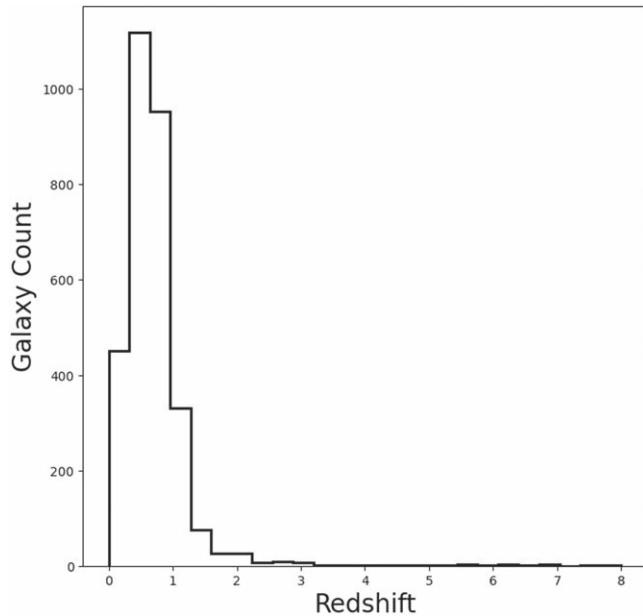
**Table 1**  
An Example of the Format of the Final Catalog for the 50 Example Images Presented in This Paper

Image No. (1)	SourceID (2)	R.A. (deg) (3)	Decl. (deg) (4)	Interaction Prediction (5)	References (6)	Status (7)
1	4001014298177	261.292845	37.162387	0.983999	No entry	Unreferenced
2	4001444190958	183.527536	33.183451	0.998016	[1994PASP..106..646K]	Referenced
3	4000809226818	93.960150	-57.813401	0.982266	[2019ApJ...878...66C]	Referenced
4	4553390202	73.581297	2.903528	0.968280	No entry	Unreferenced
5	4000907600174	259.037474	59.657617	0.999978	No entry	Unreferenced
6	4575187799	150.001883	2.731942	0.974649	[2007ApJS..172...99C]	Referenced
7	4000717342023	149.527791	2.126945	0.993912	[2007ApJS..172...99C]	Referenced
8	4001174802281	28.593114	-59.643515	0.982890	No entry	Unreferenced
9	4182689774	186.709991	21.835419	0.973232	[2016ApJS..224...1R, 2011ApJS..193...8B]	Referenced
10	4000958398690	186.719496	23.961225	0.999288	No entry	Unreferenced
11	4266881925	344.730228	-34.799824	1.000000	No entry	Unreferenced
12	4001084105393	150.128198	2.623949	0.982739	[2018ApJ...858...77H, 2007ApJS..172...99C]	Referenced
13	4000961670486	345.337556	-38.985521	0.954961	No entry	Unreferenced
14	4000719687395	338.173538	31.189718	0.974724	No entry	Unreferenced
15	4001435343326	331.771500	-27.826175	0.986885	No entry	Unreferenced
16	4001268932937	8.856781	-20.271978	0.986329	No entry	Unreferenced
17	4651336656	149.836709	2.141702	0.984389	[2007ApJS..172...99C]	Referenced
18	4000877021787	116.211231	39.462563	0.979178	No entry	Unreferenced
19	4000878525229	149.834893	2.516816	0.963694	[2007ApJS..172...99C, 2009ApJS..184..218L]	Referenced
20	6000290755870	186.774907	23.866311	0.981961	No entry	Unreferenced
21	4000806637434	210.253419	2.854869	0.960790	No entry	Unreferenced
22	4001215753971	135.898809	50.487130	0.998386	No entry	Unreferenced
23	4000813961830	163.678042	-12.776815	0.958405	[2005ApJ...630..206F]	Referenced
24	4001200639012	54.037618	-45.170026	0.991404	No entry	Unreferenced
25	4000921402261	150.417634	2.313781	0.990775	[2018ApJ...858...77H, 2012ApJ...753..121K]	Referenced
26	4001224732336	337.217339	-58.444885	0.955972	No entry	Unreferenced
27	4000781402752	216.968619	34.575819	0.974076	No entry	Unreferenced
28	4001283017901	120.202582	36.058927	0.994169	[2016ApJS..224...1R]	Referenced
29	4000833486119	116.260049	39.457642	0.971092	No entry	Unreferenced
30	4000949659908	146.342493	68.730869	0.961113	No entry	Unreferenced
31	4000982902478	53.084832	-27.765379	0.983472	[2010A&A...512A..12B]	Referenced
32	4001189505548	192.492491	2.436292	0.992574	No entry	Unreferenced
33	4001060882070	89.700725	-73.049783	0.962839	No entry	Unreferenced
34	4000889750512	151.176470	41.214096	0.962205	No entry	Unreferenced
35	6000322363510	53.149367	-27.823945	0.963889	[2016ApJ...830..51S]	Referenced
36	4000722901091	28.257843	-13.928090	0.982778	No entry	Unreferenced
37	6000198293960	264.488431	60.101798	0.986865	No entry	Unreferenced
38	4001095660911	258.587670	59.970358	0.955193	No entry	Unreferenced
39	4000972775076	330.960020	18.796346	0.989131	No entry	Unreferenced
40	4001132466571	126.545810	26.456196	0.997077	No entry	Unreferenced
41	4000933395648	312.810365	2.288410	0.976252	No entry	Unreferenced
42	4000932940918	218.066960	32.997228	0.990737	No entry	Unreferenced
43	4001048433104	93.880689	-57.754746	0.957755	No entry	Unreferenced
44	4001039919651	53.111470	-27.673717	0.994424	[2011ApJ...743..146C]	Referenced
45	4001282607544	333.765783	-14.006097	0.999520	No entry	Unreferenced
46	4000922341052	260.723839	58.849293	0.995477	No entry	Unreferenced
47	4000731518210	194.869144	14.146223	0.994651	No entry	Unreferenced
48	4001082523786	311.703084	-12.869002	0.976454	No entry	Unreferenced
49	4000767041112	149.784518	2.172233	0.991335	[2007ApJS..172...99C]	Referenced
50	4001024667142	150.661685	1.718587	0.967865	[2007ApJS..172...99C]	Referenced

(e.g., Ferreira et al. 2022)—the latter being how Zoobot identifies interacting galaxies. For example, tidal features whose initial starburst has faded may be undetected; conversely, a single galaxy with irregular star-forming clumps may appear to be multiple interacting galaxies, which we noted as a particular source of contamination during the visual inspection stage. High-redshift interacting galaxies that are detected initially by Zoobot but have unusual morphologies compared to  $z \sim 1$  sources may be removed during prediction (Section 3), given that fine-tuning is based primarily on the  $z \lesssim 1$  imagery

of Galaxy Zoo: Hubble. Therefore, the currently measured redshift distribution in Figure 11 is likely due to some combination of selection bias and training bias.

Figure 12 shows the basic parameter space sampled by the subsample of the catalog with existing photometry and redshifts. We show the distributions of redshift with the measured apparent F814W magnitude and the calculated absolute F814W magnitude. The faintest objects are, as expected, observed at approximately the limiting magnitude of the deepest observations in our catalog. Other observations



**Figure 11.** The redshift distribution of a subsample of our catalog. Of the 7583 referenced systems, 3037 of them had redshift measurements in the NED, MAST, or Simbad. This redshift distribution shows that our model confidently predicted interacting systems primarily for  $z < 1$  systems. This was anticipated, as the model was primarily trained on systems at these redshifts. There are 15 sources with a reported  $z > 5$ .

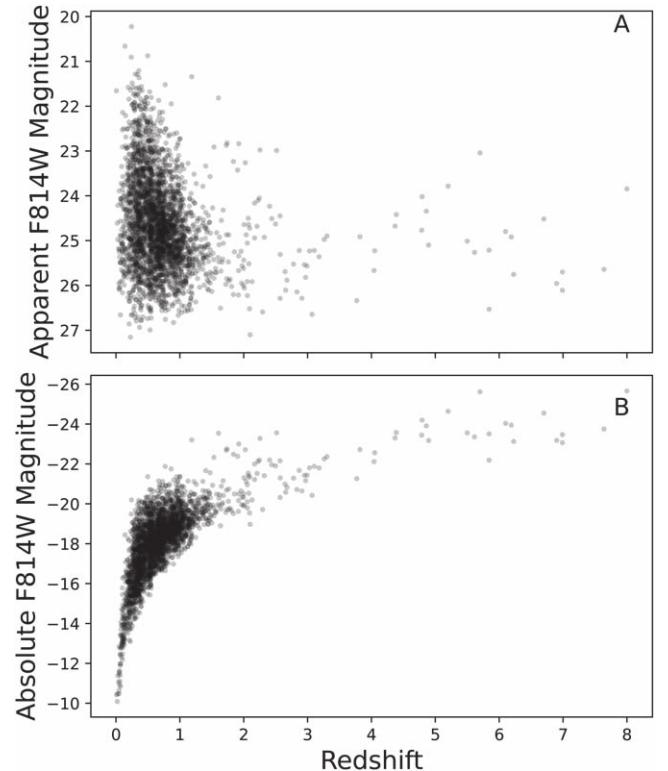
**Table 2**  
A Breakdown of Gems Found in the Visual Inspection Stage of Contamination

Category (1)	Total Found (2)	Referenced (3)	Unreferenced (4)
AGN/Quasars	35	21	14
Submillimeter Galaxies	11	8	3
Galaxy Groups	6	6	0
High-redshift Galaxies	10	7	3
Jellyfish Galaxies	18	5	13
Galaxy Jets	25	10	15
Gravitational Lenses/Lensing Galaxies	189	64	125
$\text{Ly}\alpha$ Emitters	1	1	0
Overlapping Galaxies	221	92	129
Edge-on Protoplanetary Disks	9	2	7
Radio Halos	1	1	0
Ringed Galaxies	6	1	5
Supernova Remnants	4	3	1
Transitional Young Stellar Objects	2	1	1
Unknown Objects	6	0	6
Young Stellar Clusters	2	1	1

**Note.** Each gem category has been classified based on the references associated with each object.

have brighter limits; those wishing to select a uniform or volume-limited sample from our catalog must consider the variable flux limits across the sample.

We finally investigate sources from our high-confidence sample that have multiband photometry, focusing on commonly observed filters. By construction, 100% of the sample has F814W measurements, with 45% of the catalog having F606W and only 11% having measured fluxes in F475W.



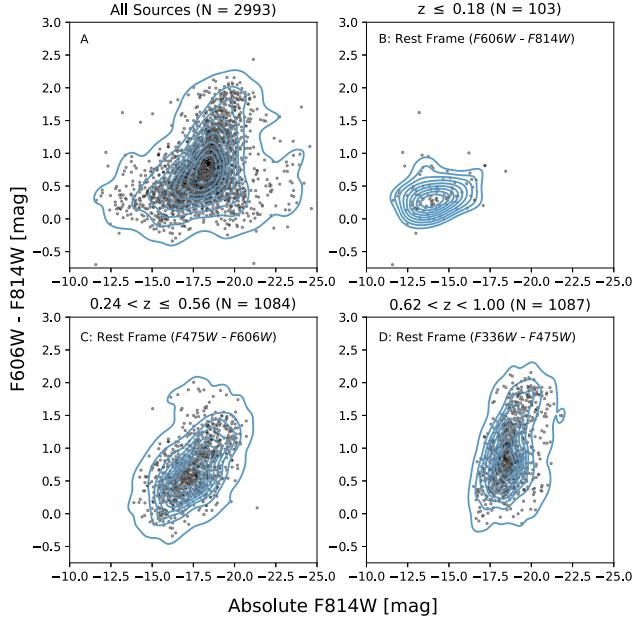
**Figure 12.** The distribution of redshift with magnitude for all sources with available data. This shows the parameter space we are sampling in this catalog. Panel (A) shows that the majority of our sources are dim, background sources at low redshift. Panel (B) shows the faintest objects we find are at the limiting magnitudes of the different surveys this data is from.

**Table 3**  
Percent of Sources in the Final Catalog That Have Observations in the Relevant Hubble Filter

Filter (s) (1)	Sources Covered (2)
F814W	100%
F606W + F814W	45.0%
F475W + F814W	11.0%
F475W + F606W + F814W	6.1%

Table 3 summarizes the filter coverage of our catalog. We found that 6.1% (1336 sources) have complete three-band photometric information in the HSC. We use these to create examples of color images from the catalog (using the algorithm of Lupton et al. 2004). We used a scaling factor  $Q = 2$  and  $\alpha = 0.75$ , with (F814W, F606W, F475W) as RGB channels and multiplicative factors of (1.25, 0.95, 2). The resultant images are shown in Appendix B.

We extract the measured magnitudes of the F606W and F814W filters, giving us two-band photometry for 9876 sources. Cross referencing with each source that had a redshift yields 2993 sources from our catalog. We calculate the color of each source and plot it against the absolute magnitude in the F814W filter. Figure 13 shows the resulting color–magnitude distribution in panel (A). The resultant distribution is very hard to interpret due to the high scatter of the sources. We extrapolate from this panel that there is little contamination from sources other than galaxies. If levels of contamination



**Figure 13.** The color–magnitude distribution of sources with a redshift measurement associated. Panel (A) shows the distribution of all galaxies, without controlling for redshift or dust extinction. The remaining panels then split these sources into distinct redshift bins where the F606W and F814W filters are observing in different rest frames. Panel (B) shows the color–magnitude distribution in the local universe, where the rest-frame observations are F606W and F814W flux. This bin reveals a blue population. Panel (C) shows the redshift bin where at 50%–100% of observed F606W and F814W flux is rest-frame F475W and F606W flux. This bin reveals a larger distribution of interacting galaxies, with a dominating population of blue systems and a minor population of red systems. Panel (D) shows the redshift bin where 50%–100% of observed F606W and F814W flux is rest-frame F336W and F475W flux. These filter bands are very sensitive to star formation, and reveal a broad distribution in color of red and blue systems.

were high, we would expect a second locus of sources with a very different color–magnitude distribution.

Plotting the color–magnitude distribution in this way captures a wide range of rest-frame wavelengths in the observed filters, which is the primary reason that panel (A) of Figure 13 is hard to interpret. In this first-look study, we do not have full spectral energy distributions (SEDs) of most sources, so  $K$ -correction of individual colors within this sample would involve assuming a template SED for each galaxy. Given that a high fraction of galaxies in our sample of mergers may deviate from standard SED templates, we wish to avoid this method. Instead, we choose redshift ranges within which to examine subsamples, such that the observed F606W and F814W bands cover consistent rest-frame colors within that subsample. Figure 13(B) shows only sources with  $z < 0.18$ , within which the observed filters can be taken to be approximately rest-frame filters, which we define as at least 50% of the flux captured in the observed band being emitted at rest-frame wavelengths covered by that band. At  $0.24 < z < 0.56$ , the observed F606W filter captures at least 50% rest-frame F475W flux, and the observed F814W filter captures at least 50% rest-frame F606W flux, so Figure 13(C) is approximately a rest-frame F475W – F606W versus F606W plot. At  $0.62 < z < 1$ , Figure 13(D) is approximately a rest-frame near-UV (NUV)–blue plot (F336W–F475W versus F475W).

The galaxies in panel (B) are observed in approximately the rest-frame F606W and F814W filters. Nearly all are blue systems (by general definitions at various redshifts; e.g.,

Kauffmann et al. 2003; Whitaker et al. 2012; Schawinski et al. 2014). This is expected for interacting systems with enough gas to fuel a starburst. The lack of many red systems is due to few gas-poor (“dry”) interactions in the (relatively) local volume (López-Sanjuan et al. 2009). In Figure 13(C), the F606W and F814W filters are still detecting rest-frame optical (F475W and F606W) emission, and we find a much broader population. There are both blue and red interacting systems, with the redder mergers occurring in more-luminous (likely higher mass) systems, broadly consistent with expectations (van Dokkum 2005; Lotz et al. 2008). The rest-frame filters approximately captured in panel (D) (F336W and F475W) sample emission across the 4000 Å break. Sensitivity to NUV means this panel effectively splits systems according to very recent star formation history (Schawinski et al. 2014; Smethurst et al. 2015). There is a significant spread in color, with equivalent red and blue systems. We, therefore, find many young blue systems undergoing star formation and brighter, elliptical, massive systems also undergoing interaction in this bin.

This initial examination of the subsample of systems with easily retrievable redshifts has revealed that the interacting galaxies in the sample broadly agree with previous studies of colors in merging systems. This demonstrates the underlying promise of the catalog. A detailed study is beyond the scope of this work, but there is considerable potential for new astrophysical insights using this high-confidence catalog with nearly an order-of-magnitude more sources than those previously published.

## 7. Conclusion

We present a large, pure catalog of 21,926 interacting galaxy systems found from the HSC. This catalog is a factor of 6 larger than previous works. Each interacting system was found using the European Space Agency’s new platform ESA Datalabs, which allowed us to directly apply an advanced CNN—Zoobot—to the entire Hubble science archive. This corresponds to predicting over 126 million sources. The compiled catalog has a contamination rate of  $\approx 3\%$  as found by bootstrapping. Table 1 shows an example of 50 entries in our new catalog, Figure 9 shows the corresponding images. The new catalog and all corresponding images can be downloaded from Zenodo: doi:10.5281/zenodo.7684876.

Each of our interacting galaxies was given a prediction score  $\geq 0.95$  by Zoobot, with such a conservative score chosen to limit contamination and maintain purity in the catalog. Contamination was removed by applying cuts in representation space (shown by Figure 8) and visual inspection. Upon visual inspection, many contaminating images were found to be objects of other astrophysical interest. These have been compiled into separate catalogs, and Table 2 shows a breakdown of the objects found. These subcatalogs have been released alongside our interacting galaxy catalog. With the priority of purity in this catalog creation, we will aim in future work to use it in the statistical analysis of interacting galaxies and begin linking the underlying parameters of interaction to the complex physical processes that occur in them. A secondary purpose of this catalog is to serve as a training set for future models that may wish to search for interacting or merging galaxies.

With the use of ESA Datalabs, this project was conducted quickly. The entire process, from creating the source cutouts, to

training *Zoobot*, to making predictions on 126 million sources, took three months to complete. Using conventional methods, such as *AstroQuery* or TAP services, downloading the data would have likely taken on this timescale. By bringing the user to the data, rather than vice versa, catalogs of a similar size—and many times larger than previous catalogs—of many different objects can be created quickly.

None of the interacting systems in this work are “new”; every one of them exists in the background of large-scale HST surveys and observations since their release. However, the method to directly search for them has been impractical until the release *ESA Datalabs*. By directly applying machine learning to existing astrophysical data repositories, a new method of creating significantly larger catalogs has been achieved.

This shows the importance of archival work, and the power that *ESA Datalabs* will bring to the field of astronomy. *ESA Datalabs* is expected to be released in Q3 and with it, the ability for large-scale exploration of archival data. It will be released with introductory tutorials, step-by-step guides, and different Python environments for ease of use for different telescopes and instruments the *ESA* is involved in. It will have a full cluster of graphical processing units (GPUs) at its disposal and a storage capability in the range of hundreds of terabytes. In future, this entire project—from training set creation to predictions—could be conducted on *ESA Datalabs*.

Such a setup as *ESA Datalabs* also allows for the creation of large observational catalogs, comparable to that we create from cosmological simulations. This is incredibly important to further constraining already existing results. In the current period of astronomy where large survey instruments are awaiting first light, or the beginning of future telescopes is uncertain, the ability to get ever more information out of the archives is paramount.

D.O.R. gratefully acknowledges the support from European Space Agencies Visitor Archival Research program, and hosting at the European Space Astronomy Centre. D.O.R. thanks Bruno Merín for supervising this project and Sarah Kendrew for aiding its creation. This project was conducted as part of D.O.R.’s PhD program supported by the UK Science and Technology Facilities Council (STFC) under grant reference ST/T506205/1. B.D.S. acknowledges support through a UK Research and Innovation Future Leaders Fellowship (grant No. MR/T044136/1). I.L.G. acknowledges support from an STFC PhD studentship (grant No. ST/T506205/1) and from the Faculty of Science and Technology at Lancaster University. M.W. gratefully acknowledges support from the UK Alan Turing Institute under grant reference EP/V030302/1. M.R.T. acknowledges the support from an STFC PhD studentship (grant No. ST/V506795/1) and from the Faculty of Science and Technology at Lancaster University.

Much of the intense computation was conducted at the High End Computing facility at Lancaster University. This publication uses data generated via the Zooniverse.org platform, and the unending enthusiasm of citizen scientists and volunteers in classifying galaxies. We also thank the many PIs whose archival data we have used to create this catalog. All data containing astrophysical objects of interest found in this work are public on MAST: [10.17909/wfke-n133](https://doi.org/10.17909/wfke-n133).

This research made use of many open-source Python packages and scientific computing systems. These included *Matplotlib* Hunter (2007), *scikit-learn* (Pedregosa

et al. 2012), *scikit-image* (van der Walt et al. 2014), *Pandas* (McKinney 2010), *Shapely* (Gillies et al. 2022), *UMAP* (McInnes et al. 2018), and *numpy* (Harris et al. 2020). This work also extensively used the community-driven Python package *Astropy* (Astropy Collaboration et al. 2018). *Zoobot* utilizes the underlying code *Tensorflow* (Abadi et al. 2016) Python package.

This project used data from the Hubble Space Telescope and stored in the archives at the European Space Astronomy Centre. These observations are obtained from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc, under NASA contract NAS 5-26555. All sources were found using v3.1 of the Hubble Source Catalog (Whitmore et al. 2016) and accessed using the *ESA Datalabs* science platform. *ESA Datalabs* is directly connected to the *ESA Hubble Science Archive*. This study makes use of data from *AEGIS*, a multiwavelength sky survey conducted with the *Chandra*, *GALEX*, *Hubble*, *Keck*, *CFHT*, *MMT*, *Subaru*, *Palomar*, *Spitzer*, *Very Large Array*, and other telescopes and supported in part by the *NSF*, *NASA*, and the *STFC*.

For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

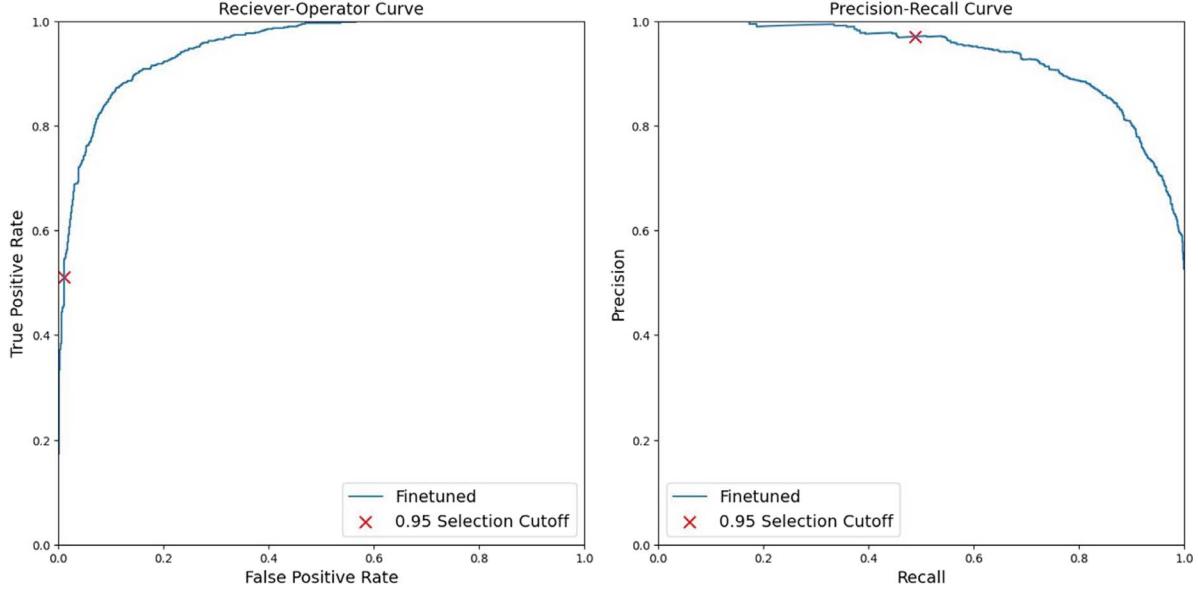
D.O.R. would like to thank those in the *ESA Traineeship* program cohort of 2022. They created a wholly welcoming environment and space of support. A special thanks must go to Karolin Frohnappel and Emma Vellard for much technical discussion. Finally, D.O.R. would like to acknowledge Aurélien Verdier.

## Appendix A Further Model Diagnostics

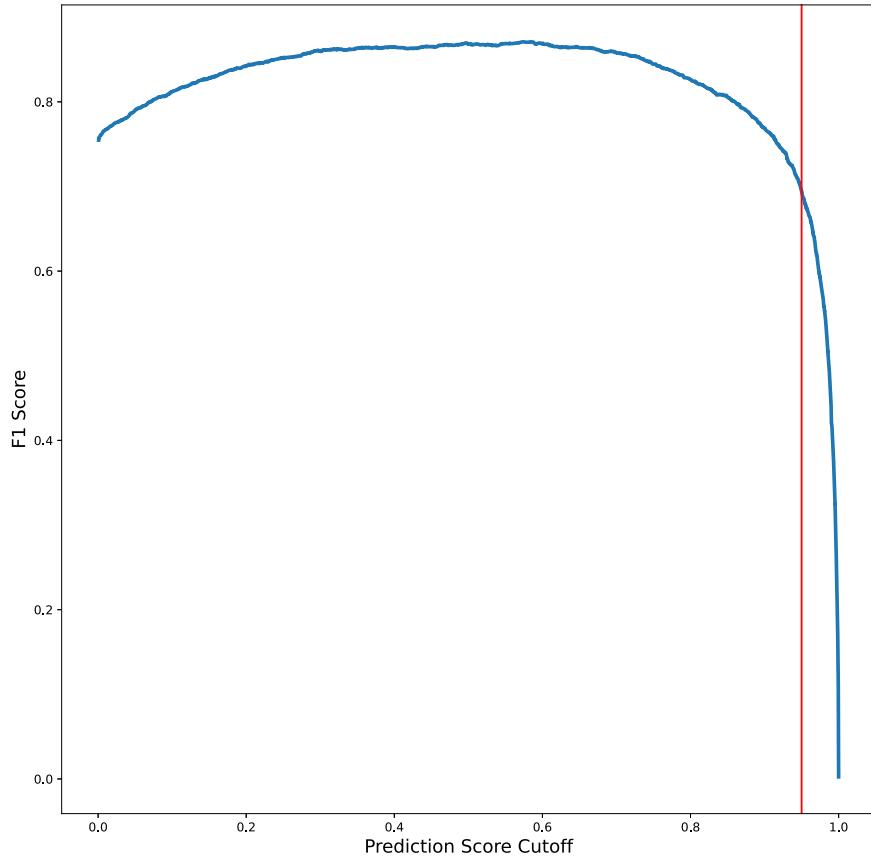
In Section 5 we presented diagnostic properties of our model. These include the accuracy measurements, purity measurements, as well as confusion matrices at different cutoffs of our model. Here, we present the receiver operating characteristic (ROC) curves, the precision-recall (PR) curves, and measures of true and false-positive rates versus the cutoff threshold.

Figure 14 shows the ROC and PR curves of the final *Zoobot* model we applied to the Hubble archives. The ROC shows the rate of change of finding true positives and false positives with changing cutoff. The PR curve shows the changes of precision against recall. Precision is the ratio of true positives (interacting galaxies correctly predicted as so) to the sum of true and false positives (noninteracting galaxies incorrectly predicted as interacting). The recall is then the ratio of true positives to the sum of true positives and false negatives (interacting galaxies that have been misclassified as noninteracting). The red crosses in both plots shows how the model was behaving when we used a cutoff of 0.95.

These are both as expected. Both curves show that the model behaves well, and are much better than a random classifier (which would have a 1:1 relation). The ROC plot shows that we are minimizing our false-positive rate when using a prediction score cutoff of 0.95. However, we are misclassifying approximately 50% of interacting galaxies as noninteracting galaxies. The contamination rate in our final catalog (false-positive rate) will be very low (close to zero in this ideal validation set). The PR curve shows a similar result. Here, we are operating with high precision (finding a pure catalog) while keeping our recall minimal.



**Figure 14.** The receiver-operator and precision-recall curve for the Zoobot model that was used to explore the Hubble archives. These curves measure the relevant rates or characteristics based on the changing cutoff applied to how Zoobot defines an interacting galaxy. The red crosses are where the prediction score cutoff is for this work. We can see in the receiver-operator curve that the prediction score cutoff we use would have an incredibly low false-positive rate, while it would be misclassifying  $\approx 50\%$  of interacting galaxies. This is also shown in the precision-recall curve where our recall is  $\approx 50\%$ .



**Figure 15.** The F1 score found during the diagnostics of the model used in this work. The F1 score is a measure combining accuracy and purity into one metric. The cutoff we use is at the point where the F1 score begins to rapidly decline. This point is shown by the red vertical line.

We also present the changing F1 score for the model used in this work, shown in Figure 15. The F1 score is twice the ratio of precision multiplied by recall upon precision summed to recall. This combines our measure of accuracy and purity into a single metric. The cutoff we use in this work is at the point

where the F1 score has begun to decline. This is because we are beginning to lose recall rapidly, but gaining significantly in precision. As discussed in Section 5, this was an acceptable trade-off in this work for a very large, pure interacting galaxy catalog.

## Appendix B

### Examples of Sources with Three-band Information

Of the full catalog of 21,926 interacting systems, only 1336 of them had obtained all three-band information. Six examples

are shown in Figure 16. These were created using the Lupton et al. (2004) algorithm, with a scaling factor  $Q=2$  and  $\alpha=0.75$ , with (F814W, F606W, F475W) as RGB channels and multiplicative factors of (1.25, 0.95, and 2).



**Figure 16.** Example of six interacting systems in the catalog with full three-band imagery.

## Appendix C

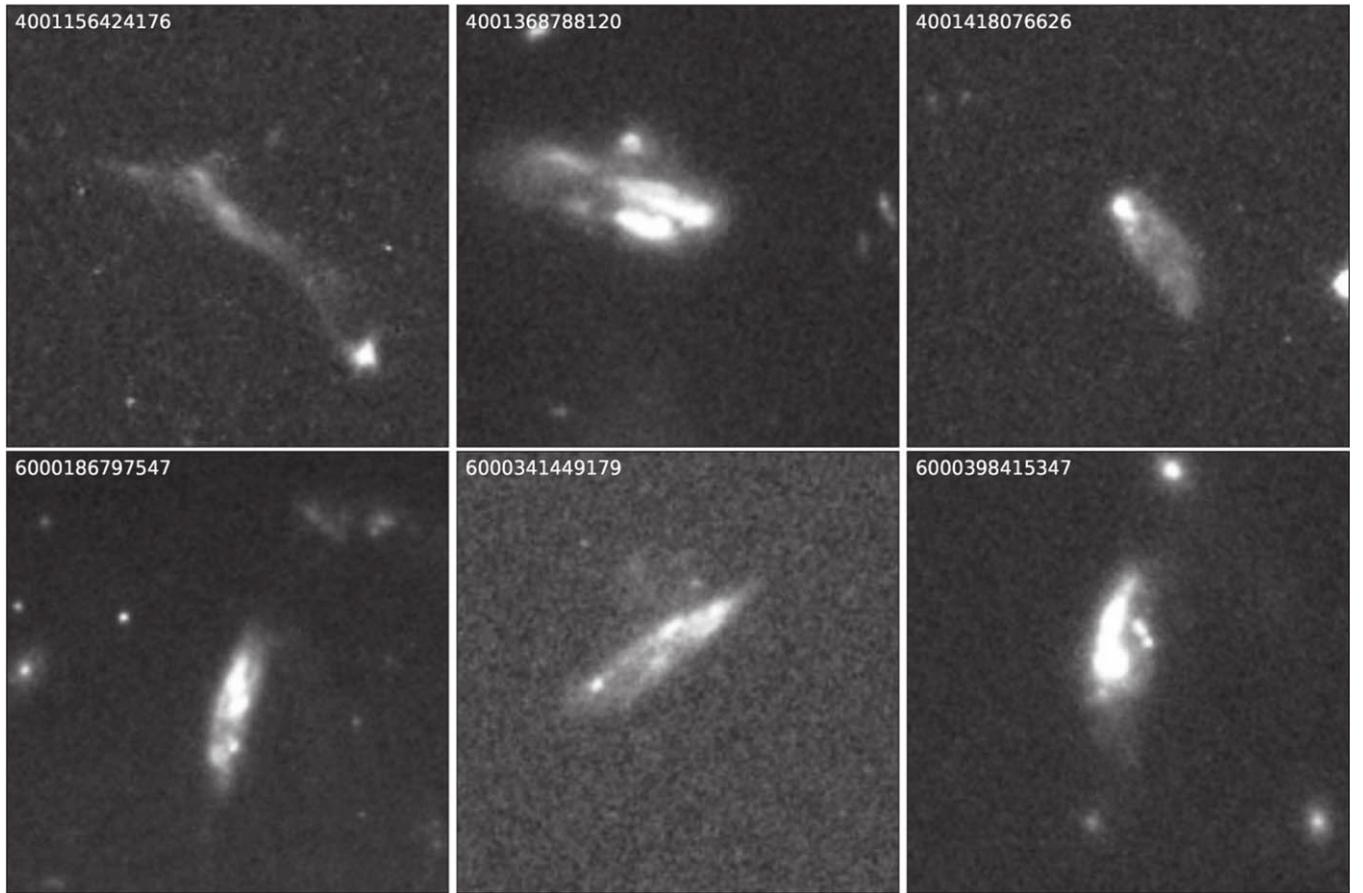
### Unknown Objects

From the final catalog, there were six sources that we could not visually identify. These objects were also not referenced anywhere in astrophysical literature. F814W cutouts of the six objects are shown in Figure 17. Their Source IDs are shown in the upper left of each image, and a separate catalog has been released of these with all other objects. This catalog can be found at the data release on Zenodo.

Four of the six objects (40001156424176, 4001368788120, 4001418076626, and 6000398415347) have a bright central

source, followed by a low-surface brightness tail. Initially, it was assumed that these were solar system objects such as comets. This, however, could not be confirmed. The first of these four sources is also thought to potentially be a highly disrupted system with a significantly elongated tidal feature. The final two unknown sources (6000186797547 and 6000341449179) have no clear central source, though there is extended structure to them. These are likely to be highly irregular galaxies, but no confirmation could be found.

These objects are released to the community for identification and investigation, as the authors cannot find definitive agreement on what they are.



**Figure 17.** The six unknown systems found in this work. These have no reference in Simbad or in NED, and their morphology could not be classified by the authors. Investigation into these six objects is presented to the community, with the authors hoping that future work and investigation can be conducted.

## Appendix D

### Acknowledging PIs

In the final section of this work, we wish to acknowledge all of the PIs whose observations we have used. A machine readable table containing the proposal IDs, the DOIs and the full observation filename accompanies this work. In this MRT, the specific filename in ESA Datalabs is also specified so that

the observation can be easily found. The Astroquery Python package can also be used with this filename to download each observation. This MRT contains details of every observation we used. Table 4 presents each observation for which a reference was recommended to use by the PI so that we could formally acknowledge them. It also states their DOI, observation date and proposal ID, while the filenames are omitted.

**Table 4**  
Twenty Examples of the Accompanying Data Table of Observations Used

Proposal ID (1)	Observation ID (2)	Observation Date (3)	DOI (4)	References (5)
8183	hst_8183_54_acs_wfc_f814w_j59l54	18/07/2002	<a href="https://doi.org/10.5270/esa-88k8vcj">https://doi.org/10.5270/esa-88k8vcj</a>	
9075	hst_9075_2a_acs_wfc_f814w_j6fl2a	24/07/2002	<a href="https://doi.org/10.5270/esa-gsxhb4b">https://doi.org/10.5270/esa-gsxhb4b</a>	
9351	hst_9351_11_acs_wfc_f814w_j8d211	31/03/2003	<a href="https://doi.org/10.5270/esa-5lba8bo">https://doi.org/10.5270/esa-5lba8bo</a>	
9361	hst_9361_03_acs_wfc_f814w_j8d503	22/07/2003	<a href="https://doi.org/10.5270/esa-ecmnqgh">https://doi.org/10.5270/esa-ecmnqgh</a>	
9363	hst_9363_09_acs_wfc_f814w_j8d809	02/07/2002	<a href="https://doi.org/10.5270/esa-ethtec5">https://doi.org/10.5270/esa-ethtec5</a>	
9367	hst_9367_02_acs_wfc_f814w_j8ds02	10/06/2003	<a href="https://doi.org/10.5270/esa-3j404ll">https://doi.org/10.5270/esa-3j404ll</a>	
9373	hst_9373_02_acs_wfc_f814w_j6la02	05/07/2002	<a href="https://doi.org/10.5270/esa-ztsq94u">https://doi.org/10.5270/esa-ztsq94u</a>	Rejkuba et al. (2005)
9376	hst_9376_02_acs_wfc_f814w_j8e302	13/07/2002	<a href="https://doi.org/10.5270/esa-h90iavd">https://doi.org/10.5270/esa-h90iavd</a>	Keel et al. (2006)
9381	hst_9381_02_acs_wfc_f814w_j8fu02	13/03/2003	<a href="https://doi.org/10.5270/esa-vlapyea">https://doi.org/10.5270/esa-vlapyea</a>	
9400	hst_9400_04_acs_wfc_f814w_j6kx04	29/05/2003	<a href="https://doi.org/10.5270/esa-39mout">https://doi.org/10.5270/esa-39mout</a>	
9403	hst_9403_02_acs_wfc_f814w_j8fp02	09/07/2002	<a href="https://doi.org/10.5270/esa-k5mv9ct">https://doi.org/10.5270/esa-k5mv9ct</a>	
9405	hst_9405_6k_acs_wfc_f814w_j8iy6k	22/05/2003	<a href="https://doi.org/10.5270/esa-zy9phm1">https://doi.org/10.5270/esa-zy9phm1</a>	
9409	hst_9409_03_acs_wfc_f814w_j6n203	29/06/2003	<a href="https://doi.org/10.5270/esa-vjngw7r">https://doi.org/10.5270/esa-vjngw7r</a>	Goudfrooij et al. (2004)
9411	hst_9411_09_acs_wfc_f814w_j8dl09	11/02/2003	<a href="https://doi.org/10.5270/esa-debpiln">https://doi.org/10.5270/esa-debpiln</a>	
9427	hst_9427_13_acs_wfc_f814w_j6m613	21/10/2002	<a href="https://doi.org/10.5270/esa-bw1b97v">https://doi.org/10.5270/esa-bw1b97v</a>	
9438	hst_9438_01_acs_wfc_f814w_j6me01	16/01/2003	<a href="https://doi.org/10.5270/esa-e5eaam5">https://doi.org/10.5270/esa-e5eaam5</a>	Gregg & West (2017)
9450	hst_9450_02_acs_wfc_f814w_j8d402	25/08/2002	<a href="https://doi.org/10.5270/esa-9ttmykz">https://doi.org/10.5270/esa-9ttmykz</a>	York et al. (2005)
9453	hst_9453_02_acs_wfc_f814w_j8f802	03/12/2002	<a href="https://doi.org/10.5270/esa-1xvyjfy">https://doi.org/10.5270/esa-1xvyjfy</a>	Brown et al. (2003)
9454	hst_9454_11_acs_wfc_f814w_j8ff11	23/03/2003	<a href="https://doi.org/10.5270/esa-xsdowj9">https://doi.org/10.5270/esa-xsdowj9</a>	

(This table is available in its entirety in machine-readable form.)

## ORCID iDs

David O’Ryan <https://orcid.org/0000-0003-1217-4617>  
 Bruno Merín <https://orcid.org/0000-0002-8555-3012>  
 Brooke D. Simmons <https://orcid.org/0000-0001-5882-3323>  
 Mike Walmsley <https://orcid.org/0000-0002-6408-4181>  
 Izzy L. Garland <https://orcid.org/0000-0002-3887-6433>  
 Tobias Géron <https://orcid.org/0000-0002-6851-9613>  
 William Keel <https://orcid.org/0000-0002-6131-9539>  
 Sandor Kruk <https://orcid.org/0000-0001-8010-8879>  
 Chris J. Lintott <https://orcid.org/0000-0001-5578-359X>  
 Kameswara Bharadwaj Mantha <https://orcid.org/0000-0002-6016-300X>  
 Karen L. Masters <https://orcid.org/0000-0003-0846-9578>  
 Rebecca J. Smethurst <https://orcid.org/0000-0001-6417-7196>  
 Matthew R. Thorne <https://orcid.org/0000-0003-1014-5839>

## References

Abadi, M., Barham, P., Chen, J., et al. 2016, arXiv:1605.08695  
 Abadi, M. G., Navarro, J. F., Steinmetz, M., & Eke, V. R. 2003, *ApJ*, **591**, 499  
 Abd El Aziz, M., Selim, I. M., & Xiong, S. 2017, *NatSR*, **7**, 4463  
 Abramson, A., Kenney, J., Crowl, H., & Tal, T. 2016, *AJ*, **152**, 32  
 Abramson, A., & Kenney, J. D. P. 2014, *AJ*, **147**, 63  
 Ackermann, S., Schawinski, K., Zhang, C., Weigel, A. K., & Turp, M. D. 2018, *MNRAS*, **479**, 415  
 Adamo, A., Ryon, J. E., Messa, M., et al. 2017, *ApJ*, **841**, 131  
 Alonso, M. S., Lambas, D. G., Tissera, P., & Coldwell, G. 2007, *MNRAS*, **375**, 1017  
 Alonso, M. S., Tissera, P. B., Coldwell, G., & Lambas, D. G. 2004, *MNRAS*, **352**, 1081  
 Ardizzone, E., Di Gesù, V., & Maccarone, M. C. 1996, *VA*, **40**, 401  
 Arp, H. 1966, *ApJS*, **14**, 1  
 Arp, H. C., & Madore, B. 1987, A Catalogue of Southern Peculiar Galaxies and Associations (Cambridge: Cambridge Univ. Press)  
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123  
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33  
 Avila, R. J., Hack, W., Cara, M., et al. 2015, in Astronomical Data Analysis Software and Systems XXIV (ADASS XXIV), ed. A. R. Taylor & E. Rosolowsky (San Francisco, CA: ASP), 281  
 Barchi, P. H., de Carvalho, R. R., Rosa, R. R., et al. 2020, *A&C*, **30**, 100334  
 Barton, E. J., Geller, M. J., & Kenyon, S. J. 2000, *ApJ*, **530**, 660  
 Bastian, N., Schweizer, F., Goudfrooij, P., Larsen, S. S., & Kissler-Patig, M. 2013, *MNRAS*, **431**, 1252  
 Bedin, L. R., Salaris, M., Anderson, J., et al. 2015, *MNRAS*, **448**, 1779  
 Bedin, L. R., Salaris, M., King, I. R., et al. 2010, *ApJL*, **708**, L32  
 Bellini, A., Grogan, N. A., Hathi, N., & Brown, T. M. 2017, The Hubble Space Telescope “Program of Last Resort,” Instrument Science Report ACS 2017-12, 28  
 Bennet, P., Sand, D. J., Crnojević, D., et al. 2022, *ApJ*, **924**, 98  
 Bickley, R. W., Bottrell, C., Hani, M. H., et al. 2021, *MNRAS*, **504**, 372  
 Blair, W. P., Chandar, R., Dopita, M. A., et al. 2014, *ApJ*, **788**, 55  
 Bottrell, C., Hani, M. H., Teimoorinia, H., et al. 2019, *MNRAS*, **490**, 5390  
 Bradač, M., Ryan, R., Casertano, S., et al. 2014, *ApJ*, **785**, 108  
 Brown, T. M., Beaton, R., Chiba, M., et al. 2008, *ApJL*, **685**, L121  
 Brown, T. M., Ferguson, H. C., Smith, E., et al. 2003, *ApJL*, **592**, L17  
 Brown, T. M., Smith, E., Ferguson, H. C., et al. 2006, *ApJ*, **652**, 323  
 Brown, T. M., Tumlinson, J., Geha, M., et al. 2014, *ApJ*, **796**, 91  
 Brown, W. R., Anderson, J., Gnedin, O. Y., et al. 2015, *ApJ*, **804**, 49  
 Buck, T., & Wolf, S. 2021, arXiv:2111.01154  
 Calzetti, D., Lee, J. C., Sabbi, E., et al. 2015, *AJ*, **149**, 51  
 Cannon, J. M., O’Leary, E. M., Weisz, D. R., et al. 2012, *ApJ*, **747**, 122  
 Carleton, T., Cooper, M. C., Bolatto, A. D., et al. 2017, *MNRAS*, **467**, 4886  
 Chen, C.-C., Smail, I., Swinbank, A. M., et al. 2015, *ApJ*, **799**, 194  
 Chen, J., Lee, S. K., Castander, F.-J., Maza, J., & Schechter, P. L. 2013, *ApJ*, **769**, 81  
 Cheng, T.-Y., Huertas-Company, M., Conselice, C. J., et al. 2021, *MNRAS*, **503**, 4446  
 Chiboucas, K., Jacobs, B. A., Tully, R. B., & Karachentsev, I. D. 2013, *AJ*, **146**, 126  
 Cohen, R. E., Goudfrooij, P., Correnti, M., et al. 2020, *ApJ*, **890**, 52  
 Comerford, J. M., Pooley, D., Barrows, R. S., et al. 2015, *ApJ*, **806**, 219  
 Cook, D. O., Lee, J. C., Adamo, A., et al. 2019, *MNRAS*, **484**, 4897  
 Correnti, M., Goudfrooij, P., Kalirai, J. S., et al. 2014, *ApJ*, **793**, 121  
 Crnojević, D., Sand, D. J., Bennet, P., et al. 2019, *ApJ*, **872**, 80  
 Dalcanton, J. J., Williams, B. F., Lang, D., et al. 2012, *ApJS*, **200**, 18  
 Darg, D. W., Kaviraj, S., Lintott, C. J., et al. 2010a, *MNRAS*, **401**, 1552  
 Darg, D. W., Kaviraj, S., Lintott, C. J., et al. 2010b, *MNRAS*, **401**, 1043  
 Das, A., Pandey, B., & Sarkar, S. 2022, arXiv:2207.03968  
 De Lucia, G., & Blaizot, J. 2007, *MNRAS*, **375**, 2  
 de Mello, D. F., Infante, L., & Menanteau, F. 1997, *ApJS*, **108**, 99  
 Dey, A., Schlegel, D. J., Lang, D., et al. 2019, *AJ*, **157**, 168  
 Ellison, S. L., Mendel, J. T., Patton, D. R., & Scudder, J. M. 2013, *MNRAS*, **435**, 3627  
 Ellison, S. L., Patton, D. R., Mendel, J. T., & Scudder, J. M. 2011, *MNRAS*, **418**, 2043  
 Ellison, S. L., Patton, D. R., Simard, L., & McConnachie, A. W. 2008, *AJ*, **135**, 1877  
 Elmegreen, D. M., Elmegreen, B. G., Gallagher, J. S., et al. 2022, *ApJ*, **941**, 157  
 Ferreira, L., Adams, N., Conselice, C. J., et al. 2022, *ApJL*, **938**, L2  
 Finner, K., HyeongHan, K., Jee, M. J., et al. 2021, *ApJ*, **918**, 72  
 Folatelli, G., Van Dyk, S. D., Kuncarayakti, H., et al. 2016, *ApJL*, **825**, L22  
 Gallart, C., Monelli, M., Mayer, L., et al. 2015, *ApJL*, **811**, L18  
 Gallart, C., Monelli, M., Ruiz-Lara, T., et al. 2021, *ApJ*, **909**, 192  
 Geha, M., Weisz, D., Grocholski, A., et al. 2015, *ApJ*, **811**, 114  
 Gennaro, M., Geha, M., Tchernyshyov, K., et al. 2018, *ApJ*, **863**, 38  
 Ghosh, A., Urry, C. M., Wang, Z., et al. 2020, *ApJ*, **895**, 112  
 Giavalisco, M., Ferguson, H. C., Koekemoer, A. M., et al. 2004, *ApJL*, **600**, L93  
 Gillies, S., van der Wel, C., van den Bossche, J., et al. 2022, Shapely: Manipulation and Analysis of Geometric Objects, doi:10.5281/zenodo.7428463  
 Girardi, L., Goudfrooij, P., Kalirai, J. S., et al. 2013, *MNRAS*, **431**, 3501  
 Golovich, N., van Weeren, R. J., Dawson, W. A., Jee, M. J., & Wittman, D. 2017, *ApJ*, **838**, 110  
 Gonzalez, A. H., George, T., Connor, T., et al. 2021, *MNRAS*, **507**, 963  
 Goudfrooij, P. 2012, *ApJ*, **750**, 140  
 Goudfrooij, P., Gilmore, D., Whitmore, B. C., & Schweizer, F. 2004, *ApJL*, **613**, L121  
 Goudfrooij, P., Girardi, L., & Correnti, M. 2017, *ApJ*, **846**, 22  
 Gregg, M., & West, M. 2017, Early Stages of Galaxy Cluster Formation, 13, Zenodo, 10.5281/zenodo.831767  
 Guo, Q., & White, S. D. M. 2008, *MNRAS*, **384**, 2  
 Hani, M. H., Gosain, H., Ellison, S. L., Patton, D. R., & Torrey, P. 2020, *MNRAS*, **493**, 3716  
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357  
 Hernández-Toledo, H. M., Avila-Reese, V., Conselice, C. J., & Puerari, I. 2005, *AJ*, **129**, 682  
 Hillis, T. J., Williams, B. F., Dolphin, A. E., Dalcanton, J. J., & Skillman, E. D. 2016, *ApJ*, **831**, 191  
 Hoekstra, H., Donahue, M., Conselice, C. J., McNamara, B. R., & Voit, G. M. 2011, *ApJ*, **726**, 48  
 Holincheck, A. J., Wallin, J. F., Borne, K., et al. 2016, *MNRAS*, **459**, 720  
 Hopkins, P. F., Cox, T. J., Hernquist, L., et al. 2013, *MNRAS*, **430**, 1901  
 Hunter, J. D. 2007, *CSE*, **9**, 90  
 Jacobs, C., Collett, T., Glazebrook, K., et al. 2019, *ApJS*, **243**, 17  
 Jee, M. J., Stroe, A., Dawson, W., et al. 2015, *ApJ*, **802**, 46  
 Jencson, J. E., Sand, D. J., Andrews, J. E., et al. 2022, *ApJ*, **930**, 81  
 Kaaret, P., & Alonso-Herrero, A. 2008, *ApJ*, **682**, 1020  
 Kauffmann, G., Heckman, T. M., White, S. D. M., et al. 2003, *MNRAS*, **341**, 33  
 Kaviraj, S. 2014a, *MNRAS*, **440**, 2944  
 Kaviraj, S. 2014b, *MNRAS*, **437**, L41  
 Keel, W. C., White, Raymond, E., Owen, I., & Ledlow, M. J. F. N. 2006, *AJ*, **132**, 2233  
 Kenney, J. D. P., Abramson, A., & Bravo-Alfaro, H. 2015, *AJ*, **150**, 59  
 Kim, M., Barth, A. J., Ho, L. C., & Son, S. 2021, *ApJS*, **256**, 40  
 Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980  
 Laudari, S., Jáchym, P., Sun, M., et al. 2022, *MNRAS*, **509**, 3938  
 Li, C., Kauffmann, G., Heckman, T. M., White, S. D. M., & Jing, Y. P. 2008, *MNRAS*, **385**, 1915  
 Lintott, C. J., Schawinski, K., Slosar, A., et al. 2008, *MNRAS*, **389**, 1179

López-Sanjuan, C., Balcells, M., Pérez-González, P. G., et al. 2009, *A&A*, **501**, 505

Lotz, J. M., Davis, M., Faber, S. M., et al. 2008, *ApJ*, **672**, 177

Lunyan, R., Chornock, R., Berger, E., et al. 2015, *ApJ*, **804**, 90

Lupton, R., Blanton, M. R., Fekete, G., et al. 2004, *PASP*, **116**, 133

Mackey, A. D., Huxor, A., Ferguson, A. M. N., et al. 2006, *ApJL*, **653**, L105

Mackey, A. D., Huxor, A., Ferguson, A. M. N., et al. 2007, *ApJL*, **655**, L85

Marian, V., Jahnke, K., Andika, I., et al. 2020, *ApJ*, **904**, 79

Massey, R., Harvey, D., Liesenborgs, J., et al. 2018, *MNRAS*, **477**, 669

Massey, R., Williams, L., Smit, R., et al. 2015, *MNRAS*, **449**, 3393

Maybhate, A., Goudfrooij, P., Chandar, R., & Puzia, T. H. 2010, *ApJ*, **721**, 893

Maybhate, A., Goudfrooij, P., Schweizer, F., Puzia, T., & Carter, D. 2007, *AJ*, **134**, 1729

Maybhate, A., Goudfrooij, P., Schweizer, F., Puzia, T. H., & Carter, D. 2009, *AJ*, **137**, 383

McCully, C., Jha, S. W., Scalzo, R. A., et al. 2022, *ApJ*, **925**, 138

McInnes, L., Healy, J., & Melville, J. 2018, arXiv:1802.03426

McKernan, B., Ford, K. E. S., & Reynolds, C. S. 2010, *MNRAS*, **407**, 2399

McKinney, W. 2010, in Proc. of the 9th Python in Science Conf., 56

McQuinn, K. B. W., Cannon, J. M., Dolphin, A. E., et al. 2014, *ApJ*, **785**, 3

McQuinn, K. B. W., Telidevara, A. K., Fuson, J., et al. 2021, *ApJ*, **918**, 23

Merín, B., Giordano, F., Norman, H., et al. 2017, arXiv:1712.04114

Mihos, J. C., & Hernquist, L. 1996, *ApJ*, **464**, 641

Mo, W., Gonzalez, A., Jee, M. J., et al. 2016, *ApJL*, **818**, L25

Moreno, J., Torrey, P., Ellison, S. L., et al. 2021, *MNRAS*, **503**, 3113

Morganti, R., Garrett, M. A., Chapman, S., et al. 2004, *A&A*, **424**, 371

Mutlu-Pakdil, B., Sand, D. J., Crnojević, D., et al. 2022, *ApJ*, **926**, 77

Nair, P. B., & Abraham, R. G. 2010, *ApJS*, **186**, 427

Helou, G., Madore, B. F., Schmitz, M., et al. 1991, Databases & On-line Data in Astronomy, Astrophysics and Space Science Library, Vol. 171 (Dordrecht: Springer), 89

Newman, A. B., Belli, S., Ellis, R. S., & Patel, S. G. 2018, *ApJ*, **862**, 125

Nielsen, F. 2016, Introduction to HPC with MPI for Data Science. Undergraduate Topics in Computer Science (Cham: Springer), 195

Ochsenbein, F., Bauer, P., & Marcout, J. 2000, *A&AS*, **143**, 23

O’Shea, K., & Nash, R. 2015, arXiv:1511.08458

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

Pearson, W. J., Wang, L., Alpaslan, M., et al. 2019, *A&A*, **631**, A51

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2012, *JMLR*, **12**, 2825

Penny, S. J., Conselice, C. J., de Rijcke, S., & Held, E. V. 2009, *MNRAS*, **393**, 1054

Rejkuba, M., Greggio, L., Harris, W. E., Harris, G. L. H., & Peng, E. W. 2005, *ApJ*, **631**, 262

Rejkuba, M., Harris, W. E., Greggio, L., et al. 2014, *ApJL*, **791**, L2

Rodney, S. A., Balestra, I., Bradac, M., et al. 2018, *NatAs*, **2**, 324

Rodney, S. A., Patel, B., Scolnic, D., et al. 2015, *ApJ*, **811**, 70

Ross, D., Lim, J., Lin, R., et al. 2008, *Int. J. Comput. Vis.*, **77**, 125

Sabbi, E., Calzetti, D., Ubeda, L., et al. 2018, *ApJS*, **235**, 23

Sahai, R. 2019, IAU Proc. 343, Why Galaxies Care About AGB Stars: A Continuing Challenge through Cosmic Time, 343 (Cambridge: Cambridge Univ. Press) 164, <https://ui.adsabs.harvard.edu/abs/2019IAUS..343..164S/abstract>

Sahai, R., Morris, M., Sánchez Contreras, C., & Claussen, M. 2007, *AJ*, **134**, 2200

Saitoh, T. R., Daisaka, H., Kokubo, E., et al. 2009, *PASJ*, **61**, 481

Sand, D. J., Seth, A. C., Crnojević, D., et al. 2017, *ApJ*, **843**, 134

Saviane, I., Momany, Y., Da Costa, G. S., Rich, R. M., & Hibbard, J. E. 2008, *ApJ*, **678**, 179

Scalco, M., Bellini, A., Bedin, L. R., et al. 2021, *MNRAS*, **505**, 3549

Schawinski, K., Urry, C. M., Simmons, B. D., et al. 2014, *MNRAS*, **440**, 889

Scoville, N., Aussel, H., Brusa, M., et al. 2007, *ApJS*, **172**, 1

Scudder, J. M., Ellison, S. L., Torrey, P., Patton, D. R., & Mendel, J. T. 2012, *MNRAS*, **426**, 549

Sharon, K., Bayliss, M. B., Dahle, H., et al. 2017, *ApJ*, **835**, 5

Sharon, K., Bayliss, M. B., Dahle, H., et al. 2020, *ApJS*, **247**, 12

Silva, A., Marchesini, D., Silverman, J. D., et al. 2021, *ApJ*, **909**, 124

Simmons, B. D., Lintott, C., Willett, K. W., et al. 2017, *MNRAS*, **464**, 4420

Simon, J. D., Brown, T. M., Drlica-Wagner, A., et al. 2021, *ApJ*, **908**, 18

Skillman, E. D., Monelli, M., Weisz, D. R., et al. 2017, *ApJ*, **837**, 102

Smethurst, R. J., Lintott, C. J., Simmons, B. D., et al. 2015, *MNRAS*, **450**, 435

Smethurst, R. J., Masters, K. L., Lintott, C. J., et al. 2018, *MNRAS*, **473**, 2679

Sohn, S. T., Besla, G., van der Marel, R. P., et al. 2013, *ApJ*, **768**, 139

Sohn, S. T., Patel, E., Besla, G., et al. 2017, *ApJ*, **849**, 93

Springel, V. 2000, *MNRAS*, **312**, 859

Springel, V., White, S. D. M., Jenkins, A., et al. 2005, *Natur*, **435**, 629

Stockton, A., McGrath, E., Canalizo, G., Iye, M., & Maihara, T. 2008, *ApJ*, **672**, 146

Tanvir, N. R., Mackey, A. D., Ferguson, A. M. N., et al. 2012, *MNRAS*, **422**, 162

Thompson, R. I., Sauvage, M., Kennicutt, R. C., et al. 2009, *ApJ*, **691**, 1068

Toomre, A., & Toomre, J. 1972, *ApJ*, **178**, 623

Tran, K.-V. H., Moustakas, J., Gonzalez, A. H., et al. 2008, *ApJL*, **683**, L17

Trombley, C., Donahue, M., Bruch, S., et al. 2006, AAS Meeting, **209**, 77.01

van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., et al. 2014, *PeerJ*, **2**, e453

van Dokkum, P. G. 2005, *AJ*, **130**, 2647

Vorontsov-Velyaminov, B. A. 1959, Atlas and Catalog of Interacting Galaxies

Vorontsov-Velyaminov, B. A. 1977, *A&AS*, **28**, 1

Wallin, J. F., Holincheck, A. J., & Harvey, A. 2016, *A&C*, **16**, 26

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

Walmsley, M., Lintott, C., Géron, T., et al. 2022b, Zoobot: Deep learning galaxy morphology classifier, Version 0.0.2, Zenodo, doi:10.5281/zenodo.6483176

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

Weisz, D. R., Skillman, E. D., Cannon, J. M., et al. 2008, *ApJ*, **689**, 160

Weisz, D. R., Skillman, E. D., Hidalgo, S. L., et al. 2014, *ApJ*, **789**, 24

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

Werk, J. K., Putman, M. E., Meurer, G. R., et al. 2008, *ApJ*, **678**, 888

Whitaker, K. E., Kriek, M., van Dokkum, P. G., et al. 2012, *ApJ*, **745**, 179

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

Willett, K. W., Galloway, M. A., Bamford, S. P., et al. 2017, *MNRAS*, **464**, 4176

Willett, K. W., Lintott, C. J., Bamford, S. P., et al. 2013, *MNRAS*, **435**, 2835

Worrall, D. M., Birkinshaw, M., Young, A. J., et al. 2012, *MNRAS*, **424**, 1346

York, T., Jackson, N., Browne, I. W. A., Wucknitz, O., & Skelton, J. E. 2005, *MNRAS*, **357**, 124

Zohren, H., Schrabback, T., Bocquet, S., et al. 2022, *A&A*, **668**, A18