





# Astronomy Protege: Discovery through Human-machine Collaboration

Michelle Lochner<sup>1,2</sup>  and Lawrence Rudnick<sup>3</sup> <sup>1</sup> Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town, 7535, South Africa; [dr.michelle.lochner@gmail.com](mailto:dr.michelle.lochner@gmail.com)<sup>2</sup> South African Radio Astronomy Observatory, Liesbeek House, River Park, Liesbeek Parkway, Mowbray 7705, South Africa<sup>3</sup> Minnesota Institute for Astrophysics, University of Minnesota, 116 Church Street SE, Minneapolis, MN 55455, USA

Received 2024 November 14; revised 2024 December 4; accepted 2024 December 7; published 2025 February 4

## Abstract

Modern telescopes generate catalogs of millions of objects with the potential for new scientific discoveries, but this is beyond what can be examined visually. Here we introduce ASTRONOMY: PROTEGE, an extension of the general-purpose machine-learning-based active anomaly detection framework ASTRONOMY. PROTEGE is designed to provide well-selected recommendations for visual inspection, based on a small amount of optimized human labeling. The resulting sample contains rare or unusual sources that are simultaneously as diverse as the human trainer chooses and of scientific interest to them. We train PROTEGE on images from the MeerKAT Galaxy Cluster Legacy Survey, leveraging the self-supervised deep learning algorithm Bootstrap Your Own Latent to find a low-dimensional representation of the radio galaxy cutouts. By operating in this feature space, PROTEGE is able to recommend interesting sources with completely different morphologies in image space to those it has been trained on. This provides important advantages over similarity searches, which can only find more examples of known sources, or blind anomaly detection, which selects unusual but not necessarily scientifically interesting sources. Using an evaluation subset, we show that, with minimal training, PROTEGE provides excellent recommendations and find that it is even able to recommend sources that the authors missed. We briefly highlight some of PROTEGE's top recommendations, which include X- and circular-shaped sources, filamentary structures, and one-sided structures. These results illustrate the power of an optimized human-machine collaboration, such as PROTEGE, to make unexpected discoveries in samples beyond human-accessible scales.

*Unified Astronomy Thesaurus concepts:* [Astronomy data analysis \(1858\)](#); [Astronomy software \(1855\)](#); [Open source software \(1866\)](#); [Extragalactic radio sources \(508\)](#); [Radio active galactic nuclei \(2134\)](#)

## 1. Introduction

Machine learning (ML) has become a critical tool in modern astronomy as datasets from new telescopes and surveys increase in size and complexity. In particular, it has the potential to enable the discovery of rare or new classes of sources, finding these “needles” in enormous “haystacks.” V. Böhm et al. (2023) used ML to discover unusual sources such as supernovae and LINERS among thousands of spectra from the Sloan Digital Sky Survey, while S. Webb et al. (2020) used clustering and anomaly detection to discover new variables and flare stars in tens of thousands of light curves. However, for these algorithms to be effective, the dimensionality of the input data must be reduced to a smaller set of meaningful numbers. This feature extraction is typically the most challenging step for most ML applications.

Deep learning (X. Zhao et al. 2024) has revolutionized ML because, given enough data, it can perform highly effective feature extraction automatically, extracting a useful lower-dimensional representation of complicated high-dimensional data such as images. In radio astronomy, R. I. J. Mostert et al. (2021), N. Gupta et al. (2022), and A. N. Vantyghem et al. (2024) used self-organizing maps, an algorithm based on neural networks, to discover rare sources in LOFAR (M. P. van Haarlem et al. 2013), ASKAP (A. W. Hotan et al. 2021), and VLASS (M. Lacy et al. 2020) data, respectively. N. O. Ralph et al. (2019) used an autoencoder to

find a representation of the radio images on which to perform unsupervised clustering.

Several recent works (M. Walmsley et al. 2022b; I. V. Slijepcevic et al. 2023; S. Andrianomena & H. Tang 2024; K. Mohale & M. Lochner 2024; S. Riggi et al. 2024) have introduced the use of modern self-supervised learning techniques to automatically learn general representations for later downstream tasks, including anomaly detection.

However, anomaly detection alone is not sufficient in the face of large datasets because many anomalies are not scientifically interesting, and what is considered “interesting” varies between scientists. M. Lochner & B. A. Bassett (2021) introduced ASTRONOMY, a general-purpose anomaly detection framework that uses active learning, a method of obtaining a small number of strategically chosen labels from a human user, to refine the anomaly detection algorithm to focus on more interesting sources. ASTRONOMY was used to discover an unusual radio galaxy in MeerKAT (J. Jonas & MeerKAT Team 2016) data, called a Steep and Uneven Ring of Nonthermal Radiation (SAURON; M. Lochner et al. 2023). An extension of the ideas of ASTRONOMY was explored in A. Vafaei Sadr et al. (2022) where they iteratively retrained a convolutional neural network with active learning, effectively turning an anomaly detection problem into a classification problem.

Our goal in this work is to use ASTRONOMY to find sources that we find scientifically interesting in the MeerKAT Galaxy Cluster Legacy Survey (MGCLS; K. Knowles et al. 2022). To achieve this goal, we first extract meaningful features from the data using the self-supervised deep learning algorithm



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.

Bootstrap Your Own Latent (BYOL; J.-B. Grill et al. 2020), designed to automatically learn good representations of images.

We then use a second phase of ML to identify, from these features, a sample of diverse, rare, and “scientifically interesting” sources for further follow-up. We introduce here an extension to ASTRONOMY called PROTEGE, based on the work of M. Walmsley et al. (2022a), which is not a blind anomaly detection algorithm. Rather, PROTEGE uses active learning to recommend sources of interest to a human trainer, based on a minimal, optimal set of user-provided scores. We will show that, although PROTEGE can be used to find very specific source classes, by training it on a broad range of sources exhibiting unusual morphology, it is capable of learning a somewhat more general concept of what is interesting.

We begin this paper with an introduction to the broad landscape of ML in astronomy in Section 2, allowing us to place our work in context for the overall astronomy community. Section 3 then provides an overview of the MGCLS data and preprocessing techniques. Section 4 discusses the feature extraction method we use. We briefly explore in Section 5 practical aspects of what is meant by “scientifically interesting.” In Section 6, we introduce ASTRONOMY: PROTEGE and evaluate its effectiveness with a small hand-scored subset of data in Section 7.

Finally, we apply PROTEGE to the full MGCLS dataset and examine the types of recommendations PROTEGE makes in Section 8, with a discussion and conclusions in Sections 9 and 10. Appendices A–D contain a mosaic of random sources to provide a comparison with PROTEGE’s performance, more detailed technical information on the hyperparameter optimization, and a numerical evaluation of PROTEGE’s performance and comparison with other techniques.

We include the BYOL code for extracting features and the PROTEGE algorithm in a new release of the ASTRONOMY package (M. Lochner et al. 2024)<sup>4</sup> and release the catalog of PROTEGE’s recommendations, with strong caveats for its use in Zenodo at doi:10.5281/zenodo.14440675 and in github at <https://github.com/MichelleLochner/mgcls.protege>. Images of all of the sources in the catalog, ordered by PROTEGE score, are also available.<sup>5</sup>

## 2. PROTEGE’s Place in Machine Learning

A dizzying array of ML algorithms have been applied to a wide variety of problems in astronomy. Both the choice of specific technique and the amount of human labeling required will depend on the scientific problem to be solved. We first discuss the common, critical step of feature extraction, and then present Figure 1 to help understand PROTEGE’s unique role in this universe of algorithms.

### 2.1. Feature Extraction

ML algorithms in general cannot work directly with complex data such as images, light curves, or spectra. These must be simplified into a lower-dimensional, but still informative, representation of the data in a process called feature extraction. The choice of feature extraction method is crucial as it dictates the success of the downstream ML algorithm. The popularity of

deep learning algorithms (e.g., Y. LeCun et al. 2015), which rely on many-layered neural networks, is due to their ability to learn useful features automatically. *Self-supervised learning* is a branch of ML that automatically learns a representation that maintains similarity between sources in the lower-dimensional space, without needing any labels for the data. It is self-supervised learning that has unlocked the potential of anomaly detection in complex image data but also necessitated the development of a new approach to anomaly detection, namely PROTEGE, discussed below.

The behaviors of the different ML algorithms can be captured succinctly in this feature space. To do this, we use a visualization called a Uniform Manifold Approximation and Projection (UMAP) plot (L. McInnes et al. 2018), described in more detail in Section 4.4. In Figure 1 we use these plots to illustrate how different science problems and algorithms select or classify different parts of feature space, where each point represents some piece of underlying complex data (for example, an image or a spectrum).

### 2.2. Types of Machine Learning Problems


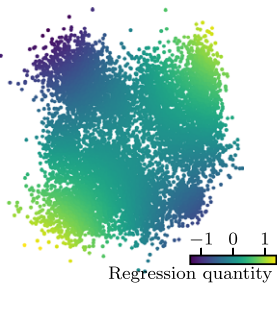

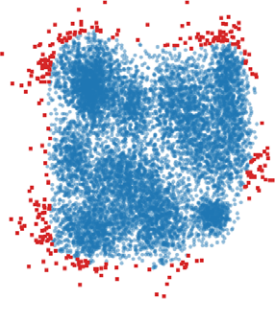
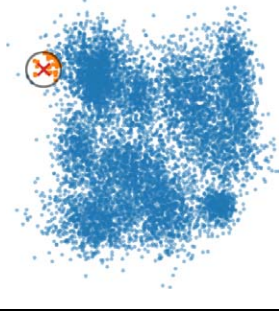
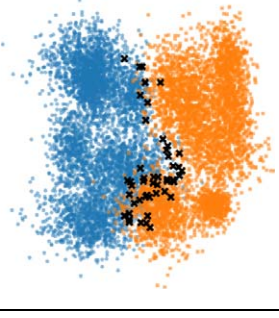
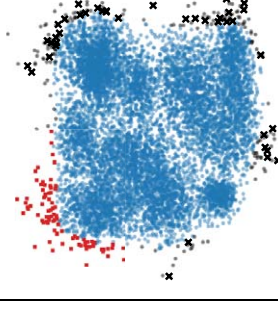
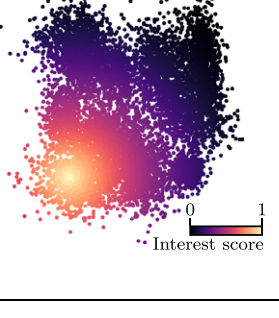
The top row in Figure 1 illustrates two common types of *supervised learning* problems where many labels are available, either from human labeling or additional data. These labels can be either discrete classes (Classification) or continuous variables (Regression). The designation of radio galaxies as either FRI or FRII (B. Becker et al. 2021) is an example of Classification. It requires the algorithm to recognize a wide range of morphologies, sizes, orientations, and brightness distributions, and group them into two discrete classes. We indicate those using two different colors in the UMAP feature space. When the algorithm is based on the human assignment of a continuous range of values, as in the determination of photometric redshifts (J. De Vicente et al. 2016), then it is called Regression, and can be seen as the smoothly varying color scale across feature space.

On the other hand, *unsupervised learning*, illustrated in the second row of Figure 1, requires no labels. Clustering can be used to break the dataset up into discrete regions in feature space, within which the objects are all close to each other, and farther away from other objects. These can later be labeled by a human, such as was done in K. Mohale & M. Lochner (2024) on optical galaxies, to create training sets or even potentially to find new patterns such as subclasses of sources. Anomaly detection algorithms are used to identify those sources that are at large distances in feature space from all other sources. Such outliers in feature space were used, e.g., by N. Gupta et al. (2022) to locate unusual galaxies in ASKAP data. Training data is unavailable, essentially by definition, for never-before-seen sources.

While obtaining a large enough number of labels for supervised learning may be prohibitive for many datasets, it is often possible to provide a few strategically chosen labels to improve algorithms, illustrated in the last two rows of Figure 1. A similarity search can be used to find more examples of a specific type of object by searching nearby feature space. This is similar to the goals of Classification, but focused on one, or a few, specific types. The remaining three problem types in Figure 1 are all examples of *active learning*, where an algorithm is iteratively improved by the introduction of a few optimally chosen labels. For instance, the RESSPECT system (N. Kenamer et al. 2020) uses Classification with active

<sup>4</sup> <https://github.com/MichelleLochner/astromy>

<sup>5</sup> <https://michellelochner.github.io/mgcls.protege>

	Problem type	Illustrative UMAP	Problem type	Illustrative UMAP
Many labels	<b>Classification</b> - Supervised learning problem where a substantial number of examples of objects of known classes are available and an algorithm is trained to classify new examples.		<b>Regression</b> - Supervised learning problem where a substantial number of objects with a known corresponding quantity are available and an algorithm is trained to predict this quantity as a function of features.	
	<b>Clustering</b> - The unsupervised learning equivalent of classification where all underlying classes are unknown and an algorithm attempts to identify classes by clustering similar objects together. Researchers must label classes after they have been identified.		<b>Anomaly detection</b> - Unsupervised learning problem which aims to identify anomalies - objects which appear substantially different from the norm in feature space. Many of these will not be scientifically interesting.	
Few labels	<b>Similarity search</b> - Technique to rapidly identify similar objects to a known exemplar by locating nearby objects in feature space.		<b>Classification with active learning</b> - A small number of difficult-to-classify objects are strategically selected for human labeling in order to improve the classifier.	
	<b>Active anomaly detection</b> - Uses a small number of human-supplied labels (black crosses) to improve anomaly detection by excluding regions of feature space considered by the user to be uninteresting (gray points), leaving only the most interesting anomalies (red squares). This is the basis for ASTRONOMALY.		<b>Astronomy: Protege</b> - Uses active learning to recommend sources which may be of interest to the user through iterative labeling and retraining of a regression algorithm, without running any initial anomaly detection algorithms.	

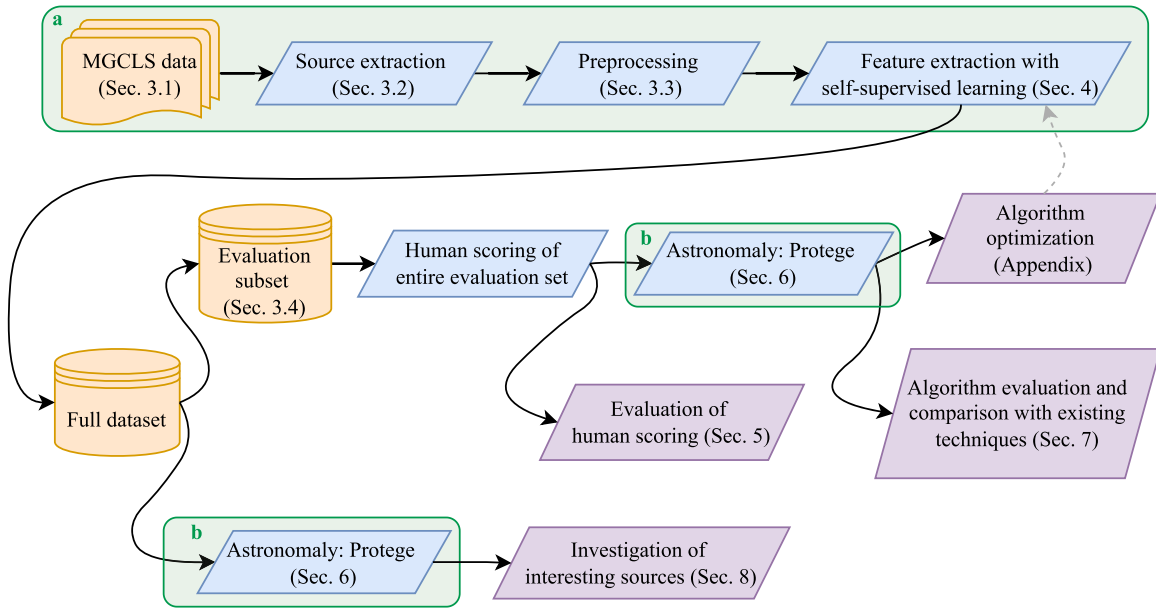
**Figure 1.** A summary of different problems in astronomy that can be solved with different machine learning techniques. Each example is illustrated with an example Uniform Manifold Approximation and Projection (UMAP) plot (see Section 4.4), which is a low-dimensional embedding of a high-dimensional data space where each point represents a different image or sample in the dataset.

learning to select the most useful optical transients to follow up spectroscopically in order to improve a photometric light-curve classifier.

Active learning can also solve a key problem with pure anomaly detection: many anomalous sources are not scientifically interesting and must be rejected or ignored. ASTRONOMALY (M. Lochner & B. A. Bassett 2021) is a general-purpose active anomaly detection framework that combines

unsupervised learning with minimal human labeling. ASTRONOMALY was successfully applied to  $\sim 4$  million optical galaxies in V. Etsebeth et al. (2024), discovering a large number of unusual sources, including mergers, strong lens candidates and several unidentified classes.

The anomaly detection algorithms that ASTRONOMALY relies on can be limited by the feature extraction method used, as alluded to in Section 2.1. This is illustrated in the bottom right of



**Figure 2.** High-level flowchart of the full methodology with references to relevant Sections. The section of the flowchart indicated by “a” outlines the procedure for extracting features from the raw data and is further expanded in Figure 3. After feature extraction, the analysis is divided in two, making use of a fully labeled subset of the data to evaluate performance and optimize algorithm hyperparameters and subsequently applying the methodology to the full dataset to locate and follow up high-scoring sources. In both cases, PROTEGE is used with the same settings and is described in more detail in Figure 8.

Figure 1 where the interesting sources occupy an extended region within feature space, rather than forming isolated patches near the edge. To solve this problem, we introduce ASTRONOMY: PROTEGE, heavily based on the algorithm from M. Walmsley et al. (2022a), which, instead of pure anomaly detection, leverages active learning to rapidly learn what the human trainer finds interesting. PROTEGE acts, to some degree, like a *recommendation engine*, similar to algorithms used by popular streaming services to recommend content to users based on their interests (see D. Roy & M. Dutta 2022, for a recent review). This form of active regression, when combined with the self-supervised features used in this paper, is highly effective at discovering truly scientifically interesting sources without requiring significant human time.

PROTEGE’s goal is to identify relatively rare objects that the trainer finds scientifically interesting, but span a wide range of morphologies. It is thus broader than the similarity search, which is targeted toward specific morphologies, and also differs from classification or clustering, which are also focused on finding common and related morphologies. The only criterion for PROTEGE is that the source in question is of interest to the trainer. Its power lies in its ability to efficiently select interesting sources, as will be shown in Figure 10. Most exciting, however, is that PROTEGE can make new discoveries, i.e., to find sources that the human trainer will recognize as interesting, but that do not appear to resemble those used in the training. Examples of this will be shown in Section 7.5.

PROTEGE depends on several algorithms working in concert. As we describe each of these steps, it is useful to consult the high-level flowchart in Figure 2, which includes references to the relevant text Sections for each functional block.

### 3. Data and Preprocessing

#### 3.1. The MGCLS Survey

The MeerKAT Galaxy Cluster Legacy Survey (MGCLS; K. Knowles et al. 2022) consists of 6–10 hr, full polarization

observations of 115 galaxy clusters made using the MeerKAT telescope in the  $L$  band (900–1670 MHz). The images have a typical sensitivity of  $\sim 3\text{--}5 \mu\text{Jy beam}^{-1}$  and are publicly available (K. Knowles et al. 2021). We make use of the high-resolution total intensity images for this work, which have a typical synthesized beam size of  $\sim 7''\text{--}8''$ . Based on the work of K. Knowles et al. (2022), we knew that this MGCLS dataset would contain a number of new and interesting sources. Follow-up work highlighted additional unusual radio galaxies, diffuse plasmas, and magnetized threads in the intracluster medium (L. Rudnick et al. 2021, 2022; F. de Gasperin et al. 2022; C. J. Riseley et al. 2022).

#### 3.2. Extracting Extended Sources

PYBDSF (N. Mohan & D. Rafferty 2015) is one of the most commonly used source extractors in radio astronomy. The algorithm works by first determining the rms of an image and then identifying contiguous islands of emission that are a certain threshold above the noise (typically  $3\sigma$ ). Then, PYBDSF decomposes the islands into Gaussian components, which are typically used as a source catalog. Alternatively, the islands can be decomposed using shapelets to better fit extended sources.

Automatically extracting extended sources is an active field of research and remains challenging for most source-finding algorithms (see M. M. Boyce et al. 2023 for an in-depth comparison). Often, the Gaussian components alone do not adequately describe extended sources. We propose an alternative approach by making use of the islands themselves as the input catalog, with the Gaussian components used as a proxy for complexity. Islands that require more Gaussians are more likely to be extended sources. The position of each source is defined as the center of mass of the island, which may or may not coincide with the true visual center of the source.

As the vast majority of sources in the original dataset are unresolved or only slightly resolved, we applied a cut requiring

each island to have at least four Gaussian components to be included in the catalog. Visual inspection of individual images suggested this cut included the majority of extended sources with complex morphologies.

For a more quantitative check of this procedure, we made use of the extended source catalog from K. Knowles et al. (2022) provided for two of the clusters: A209 and AS295. For A209, 20 out of 33 extended sources are recovered by our method, while for AS295, 19 out of 26 are recovered. Visual inspection confirms that the missed sources fall into only one of two categories. First, there were highly separated FRII galaxies where each component was described by only a single Gaussian; we did not consider such sources as “scientifically interesting,” so their exclusion is not a problem for this experiment.<sup>6</sup> The second category included large sources for which the human-identified center differs significantly from that detected by PYBDSF, so they were misidentified, but still included.

Examination of a larger sample of objects shows that occasionally, sources that we would consider scientifically interesting can be missed by the above procedure. The star-forming ring in the A548 field is a near perfect circle (see Figure 25 of K. Knowles et al. 2022) and worthy of further scientific investigation. However, it was missed because its circular nature meant PYBDSF could fit it with fewer than four Gaussian components. While the extended source extraction technique can certainly be improved, there will always be a trade-off between losing target sources and managing an increased number of contaminants.

### 3.3. Preprocessing

After obtaining a catalog using PYBDSF, we then extracted image cutouts of each source. One challenge is that source size varies significantly, and this property can inadvertently dominate ML algorithms (as found in I. V. Slijepcevic et al. 2023). As our main interest is in radio galaxy morphology, the apparent size of the source is unimportant. We thus used the islands derived from PYBDSF to estimate the minimum side of a square that fully contains the source. The cutouts were then resized using SKIMAGE (S. van der Walt et al. 2014) to a common size of  $128 \times 128$  pixels. Most sources are smaller than this and so are interpolated, but a few are larger than the window size and are instead downsampled.

As in M. Lochner & B. A. Bassett (2021), we applied sigma-clipping to the cutouts at  $3\sigma$  to remove potentially confusing background sources. We used the package ASTROPY to estimate the standard deviation of the image and then thresholded the image at  $3\sigma$ . We then fit a contour around the central source using the package OPENCV (since our sources are always in the center of the image by definition) and removed all other background sources. We investigate the efficacy of sigma-clipping in Appendix C.

Radio images tend to have high dynamic range, and astronomers often resort to scaling functions, such as  $\log$  or  $\text{asinh}$ , restricting the brightness range in an image to some fraction of the maximum or a combination of these techniques to better visualize faint sources. As we have already made use of sigma-clipping to remove bright background sources, we do not apply any transforms to the cutouts provided to the feature

<sup>6</sup> Other users could well determine that these were of interest, and have to tailor their extraction parameters accordingly.

**Table 1**  
Human Scores Used for the Evaluation Subset

Score	Example Sources	Number of Sources
1	Unresolved, artifacts	384
2	“Normal” FRI or FRII	439
3	WAT, double-double, disturbed	122
4	NAT, S-shaped, spiral, diffuse	62
5	X-shaped, turbulent, unknown	24

extractor. However, we do apply an  $\text{asinh}$  scaling to the images viewed by the human user and do not include sigma-clipping in these images. The larger context is often helpful to humans for source identification, but can negatively impact performance of ML algorithms due to the introduction of background sources. The use of transforms is discussed further in Section 9, and our choices are supported by the analysis in Appendix C.

Finally, we removed any source that results in a cutoff image due to its bounding box intersecting the edge of the original MGCLS image. This results in a final catalog of 6161 sources.

### 3.4. Evaluation Subset

From this catalog, we created a fully labeled subset of sources. We use this evaluation subset to compare PROTEGE against other approaches in Section 7. In the process of evaluating performance, we also used this subset to explore the impact of hyperparameters for both PROTEGE and traditional ASTRONOMALY (see Section 4.2 and Appendix C).

To create the evaluation subset, we selected the two fields that have an extended source catalog (A209 and S295), and then 18 other randomly selected fields in the MGCLS. After source extraction and preprocessing, this results in a catalog of 1031 sources, which were then scored by M.L. on a scale of 1 to 5. Section 5 discusses further what is meant by an interesting source and the subjective scoring procedure used throughout this paper. Table 1 shows a breakdown of the number of sources per score and a rough description of the types of sources corresponding to that score. We define any object with a score of 4 or 5 to be of high interest, which corresponds to 8.34% of the evaluation subset. The evaluation subset human scores are included as a separate column in the catalog<sup>7</sup> provided with this paper.

## 4. Self-supervised Learning for Feature Extraction

After preprocessing, the next step is to simplify the images in a process known as feature extraction. Features should be low-dimensional but still capture key information in the images to ensure successful application of the downstream algorithm. The choice of feature extractor can dictate the types of sources the algorithm will be sensitive to.

M. Lochner & B. A. Bassett (2021) used a simple approach to feature extraction that focused on shape parameters. Although initially designed for optical galaxies, this works remarkably well for radio galaxies (see Section 7). However, this simple approach is insufficient for our more ambitious goal of using active learning to group sources of interest in feature space, and so we make use of a more sophisticated feature extraction scheme.

<sup>7</sup> <https://github.com/MichelleLochner/mgcls.protege>

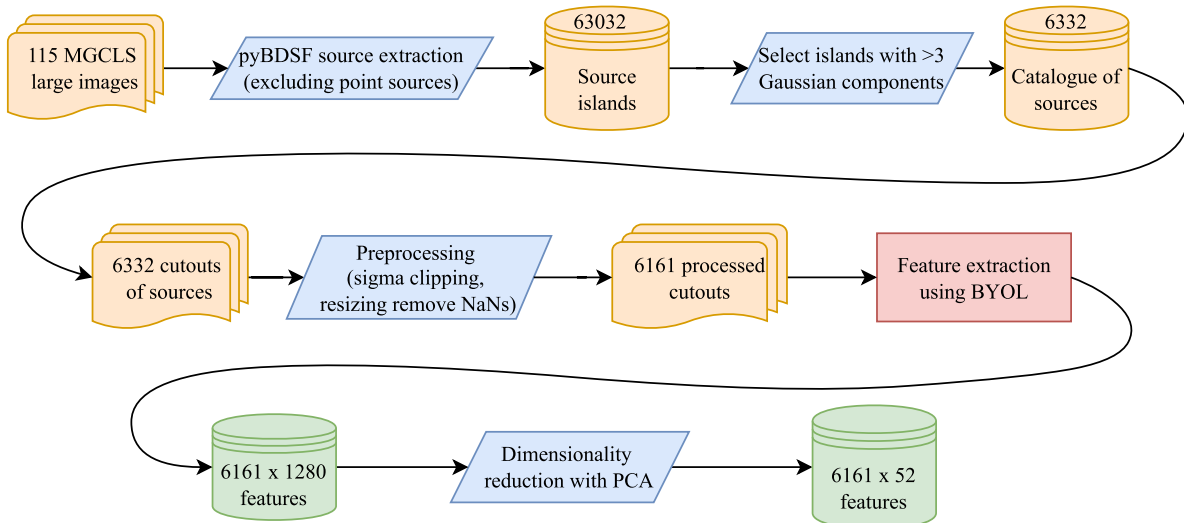


Figure 3. Flowchart of the preprocessing and feature extraction process.

The deep learning algorithm BYOL (J.-B. Grill et al. 2020) has emerged as a powerful method to automatically learn a low-dimensional feature space such that similar looking objects should have similar features. K. Mohale & M. Lochner (2024) showed that BYOL can be used to extract meaningful representations of both optical and radio galaxies. These features were used to automatically cluster together galaxies with similar morphology, as well as to detect anomalous galaxy types such as mergers. We used the exact same technique in this work to extract features for the MGCLS sources.

The full procedure required to extract features from the MGCLS images is fairly complex, involving many steps. Figure 3 shows a flowchart of the process, including the preprocessing steps described in Section 3, as well as the feature extraction methodology described in the Sections below.

#### 4.1. BYOL

BYOL is a self-supervised learning technique that aims to automatically learn useful features for a dataset without requiring any labels for the data. BYOL operates on the key insight that if an image is *augmented*, in other words changed in some trivial way such as applying rotation, zooming, or blurring, the features that describe that image should not change significantly. The goal is therefore to train a neural network to predict features such that augmented pairs produce similar values.

BYOL consists of a pair of identical neural networks: an online and a target network. The goal of the online network is, given an augmented view of an image, to predict the output of the target network that has a different augmented view of the same image. The two networks learn in such a way that they should converge to the same representation of the data. The second-to-last layer of weights of the online network is used as features.

Any network architecture can be used for the online and target networks. For this work, we use EfficientNet-B0 (M. Tan & Q. Le 2019) due to its lightweight architecture, which makes it computationally efficient and suitable to the relatively small dataset size.<sup>8</sup> K. Mohale & M. Lochner (2024) showed that initializing the network with weights learned from training on

the terrestrial dataset ImageNet (O. Russakovsky et al. 2015), instead of random weights as is commonly used, results in much faster convergence and better performance. We thus initialize the EfficientNet network with these weights and fine-tune the network to our data using the packages PYTORCH (A. Paszke et al. 2019) and BYOL-PYTORCH (X. Chen & K. He 2020).

It is the second-to-last layer of EfficientNet, called `avg_pool`, that is used as features. In normal supervised learning, the last layer would be a final fully connected layer that aims to predict the class or quantity the network is being trained to predict. For self-supervised learning, which aims to extract useful features, this layer is ultimately discarded, and the output of the preceding layer is used. For EfficientNet-B0, this results in 1280 features.

#### 4.2. Hyperparameters

BYOL has several hyperparameters and training choices that can impact performance. In addition, the choice of augmentation strategy is critical for good results and is expected to be different for astronomical data compared to those used for terrestrial image datasets. Fortunately, I. V. Slijepcevic et al. (2023) performed an extensive study of these choices on a radio astronomy dataset. Although I. V. Slijepcevic et al. (2023) were primarily interested in classification performance, it is reasonable to expect similar choices will work well for general feature extraction and other downstream tasks. We thus do not perform an extensive analysis of hyperparameters and augmentation choices, and rather use a similar training procedure as I. V. Slijepcevic et al. (2023), highlighting variations where appropriate. Appendix C has a detailed discussion on the impact of hyperparameter choices, using the evaluation subset (Section 3.4).

We use Adam (D. P. Kingma & J. Ba 2014) as an optimizer, with a batch size of 32 and a base learning rate of 0.0005 (which is scaled by batch size/256). We find that training for just 100 epochs is sufficient for excellent features (see Figure 24 in Appendix C). We use almost identical augmentations as those of I. V. Slijepcevic et al. (2023), as shown in Table 3. The only difference is the size of the center crop, since

<sup>8</sup> While investigating other architectures would be useful, the high computational cost made such an analysis prohibitive for this work.

our images are rescaled such that the source should occupy most of the cutout.

#### 4.3. Dimensionality Reduction

As described in Section 4.1, the number of features from the network is 1280. This is generally too high-dimensional for most downstream applications so must be reduced. While there are many approaches to this problem, a common solution is principal component analysis (PCA; H. Hotelling 1933). PCA decomposes a dataset into a new coordinate system such that each orthogonal component aligns with directions of maximal variance. After the decomposition, the majority of components can be discarded, retaining only those that describe the most variance in the dataset: the principal components. Dramatic reductions in dimensionality are possible while losing only a small amount of information. We apply a threshold of 95% variance, reducing the original 1280 features to just 52 for the full dataset.

#### 4.4. Visualization

For unsupervised learning tasks, it is useful to visualize the extracted feature space to gain intuition about the effectiveness of the method and where classes of objects may lie in relation to each other. Naturally, this is very challenging with a high-dimensional space. As is standard practice in the field, we make use of the algorithm UMAP (L. McInnes et al. 2018), implemented using the UMAP-LEARN python package (T. Sainburg et al. 2021). The aim of visualization algorithms like UMAP is to embed the high-dimensional feature space into a two-dimensional one, while preserving relative distances between data points. This means that after extracting features from the images, reducing them using PCA and further embedding them using UMAP, similar looking galaxies *should* appear close together in the UMAP plot.

UMAP achieves this by constructing a fuzzy topological representation of the features and then learning a lower-dimensional approximation to this via an optimization algorithm. While UMAP plots can be useful tools for visualization and can actually work on high-dimensional spaces, the algorithm is not usually used as a dimensionality reduction algorithm directly. This is because a linear algorithm like PCA better preserves global structure than embedding algorithms like UMAP, which can occasionally create artificial clusters during the embedding process. This can be mitigated to some degree by careful choice of hyperparameters. Following K. Mohale & M. Lochner (2024), we set the number of neighbors to 15 and the minimum distance to 0.01 for all UMAP plots.

### 5. What Makes a Source Interesting?

Since our goal is to identify scientifically interesting sources, i.e., those worthy of further study (hereinafter “interesting”), it is worth considering what that really means in practice. In the simplest terms, and with no intent to be frivolous, an “interesting” source is whatever an individual researcher or community decides it is. This definition is not trivial because it strongly influences awards of telescope time, funding proposals, reviews of journal articles, and the focus of theoretical and simulation work. In the case of large surveys, defining what is “interesting” is imperative, so that limited time and resources can be focused on the most promising sources and scientific questions. Since objects of scientific interest will differ between

different groups and over time, it is not useful to search for algorithms to find “objectively” interesting sources, only those that are relevant to the current effort.

here are currently two complementary approaches to this challenge. First is to explore algorithms that select samples of sources from a larger group, and to then judge how well each algorithm has preferentially selected interesting sources. The use of image complexity (G. Segal et al. 2019, 2023) is one such successful method. The second approach is to predefine what is “interesting”, in terms of identifying sources with a specific morphology, and then use *similarity* algorithms (e.g., M. Walmsley et al. 2022a) to extract them from a large dataset.


An intermediate approach is used by ASTRONOMALY (TM. Lochner & B. A. Bassett 2021), which combines anomaly detection with some active learning based on human input to reject sources the user identifies as not “interesting.” However, the combination of traditional anomaly detection and subsequent active learning cannot handle the situation where “interesting” sources are buried in feature space (see Section 2).

PROTEGE is intended to take the approach of ASTRONOMALY to a different level—to recognize the broad range of morphologies that its scientist trainer finds “interesting,” and then select those rare cases from the enormous datasets produced in modern surveys. This breadth in morphology is the essential challenge—can a machine find similarities in feature space that allow it to mirror the interests of its trainer? Can it go even further and find sources with completely different morphologies from those in its training, that are also “interesting”? This is the challenge for which PROTEGE, described in Section 6, is designed to address.

To accomplish this, the trainer (or “oracle,” as used in the active learning literature; E. Mosqueira-Rey et al. 2023) provides input by assigning a numerical score of 1–5 on how “interesting” each of a given group of sources appears. The sources are therefore not labeled, but given a numerical score that can be used by the algorithm. Figure 4 shows example sources from this work, along with their scores.

Here, author M.L. gave some of the unusual morphologies a score of 5, because they were considered worthy of further study. These included SAURON, a pair of galaxies with extremely bent jets, and some filamentary structures. Other sources shown in Figure 4 appeared rare and interesting, but were of known types and thus given a score of 4. These included spiral galaxies, head-tail galaxies and diffuse emission. Radio galaxies with slightly disturbed morphologies were scored a 3, typical radio galaxies a 2, and pointlike sources a 1.

Independent of the success of the algorithm, in order to be scientifically useful, there must also be some agreement on “interesting” sources among multiple users. To investigate the sense to which this is true in the current case, we performed a small experiment comparing the scores of the two authors of this paper, M.L. and L.R. Note that M.L. is relatively new to the field of radio galaxies, while L.R. has decades of experience. They had several conversations, on different projects, of which sources appeared most “interesting.” Thus, there was some common understanding, although not a prescription for finding specific morphologies. They each assigned a score to a random subset of 100 sources, and their scores were compared, as shown in Figure 5. Perfect agreement is neither expected (see E. Savary et al. 2022 and K. Rojas et al. 2023 for differences in expert labeling of strong lens candidates) nor necessarily desirable. We found a reasonable agreement between the two scorers, with a




**Figure 4.** Examples of sources with their human-assigned score. Each image has been normalized using an *asinh* scaling with a threshold at 90% of the maximum brightness to highlight faint features. The beam size for each image ( $\sim 7''6$ ) is indicated with a bar in the lower left.

nonparametric Pearson correlation coefficient of 0.70, and only 12% of the source scores differing by more than one point. This gives us confidence that there is sufficient agreement on what constitutes an “interesting” source to make the approach meaningful.


However, it is also instructive to explore where there was significant disagreement between the two scorers, as shown in Figures 6 and 7. The basis for scoring was the same as described for Figure 4, but here the scorers disagreed on how well further study was justified. For example, if one of the scorers thought that an unusual morphology was due simply to projection effects, they would assign a lower score. Such disagreements form an excellent basis for scientific discussion, as the community seeks to reach a common understanding. Although not pursued further here, it would be interesting to have a variety of different users each train PROTEGE, and compare the resulting samples. Another opportunity, as discussed below, is to think about PROTEGE’s trained scoring as an oracle in its own right, and look for novel examples of what may be “interesting” to the human users, that they may have missed.

## 6. Finding “Interesting” Sources

ASTRONOMALY’s flexibility as a general framework makes it ideal for application on a novel and unlabeled dataset



**Figure 5.** Comparison between the two lead authors’ scoring for a random sample of 100 sources. The error bars represent the standard deviation of author L.R.’s scores for a given value of M.L.’s scores.




**Figure 6.** Examples of sources author L.R. scored high while M.L. scored low. The images have been normalized with *asinh* scaling.

like MGCLS. By relying on an initial feature extraction step, ASTRONOMALY can work with essentially any type of astronomical data. We use the extracted features from a trained BYOL model, as described in Section 4, with the aim of finding sources that we, the authors, find “interesting.”

As discussed in Section 2, M. Walmsley et al. (2022a) found that for deep features, the traditional anomaly detection algorithms used in M. Lochner & B. A. Bassett (2021) did not perform well. They effectively bypassed anomaly detection entirely and used a regression algorithm, specifically Gaussian processes (C. E. Rasmussen & C. K. Williams 2006), coupled with a traditional active learning approach, to rapidly locate the objects of interest to the user.

While we did explore the use of traditional anomaly detection algorithms in this work (Section 7), we found that the approach developed in M. Walmsley et al. (2022a) has the highest performance for this dataset and so is what we will focus on. Because this approach is philosophically different



**Figure 7.** Examples of sources author M.L. scored high while L.R. scored low. The images have been normalized with *asinh* scaling.

from traditional anomaly detection, we give it a new name to avoid confusion, including it as a module within ASTRONOMALY. Since the method aims to learn what is “interesting” to the user (not strictly speaking anomalous), and does this by iteratively requesting what sources are “interesting,” we name it PROTEGE.

It is worth reiterating that, unlike anomaly detection algorithms, PROTEGE can be trained to find any source of interest, even if it is common (for example, FR II radio galaxies). However, for a class that has many examples, a trained supervised learning algorithm will generally outperform PROTEGE. The goal of PROTEGE is rather to act like a recommendation engine for new discoveries, fine-tuning its search for sources a user may find “interesting” based on that user’s input.

### 6.1. PROTEGE

PROTEGE essentially aims to learn a function across feature space that predicts what score a user will give each source. As with other active learning algorithms, it uses an acquisition function to prioritize which sources, if scored by the trainer, would be most informative in training the algorithm. Different active learning approaches use different acquisition functions.

Figure 8 shows a flowchart broadly outlining the algorithm. The first step is to select an initial set of sources to query. These will be shown to the user (or “oracle”), using ASTRONOMALY’s front end interface, to be scored from 1 to 5 based on how “interesting” the user thinks the source is. While the range of scores is arbitrary, a breakdown of five levels is sufficient to allow a clear gradient for an algorithm to learn.

The goal of the initial set of sources presented in the queries for human scoring is to broadly sample the full feature space. We take advantage of PCA by selecting the 10 sources equally spaced from a list ordered by the first component (i.e., the first feature). This differs from the approach in M. Walmsley et al. (2022a), where the initial sample was selected randomly. We found a PCA-based approach higher performing and more consistent when the training was repeated.

In the next step, PROTEGE predicts the score of the entire dataset based on the training data provided so far, and selects the next set for human scoring. We use a Gaussian process

(GP) for predicting the user score because it provides an estimate of the uncertainty, which can then be used to calculate the acquisition function. A Gaussian process effectively models the data using a Gaussian probability distribution in function space, predicting both the mean and standard deviation of the target quantity. We use the same acquisition function used in M. Walmsley et al. (2022a): the *expected improvement* (EI, J. B. Mockus & L. J. Mockus 1991; D. R. Jones et al. 1998). This function, described in detail in M. Walmsley et al. (2022a), includes a hyperparameter  $\epsilon$ , which can be used to tune the trade-off between exploitation (low values of  $\epsilon$ ) and exploration (high values of  $\epsilon$ ). We used a value of  $\epsilon = 3$  (halfway between our minimum and maximum expected user scores), which encourages the algorithm to explore regions of high uncertainty. This makes it less likely to get fixated on a particular region, and increasing the chances of new discoveries. An analysis described in Appendix C showed that as long as this parameter was higher than 0.5, it did not make a significant difference to which sources were discovered.

After predicting the user scores for the entire dataset and calculating the acquisition function, PROTEGE then selects the 10 sources with the highest values of the acquisition function and presents them as the next query. While any number of sources can be scored in each iteration, we find that 10 is a good balance between providing a reasonable amount of training data for the GP and updating rapidly enough to locate high human-scored sources quickly.

This process is repeated until the trainer determines that the algorithm is trained well enough and is presenting many more sources of interest to them than uninteresting sources. The list of sources can then be sorted by the most recent predicted score from the GP and is now ready for user inspection. Figure 9 illustrates an example of the interactive approach of PROTEGE. Appendix D investigates the performance of the GP algorithm in predicting user scores.

## 7. Performance Evaluation

We use the evaluation subset described in Section 3.4 to compare our method with alternative approaches. The goal was to determine whether PROTEGE, traditional ASTRONOMALY or a non-ML approach would best recover “interesting” sources with minimal human labeling.

### 7.1. Complexity Comparison

Several recent works (G. Segal et al. 2019, 2023) have shown that anomalies in radio images, and thus sources that might be considered “interesting,” can be effectively detected using measures of complexity. Here, we use a simplified measure of complexity, described below, as a demonstration of what can be achieved without additional processing steps.

The number of Gaussian components in each island identified with PYBDSF was taken as a measure of source complexity. Sources with the most components are given higher scores, extending the idea behind the minimum four-component cut used to create the catalog (Section 3.2). This simple procedure provides a useful baseline to judge whether the ML approach of PROTEGE is sufficiently superior to a simple catalog ordering to warrant the additional computational cost.

Our simple complexity measure does not, however, provide features that can be used by active learning algorithms to target

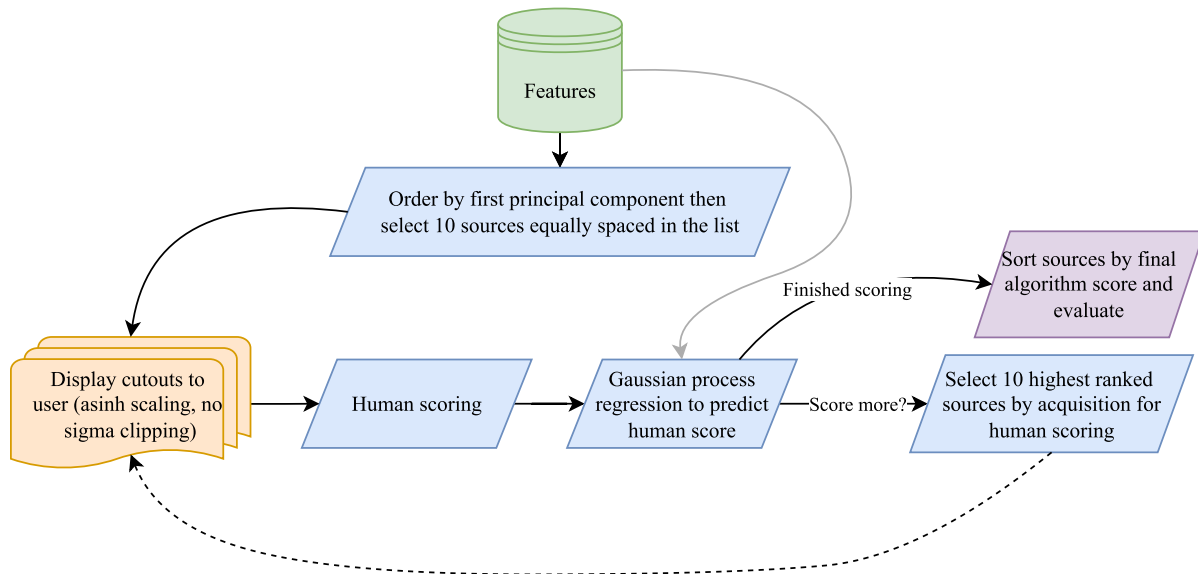


Figure 8. Flowchart of PROTEGE.

specific morphologies and cannot be used for downstream tasks such as disregarding artifacts, applying classification or similarity searches.

### 7.2. Ellipse-based Feature Comparison

In M. Lochner & B. A. Bassett (2021), ASTRONOMY originally used hand-crafted morphology features based on fitting ellipses to isophote contours. This was followed by a traditional anomaly detection algorithm (isolation forest; F. T. Liu et al. 2008) combined with active learning. Although originally applied to optical galaxies, it is reasonable to assume a similar approach would work for radio sources, so it was the first approach we used. However, inspired by the high performance of deep-learning-based feature extraction used in M. Walmsley et al. (2022a), we moved on to the methodology described in Section 4. The ellipse-based approach, described in more detail in Appendix B is included here only as another baseline against which to compare. As there is some randomness in the isolation forest algorithm, we run the full ASTRONOMY pipeline 10 times to estimate the variance in performance on the evaluation subset.

### 7.3. BYOL-based Feature Comparison

We compared the previous two approaches to our new method, BYOL-based feature extraction (Section 4.1) combined with PROTEGE (Section 6.1).

One advantage of a self-supervised algorithm like BYOL is that a larger, unlabeled dataset can be used to train the algorithm, improving the features extracted. To investigate the effect of training set size on performance, we considered three different subsets:

1. The evaluation subset (1031 sources).
2. A random subset of 3500 sources, excluding any source in the evaluation subset.
3. The full dataset (6161 sources).

Since the augmentations BYOL relies on are applied randomly, we ran trained BYOL 10 times on each subset, allowing an estimate of the model variance.

BYOL trains by minimizing the loss—the mean-squared error between the online and target representations. Figure 24 in Appendix C shows the loss as a function of epoch for each of these three subsets. Not surprisingly, increasing the amount of data improves the loss significantly.


Each iteration produces a different feature set, trained using the same subset of data but different random augmentations. The features for the evaluation subset are then passed to PROTEGE and its performance evaluated. PROTEGE is trained using the procedure outlined in Section 6.1, in batches of 10, until 100 sources have been given human scores.

Appendix C explores using the original ASTRONOMY approach with BYOL features, which performed poorly, so we exclude it here.

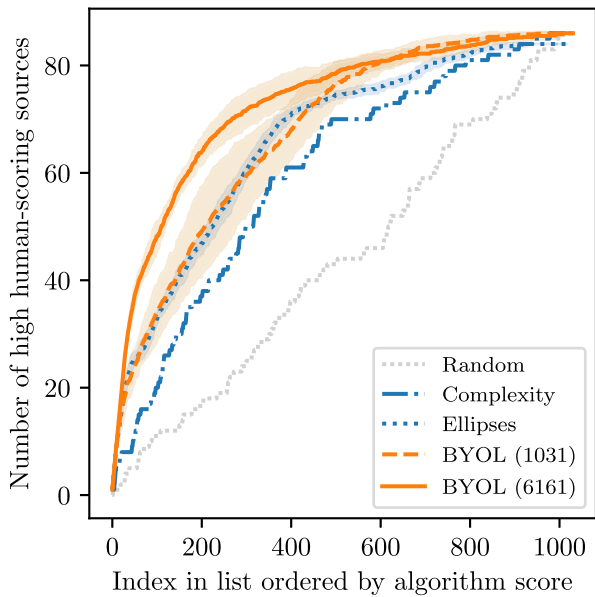
### 7.4. Results from the Evaluation Subset

Figure 10 shows the recall as a function of index in an ordered dataset. In other words, for any given number of sources ( $N$  on the  $x$ -axis) starting from the highest ranked, the  $y$ -axis reports the number of those sources with a human score of 4 or 5. The steeper this plot, the better the algorithm is at recovering “interesting” sources, without needing a user to look through a vast number of objects. A random ordering of sources results in a linear growth in the number of sources with human scores of 4 or 5, as expected, up to the full number (86) of such sources. Ordering by our complexity proxy performs surprisingly well, indicating that such methods have merit in improving the efficiency of identifying sources of interest.

The ellipse-based method described in Section 7.2 also performs well, despite being developed for optical galaxies. It is actually comparable to the algorithm using BYOL features combined with PROTEGE in our initial training experiment. However, BYOL + PROTEGE improve significantly as the BYOL algorithm is trained on more data. The random sample of 3500 was excluded from the plot for clarity, as the curve is almost identical to that of the full dataset. This implies that these results are not due to “data leakage,” as BYOL was trained on a subset excluding the evaluation subset. However, the performance improves slightly and the variance decreases when using all available data. We thus expect the algorithm to



**Figure 9.** Demonstration of the interactive training procedure. In a “query,” a set of objects is shown to the “oracle” (the user), who provides an interest level score for each object (“1” for compact sources or pure noise, to “5” for highly unusual morphology). Images are normalized using *asinh* scaling with the brightness truncated to 90% of maximum. The active learning algorithm described in Section 6 is then applied, the objects are sorted by acquisition score (which prioritizes which objects will provide the algorithm with the most useful information), and the next set of objects are queried. Within just a few iterations, the algorithm hones in on high human-scoring sources.



**Figure 10.** Number of sources author M.L. scored 4 or 5 in the evaluation subset when ordered by the predicted score from several algorithms: random ordering (*Random*), ordering by number of components in an island from PYBDSF (*Complexity*), ordering by ASTRONOMALY using ellipse-based features with an anomaly detection algorithm and active learning (*Ellipses*), and two cases of ordering using the BYOL features combined with PROTEGE for active learning (*BYOL*). The numbers in brackets for BYOL refer to the number of sources used to train the algorithm. The 3500 source subset is excluded due to the curve’s similarity to the full dataset curve. The envelope around the lines indicates the standard deviation from 10 repetitions. As it results in the steepest curve, BYOL combined with PROTEGE outperforms all other methods when sufficient training data is used.

perform best on the full dataset when BYOL is also trained using the full dataset. Note that there is no way to similarly scale the ellipse-based method.

To make the practical application of Figure 10 clearer, we consider a simple thought experiment. If a user were willing to look through all 1031 sources, they would naturally find all of the high-scoring sources. However, looking through all of the sources is already tedious and will quickly become impossible, as dataset sizes increase. If instead, a user only had the time and energy to view 100 sources, they would discover only eight high-scoring sources if the sources were chosen at random, reaching 20 using a list ranked by complexity measure or 32 using traditional ASTRONOMALY with ellipse features. But using PROTEGE with BYOL, an impressive 56% (48 out of 86) sources were found, even though  $<10\%$  (100/1031) of the sources had been checked. Given that the original full catalog consists of 608,595 sources, the chance of finding a high human-scoring source looking through just 100 examples from it is extremely small ( $\sim 0.084$  in a random sample of 100).

Figure 11 shows two UMAP visualizations (described in Section 4.4) of the BYOL features (using one of the full dataset runs) for the evaluation subset. The top panel reveals that “boring” sources (given a human score of 1) are grouped together near the edges, making it very easy for PROTEGE to discard them. There are also fairly pure regions containing sources of “medium interest” (human scores of 2 or 3) that can be ignored. While some sources with a human score of 4 or 5 are embedded among others, most are found clustered in a few key regions. These correspond to the high-scoring regions in the lower panel, where points are colored by PROTEGE score. These plots demonstrate two key points: BYOL is

effective, but not perfect, at grouping together sources the authors find “interesting,” and PROTEGE is able to find the regions where these sources tend to reside, even with only 100 training labels provided.

### 7.5. Sources PROTEGE Finds, but Humans Missed

A key question to explore is whether PROTEGE can find sources that are “interesting,” but were initially missed by humans. Figure 12 shows four examples that were scored high by PROTEGE but low by the authors. Some of these sources contain subtle diffuse structures that were missed, which can happen when viewing large samples of images. Source “a” in Figure 12 contains a faint, diffuse structure (halo), located in the center of cluster AS295, with a remarkable long filament. Source “b” is a radio galaxy that appears typical at first glance, but upon closer inspection reveals transverse filamentary structures similar to the “ribs” described in L. Rudnick et al. (2021). The asymmetry of source “c” warrants a second look from a science perspective.

These examples demonstrate the power of ML to not only locate rare sources in large datasets, but also to locate sources that may have been missed by humans and deserve a second look. The less-biased nature of ML algorithms, coupled with their ability to detect subtle structures not initially obvious to human users, provides a powerful tool for scientific discovery.


We also note that the unusual and scientifically interesting source SAURON (originally detected using the ellipse-based method) appears second in the list when ordered by PROTEGE score (using BYOL trained on the full dataset as feature extractor), placing it in the top 0.2% of the data.

The current algorithm is not perfect. Source “d” Figure 12 is a high-scoring source that we considered uninteresting, even upon a second look. It is unclear why this source appears in the same area of feature space as the high human-scoring sources. A similarity search (Section 2) revealed that all nearby sources indeed look similar and would also be given a low human score. This is an indication of an imperfect feature space, which could perhaps be improved with more training data and better preprocessing. However, there were very few examples of such sources (in fact, it was a challenge to find one), and we agreed with the majority of PROTEGE’s top recommendations.

### 7.6. Sources with High Human but Low PROTEGE Scores

Figure 11 shows that some sources with high human scores are missed by PROTEGE, primarily because they end up with similar features to sources the authors were not interested in. A closer look at a high-human-scoring source that PROTEGE ranked low can be seen in Figure 13, with its eight nearest neighbors, found using a similarity search. This source has faint diffuse emission, which is largely invisible with the linear scaling used by PROTEGE, but obvious with *asinh* scaling used for reinspection. Thus, it is the preprocessing that is ineffective. The majority of “missed” sources fall into this category, indicating that improving the preprocessing could be key to improving the overall method (see Section 9 for further discussion).

A similarity search for the source in Figure 14 shows that its neighbors have a similar morphology to other sources that are mostly slightly resolved. This is one of only a few examples where the source of interest is simply not very well-resolved and easy to mistake for other types.



**Figure 11.** UMAP visualization of the BYOL feature space. Each point represents a distinct source in the evaluation subset. Upper panel: points are colored by their underlying human score, showing that the sources with a score of 1 collect toward the edge of feature space while the sources with scores of 4 and 5 form clusters embedded throughout feature space. Lower panel: the same sources colored by the predicted PROTEGE score, after training on 100 labels selected by the acquisition function. It can be seen that, even with only a few labels, the algorithm learns which regions of feature space to focus on, although naturally a few high human-scoring sources are missed.

On balance, the overall high efficiency of PROTEGE in identifying the most “interesting” sources (Figure 10), and the discoveries that it makes that are missed by humans, shows that it is a very important tool as datasets become extremely large. At the same time, it is important to acknowledge that it is imperfect, missing some sources that humans would rank as “interesting.”

### 8. Recommendations from PROTEGE


So far, we have demonstrated, on a sample of 1031 sources, how PROTEGE can be an effective tool for selecting scientifically “interesting” sources from a large dataset with minimal human training. We now apply PROTEGE to the entire dataset of 6161 extended sources in the MGCLS and look at the sources it considers most interesting.

We applied exactly the same procedure described in Section 6.1, using the BYOL features trained on all 6161 extended sources.<sup>9</sup>


<sup>9</sup> Note that this includes the 1031 sources in the evaluation subset, but those human scores were not used. PROTEGE was trained fresh on the full dataset, and there is no separate testing or validation phase.

Author M.L. iteratively trained PROTEGE in batches of 10 until 400 sources had been given human scores. PROTEGE then gave final scores to the entire dataset. After this point, the proportion of sources PROTEGE was showing that M.L. found “interesting” was relatively low. This illuminates an important aspect of how PROTEGE training works. In the beginning stages of training, for the first 200 sources, ~80% of the objects PROTEGE presented were found to be “interesting” (human score of 4 or 5). This is much higher than for the evaluation subset simply because there are more “interesting” sources to find. For the next 100 sources, only 37% were interesting, and that number dropped dramatically to 9% for the last 100 sources that M.L. scored. This is an important byproduct of the training process. If the user finds that the fraction of “interesting” sources in each iteration decreases significantly, then further training is likely of marginal value.

The primary goal of PROTEGE is to dramatically reduce the number of sources that must be examined in a very large database to find the ones that are scientifically “interesting.” After training, PROTEGE’s performance on this task can be evaluated in the current experiment by comparing images of the




**Figure 12.** Four examples of sources given a high score by PROTEGE and a low score by humans. Upon reinspection, we agree that the sources in (a), (b), and (c) are worth further consideration, as described in the text, but not (d). All images are produced with *asinh* scaling. To reveal faint structure, the maximum brightness is truncated to 10%, 10%, 90%, and 90% of the maximum for panels (a), (b), (c), and (d), respectively.



**Figure 13.** Similarity search of feature space for a high human-scoring source for which PROTEGE gave a low score. The source of interest is displayed in the center, with its eight nearest neighbors (by Euclidean distance) around it. These images have not been rescaled and are identical to the input to the feature extraction algorithm. Viewing the central source with *asinh* scaling reveals the faint structure around the central core, so this source was given a low score purely because of the preprocessing scaling. Nearly all high human-scoring sources that were missed by PROTEGE fall into this category.

top 260 (4.2%) scorers (Figure 15) with 260 randomly selected sources shown in Appendix A (Figure 21) for comparison. The differences are dramatic, showing that PROTEGE has been very successful in its primary goal. A catalog (see footnote 7) of coordinates, PROTEGE scores, and human scores (where available), as well as an additional version (see footnote 5) of Figure 15 for all 6161 sources can be found at the links provided in Section 1, for follow-up work.

PROTEGE’s high-scoring sources show a very diverse set of morphologies, reflecting the diversity of the human training



**Figure 14.** Similarity search of feature space for a second high human-scoring source (center) for which PROTEGE gave a low score, showing its eight nearest neighbors. This source is similar to other double core sources, so it is difficult to see how the algorithm could have distinguished between them.

and demonstrating that PROTEGE achieved its key design goal. In the rest of this discussion, numbers in parentheses refer to the high-scoring sources in Figure 15. Numbers in bold refer to sources highlighted in K. Knowles et al. (2022), as discussed further below.


Prominent among the high-scoring sources from PROTEGE are spiral galaxies (4, 16, 31, 47), head-tail or narrow-angle tailed sources (26, 44, 45, 106). A variety of radio galaxies exhibiting diffuse and disturbed lobes (0, 2, 6, 8) are also seen, some far from the centers of the target clusters (32, 83, 91). Particularly dramatic are the “kissing WATs” featuring apparently two wide-angle or narrow-angle tailed galaxies with extremely bent jets (60).

PROTEGE is also fairly sensitive to diffuse emission, with several high-scoring sources at their cluster centers (1, 10, 19, 22), as well as sources significantly displaced from the centers (11, 21, 25). These sources, with the exception of source 11, were seen and noted in Table 4 of K. Knowles et al. (2022), although they were not among the highlights of the paper. A newly discovered diffuse source (48), was found close to the center of cluster J0616.8–4748. It has an unusual structure, with an apparent hole on one side. The utility of PROTEGE retrained to specifically detect diffuse emission will be explored in an upcoming paper (V. Etsebeth et al. 2025, in preparation).

The enormous diversity among the high-scoring sources is also reflected in UMAP feature space, shown in Figure 18, where the high-scoring sources span a large region. The ability to recover such diverse sources distinguishes PROTEGE from other algorithms, as shown in Figure 1.

### 8.1. Case Studies

To explore in more detail the selections that PROTEGE has made, we present case studies of two different source types frequently considered to be “anomalous”—X-shaped, and circular shaped.



**Figure 15.** Highest-ranked sources according to PROTEGE. Images have been normalized using an *asinh* scaling with a threshold at 90%, and the beam size ( $\sim 7''/6$ ) is indicated with a bar in the lower left. Known sources that have published Figures from K. Knowles et al. (2022) and related literature are highlighted with a yellow border, while new recommendations discussed in the text have a red border.



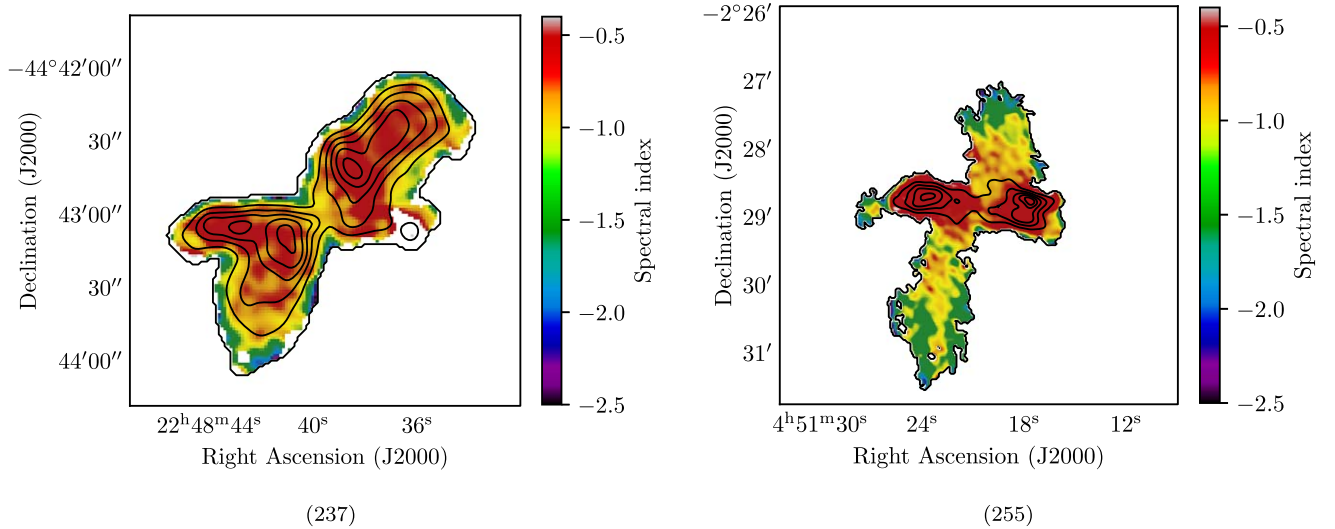


Figure 15. (Continued.)



**Figure 16.** X-shaped radio galaxies scored highly by PROTEGE. The flux density in these images is truncated to 99% of maximum for the left and right images and 95% for the central image, to highlight faint features. The potential host, where one is identified, is shown as a white circle, and the synthesized beam ( $\sim 7''$ ) is shown as a gray circle in the lower left of each image.



**Figure 17.** Spectral index for source 237 (left) and 255 (right), only including pixels with flux density at least 10 times higher than the rms for the image ( $4.00 \mu\text{Jy beam}^{-1}$  and  $7.36 \mu\text{Jy beam}^{-1}$ , respectively) and for which  $\sigma_\alpha/\alpha < 0.2$ . The spectral index is estimated in-band using least-squares fitting applied to the 12 sub-band images. Contours are drawn at five evenly spaced flux density levels between the minimum and the maximum flux, with background sources excluded.

### 8.1.1. Case Study: X-shaped Sources

X-shaped radio galaxies (XRGs) are a fairly rare class of radio galaxies (see G. Giri et al. 2024 for a recent review). PROTEGE was not trained to find this specific morphology, so the sources shown here (237, 255, 110) were identified visually among the high-scorers.

The most important lesson from this case study is that even among the subset of X-shaped sources, there are significant differences in morphology, and probably the corresponding physical processes at play. This likely reflects the fact that multiple mechanisms can lead to a general X-shaped source (see, e.g., D. Patra et al. 2023).

Source 237 (Figure 16, left panel) is a typical XRG, similar to PKS2014-55 in W. D. Cotton et al. (2020), with long straight lobes, indicating the likely orientation of the jets, and short transverse extensions/wings at the back of the lobes, forming an X-shape.<sup>10</sup> Using the observed magnetic field structure, W. D. Cotton et al. (2020) showed that PKS2014-55 was best understood in terms of hydrodynamic backflows from the ends of the jets, back toward the host galaxy, where they are deflected into an X-shape by the large, oblique hot gas halo of the host galaxy. The structure of 237 appears consistent with

<sup>10</sup> A possible host for 237, has a DELVE (A. Drlica-Wagner et al. 2022) photo- $z$  of  $0.239 \pm 0.006$ , placing it outside the target cluster for this field, AS1063, with a redshift of 0.348.

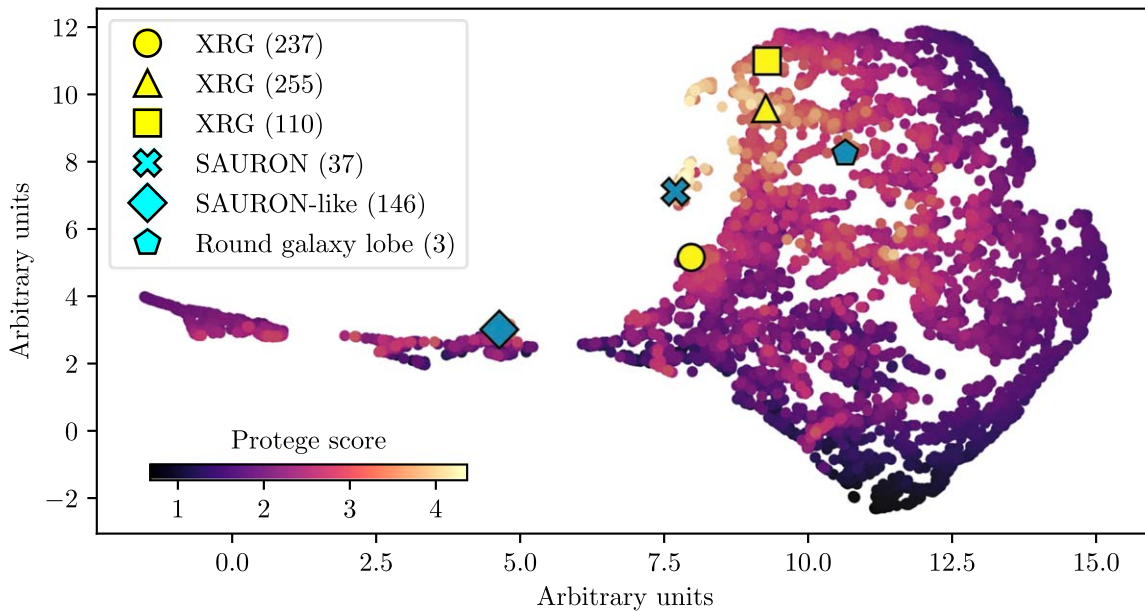
this picture. We also find that 237's radio spectra are fairly uniform throughout the source, similar to that of PKS2014-55, with an average value of  $-0.79 \pm 0.02$  (see Figure 17).

In contrast, source 255 in the middle panel of Figure 16 has a quite different morphology. One set of structures, in the north-south direction, are significantly fainter and longer than the east-west wings. The long, faint structures are not perfectly aligned with the central source, and the west wing displays a very unusual double structure in its lobe. The radio spectra are also quite different than those of 237 (Figure 17). The faint structures steepen away from the host, and on average have much steeper spectra than those of the short wings ( $-1.0$ , compared to  $-0.7$ ). These findings are all consistent with a physical model where the jets were originally oriented north-south, and have steepened through synchrotron losses, and then were re-oriented in the east-west direction. Such reorientations could occur, e.g., during a binary black hole merger (G. Giri et al. 2024).<sup>11</sup>

Source 110 (Figure 16, right panel) has an overall X-shaped morphology, but is unusual in the asymmetry in length and brightness of its opposing structures.<sup>12</sup> It bears some resemblance to the peculiar XRG of S. Kumari et al. (2024), which has a

<sup>11</sup> Its apparent host has a DES Y3 GOLD (I. Sevilla-Noarbe et al. 2021) photo- $z$  of  $0.271 \pm 0.025$ , also likely placing it outside the target cluster, PLCK G200.9–28.2 cluster ( $z \sim 0.220$ ).

<sup>12</sup> The host is unclear, as the bright source near the center could be a foreground contaminant.



**Figure 18.** UMAP plot of the BYOL features for the full dataset, colored by the predicted score from PROTEGE. Candidate XRGs (source 237, 255, and 110) and unusually circular sources (37, 146, and 3) are highlighted. The remarkable diversity even within a single “class” of “interesting” source type is reflected by the large distance between them in feature space. Nonetheless, they all lie in regions of interest to PROTEGE.

one-sided secondary lobe. It is not clear whether either of the two models mentioned above would be sufficient in this case. It is likely that additional physical processes are involved. The source was too faint to derive spectral maps.

The differences in morphology and spectra between these examples are also reflected in the UMAP feature space (Figure 18). Thus, if one were interested in finding XRGs, the diversity within this group means that no single similarity search would be sufficient. In addition, since these sources are not isolated or on the edges of the UMAP distribution, they would not be identified through pure anomaly searches in this feature space.

### 8.1.2. Case Study: Circular Shaped Sources

Circular radio sources, such as supernova remnants, planetary nebulae, or face-on spiral galaxies, are well known. However, a newly discovered class of extragalactic circular radio sources, dubbed odd radio circles (ORCs; B. S. Koribalski et al. 2021; R. P. Norris et al. 2021; R. P. Norris et al. 2022) are not compatible with standard models of jetted radio galaxies. Some ORCs have central galaxies, and so must originate in a spherically symmetric outflow from the AGN (R. P. Norris et al. 2021). Other ORCs appear to be associated with diffuse radio lobes from an adjacent AGN (e.g., S. S. Shabala et al. 2024). Yet another type of circular object, SAURON, was discovered using ASTRONOMALY (M. Lochner et al. 2023)—a circular structure with quadrilateral brightness asymmetries. Given this new and interesting phenomenon, we looked for circular sources among PROTEGE’s high-scorers.

Several high-scoring circular structures were found, including SAURON (37, left panel of Figure 19) and another source (146, middle panel of Figure 19) with a similar morphology and steep spectrum ( $-1.40 \pm 0.07$ ). It is not clear whether source 146 has a host galaxy.<sup>13</sup> However, if it does, then source 146 has a

comparable luminosity to SAURON ( $\sim 10^{25} \text{ W Hz}^{-1}$ ). No other candidate ORCs were found in PROTEGE’s top 260 recommendations.

A final, quite unusual high-scoring circular source (3) is shown in the right panel of Figure 19. A zoomed-out view (Figure 20) shows that this source may be the southern lobe of a large, highly disturbed radio galaxy. S. W. Randall et al. (2010) hypothesized that this could be a background giant radio galaxy at redshift 0.293. Its northern lobe would in that case coincide with a radio relic at the end of an X-ray plume (Y. Fujita et al. 2004), which is located in the foreground target cluster, A133 at  $z = 0.057$ .


These three circular sources are also shown in the UMAP of Figure 18, again illustrating their diversity as they lie far apart in feature space. This seems odd given how visually similar sources 37 and 146 appear, but visual inspection of a similarity search and these sources’ neighbors indicates that the BYOL algorithm has picked up on the diffuse features of source 37, which differ from the fairly smooth, less resolved features of source 146. Again, neither a similarity search nor anomaly selection can replicate the ability of PROTEGE to identify this diverse, “interesting” class of sources.

## 8.2. Previously Identified Interesting Sources


Above, we have shown that PROTEGE efficiently selected a very diverse set of sources that are of high human interest, as was the goal. We now turn to a related challenge, i.e., how effective is PROTEGE in recommending sources that were previously identified as of high interest, completely separate from its training?

For this task, we looked at the PROTEGE rank for sources that were highlighted in K. Knowles et al. (2022). A number of the more extraordinary sources discussed there appear in Figure 15 including a “phoenix” (7), the “double-scythe” (20; see J. O. Chibueze et al. 2021), and an unusual source demonstrating ribs and tethers (149; L. Rudnick et al. 2021). A highly morphologically disturbed source at cluster center

<sup>13</sup> The optical source in the middle is classified as a star in the DES Y3 GOLD catalog. However, it has no measured parallax in Gaia (Gaia Collaboration et al. 2016, 2023) and DELVE’s algorithm instead flags it as a galaxy, with a photo- $z$  of  $0.83 \pm 0.04$ .



**Figure 19.** Sources detected by PROTEGE that are similar to odd radio circles (ORCs). Source 37 (left) is SAURON, and source 146 (middle) is a morphologically similar source discovered by PROTEGE. While not a candidate ORC, source 3 (right) is an unusually rounded source. The flux density in these images is truncated to 99% of maximum to highlight faint features, and the potential host galaxy (where identified) is indicated with a white circle. The synthesized beam ( $\sim 7''$ ) is shown as a gray circle in the lower left of each image.



**Figure 20.** A zoomed-out image of source 3 revealing it as the likely lobe of a large (possibly giant) radio galaxy. The flux density in this image is truncated to 99% of maximum to highlight faint features, and the synthesized beam ( $\sim 7''$ ) is shown as a gray circle in the lower left.

exhibits remarkable filaments (54; L. Rudnick et al. 2022), and we find a new source (195), which could either be another highly disturbed source or even two galaxies, with clear filaments connecting lobes of the source.

In another example, Figure 14 of K. Knowles et al. (2022) shows a faint radio galaxy exhibiting lateral-edge enhancement (135), a puzzling finding not explained in current radio galaxy models. PROTEGE found another example, ranked extremely high at (5). Equally intriguing are two one-sided sources that appear in the high-scoring sources. One (133), has been suggested to result from one of the jets becoming trapped and shocked by the intergalactic medium as two sources merge (A. Hota et al. 2023). The second source (140) appears to show a normal jet and bright extended lobe, but nothing on the other side (outside of the cutout seen here and by PROTEGE). K. Kurahara et al. (2023) found one lobe to have an ultrastep spectrum, opening the possibility that the other lobe has faded below detectability.

PROTEGE successfully gave high scores even to some of the diffuse sources, for which PYBDSF is not optimized, including a halo + double relic system (29), “boomerang” relics (187), and some dying radio galaxies (27, 36).

To look at this more quantitatively, we checked the 23 examples chosen by K. Knowles et al. (2022) to highlight in their paper. These were chosen from hundreds of sources originally suggested for inclusion by team members, of the

many thousands of individual sources that were present in the images. Thus, the highlighted sources represent a set selected to be of the highest interest.

Of these 23, nine were either too big or too faint to be captured by PYBDSF and/or extended beyond the images used as inputs to PROTEGE. These included cluster-wide sources such as halos and large peripheral relics. This illustrates a limitation of our current selection. If the objective were to look for faint sources larger than several arcminutes, e.g., a different type of initial source selection than used here would be needed.

Of the remaining 14, the median PROTEGE rank was 70, i.e., among the top 1% of the sample. 13 of these 14 sources were in the top 2%, with the remaining source that had been selected by K. Knowles et al. (2022) for its interesting, but extremely faint diffuse emission. This very high correspondence demonstrates PROTEGE’s ability to identify those sources that will be of the highest interest to a large team of people. Author M.L., who did the full dataset training here, was not involved in the selection of the sources to highlight in K. Knowles et al. (2022), reinforcing this point.

## 9. Discussion

We have demonstrated that PROTEGE is a powerful tool for discovering “interesting” sources in the MGCLS dataset. However, there are several avenues for improving PROTEGE in order to prepare it for the truly large datasets of the near future. In this discussion, we highlight some of the strengths and weaknesses of PROTEGE with future applications in mind.

In Section 9.1 we discuss how the self-supervised deep features derived from BYOL have incredible potential for the field of unsupervised learning, but introduce new challenges that PROTEGE is uniquely positioned to address. Ultimately, this self-supervised methodology coupled with small dataset size did not allow for a separate test set. This is unusual for ML applications, as discussed in Section 9.2. Preprocessing is another area where improvement is needed, as discussed in Section 9.3, since it led to most of PROTEGE’s “failures.” Finally, in Section 9.4 we speculate how increasingly popular multimodal approaches, combining different data types, could further improve PROTEGE.

### 9.1. Unsupervised Learning with Deep Features

As discussed in Section 2.1, feature extraction is generally the limiting factor for ML applications, a problem that deep

learning has solved for many data types. M. Walmsley et al. (2022a) and K. Mohale & M. Lochner (2024) demonstrated deep-learning-derived features can unlock the potential of unsupervised learning for astronomical data, but also revealed a key problem for traditional anomaly detection. UMAP plots in those papers and in Figure 11 show that while interesting sources group together, they rarely fall on the boundary of feature space, and the features do not form distinct clusters. This is possibly because the loss functions in deep learning models do not include any specific penalty term for encouraging separation between groups of similar objects. It could also be related to the large number of parameters in deep learning models.

Traditional anomaly detection methods like isolation forest will fail to detect “interesting” sources that do not fall on the boundary. PROTEGE, on the other hand, enables the discovery of “interesting” sources even if they are buried deep in feature space. The algorithm only has to find one example to quickly locate any other sources nearby. This is a critical advantage of PROTEGE over traditional anomaly detection or the original form of ASTRONOMALY, which used anomaly detection as a first step (as demonstrated in Figure 10).

However, PROTEGE could easily still miss some truly rare sources buried among uninteresting ones and is reliant on the quality of the feature extractor. Therefore, improving the self-supervised methodology, particularly by using foundation models (models trained as general-purpose feature extractors using large datasets), would naturally improve PROTEGE’s recommendations.

### 9.2. Validation and Test Sets

It is common in ML applications to separate data into subsets for training (to train the algorithm), validation (to help optimize hyperparameters), and testing (to evaluate final performance). This separation prevents overfitting and ensures that models are generalizable. However, our application is not a standard supervised learning problem. We have a two-stage process: feature extraction with self-supervised learning followed directly by an active learning process. This structure, combined with a small dataset, does not easily accommodate the standard training-test data splitting approach.

We ultimately trained BYOL on the full dataset, because more data resulted in better features for the subsequent active learning stage (see Figure 10). At the same time, only the evaluation subset was used to evaluate performance. Although we attempted some hyperparameter optimization and checks for overfitting using validation sets in Appendix C, the use of the full dataset in our final application means we cannot claim that the trained model would generalize to other radio astronomy datasets.

However, generalization is not the intention of this work. The goal is to detect as many “interesting” sources as possible in the relatively unexplored and completely unlabeled MGCLS dataset. For use with another dataset, the entire procedure, including retraining of the model, would need to be repeated.

This raises a difference in philosophy between the current experiment and more traditional ML approaches. Even if we had a large enough dataset to make separate training and test sets, it would not be advantageous to do so. Use of such a hold-out test set would mean ignoring the training set in the search for “interesting” sources, which is our goal. The approach of combining self-supervised and active learning methods is still a

new field, and understanding the role of validation and test sets in this regime is an important avenue of future research.

### 9.3. Preprocessing

Figure 14 shows a high human-scoring source that PROTEGE ranked low. This happens because of the high dynamic range present in these radio images, causing faint emission to be close to invisible. While many transforms have been introduced to mitigate this problem (such as *asinh* used frequently throughout this paper), no universal transform has yet been identified that scales an image perfectly to maintain the necessary detail on multiple flux levels. As found in V. Etsebeth (2020), ML performance is often higher when using a transform that is suboptimal for human identification, leading us to use separate transforms for human and machine tasks. This counterintuitive result further highlights the importance of understanding the impact of preprocessing in deep learning applications.

In Appendix C, we found that providing BYOL with linear images without any scaling slightly outperformed *asinh* for this dataset, but that is unlikely to be a general result. An alternative may be to include multiple views of the same image, either with different scaling algorithms applied or in multiple “slices” of flux level. As CNNs frequently work with RGB images, the architecture is capable of ingesting this type of data and ensuring that correlations between views of the same image are preserved. This is a useful avenue to explore for any ML application for radio or other wavelength data.

### 9.4. Adding Information to Improve PROTEGE

Introducing additional information about each cutout could improve PROTEGE’s ability to recommend sources the user finds “interesting.” When we further investigated sources of interest identified by PROTEGE, we frequently used optical or infrared images. This helped, for example, to distinguish between a spiral galaxy or an unusual source like SAURON. Frequently, the larger context around the source is necessary for understanding its origins (Figure 20 for example). Figure 17 shows how useful spectral information can be, if it is available, in determining the physics of a particular system. In theory, if multiwavelength, spectral, or polarization data are available, as well as larger cutouts for each source, these can all be introduced as additional channels for each image before training BYOL. Alternatively, separate feature extractors can be trained and then later combined before passing the features to PROTEGE (this is particularly useful if additional data is not available for all sources). In general, multimodal foundation models, deep learning algorithms trained using self-supervised learning on large amounts of different types of data, have proven to be incredibly effective for a number of complex tasks (P. P. Liang et al. 2024). This is thus a critical avenue of research for the era of modern astronomy, dominated by multiwavelength surveys.

## 10. Conclusions

In this paper, we have made use of a new implementation of ML, ASTRONOMALY: PROTEGE, to make discoveries in the MeerKAT Galaxy Cluster Legacy Survey (MGCLS). We explored the concept of how “interesting” a source is, numerically comparing the results of two authors, one considerably more experienced with radio galaxy work than

the other. We find good agreement between authors, and with the highly selected “interesting” sources from K. Knowles et al. (2022). This indicates that the concept of what is scientifically “interesting,” based on the visual morphology, is a useful one.

We find that ASTRONOMALY: PROTEGE outperforms other ML approaches, as long as the deep-learning based feature extractor (BYOL) has access to sufficient data. The key metric is the recall plot (Figure 10), which asks how much data a human must sift through before detecting a significant number of high human-scoring sources.

With PROTEGE, a user only has to search a small amount of data in order to make potentially significant discoveries. We also found examples of sources that PROTEGE recommends and humans missed, which highlight that the algorithm has learned something somewhat more fundamental about what makes a source interesting to us than the specifics of the human training.

Many of the sources featured in the original MGCLS paper (K. Knowles et al. 2022) are ranked high by PROTEGE, but it has discovered many more. We selected a few specific sources to examine in more detail: three candidate X-shaped radio galaxies, SAURON (discovered using a previous version of ASTRONOMALY and rediscovered here), a second candidate SAURON, and an unusual ring-shaped lobe of a galaxy. These sources, which would be considered “interesting” anomalies by most scientists, demonstrate the diversity of sources recommended by PROTEGE and also the richness and discovery potential of modern radio datasets like MGCLS.

Aside from performance, PROTEGE has several conceptual advantages over other methods:

1. Blind methods aimed at finding unusual sources, such as complexity measures and traditional anomaly detection, inevitably include many uninteresting sources such as artifacts. The active learning step introduced in ASTRONOMALY mitigates this contamination.
2. Even with active learning, anomaly detection methods can fail in complex feature spaces such as Figure 11. By instead reframing the problem in terms of learning to make recommendations, PROTEGE can locate sources of interest buried among other sources.
3. While we focused on developing a general recommender, PROTEGE has the flexibility to be tuned to any source of interest including specific classes of source. Although it cannot compete with supervised algorithms if large amounts of training data are available, PROTEGE has the advantage of being able to learn from very few examples.
4. The idea of a recommender rather than an anomaly detector naturally lends itself to collaborative and citizen-science applications, where multiple users’ scores could be combined and used as training data to improve PROTEGE and provide recommendations to a community of users.
5. The combination of BYOL-based features and PROTEGE can also be applied to image data of any wavelength, and wavelengths can even be combined before training BYOL.

There are several future research directions that could improve PROTEGE. The first and simplest is to increase the size of the training set for the feature extractor, perhaps exploring the use of radio-specific foundation models such as

those described in I. V. Slijepcevic et al. (2023), fine-tuned onto the dataset of interest.

We also discussed, in Section 9, the need to improve preprocessing, specifically in scaling images with high dynamic range, as well as the rich potential in combining other sources of information, such as multiwavelength data, spectral index maps, polarization, and zoomed-out images for context, in improving PROTEGE’s recommendations.

A number of sources highly ranked by PROTEGE appear immediately to the authors to be worthy of further science investigation. We encourage readers to examine Figure 15 to identify their own sources to pursue, with coordinates and MGCLS target fields also available (see footnote 7). We further encourage the use of the PROTEGE software (see footnote 4), to facilitate discoveries in other large and rich astronomical datasets.

### Acknowledgments

The authors thank Mike Walmsley, Micah Bowles, and the anonymous referee for providing useful comments on the draft. M.L. acknowledges support from the South African Radio Astronomy Observatory and the National Research Foundation (NRF) toward this research. Opinions expressed and conclusions arrived at are those of the authors and are not necessarily to be attributed to the NRF.

MGCLS data products were provided by the South African Radio Astronomy Observatory and the MGCLS team and were derived from observations with the MeerKAT radio telescope. Access to the variety of MGCLS data products is described in K. Knowles et al. (2022). The MeerKAT telescope is operated by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation.

This work has made use of data from the European Space Agency (ESA) mission Gaia (<https://www.cosmos.esa.int/gaia>), processed by the Gaia Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement.

This research has made use of ESASky, developed by the ESAC Science Data Centre (ESDC) team and maintained alongside other ESA science mission’s archives at ESA’s European Space Astronomy Centre (ESAC, Madrid, Spain).

This research has made use of the SIMBAD database, CDS, Strasbourg Astronomical Observatory, France. This research has made use of Aladin sky atlas, CDS, Strasbourg Astronomical Observatory, France.

This research uses services or data provided by the Astro Data Lab, which is part of the Community Science and Data Center (CSDC) Program of NSF NOIRLab. NOIRLab is operated by the Association of Universities for Research in Astronomy (AURA), Inc. under a cooperative agreement with the U.S. National Science Foundation.

This project used public archival data from the Dark Energy Survey (DES). Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute

of Cosmological Physics at the University of Chicago, the Center for Cosmology and Astro-Particle Physics at the Ohio State University, the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Científico e Tecnológico and the Ministério da Ciência, Tecnologia e Inovação, the Deutsche Forschungsgemeinschaft, and the Collaborating Institutions in the Dark Energy Survey. The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgenössische Technische Hochschule (ETH) Zürich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ciències de l'Espai (IEEC/CSIC), the Institut de Física d'Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig-Maximilians Universität München and the associated Excellence Cluster Universe, the University of Michigan, the National Optical Astronomy Observatory, the University of Nottingham, The Ohio State University, the OzDES Membership Consortium, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, and Texas A&M University. Based in part on observations at Cerro Tololo Inter-American Observatory, National Optical Astronomy Observatory, which is

operated by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.


*Facilities:* MeerKAT, Blanco, Gaia, Astro Data Lab.

*Software:* Numpy, SciPy (E. Jones et al. 2001), Matplotlib (J. D. Hunter 2007), Seaborn (M. L. Waskom 2021), scikit-learn (F. Pedregosa et al. 2011), scikit-image (S. van der Walt et al. 2014), Pandas (W. McKinney 2010), Astropy (The Astropy Collaboration et al. 2013, 2018), umap-learn (T. Sainburg et al. 2021), PyTorch (A. Paszke et al. 2019) and BYOL-PyTorch (X. Chen & K. He 2020).

## Appendix A Random Sample of Sources

The primary purpose of PROTEGE was to dramatically reduce the number of sources needing visual inspection to find the scientifically interesting ones. To see how PROTEGE performed on this task, we present in Figure 21 cutout images of 260 random sources from the dataset, for direct comparison with Figure 15. Note that these random sources are themselves only a small fraction of sources found by PYBDSF, viz., those with at least four Gaussian components. The results are dramatic, with PROTEGE's top scorers filled with interesting sources, while the random sources are mostly slightly extended blobs, showing very few features.

Without PROTEGE, a user would have to search through thousands of these sources to make such discoveries as described in Section 8.



**Figure 21.** A random sample of sources from the full dataset. Images have been normalized using an *asinh* scaling with a threshold at 90%, and the beam size ( $\sim 7''6$ ) is indicated with a bar in the lower left.




Figure 21. (Continued.)

## Appendix B Comparison with Previous Techniques

In Section 7, we compared PROTEGE, using the BYOL features, with the approach introduced in M. Lochner & B. A. Bassett (2021), which uses hand-crafted features and more traditional “true” anomaly detection techniques. PROTEGE significantly outperforms this approach for the MGCLS dataset, but we include it here for completeness. Figure 22 shows a flowchart of the method.

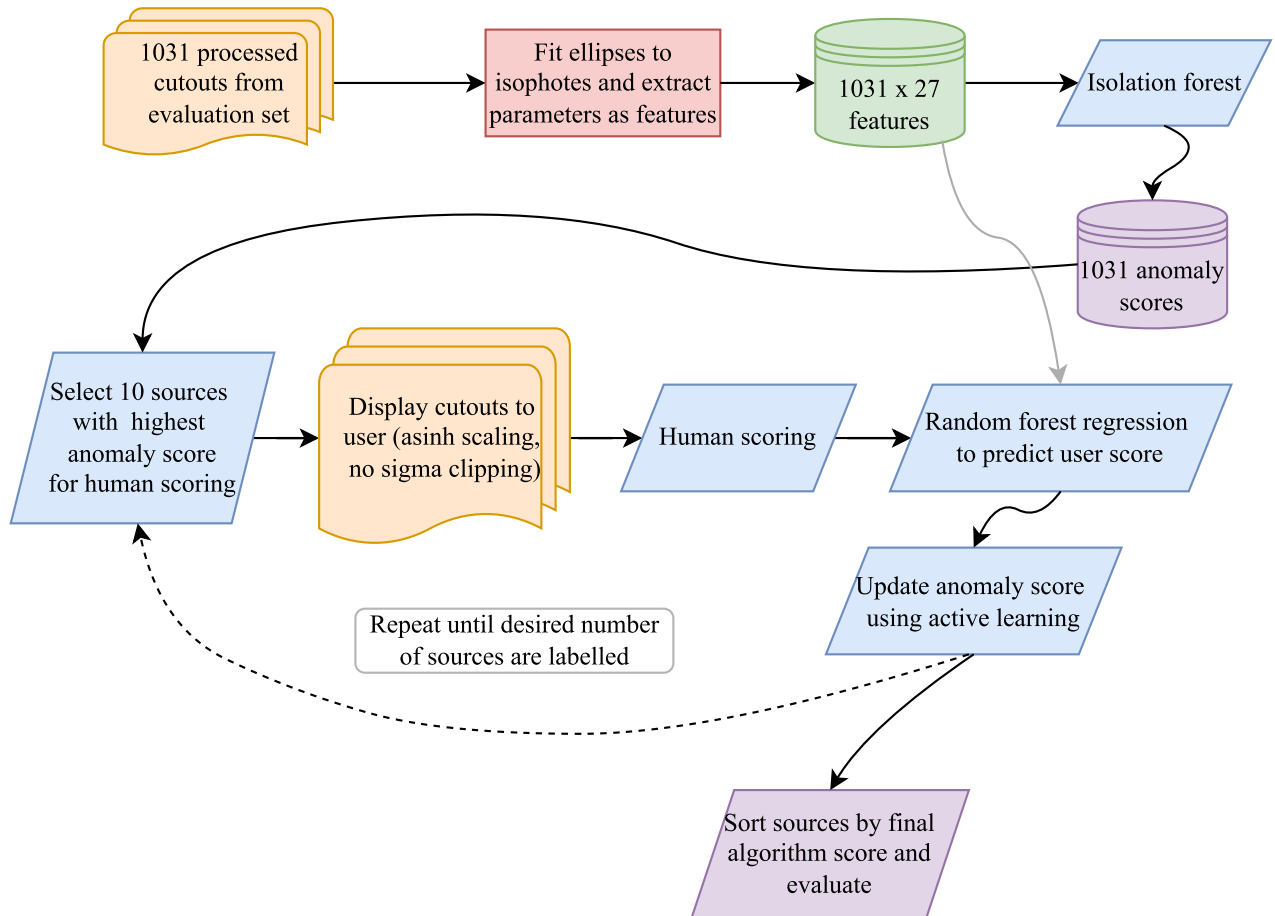
We used the same preprocessing steps described in Section 3.3 to extract cutouts of sources. Then, following M. Lochner & B. A. Bassett (2021), we use OPENCV to extract isophotes—contours of equal brightness—corresponding to 90th, 80th, 70th, 60th, 50th, and 0th percentiles of the brightness distribution of the pixels. We then fit ellipses to these contours, resulting in a set of parameters for each ellipse. The intuition behind using these as features is that more morphologically complicated (and therefore “interesting” to us) sources will have significant differences in parameters such as the rotation and offset from the center of each ellipse, while simpler sources will have ellipses with similar parameters. There are 21 features in the final feature set, described in more detail in M. Lochner & B. A. Bassett (2021).

We then apply the anomaly detection algorithm isolation forest (iforest; F. T. Liu et al. 2008), which is based on the popular random forest algorithm (L. Breiman 2001). iforest works by applying random cuts in feature space, using decision trees, in an attempt to isolate individual objects in a dataset. The idea is that more anomalous sources require fewer cuts to isolate. iforest allows for the sorting of the evaluation subset according to how anomalous each source is.

Finally, we apply the novel active learning technique introduced in M. Lochner & B. A. Bassett (2021), which we call the “neighbor score.” This is also an iterative process, where the top 10 sources are displayed to a user to score, and then active learning is applied. The process is repeated until 100 sources have been scored. The neighbor score approach first uses a regression algorithm, in this case random forest, to predict the user score for the full dataset. Then the new score is calculated as:

$$\hat{S} = S \tanh(\delta - 1 + \operatorname{arctanh}(\tilde{U})), \quad (\text{B1})$$

where  $S$  is the raw anomaly score returned from iforest (normalized from 0 to 5),  $\hat{S}$  is the updated active learning anomaly score,  $\delta$  is a distance penalty term, which is large for regions where the user score is not well known, and  $\tilde{U}$  is the normalized “relevance” score given by user labeling,



**Figure 22.** Flowchart describing the original approach taken in M. Lochner & B. A. Bassett (2021) as applied to the MGCLS data. In contrast to the method described in the main text, features are extracted by fitting ellipses to isophotes (contours of constant brightness), and the standard machine learning algorithm isolation forest is used to indicate anomalies. The active learning process to refine the anomaly scores is a reweighting by combining the isolation forest score, predicted human labels, and an uncertainty penalty term. The process is fundamentally different to PROTEGE, which is trained to focus on regions of interest rather than first filtering through an anomaly detection algorithm.

**Table 2**

Effect of Preprocessing Choices on the Ability of PROTEGE to Recover High Human-scoring Sources in the Evaluation Subset

Scaling	Sigma-clipping	Recall <sub>100</sub> Mean	Recall <sub>100</sub> Std Dev
linear	true	42.88	5.30
linear	false	38.22	6.41
asinh	true	31.00	8.26
asinh	false	36.14	8.03

**Note.** The number of sources given a human score of 4 or 5 in the top 100 sources is referred to as recall<sub>100</sub>. Generally, linear scaling and sigma-clipping are slightly preferred.

computed as:

$$\tilde{U} = \epsilon_1 + \epsilon_2 \left( \frac{U}{U_{\max}} \right), \quad (\text{B2})$$

where  $U$  is the predicted user score,  $U_{\max} = 5$  is the maximum possible user score, and  $\epsilon_1 = 0.1$ ,  $\epsilon_2 = 0.85$  are normalization constants.

The distance penalty term,  $\delta$ , is given by:

$$\delta = \exp\left(\alpha \frac{d}{d_0}\right), \quad (\text{B3})$$


where  $d$  is the Euclidean distance, in feature space, of the object in question to its nearest human-labeled neighbor,  $d_0$  is the mean distance to a labeled neighbor in the dataset, and  $\alpha$  is a tuning parameter that essentially dictates the trade-off between the anomaly detection scores and the user scores. We set  $\alpha$  to 0.1 for this work, significantly smaller than what was used in M. Lochner & B. A. Bassett (2021), as we found it has slightly higher performance. This suggests that the user scores are more informative than the raw anomaly scores.

Figure 10 shows that although the ellipse-based features, combined with iforest and the neighbor score active learning approach, perform fairly well, they are outperformed by BYOL and PROTEGE.

### Appendix C Hyperparameter Optimization

In supervised ML, hyperparameter optimization is typically done using a validation set (or cross-validation) and evaluating a clearly defined metric. In this paper, we have introduced a number of algorithms with a dizzying array of hyperparameters. Simultaneously, we have the problems of a small dataset and a fairly ill-defined metric (finding “interesting” sources). Truly optimizing all hyperparameters is extremely difficult due to the high variance inherent in the problem. While we spent considerable effort on hyperparameters following best practices, we focused on understanding the impact of the hyperparameters and selecting values likely to be appropriate for this specific dataset.

We used the evaluation subset introduced in Section 3.4 to perform several experiments on the effect of various hyperparameters. We will not go into detail within what is already a lengthy paper, but we recorded the results of these experiments at <https://github.com/MichelleLochner/mgcls.protege> for an interested reader to explore. We will describe the broad conclusions drawn from these experiments, and repeat the caution that due to the small dataset size, these are not likely to



**Figure 23.** The recall<sub>100</sub> as a function of loss (in BYOL), for a series of experiments with the evaluation subset. There is a very broad correlation between them, suggesting that better features result in a better recovery of sources that the user finds “interesting.” At the same time, the scatter shows the limitations of our small dataset.

generalize to other datasets. All experiments were performed on an ordinary desktop computer, making use of an NVIDIA RTX 3060 GPU. Training BYOL on the 1031 sources of the evaluation subset took around 10 minutes.

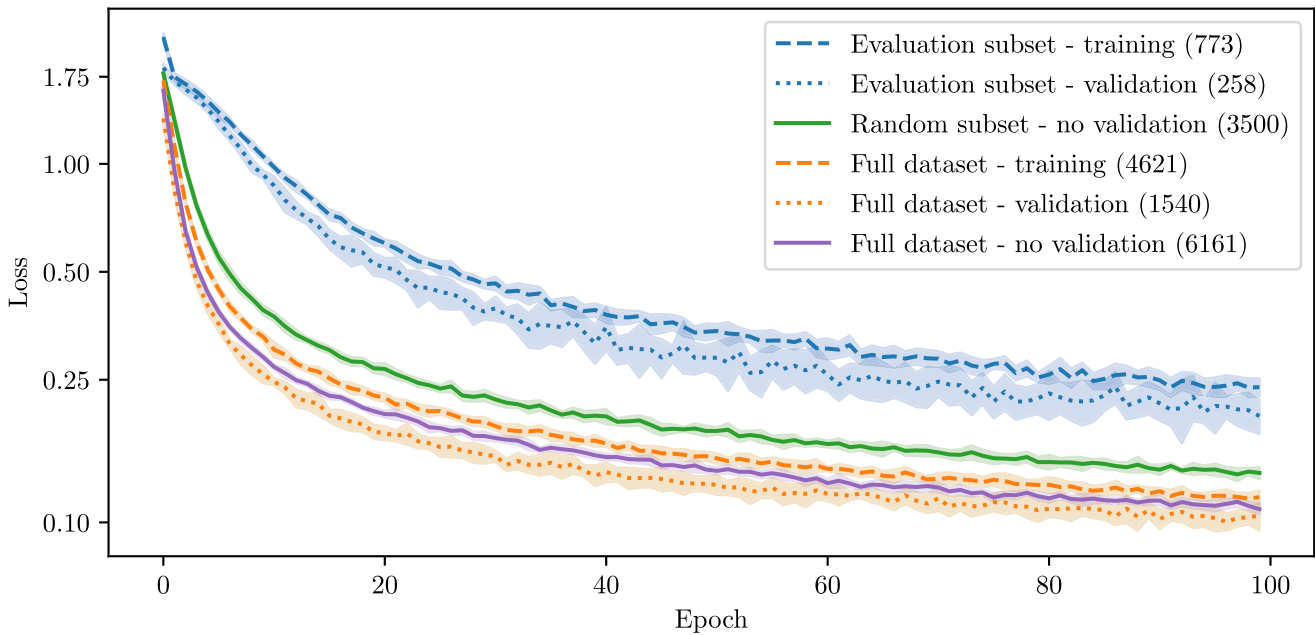
#### C.1. A Preliminary End-to-end Pipeline Investigation

We initially ran the full pipeline from feature extraction to the detection of high human-scoring sources to examine the impact of several hyperparameter choices. To evaluate performance, we use the metric recall<sub>100</sub>, which is simply defined as the number of sources with a human score of 4 or 5 in the top 100 most highly ranked sources. Each experiment was performed 50 times, and the mean recall<sub>100</sub> was recorded, as well as its standard deviation, minimum, and maximum values. This metric could vary considerably from experiment to experiment, highlighting the challenge of optimizing these hyperparameters.

First we investigated, for the preprocessing stage, whether to apply linear scaling or *asinh* scaling and whether or not to include sigma-clipping. Table 2 shows that using linear scaling with sigma-clipping produces slightly better results than the other options.

It is standard practice in deep learning to normalize the images to the same mean and standard deviation of the ImageNet (O. Russakovsky et al. 2015) dataset. We find this very slightly improves performance, but it is within the noise. We also investigated the impact of the batch size, finding a batch size of 32 significantly outperforms that of 64 or 128.

We explored some amount of variation in the augmentation choices used for BYOL. These were not exhaustive, as it was computationally intensive to perform a full ablation study (already done in I. V. Slijepcevic et al. 2023). However, in the process, we were able to establish that the recall<sub>100</sub> broadly correlates with the BYOL loss (Figure 23). This implies, somewhat obviously, that better features tend to be more reliable in recovering high human-scoring sources. To separate



**Figure 24.** Loss as a function of epoch for various subsets of the data. The solid lines correspond to the subsets used to evaluate performance in Section 7 while the dashed and dotted lines split either the evaluation subset or full dataset into training and validation sets to check for overfitting. The number of sources in each subset is given in brackets in the legend. It can be seen that the validation loss shows a similar trend to that of the training, indicating no evidence of significant overfitting. Note that the validation loss is slightly lower than the training loss because dropout is disabled for evaluation. Introducing more data naturally improves the loss function.

out the impact of BYOL hyperparameters from those of PROTEGE, we then proceeded to explore the impact of augmentations on the loss directly.

### C.2. BYOL Hyperparameters

For each combination of hyperparameters, we trained BYOL 10 times, selecting a different random 75% of the evaluation subset for training each run. We evaluated performance on the validation set loss. Unlike for the experiment in Section C.1, we found that a batch size of 64 performs slightly better than 32. However, since the difference in recall in the first experiment was far more pronounced, we elected instead to use a batch size of 32.

We explored base learning rates from  $10^{-5}$  to  $10^{-3}$ , finding our selected base learning rate of 0.0005 (then scaled by batch size/256) provides a mean validation loss of 0.2 compared to 0.7 for the extreme ends of the search range, justifying our selection of this value.

We used the KORNIA package (E. Riba et al. 2020) to apply augmentations. We did not repeat the augmentation ablation study of I. V. Slijepcevic et al. (2023). However, to ensure the same augmentation choices are applicable for our dataset, we varied some of the individual augmentations and their hyperparameters. We broadly find that for all augmentation hyperparameters explored, the mean validation loss varies by about 25% with our final choice of augmentations minimizing the loss, with one notable exception. When we removed the augmentation `ColorJiggle` (the KORNIA equivalent of `ColorJitter`), the loss reduced by a factor of 3. `ColorJiggle` randomly varies the brightness, contrast, saturation, and hue of an image. To test the impact of this choice on the ability of PROTEGE to detect high human-scoring sources, we produced two independent feature sets, one with and one without `ColorJiggle`, using the selected values for all other hyperparameters. Figure 24 shows the loss function for various subsets of data,

using the final set of augmentations selected. It can be seen that there is no evidence for overfitting (the validation loss does not increase above the training loss), and more training data improves the loss overall. This motivates our choice of using the full dataset to train the BYOL feature extractor when applying PROTEGE to the full dataset in Section 8.

### C.3. Optimizing the Search for Interesting Sources

For this experiment, we used two feature sets, switching `ColorJiggle` on and off, derived using the full dataset and tested only on the evaluation subset. Each experiment is repeated 10 times, with the mean and standard deviation of the  $\text{recall}_{100}$  reported.

We first considered using the active anomaly detection approach of ASTRONOMY introduced in M. Lochner & B. A. Bassett (2021) with the BYOL features. For a fair comparison with PROTEGE, we explored the hyperparameters of isolation forest, local outlier factor, random forest (L. Breiman 2001; used in the neighbor score algorithm of ASTRONOMY), and Gaussian processes (C. E. Rasmussen & C. K. Williams 2006). We ended up using the default parameters from the SCIKIT-LEARN implementations of these algorithms with the following exceptions:

1. For isolation forest and random forest, we set `n_estimators` to 200.
2. For local outlier factor, we set `n_neighbors` to 50.
3. For Gaussian processes, we use a `WhiteKernel` and `Matern` kernel and allow SCIKIT-LEARN to optimize the hyperparameters.

We found that when using the BYOL features, the standard anomaly detection approach does not have high performance. Using the neighbor score active learning method results in a mean  $\text{recall}_{100}$  of 24 when using isolation forest (F. T. Liu et al. 2008) for anomaly detection and just 14 when using local

**Table 3**

Final Set of Augmentations Used for Training BYOL, from the `kornia`. augmentation Library, the Associated Probability of Being Applied, and Any Relevant Hyperparameters

Augmentation	Probability	Hyperparameters
RandomRotation	1.0	0°–360°
RandomHorizontalFlip	0.5	N/A
RandomVerticalFlip	0.5	N/A
CenterCrop	1.0	110 pixels
RandomResizedCrop	1.0	80%–100% of image
RandomGaussianBlur	0.1	kernel = 15 pixels, $\sigma = 10$ –15 pixels
ColorJiggle	0.8	All parameters set to 0.5

**Note.** A range is indicated if the hyperparameter is randomly selected from this range at each iteration.

outlier factor (M. M. Breunig et al. 2000). The best-performing run from PROTEGE in comparison results in a recall<sub>100</sub> of 44. This is why we did not attempt to use standard anomaly detection with these features.

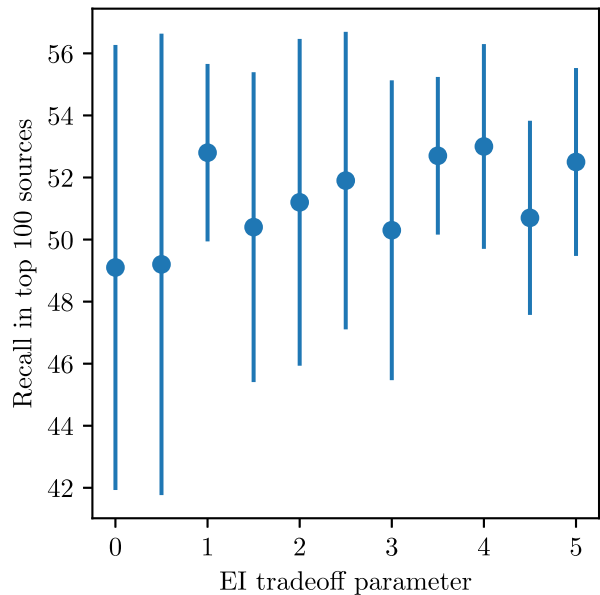
When considering the BYOL features with PROTEGE instead, we discovered that retaining the `ColorJiggle` augmentation slightly improves the recall<sub>100</sub> (39 against 36). We thus retained this augmentation, and Table 3 summarizes the final set of augmentations used in this work.

We explored different numbers of PCA components finding the impact surprisingly small (recall<sub>100</sub> of 41 against 44). This is why we elected to simply use the variance to dictate the number of PCA components, especially as this will vary depending on which subset of data is used to train BYOL.

Although best practices were followed in hyperparameter optimization and the final set of hyperparameters chosen had the highest performance, we conclude that for the majority of hyperparameters, the dataset is too small and the variance too high for them to have significant impact on the results.

#### C.4. Effect of the Trade-off Parameter

The trade-off parameter  $\epsilon$  used in the acquisition function of PROTEGE (discussed in Section 6.1) could potentially have an impact on performance. In this work, we set this parameter to 3 (halfway between 1 and 5, the minimum and maximum user scores). However, it seemed useful to test for sensitivity to the parameter. We used our final choice of hyperparameters described in Appendix C.3 and trained BYOL 10 times on 75% of the full dataset, retaining 25% for validation (as done in Appendix C.2). We then ran PROTEGE with varying values for  $\epsilon$ , recording the mean and standard deviation in recall<sub>100</sub>. The results, shown in Figure 25, indicate that values of low  $\epsilon$  (less exploration) result in more variance. Beyond that, all other values are within the variance, and it is difficult to select an




**Figure 25.** Number of “interesting” sources (human score of 4 or 5) detected in the top 100 sources with the highest PROTEGE score, after training, as a function of different trade-off parameter values. The results are shown for the evaluation subset only. In this optimization exercise, the features are extracted using 75% of the full dataset for training BYOL with 25% kept for validation. The error bars are the standard deviation from 10 different training-validation set splits. While the trade-off parameter does not significantly affect performance on average, the variance is far lower for values above 0.5.

optimal value. We thus continued to use an  $\epsilon$  of 3 for this work. This parameter could have greater impact for larger datasets or different feature sets.

## Appendix D Evaluation of User Score Regression

A key component of PROTEGE is its ability to predict the user score for the entire feature space based on a very small amount of training data. We used the evaluation subset (Section 3.4) to test the performance of the Gaussian process regression algorithm used in this work. Figure 26 compares the user scores provided by author M.L. with the predictions from PROTEGE after training on 100 user-provided scores, in batches of 10 per iteration (as described in Section 6.1). Despite a large scatter, there is a noticeable correlation between user score and the score from PROTEGE. The scatter is likely due to the very small training sample, e.g., with just 12 examples of sources with a score of 5 available for the algorithm to learn from. The predictions improve with larger training sets, but even these poor predictions are sufficient to rapidly prioritize sources the user finds “interesting.”



**Figure 26.** PROTEGE predictions vs. author M.L.'s scores for the entire evaluation subset, after training on 100 labels. The data are illustrated using a box and whisker plot, where the median is indicated by a horizontal line in the box, and the boundaries of the box are placed at the quartiles of the distribution, with the whiskers showing the maximum and minimum values barring outliers, given by circles. While the PROTEGE scores do not perfectly match the user scores and have high scatter due to a small training set, they nonetheless correlate well enough to quickly locate sources of interest to the user.

### ORCID iDs

Michelle Lochner  <https://orcid.org/0000-0003-2221-8281>

Lawrence Rudnick  <https://orcid.org/0000-0001-5636-7213>

### References

- Andrianomena, S., & Tang, H. 2024, *JCAP*, 2024, 034
- Becker, B., Vaccari, M., Prescott, M., & Grobler, T. 2021, *MNRAS*, 503, 1828
- Böhm, V., Kim, A. G., & Juneau, S. 2023, *MNRAS*, 526, 3072
- Boyce, M. M., Hopkins, A. M., Riggi, S., et al. 2023, *PASA*, 40, e027
- Breiman, L. 2001, *Mach. Learn.*, 45, 5
- Breunig, M. M., Kriegel, H.-P., Ng, R. T., & Sander, J. 2000, *SIGMOD Rec.*, 29, 93104
- Chen, X., & He, K. 2020, in *IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, ed. Lisa O'Conner (New York: IEEE)
- Chibueze, J. O., Sakemi, H., Ohmura, T., et al. 2021, *Natur*, 593, 47
- Cotton, W. D., Thorat, K., Condon, J. J., et al. 2020, *MNRAS*, 495, 1271
- de Gasperin, F., Rudnick, L., Finoguenov, A., et al. 2022, *A&A*, 659, A146
- De Vicente, J., Sánchez, E., & Sevilla-Noarbe, I. 2016, *MNRAS*, 459, 3078
- Drlica-Wagner, A., Ferguson, P. S., Adamów, M., et al. 2022, *ApJS*, 261, 38
- Etsebeth, V. 2020, Master's thesis, University of the Western Cape, <http://hdl.handle.net/11394/9028>
- Etsebeth, V., Lochner, M., Walmsley, M., & Grespan, M. 2024, *MNRAS*, 529, 732
- Fujita, Y., Sarazin, C. L., Reiprich, T. H., et al. 2004, *ApJ*, 616, 157
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, 595, A1
- Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, *A&A*, 674, A1
- Giri, G., Fendt, C., Thorat, K., Bodo, G., & Rossi, P. 2024, *FrASS*, 11, 1371101
- Grill, J.-B., Strub, F., Alché, F., et al. 2020, *Advances in Neural Information Processing Systems*, Vol. 33 (New York: Curran Associates, Inc.), 21271
- Gupta, N., Huynh, M., Norris, R. P., et al. 2022, *PASA*, 39, e051
- Hota, A., Dabhade, P., & Vaddi, S. 2023, in *IAU Symp. 375, The Multimessenger Chakra of Blazar Jets*, ed. I. Lioudakis (Cambridge: Cambridge Univ. Press), 35
- Hotan, A. W., Bunton, J. D., Chippendale, A. P., et al. 2021, *PASA*, 38, e009
- Hotelling, H. 1933, *J. Educ. Psychol.*, 24, 417
- Hunter, J. D. 2007, *CSE*, 9, 99

- Jonas, J. & MeerKAT Team 2016, in *MeerKAT Science: On the Pathway to the SKA*, ed. R. Taylor et al. (Trieste: SISSA), 1
- Jones, D. R., Schonlau, M., & Welch, W. J. 1998, *J. Glob. Optim.*, 13, 455
- Jones, E., Oliphant, T., Pearu, P., et al. 2001, *SciPy: Open Source Scientific Tools for Python.*, <http://www.scipy.org/>
- Kennamer, N., Ishida, E. E. O., Gonzalez-Gaitan, S., et al. 2020, in *IEEE Symp. Ser. Computational Intelligence*, ed. H. Abbass (New York: IEEE)
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Knowles, K., Cotton, W., Rudnick, L., et al. 2021, *The MeerKAT Galaxy Cluster Legacy Survey—Survey Overview and Highlights*, South African Radio Astronomy Observatory, doi:10.48479/7EPD-W356
- Knowles, K., Cotton, W. D., Rudnick, L., et al. 2022, *A&A*, 657, A56
- Koribalski, B. S., Norris, R. P., Andernach, H., et al. 2021, *MNRAS*, 505, L11
- Kumari, S., Pal, S., Hardcastle, M. J., & Horton, M. A. 2024, *A&A*, 689, A301
- Kurahara, K., Akahori, T., Kale, R., et al. 2023, *PASJ*, 75, S138
- Lacy, M., Baum, S. A., Chandler, C. J., et al. 2020, *PASP*, 132, 035001
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Natur*, 521, 436
- Liang, P. P., Zadeh, A., & Morency, L.-P. 2024, *ACM Comput. Surv.*, 56, 1
- Liu, F. T., Ting, K. M., & Zhou, Z. 2008, in *2008 Eighth IEEE Int. Conf. Data Mining (New York: Curran Associates, Inc.)*, 413
- Lochner, M., & Bassett, B. A. 2021, *A&C*, 36, 100481
- Lochner, M., Ramonyai, M., Etsebeth, V., & Gomez, P. 2024, *MichelleLochner/astrometry: Astrometry, v2.0*, Zenodo, doi:10.5281/ZENODO.14441058
- Lochner, M., Rudnick, L., Heywood, I., Knowles, K., & Shabala, S. S. 2023, *MNRAS*, 520, 1439
- McInnes, L., Healy, J., & Melville, J. 2018, *JOSS*, 3, 29
- McKinney, W. 2010, in *Data Structures for Statistical Computing in Python (Austin, TX: SciPy)*, <http://conference.scipy.org/proceedings/scipy2010/mckinney.html>
- Mockus, J. B., & Mockus, L. J. 1991, *J. Optim. Theory Appl.*, 70, 157
- Mohale, K., & Lochner, M. 2024, *MNRAS*, 530, 1274
- Mohan, N., & Rafferty, D. 2015, *PyBDSF: Python Blob Detection and Source Finder*, Astrophysics Source Code Library, ascl:1502.007
- Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Rios, D., Bobes-Bascarán, J., & Fernández-Leal, A. 2023, *Artif. Intell. Rev.*, 56, 3005
- Mostert, R. I. J., Duncan, K. J., Röttgering, H. J. A., et al. 2021, *A&A*, 645, A89
- Norris, R. P., Collier, J. D., Crocker, R. M., et al. 2022, *MNRAS*, 513, 1300
- Norris, R. P., Intema, H. T., Kapińska, A. D., et al. 2021, *PASA*, 38, e003
- Paszke, A., Gross, S., Massa, F., et al. 2019, in *Advances in Neural Information Processing Systems*, ed. H. Wallach et al. (New York: Curran Associates Inc)
- Patra, D., Joshi, R., & Gopal-Krishna 2023, *MNRAS*, 524, 3270
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, 12, 2825
- Ralph, N. O., Norris, R. P., Fang, G., et al. 2019, *PASP*, 131, 108011
- Randall, S. W., Clarke, T. E., Nulsen, P. E. J., et al. 2010, *ApJ*, 722, 825
- Rasmussen, C. E., & Williams, C. K. 2006, *Gaussian Processes for Machine learning* (Cambridge, MA: MIT Press)
- Riba, E., Mishkin, D., Ponsa, D., Rublee, E., & Bradski, G. 2020, in *Winter Conf. Applications of Computer Vision (New York: IEEE)*
- Riggi, S., Ceconello, T., Palazzo, S., et al. 2024, *PASA*, 41, e085
- Riseley, C. J., Bonnasieux, E., Vernstrom, T., et al. 2022, *MNRAS*, 515, 1871
- Rojas, K., Collett, T. E., Ballard, D., et al. 2023, *MNRAS*, 523, 4413
- Roy, D., & Dutta, M. 2022, *J. Big Data*, 9, 59
- Rudnick, L., Brüggem, M., Brunetti, G., et al. 2022, *ApJ*, 935, 168
- Rudnick, L., Cotton, W., Knowles, K., & Kolokythas, K. 2021, *Galax*, 9, 81
- Russakovsky, O., Deng, J., Su, H., et al. 2015, *IJCV*, 115, 211
- Sainburg, T., McInnes, L., & Gentner, T. Q. 2021, *Neural Comput.*, 33, 2881
- Savary, E., Rojas, K., Maus, M., et al. 2022, *A&A*, 666, A1
- Segal, G., Parkinson, D., Norris, R. P., & Swan, J. 2019, *PASP*, 131, 108007
- Segal, G., Parkinson, D., Norris, R. P., et al. 2023, *MNRAS*, 521, 1429
- Sevilla-Noarbe, I., Bechtol, K., Carrasco Kind, M., et al. 2021, *ApJS*, 254, 24
- Shabala, S. S., Yates-Jones, P. M., Jerrim, L. A., et al. 2024, *PASA*, 41, e024
- Slijepcevic, I. V., Scaife, A. M. M., Walmsley, M., et al. 2023, *RASTI*, 3, 19
- Tan, M., & Le, Q. 2019, in *Proc. Machine Learning Research*, Vol. 97, Proc. 36th Int. Conf. Machine Learning, ed. K. Chaudhuri & R. Salakhutdinov (Long Beach, CA: MLR Press), 6105
- The Astropy Collaboration, Price-Whelan, A. M., Sipocz, B. M., et al. 2018, *AJ*, 156, 123
- The Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33
- Vafaei Sadr, A., Bassett, B. A., & Sekyi, E. 2022, arXiv:2210.16334

- van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., et al. 2014, [PeerJ](#), 2, e453
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, [A&A](#), 556, A2
- Vantyghem, A. N., Galvin, T. J., Sebastian, B., et al. 2024, [A&C](#), 47, 100824
- Walmsley, M., Scaife, A. M. M., Lintott, C., et al. 2022a, [MNRAS](#), 513, 1581
- Walmsley, M., Slijepcevic, I., Bowles, M. R., & Scaife, A. 2022b, in ICML 2022 Workshop on Machine Learning for Astrophysics (Baltimore, MD: ICML)
- Waskom, M. L. 2021, [JOSS](#), 6, 3021
- Webb, S., Lochner, M., Muthukrishna, D., et al. 2020, [MNRAS](#), 498, 3077
- Zhao, X., Wang, L., Zhang, Y., et al. 2024, [Artif. Intell. Rev.](#), 57, 57