Detecting Low Surface Brightness Galaxies with Mask R-CNN

Caleb Levy

Abstract

Low surface brightness galaxies (LSBGs) inhabit the ultra-faint parameter space of the galactic population and are infamous for their difficulty to detect. Studying these galaxies is essential for understanding of dark matter's role in galaxy formation at low-halo-mass scales. In this project, we train a deep learning model using the Mask R-CNN framework on a set of simulated LSBGs inserted into a small fraction of the Dark Energy Survey (DES) Data Release 2 (DR2). This model is used in a pipeline for the detection of LSBGs. All 10,169 DES DR2 coadded tiles were ran through the pipeline, and preliminary results show the detection of 22 large, high-quality LSBG candidates that went undetected by conventional algorithms and a high prominence of Galactic cirrus in the model.

Keywords: Machine Learning – Deep Learning – Low-Surface-Brightness Galaxies

1. Introduction

The limited sensitivity of astronomical observations leads to an inability to detect systems that are relatively faint (Disney, 1976). In particular, low surface brightness galaxies, galaxies with central brightness fainter than the night sky, prove very difficult to detect due to observational selection effects. To address this issue, efficient and effective ways of identifying and cataloging very faint galaxies is necessary. Low surface brightness galaxies are traditionally defined as galaxies with a central brightness lower than the night sky (Bothun et al., 1997), and have historically proven very difficult to detect. These systems are extremes of the galactic population, and thus provide a good test for the standard models of galaxy formation. Thus, it is essential that efforts be made to improve the detection efficiency of LSBGs as studying them could give valuable insights into cosmology and galaxy formation.

Recently, a catalog of 23,790 LSBGs has been identified in the first three years of data from the Dark Energy Survey (DES; DES Collaboration, 2005; DES Collaboration et al., 2018) using results from `SourceExtractor` and a classifier machine learning algorithm (Tanoglidis et al., 2021). This labelled dataset has subsequently been used to train a Convolutional Neural Network (CNN) for separating LSBGs from artifacts (Tanoglidis et al., 2020), showing the promise of machine vision for studying the low-surface-brightness universe. The focus of this work is to leverage the power of machine vision, specifically through the Mask R-CNN framework (He et al., 2018), to train a model for detecting even larger and fainter LSBGs in DES data.

2. Background

2.1. Galaxy–Halo Connection

Modern cosmological models attribute $\sim 5/6^{\text{th}}$ of the mass budget of the universe to dark matter. In the inflationary period of the early universe, small fluctuations in the matter distribution of the universe emerged, planting the seeds for dark matter's gravitational collapse into halo structures. Baryonic matter eventually followed, cooling and falling into the centers of dark matter halos. In the most massive halos, where the gas was able to cool sufficiently, star formation began and the first protogalaxies were born. Over time, the star formation rate increased and halos merged to form larger structures that host the galaxies we see today. The corollary of this model of galaxy formation is a close relationship between dark matter halo properties and galaxy properties, known as the *galaxy-halo connection* (Wechsler and Tinker, 2018).

Understanding the relationship between galaxies and their dark matter halos is a rich sub-field and

features a wide variety of research. On the empirical side, one way to constrain the galaxy-halo connection is the stellar mass/halo mass relation (SHMR). A wide array of observations have converged on the form of this relation for halo masses $M_h \gtrsim 10^{11}$ $M_\odot$ because the galaxies that inhabit these halos are massive and bright. Below this mass, most models diverge due to the difficulty in detecting the faint galaxies that reside in low-mass halos. Better understanding this relationship in the low-halo-mass region is imperative for a complete picture of the galaxy-halo connection. There thus arises a necessity for novel ways of detecting the low-surface-brightness universe.

2.2. Low Surface Brightness Galaxies

Low surface brightness galaxies are historically defined as galaxies fainter than the night sky (Bothun et al., 1997). Theoretical calculations within the standard cosmological model, ΛCDM, predict that these extremes form naturally from proto-galaxies with high angular momentum and low mass (Dalcanton et al., 1997). However, challenges to this model has recently emerged upon the discovery of LSBGs with anomalously low dark matter content (van Dokkum et al., 2018). These discoveries pose a challenge to our understanding of galaxy formation and may require modified version of these models to match the observed universe. It has thus become even more imperative to develop efficient ways to detect LSBGs to gain a complete picture of galaxy formation in ΛCDM.

2.3. Dark Energy Survey

In recent years, the emergence of wide-area sky surveys, such as the Sloan Digital Sky Survey[1] and the Dark Energy Survey[2], has allowed for large-scale, untargeted searches for LSBGs, (Zhong et al., 2008; Tanoglidis et al., 2021). In contrast to the targeted searches of telescopes optimized for low-surface-brightness detection in the near-field (Eg. the Dragonfly Telephoto Array; Abraham and van Dokkum, 2014), these untargeted searches allow us to gain a statistical understanding of the distribution of LSBGs to better understand their place in the galaxy-halo connection.

[1]https://www.sdss.org/

[2]http://www.darkenergysurvey.org/

The Dark Energy Survey is a multi-band ($grizY$) imaging survey covering ~ 5000 deg^2 of the Southern sky. Imaging is done at the Cerro Tololo Inter-American Observatory (CTIO) 4m Blanco telescope in Chile using the Dark Energy Camera (DECam), a 570 megapixel wide-field camera. The objective of DES is to gain a more complete understanding of our universe through the study of dark matter and dark energy. In this work, we use coadded DES image tiles of size 0.534 deg^2 from DES Data Release 2 (DR2). An initial search for LSBGs in the first three years of DES observation has been done using a combination of `SourceExtractor` modelling and a machine learning classifier (Tanoglidis et al., 2021). The resulting catalog contains 23,790 visually-confirmed LSBGs with mean surface brightness $\bar{\mu}_{\rm eff}(g) > 24.2$ mag arcsec^{-2}. In addition, this catalog has been used to develop the deep-learning algorithm *DeepShadows*, a Convolutional Neural Network (CNN) trained to separate LSBGs from artifacts like Galactic cirrus and tidal ejecta from bright galaxies (Tanoglidis et al., 2020). The model achieved a test accuracy of 92%, demonstrating the power of machine vision for understanding the low-surface-brightness universe.

3. Mask R-CNN Model

In this work, we utilize the instance segmentation and object detection algorithm Mask R-CNN. We begin this section with a brief overview of how the model works, and refer the interested reader to the original paper for a more technical discussion (He et al., 2018).

A high-level depiction of the algorithm is shown in Fig. 1. The input image is first fed into the "backbone" of the network, a pre-trained CNN, in this case ResNet, whose output is a feature map. In a traditional CNN used for object detection, this feature map would be passed to a fully connected classification layer. However, in the Mask R-CNN implementation, the feature maps are fed into a Region Proposal Network (RPN) whose output is a set of Regions of Interest (RoI) where potential objects are located. These regions are identified from a CNN trained to identify *anchor boxes* of various sizes at each position. Doing this is a two-fold process involving binary classification and regression. The classification aspect involves assigning a probability to each box based on whether or not it contains an object of *any* class, while the regression component is associated with finding the optimal

box coordinates. The regions with the highest probabilities of containing an object are then passed into the RoIAlign method along with the feature maps from the backbone to produce RoI feature maps of the same size. This step is crucial as fully connected networks implemented for prediction require inputs of the same size.

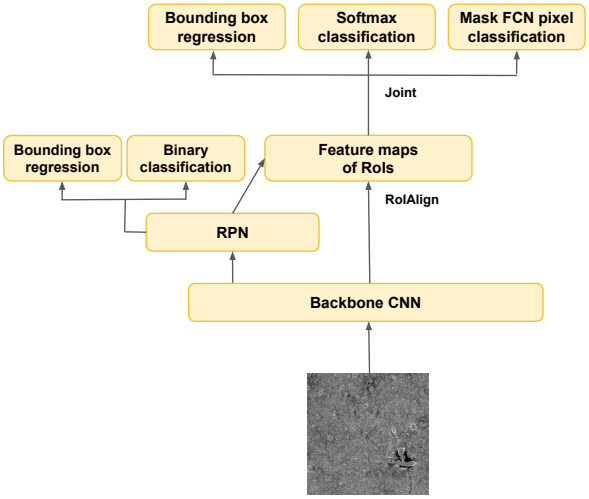

Figure 1: Schematic description of the Mask R-CNN algorithm.

The final part of the algorithm is a simultaneous three-fold process: classification of the object in the RoI with a softmax classifier, regression to learn the best bounding box coordinates, and per-pixel classification with a Fully Convolutional Network (FCN) to create an object mask. The final output of the algorithm is thus, for each object detected, a bounding box region, a per-pixel mask, an object class, and a confidence score.

3.1. Image Processing and Training

Training on the Mask R-CNN framework requires a set of images along with per-pixel masks of training objects in each image and their corresponding class. In order to train a model for detecting LSBGs, we rely on simulated galaxies injected into 94 random tiles of DES. A simulated dataset provides the distinct advantage of having pre-labeled ground truth masks, bypassing the requirement of manually labelling a dataset. Artificial LSBGs are modeled with GALSIM (Rowe et al., 2015), an open-source simulation toolkit. GALSIM can simulate

Sersic Index	0.3–4.5
Ellipticity	0–0.8
Effective Radius (arcsec)	2–150
Position Angle (°)	0–180
Surface Brightness (mag/arcsec^2)	24–27

Table 1: Low surface brightness galaxy simulation parameters for training data.

galaxies from a variety of simple parametric models, and we used a Sersic profile for our artificial LSBGs. Our input parameters are displayed in Table 1. We used GALSIM features that allowed us to convolve model galaxies with the measured Point Source Functions (PSFs) of each DES coadded image and add noise drawn from a Poisson distribution. We generated artificial LSBGs in g, r, i, and z bands, and then injected them into DES coadd images.

After injecting 16 LSBGs per tile, each tile was divided into sixteen 2500 $\times$ 2500 quadrants with each containing one LSBG. In addition to the LSBG class, we also included a training class for very large objects identified from the `SourceExtractor` catalog from DES DR2 (Collaboration et al., 2021). Objects in this class were identified as those whose masked pixel count exceeded 10,000 pixels ($\sim$ 0.2 deg^2) in angular size. The purpose of this was to remove contamination of LSBG detection by training the algorithm to specifically detect this class of object as distinct from LSBGs. Initial models without this training class showed a high percentage of false positives centered on these objects and thus it was necessary for the model to learn the difference between LSBGs and these large objects, which is most effectively done by training the model with an additional class.

After injecting simulated data into the DES dataset, the next step is to pre-process images before training in order to accentuate low surface brightness features. This process occurs in three steps:

1. Replace pixels associated with all sources detected in the DES DR2 `SourceExtractor` catalog with noise.
2. Convolve the image with a Gaussian kernel.
3. Bin the image by a factor of 10 $\times$ 10.

The first step is to remove all objects identified in the DES pipeline with `SourceExtractor` and

replace with a normally-distributed background. Firstly, all pixels identified as background in the catalog are averaged and their standard deviation computed. Then, the masks of the objects we wish to remove are binary-dilated with a 3×3 matrix of ones over five iterations using the binary dilation function of `SciPy`'s multidimensional image processing module, `SciPy.ndimage` (Virtanen et al., 2020). This is to ensure that pixels surrounding the mask that may be associated with the object, but are not masked, are removed. The final step is to replace the object masks with random samples from a normal distribution of the background.

The next pre-processing step is to perform a convolution operation on the image with a Gaussian kernel. This is known as a 'Gaussian blur,' and is an effective way to accentuate low surface brightness features in an image if the scale of the kernel coincides with the scale of the features of interest. In this work, we utilize a Gaussian kernel defined over a 20×20 pixel window (angular size ~ 30 arcsec^2) given by the following equation:

$$G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}, \tag{1}$$

where (x, y) is the distance measured from the center of the image[3], and σ is the standard deviation, which we took as $\sigma = 5$.

The final pre-processing step is to reduce the pixel count of the image by binning it with a mean filter. This is also a convolutional operation and involves taking the average of a group of pixels (whose size is determined by the binning factor) and mapping it to one pixel in the binned image. This was done with a bin factor of ten, reducing the image size from 10000×10000 to 1000×1000. In addition to binning the image itself, the masks containing the ground truth locations of the training objects must also be reduced by the same factor for training. To do this, the same process is done, but with a max filter which, instead of taking the average of a group of pixels, takes the maximum value. This ensures that the masks remain suitable as ground-truth representations as each pixel in the mask must take a boolean value.

In Fig. 2 we plot the total loss of the model after training 94 tiles with 1034 simulated LSBGs over 60 epochs. The total number of LSBGs injected

[3]In actuality, the Gaussian is normalized such that the central pixels take a value of 1.

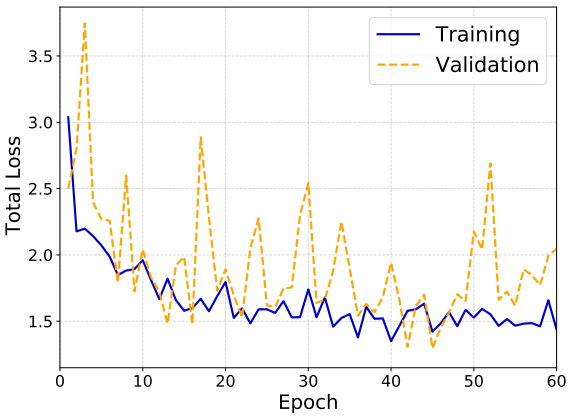

Figure 2: Total training and validation loss of Mask-RCNN model trained on 1034 LSBGs in 94 DES tiles.

was 1504, however $\sim 70\%$ of these were assigned to a training set, with $\sim 15\%$ left for validation and $\sim 15\%$ for testing. Each epoch consisted of 150 steps, with 15 steps at the end of each for validation. In the Mask R-CNN model, the total loss is a combination of all the network losses and is given by:

$$L_{\text{tot}} = L_{\text{RPN-bbox}} + L_{\text{RPN-cls}} + L_{\text{cls}} + L_{\text{bbox}} + L_{\text{mask}}, \tag{2}$$

where $L_{\text{RPN-bbox}}$ is the RPN bounding box loss associated with locating objects in the image, $L_{\text{RPN-cls}}$ is the classification loss associated with labelling anchor boxes in the RPN as foreground or background, L_{cls} is the classification loss of objects identified, L_{bbox} is the bounding box loss associated with precisely identifying the object's location after region proposal, and L_{mask} is the loss in masking the object.

The final step in the training phase was to test the model on the reserved test set of 282 LSBGs. While examining the quantitative results of this process is useful, we focused our efforts more on a qualitative analysis of the model. After looking at the model's output on the entire test set, we were confident in the model's performance and decided to apply the model to the DES dataset.

3.2. Application to DES Data

We began by applying the model to $\sim 1\%$ of the DES dataset that the model had yet to see. Our pipeline took the following form:

1. Pre-process tiles on the DES cluster.

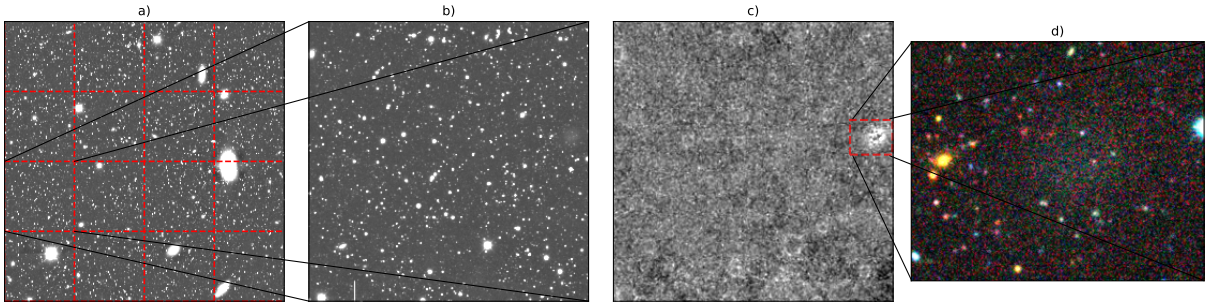

Figure 3: Example of a successful LSBG detection using the pipeline described in Section 3.2. Panel a) represents the full, unprocessed image along with a red gird outline indicating the cutouts made. Panel b) shows a zoomed in version of the indicated cutout, unprocessed. Panel c) represents the fully processed version of panel b), showing the accentuation of the LSBG in the right of the image. The region identified by the Mask R-CNN is outlined in red. Panel d) shows the bounded detection region zoomed in from the DESI Legacy Imaging Sky Viewer.

2. Transfer pre-processed images to Google Drive.

3. Run detection with Mask R-CNN model in the Google Colaboratory Pro environment.

4. Collect results and display pre-processed image along with cutouts from DESI Legacy Imaging Sky Viewer[4] on a website.

5. Visually inspect results and select potential LSBGs.

Figure 3 schematically shows this pipeline in action for a successful LSBG detection. We initially ran detection on 100 tiles and got 109 claimed LSBG detections, 1 of which showed real promise of being a true LSBG. This was sufficient information to estimate the time and human effort it would take to run this pipeline on the entire dataset as well as what to expect in terms of LSBG detection efficiency. Given the limited time in the internship, we decided that rather than further refining the model, it would be more fruitful to run the entire DES dataset.

4. Results

In this section we summarize the results of running the detection pipeline on all 10,169 DES tiles in DES DR2. The model returned 13,336 results for the LSBG class of which $\sim$ 11,000 have been visually inspected. We divided LSBG detections into two tiers after visual inspection based on our confidence that they are truly LSBGs: tier 1, containing the most promising candidates, and tier 2, the candidates we suspect are LSBGs but are unsure. We found 22 tier 1 LSBGs and 19 tier 2 LSBGs. Figure A.5 in Appendix A shows the zoomed in bounding box of 20 out of the 22 tier 1 LSBG candidates in the DESI Legacy Imaging Surveys Sky Viewer (the same as panel (d) in Fig. 3). These galaxies were generally larger (with most having angular sizes > 20 arcsec) and lower surface brightness than those detected in previous DES LSBG catalogs (e.g., Tanoglidis et al., 2020).

The vast majority of the model's output labelled as LSBGs were in actuality artifacts. The most prominent artifact class consisted of Galactic cirrus, starlight scattered off of the interstellar dust surrounding our Galaxy. Contamination from Galactic cirrus is a well-known barrier for detecting the low surface brightness universe (Eg. being confused with faint tidal features; Cortese et al., 2010). The model also detected a number of tidal features associated with stars, galaxies, and mergers. After visually inspecting the images, we conclude that this is likely because these features are not completely characterized by `SourceExtractor` in the DES pipeline, and thus when the compact sources are removed, their tidal remnants remain and become amplified by the pre-processing. We hoped to prevent this from occurring with the binary dilation step in the image pre-processing, which was implemented to cover as much of the compact sources in the mask as possible. However, we were not able to completely remove all features from these compact sources, hence their prominence in the model's output. The same is true for very large, bright stars and galaxies. Our initial models classified a high volume of these artifacts as LSBGs; however, they

[4] http://legacysurvey.org/

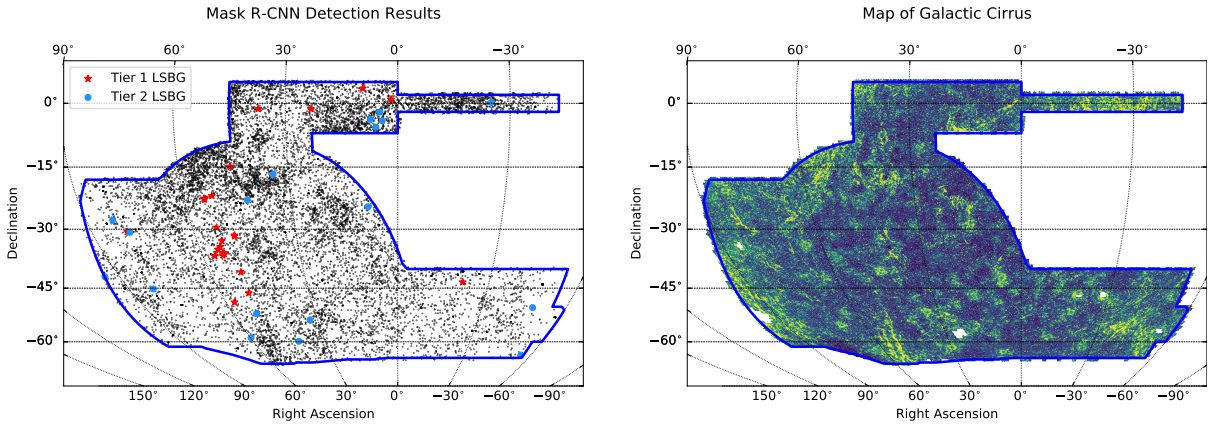

Figure 4: Left panel: Map of all Mask R-CNN model detections (black dots) along with tier 1 and tier 2 LSBG detections. Right panel: map of Galactic cirrus, where brighter regions represent higher detection densities. In both panels, the blue region represents the region of observation of the Dark Energy Survey.

were much less prominent in the pipeline as their addition as a training class greatly improved the purity of the results.

In Fig. 4, we plot the spatial distribution of all our detections (with the tier 1 and tier 2 LSBGs highlighted) as compared to a map of the Galactic cirrus in the DES footprint.[5] As mentioned previously, our visual inspection showed a high incidence of Galactic cirrus in the pipeline output. We therefore decided to compare our results to a map of the Galactic cirrus which showed a very high visual correlation in most areas, confirming the model's apparent propensity for detecting cirrus. It is interesting to note that some regions appear to visually lack a correlation to the map. Specifically, the left and lower right of the DES footprint appear to be lacking in our detection results while the map of Galactic cirrus seems to suggest high densities in this region. This is likely a result of the fact that these regions are closer to the Galactic plane and thus the model's detections are most likely dominated by artifacts of very bright, nearby stars.

Another intriguing feature of Fig. 4 is the cluster of tier 1 LSBG detections in the region around (RA, DEC) $= (55°, -35°)$. This is approximately the location of the nearby Fornax cluster of galaxies. The previous DES LSBG study in Tanoglidis et al. (2021) showed prominent clustering of red LSBGs around the Fornax cluster, a fact our model verifies.

[5]Cirrus map courtesy of the DES Collaboration (private communication).

5. Conclusions

In this work we applied the deep learning framework Mask R-CNN to the problem of detecting large low surface brightness galaxies in the Dark Energy Survey Data Release 2 using a simulated, labelled dataset of LSBGs generated with GALSIM with a range of physical parameters. We trained on 94 tiles with 16 LSBGs simulated per tile, with $\sim 70\%$ for training, $\sim 15\%$ for validation and $\sim 15\%$ for testing. We then developed a detection pipeline which we tested on $\sim 1\%$ of the DES dataset (100 tiles). We then ran the pipeline on the entire DR2 dataset (10,169 tiles), identifying 22 high-quality LSBGs. These galaxies were generally larger and lower surface brightness than those detected by conventional algorithms. In addition, we found the model to be very efficient at detecting Galactic cirrus and tidal features. The spatial distribution of all the detections appears to have a strong visual correlation with maps of Galactic cirrus. We have thus demonstrated that the Mask R-CNN framework could be very beneficial for detecting low surface brightness features.

Appendix A. LSBG Candidates

In this appendix, we present 20 high probability LSBG detections from the pipeline in Section 3.2. Figure A.5 contains images from the DESI Legacy Imaging Sky Viewer along with the right ascension and declination of 20 tier 1 LSBG candidates out of the 22. These were identified as high quality LSBG candidates upon visual inspection after being detected in the pipeline. Note that the brightness and contrast has been increased to accentuate LSBGs.

References

Abraham, R.G., van Dokkum, P.G., 2014. Ultra-Low Surface Brightness Imaging with the Dragonfly Telephoto Array. PASP 126, 55. 1401.5473.

Bothun, G., Impey, C., McGaugh, S., 1997. Low-surface-brightness galaxies: Hidden galaxies revealed. Publications of the Astronomical Society of the Pacific 109, 745.

Collaboration, D., Abbott, T.M.C., Adamow, M., Aguena, M., Allam, S., Amon, A., Annis, J., Avila, S., Bacon, D., Banerji, M., Bechtol, K., Becker, M.R., Bernstein, G.M., Bertin, E., Bhargava, S., Bridle, S.L., Brooks, D., Burke, D.L., Rosell, A.C., Kind, M.C., Carretero, J., Castander, F.J., Cawthon, R., Chang, C., Choi, A., Conselice, C., Costanzi, M., Crocce, M., da Costa, L.N., Davis, T.M., Vicente, J.D., DeRose, J., Desai, S., Diehl, H.T., Dietrich, J.P., Drlica-Wagner, A., Eckert, K., Elvin-Poole, J., Everett, S., Evrard, A.E., Ferrero, I., Ferté, A., Flaugher, B., Fosalba, P., Friedel, D., Frieman, J., García-Bellido, J., Gelman, L., Gerdes, D.W., Giannantonio, T., Gill, M., Gruen, D., Gruendl, R.A., Gschwend, J., Gutierrez, G., Hartley, W.G., Hinton, S.R., Hollowood, D.L., Honscheid, K., Huterer, D., James, D.J., Jeltema, T., Johnson, M.D., Kent, S., Kron, R., Kuehn, K., Kuropatkin, N., Lahav, O., Li, T.S., Lidman, C., Lin, H., MacCrann, N., Maia, M.A.G., Manning, T., March, M., Marshall, J.L., Martini, P., Melchior, P., Menanteau, F., Miquel, R., Morgan, R., Myles, J., Neilsen, E., Ogando, R.L.C., Palmese, A., Paz-Chinchón, F., Petravick, D., Pieres, A., Plazas, A.A., Pond, C., Rodriguez-Monroy, M., Romer, A.K., Roodman, A., Rykoff, E.S., Sako, M., Sanchez, E., Santiago, B., Serrano, S., Sevilla-Noarbe, I., Smith, J.A., Smith, M., Soares-Santos, M., Suchyta, E., Swanson, M.E.C., Tarle, G., Thomas, D., To, C., Tremblay, P.E., Troxel, M.A., Tucker, D.L., Turner, D.J., Varga, T.N., Walker, A.R., Wechsler, R.H., Weller, J., Wester, W., Wilkinson, R.D., Yanny, B., Zhang, Y., Nikutta, R., Fitzpatrick, M., Jacques, A., Scott, A., Olsen, K., Huang, L., Herrera, D., Juneau, S., Nidever, D., Weaver, B.A., Adean, C., Correia, V., de Freitas, M., Freitas, F.N., Singulani, C., Vila-Verde, G., 2021. The dark energy survey data release 2. 2101.05765.

Cortese, L., Bendo, G., Isaak, K., Davies, J., Kent, B., 2010. Diffuse far-infrared and ultraviolet emission in the ngc 4435/4438 system: Tidal stream or galactic cirrus. Monthly Notices of the Royal Astronomical Society: Letters 403.

Dalcanton, J.J., Spergel, D.N., Summers, F.J., 1997. The Formation of Disk Galaxies. ApJ 482, 659–676. astro-ph/9611226.

DES Collaboration, 2005. The Dark Energy Survey. arXiv e-prints , astro–ph/0510346astro-ph/0510346.

DES Collaboration, Abbott, T.M.C., Abdalla, F.B., Allam, S., Amara, A., Annis, J., Asorey, J., Avila, S., Ballester, O., Banerji, M., Barkhouse, W., Baruah, L., Baumer, M., Bechtol, K., Becker, M.R., Benoit-Lévy, A., Bernstein, G.M., Bertin, E., Blazek, J., Bocquet, S., Brooks, D., Brout, D., Buckley-Geer, E., Burke, D.L., Busti, V., Campisano, R., Cardiel-Sas, L., Carnero Rosell, A., Carrasco Kind, M., Carretero, J., Castander, F.J., Cawthon, R., Chang, C., Chen, X., Conselice, C., Costa, G., Crocce, M., Cunha, C.E., D'Andrea, C.B., da Costa, L.N., Das, R., Daues, G., Davis, T.M., Davis, C., De Vicente, J., DePoy, D.L., DeRose, J., Desai, S., Diehl, H.T., Dietrich, J.P., Dodelson, S., Doel, P., Drlica-Wagner, A., Eifler, T.F., Elliott, A.E., Evrard, A.E., Farahi, A., Fausti Neto, A., Fernandez, E., Finley, D.A., Flaugher, B., Foley, R.J., Fosalba, P., Friedel, D.N., Frieman, J., García-Bellido, J., Gaztanaga, E., Gerdes, D.W., Giannantonio, T., Gill, M.S.S., Glazebrook, K., Goldstein, D.A., Gower, M., Gruen, D., Gruendl, R.A., Gschwend, J., Gupta, R.R., Gutierrez, G., Hamilton, S., Hartley, W.G., Hinton, S.R., Hislop, J.M., Hollowood, D., Honscheid, K., Hoyle, B., Huterer, D., Jain, B., James, D.J., Jeltema, T., Johnson, M.W.G., Johnson, M.D., Kacprzak, T., Kent, S., Khullar, G., Klein, M., Kovacs, A., Koziol, A.M.G., Krause, E., Kremin, A., Kron, R., Kuehn, K., Kuhlmann, S., Kuropatkin, N., Lahav, O., Lasker, J., Li, T.S., Li, R.T., Liddle, A.R., Lima, M., Lin, H., López-Reyes, P., MacCrann, N., Maia, M.A.G., Maloney, J.D., Manera, M., March, M., Marriner, J., Marshall, J.L., Martini, P., McClintock, T., McKay, T., McMahon, R.G., Melchior, P., Menanteau, F., Miller, C.J., Miquel, R., Mohr, J.J., Morganson, E., Mould, J., Neilsen, E., Nichol, R.C., Nogueira, F., Nord, B., Nugent, P., Nunes, L., Ogando, R.L.C., Old, L., Pace, A.B., Palmese, A., Paz-Chinchón, F., Peiris, H.V., Percival, W.J., Petravick, D., Plazas, A.A., Poh, J., Pond, C., Porredon, A., Pujol, A., Refregier, A., Reil, K., Ricker, P.M., Rollins, R.P., Romer, A.K., Roodman, A., Rooney, P., Ross, A.J., Rykoff, E.S., Sako, M., Sanchez, M.L., Sanchez, E., Santiago, B., Saro, A., Scarpine, V., Scolnic, D., Serrano, S., Sevilla-Noarbe, I., Sheldon, E., Shipp, N., Silveira, M.L., Smith, M., Smith, R.C., Smith, J.A., Soares-Santos, M., Sobreira, F., Song, J., Stebbins, A., Suchyta, E., Sullivan, M., Swanson, M.E.C., Tarle, G., Thaler, J., Thomas, D., Thomas, R.C., Troxel, M.A., Tucker, D.L., Vikram, V., Vivas, A.K., Walker, A.R., Wechsler, R.H., Weller, J., Wester, W., Wolf, R.C., Wu, H., Yanny, B., Zenteno, A., Zhang, Y., Zuntz, J., DES Collaboration, Juneau, S., Fitzpatrick, M., Nikutta, R., Nidever, D., Olsen, K., Scott, A., NOAO Data Lab, 2018. The Dark Energy Survey: Data Release 1. ApJS 239, 18. 1801.03181.

Disney, M.J., 1976. Visibility of galaxies. Nature 263, 573–575.

He, K., Gkioxari, G., Dollár, P., Girshick, R., 2018. Mask r-cnn. 1703.06870.

Rowe, B.T.P., Jarvis, M., Mandelbaum, R., Bernstein, G.M., Bosch, J., Simet, M., Meyers, J.E., Kacprzak, T., Nakajima, R., Zuntz, J., Miyatake, H., Dietrich, J.P., Armstrong, R., Melchior, P., Gill, M.S.S., 2015. GALSIM: The modular galaxy image simulation toolkit. Astronomy and Computing 10, 121–150. 1407.7676.

Tanoglidis, D., Drlica-Wagner, A., Wei, K., Li, T.S., Sánchez, J., Zhang, Y., Peter, A.H.G., Feldmeier-Krause,

A., Prat, J., Casey, K., Palmese, A., Sánchez, C., DeRose, J., Conselice, C., Gagnon, L., Abbott, T.M.C., Aguena, M., Allam, S., Avila, S., Bechtol, K., Bertin, E., Bhargava, S., Brooks, D., Burke, D.L., Rosell, A.C., Kind, M.C., Carretero, J., Chang, C., Costanzi, M., da Costa, L.N., De Vicente, J., Desai, S., Diehl, H.T., Doel, P., Eifler, T.F., Everett, S., Evrard, A.E., Flaugher, B., Frieman, J., García-Bellido, J., Gerdes, D.W., Gruendl, R.A., Gschwend, J., Gutierrez, G., Hartley, W.G., Hollowood, D.L., Huterer, D., James, D.J., Krause, E., Kuehn, K., Kuropatkin, N., Maia, M.A.G., March, M., Marshall, J.L., Menanteau, F., Miquel, R., Ogando, R.L.C., Paz-Chinchón, F., Romer, A.K., Roodman, A., Sanchez, E., Scarpine, V., Serrano, S., Sevilla-Noarbe, I., Smith, M., Suchyta, E., Tarle, G., Thomas, D., Tucker, D.L., Walker, A.R., DES Collaboration, 2021. Shadows in the Dark: Low-surface-brightness Galaxies Discovered in the Dark Energy Survey. ApJS 252, 18. 2006.04294.

Tanoglidis, D., Ćiprijanović, A., Drlica-Wagner, A., 2020. Deepshadows: Separating low surface brightness galaxies from artifacts using deep learning. 2011.12437.

van Dokkum, P., Danieli, S., Cohen, Y., Merritt, A., Romanowsky, A.J., Abraham, R., Brodie, J., Conroy, C., Lokhorst, D., Mowla, L., O'Sullivan, E., Zhang, J., 2018. A galaxy lacking dark matter. Nature 555, 629–632.

Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., SciPy 1.0 Contributors, 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. Nature Methods 17, 261–272.

Wechsler, R.H., Tinker, J.L., 2018. The Connection Between Galaxies and Their Dark Matter Halos. ARA&A 56, 435–487. 1804.03097.

Zhong, G.H., Liang, Y.C., Liu, F.S., Hammer, F., Hu, J.Y., Chen, X.Y., Deng, L.C., Zhang, B., 2008. A large sample of low surface brightness disc galaxies from the SDSS - I. The sample and the stellar populations. MNRAS 391, 986–999. 0809.3099.

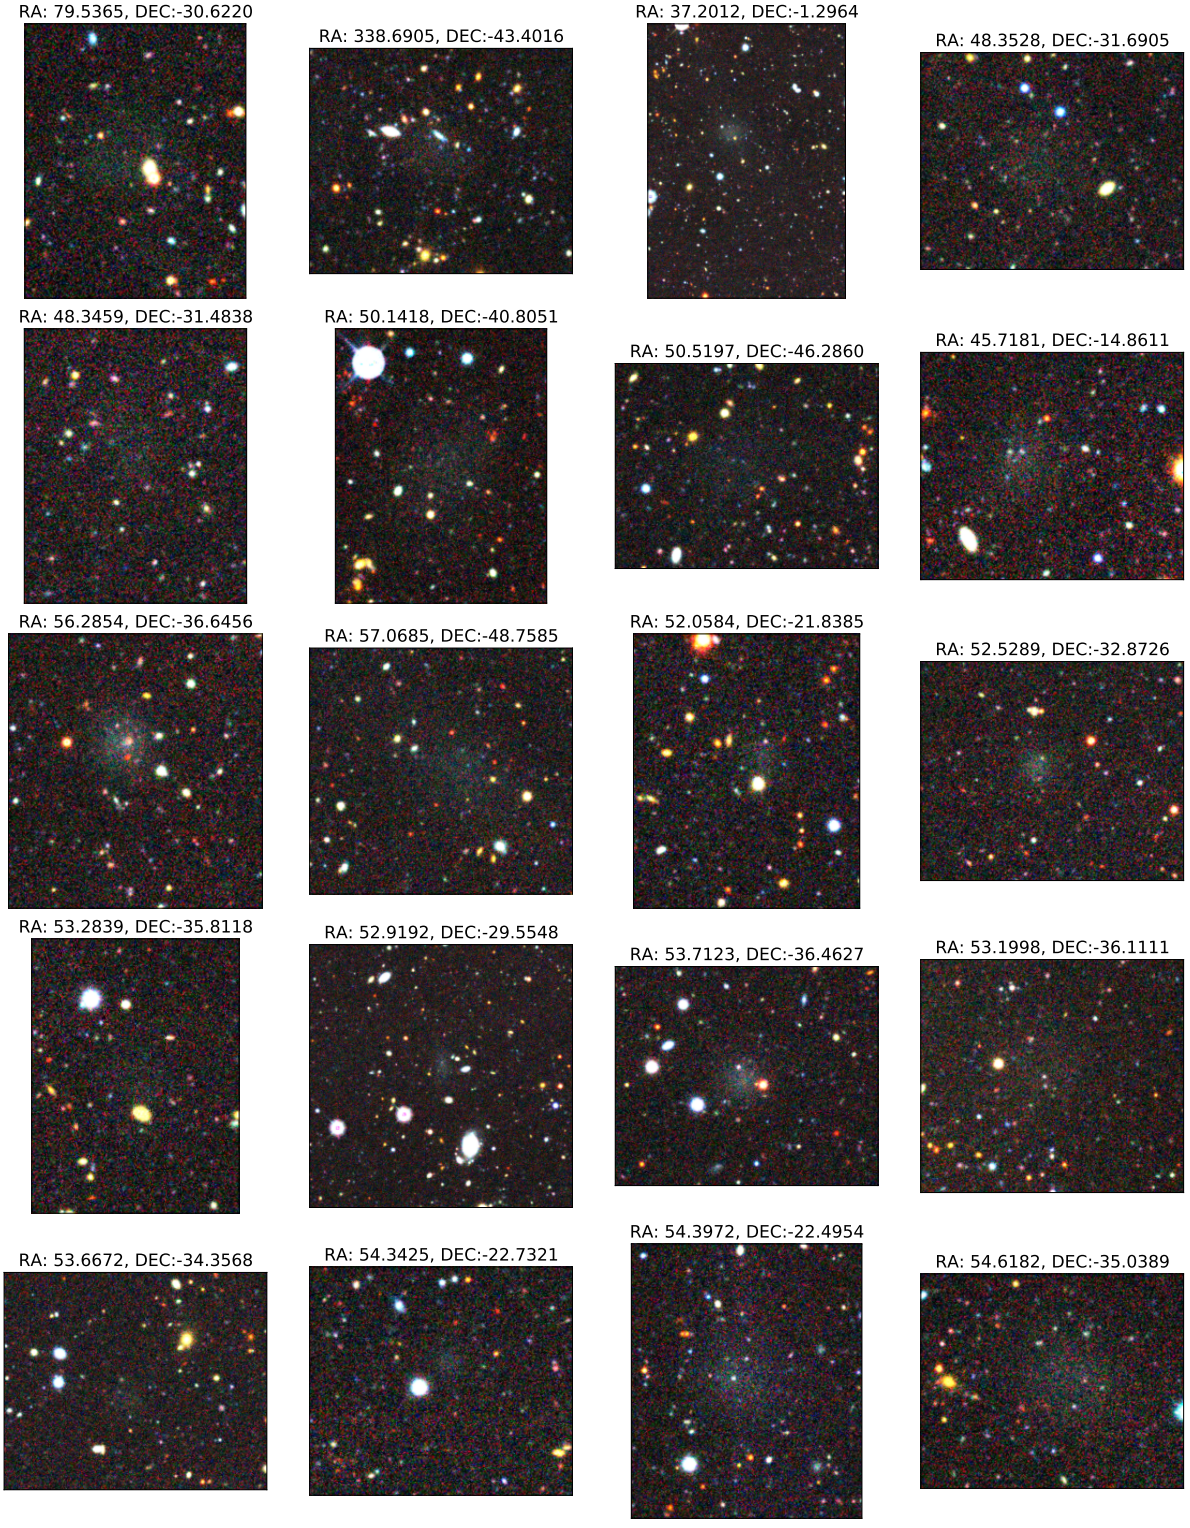

Figure A.5: Twenty tier 1 LSBG candidates from detection ran on 10,169 tiles of DES DR2 using the Mask R-CNN Model. The right ascension and declination coordinates are provided at the top of the image.