

Comparison of LSBGnet to MTO for Finding Low-Surface-Brightness Galaxies

Belu Oana Miruna
Student-number: S5061121

Bachelor Project, Computing Science
Course code: WBCS13000

GENERAL INFORMATION

First supervisor: Michael Wilkinson
Second supervisor: Kerstin Bunte

1 ABSTRACT

Low-surface-brightness galaxies (or LSBGs) present a significant challenge for detection due to their low central surface brightness. Nevertheless, discovering the most effective methods to detect them is crucial for our advancement in understanding our Universe's composition and dynamics.

This project presents a thorough comparison between LSBGnet, a neural network that specializes in the automatic LSBG detection and MTO, a Max-Tree-based approach for detecting astronomical sources, the latter one being able to detect all astronomical objects, not just low-surface-brightness galaxies. The performance of both algorithms is evaluated based on their detection rates across key galaxy parameters, including central surface brightness, effective radius, ellipticity, halo-to-disk ratio, and background density.

We find that LSBGnet has an overall detection completeness of 97%, compared to 82% for MTO - with the highest gap at $\mu_0 = 26 \text{ mag arcsec}^{-2}$ - the faintest end at which we tested, while MTO recovers the injected galaxies more faithfully. We also observe that both detection pipelines have zero False Positives when feeding them pure-noise images.

2 INTRODUCTION AND MOTIVATION

Low-surface-brightness galaxies (LSBGs) are galaxies whose central surface brightness is fainter than the one of a "normal" galaxy, often by several magnitudes, making them among the most diffuse galaxies known. LSBGs present a significant observational challenge due to their diffuse luminosity. However, they are estimated to constitute between 30% and 60% of the local galaxy number density and contribute up to 20% of the total dynamical mass [18]. Moreover, they are believed to be one of the main sources of baryons in the Universe - the "ordinary" matter, made of protons and neutrons [12] - and play an important role in the process of understanding Dark Matter, according to [18]. Therefore, identifying the most effective detection method for LSBGs is crucial to advance our understanding of the composition and dynamics of the Universe.

Before assessing any detection algorithms, we must first establish a precise definition of a LSBG. The main criterion is having a central surface brightness notably fainter than the Freeman value of $\mu_0 = 21.65 \pm 0.30$ B -mag arcsec $^{-2}$ [3]. Freeman showed that even though galaxies may differ in total brightness, the central brightness of their disk regions is nearly the same - around the Freeman value [8]. In simpler terms, this means that most galaxies have a “standard” level of brightness at their centers and any galaxy with a significantly fainter center can be considered a low-surface-brightness galaxy. The faintness is a consequence of their low star formation rate, which arises from a gas surface density below the critical threshold required for sustained star formation [3]. Another important characteristic of low-surface-brightness galaxies is their ellipticity, a study across different LSBG catalogs showed that their ellipticities span a broad range, from nearly circular ($\epsilon \approx 0$) to very elongated ($\epsilon \approx 0.8$)[13]. Their effective radius, which is defined as the radius encompassing half of the galaxy’s total light has been shown in [10] to have values between 2 and 14 arcsecs, the majority spanning between 3 and 9 arcsecs. Furthermore, studies had demonstrated that LSBGs are best described by an exponential disk, coupled with a faint and extended halo [2]. Photometric decompositions in recent studies show that the disks are accurately described by Sérsic profiles with n between 0.5 and 1.5, the majority having $n < 1$ [17]. The Sérsic profile is a method to parametrize the surface-brightness distribution in galaxies, where n controls the concentration of the light [1].




Figure 1: Sérsic profiles for various n . [4]

This thesis aims to provide a comprehensive comparison of two methods for detecting low-surface-brightness galaxies: Low Surface Brightness Galaxy net (LSBGnet), a neural network specifically designed for LSBG detection [18], and Max-Tree Objects (MTO), a Max-Tree-based approach capable of identifying all faint astronomical objects [19]. The analysis will evaluate these methods in terms of detection efficiency, false positive rates, and recovery fractions. The outcome of this study is expected to improve future detection strategies for LSBGs, subsequently refining our understanding of galaxy formation and evolution.




Figure 2: Examples of LSBGs [3]

Most LSBGs exhibit blue optical colors - an indicative of scattered star formation - while a smaller red subset shows more uniform star formation [23].

3 STATE OF THE ART

3.1 LSBGNET

One of the two technologies that will be compared using several quality criteria is LSBGnet. LSBGnet is a recently published (2024) neural network specialized in the automatic detection of low-surface-brightness galaxies [18]. Its framework consists of four main steps: image data augmentation, building an LSBG detection network, defining a loss function, and optimizing network parameters and model performance [18]. The data augmentation used consists of scaling, flipping, color gamut transformation (the process of changing colors from one gamut to another, while maintaining accuracy), and mosaic data augmentation, which stitches four images together, in order to make the location of LSBGs more random [18]. The resulting image after this process is then overlapped with a gray background image [18].

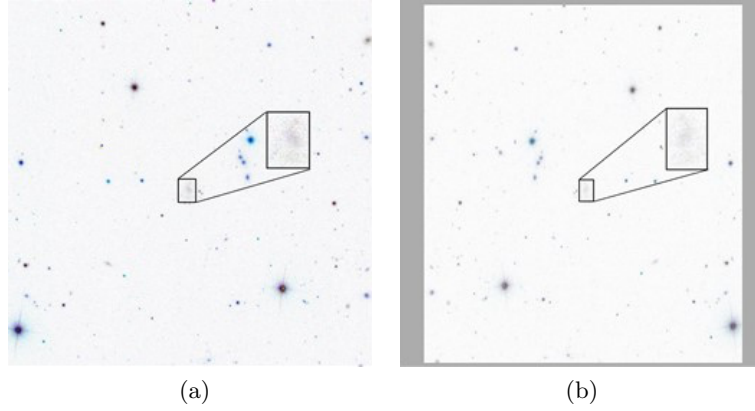


Figure 3: Comparison between original image (a) and image after data augmentation (b). The object in the box is the LSBGs [18].

The detection network is structured in four modules, as it can be observed in Figure 3. [18].

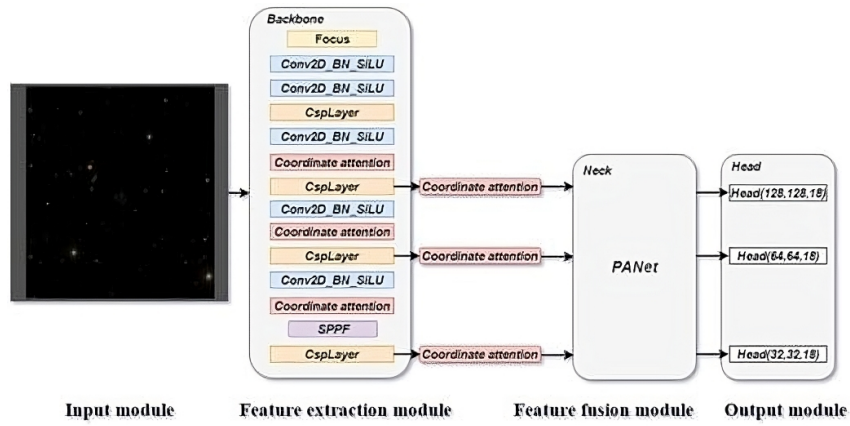


Figure 4: The network structure of LSBGnet [18].

The **Input module** is an RGB (red, green, and blue) image, that is enhanced through the data augmentation process [18]. The **Feature extraction module** (Backbone) is an improved version of CSPDarknet [18], the backbone of YOLO-v4 (the fourth version in the You Only Look Once family of models) [5].

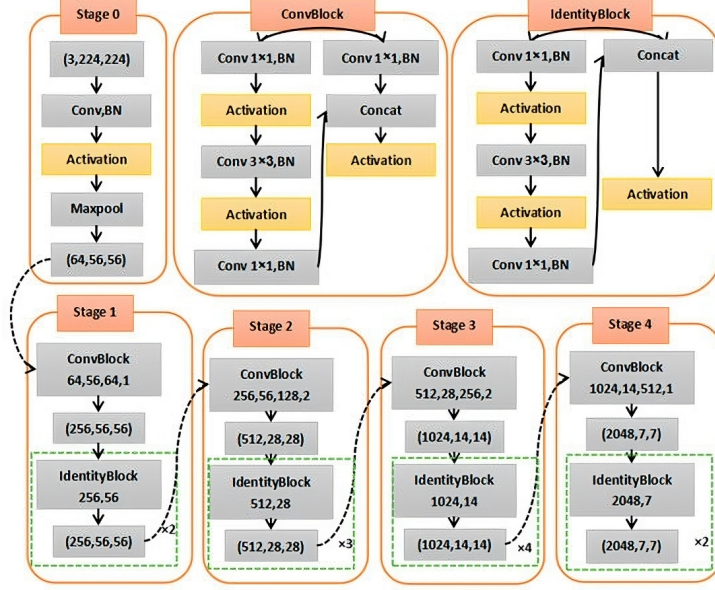


Figure 5: The structure of CSPDarknet [5].

Cross-Stage-Partial-Network, or CSPDarknet, can enhance the learning capacities of convolutional neural networks (CNNs), increasing their efficiency in image detection [21]. After passing through the Feature Extraction module, it outputs three feature layers, of sizes 128 x 128, 64 x 64 and 32 x 32 [18].

The **Feature fusion module** (Neck) is used for fusing information from the Backbone, before sending it further [18]. This is done using PANet (Path Aggregation Network), [18], an enhancer for the representation capabilities of the backbone, by fusing the bottom-up and top-down paths, therefore managing to avoid rigid matching of target size and network depth [6].

The **Output module** (Head) interprets the network’s predictions and generates detection outputs at three distinct scales: (128,128,18), (64,64,18), and (32,32,18) [18]. These scales refer to the partitioning of the original 1024 x 1024 input image into grids of sizes 128 x 128, 64 x 64, and 32 x 32 cells, respectively. Each cell in these grids corresponds to a specific region of the input image and it captures objects of different sizes: small, medium, and large [18]. In each grid cell, there are three anchor boxes [18]. Each anchor box encodes six attributes: class, confidence, center coordinates, width, and height of the box [18] - thus, there will be 18 dimensions for each grid (3 anchor boxes \times 6 attributes). The Head maps these outputs back to the original image to generate detection results [18]. The loss function of LSBGnet consists of confidence, classification, and localization losses [18]. BCELoss (Binary Cross-Entropy Loss):

$$\text{BCE} = -\frac{1}{N} \sum_{i=1}^N \left[y_i \log(p_i) + (1 - y_i) \log(1 - p_i) \right],$$

where:

N = number of observations,

$y_i \in \{0, 1\}$ is the actual binary label for observation i ,

$p_i \in (0, 1)$ is the predicted probability that $y_i = 1$.

is used for confidence and classification, while GIoU (Generalized intersection over union):

$$\text{GIoU}(A, B) = \text{IoU}(A, B) - \frac{|C \setminus (A \cup B)|}{|C|},$$

where:

$$C = \text{ConvexHull}(A \cup B)$$

(often the smallest enclosing box) improves bounding box quality [18]. The network was trained using the Adam optimizer to detect LSBGs [14]. The model outputs class, confidence, and bounding box coordinates [18]. A confidence threshold was set based on the distribution of LSBG candidates, and non-maximum suppression was applied to finalize detections [18].

3.2 MTO

The LSBGnet model will be compared to MT-Objects (MTO), which outperformed other methods (SExtractor, ProFound, NoiseChisel) in a recent study [11]. MTO is a Max-Tree-based approach for detecting faint extended astronomical sources without the need for strong image smoothing [19]. Unlike traditional methods such as SExtractor, which struggle with low-intensity sources, MTO utilizes a hierarchical Max-Tree representation, where an image is decomposed into nested connected components at different intensity levels [19]. Object detection is performed by analyzing statistical distributions within the tree nodes, distinguishing real sources from noise [19]. Figure 6 showcases the core of MTO, (a) shows the original 2D image, (b) extracts the peak connected components at a chosen threshold, and (c) organize the found components into a Max-Tree.

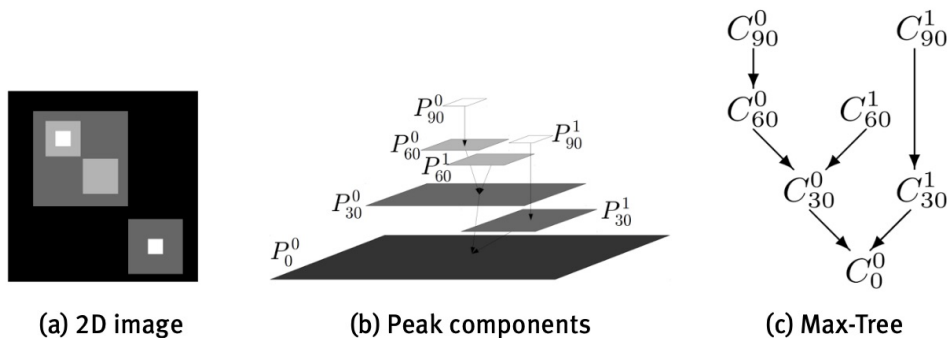


Figure 6: (a) - 2D image with intensities from 0 to 90, (b) - its peak components P_h^k at intensity h , and (c) - the Max-Tree nodes C_h^k [19].

In this framework, the tool takes as input an image, assumed to be flat, with a gain of g , that has three different components: *Gaussian background noise*, *Poisson distributed signal*, and *other source noises* [20]. The algorithm consists of five steps[20]. **The first one** is background estimation: the image is divided into tiles of 64x64 pixels [20]. Each tile is tested for Gaussianity and flatness [20]. The tiles are iteratively adjusted by doubling or halving [20]. The tiles that pass the tests are selected to represent the background and the mean value of the tiles is removed from the image [20]. The background becomes the root of the tree [20]. **The second one** is identifying the significant branches by incrementally thresholding the image and higher intensity level and at each level the maximum connected structures are recognized and then incorporated as branches in the tree [20]. **The third one** is labeling the branches so that all that belong to the same object get the same label [20]. **The fourth one** is projecting the hierarchical tree into a plane with the same size as the input image, in order to create a segmentation map [20]. **The fifth one** is parameter extraction, which computes structural parameters using the labeled pixels from the segmentation map [20]. For nested objects, to prevent bias in the surface brightness that we measure, the parent's node intensity is subtracted before calculating the parameters [20]. The output of the MTO is a segmentation map and an object list that includes the location, magnitude, size, and median surface brightness of the detected objects [20].

A recent advancement of the MTO method is MMTO (multi-band max-tree objects) [7], a new algorithm which processes multi-spectral images. Unlike MTO, which processes gray-scale images [19], MMTO incorporates cross-band correlations by structuring multiple Max-Trees

into a semantically-informed graph [7]. Given that LSBGnet also operates on color images [18], MMTO is a potential alternative for multi-band low-surface-brightness galaxy detection.

4 SOURCE CODE REPOSITORIES

We obtained the LSBGnet code from its author, and the MTO code from the public GitHub repository at <https://github.com/CarolineHaigh/mtobjects>.

5 METHODS

5.1 DATA SELECTION

This section describes the technique used for the selection and retrieval of background fields from the Sloan Digital Sky Survey (SDSS) in which synthetic LSBGs will be injected. The procedure ensures sampling across a range of sky densities and produces images that are free of gross artifacts, missing data or defects - such as blurred corners or visible seams. To guarantee high quality images, all candidate fields were visually inspected, and the two best cutouts from each density category were retained.

5.1.1 SDSS FIELD SAMPLING

We randomly selected 400 sky positions uniformly in right ascension $0^\circ \leq \alpha \leq 360^\circ$ and declination $-60^\circ \leq \delta \leq +60^\circ$ (celestial coordinates) using a random number generator with a fixed seed to ensure reproducibility. For each candidate position, we queried the SDSS PhotoObj catalog - that provides data on all detected sources - to retrieve all detections (stars, galaxies, artifacts) in that region. The density of the area was calculated as $\rho = \frac{N_{obj}}{\pi * (R/60)^2}$, where N_{obj} represents the count of objects detected and R represents the radius in which we searched for the objects.

5.1.2 DENSITY THRESHOLDS

To explore the detection performance of MTO and LSBGnet across different background densities, we sorted the 400 fields by increasing density and defined our "low" density as below the 25th percentile, the "medium" density below the 75th percentile and the "high" density above the 75th percentile. The computed percentiles were $\rho_{25} = 36160.003$ objects/deg² and $\rho_{75} = 121148.742$ objects/deg².




Figure 7: Distribution of SDSS Field Densities

We then attempted to select up to 20 candidates per field density. However, not all candidate fields yielded valid images, as some of them had missing bands. Fields without successful retrievals in all three bands (g, r, i) were excluded. This filtering step limited the number of usable steps per bin to fewer than originally requested.

5.1.3 VISUAL SELECTION

From the valid subset, we visually inspected the quality of each field, by generating RGB previews with the Lupton method. We excluded all fields containing: data gaps, truncation, strong gradients or visible seams. This step ensured that the backgrounds are free of anomalies that could bias detection outcomes.

Table 1: Summary of the six selected SDSS background fields used for LSBG injection.

Bin	Field ID	RA (deg)	DEC (deg)	Density (objects/deg ²)
Low	low_00	56.568	5.236	14897
	low_03	211.900	50.875	26738
Medium	mid_01	227.689	-1.593	50293
	mid_06	12.516	-2.260	55131
High	high_00	240.208	30.680	211230
	high_06	15.670	-0.083	790936

5.2 SYNTHETIC LSBG INJECTION

In order to evaluate the detection performance of MTO and LSBGnet under different conditions, we inject the LSBGs in the selected SDSS backgrounds. The injection pipeline is implemented using AstroPy for I/O, GalSim for galaxy modeling and convolution, and other standard libraries for image manipulation and annotation.

5.2.1 COMPOSITION AND EFFECTIVE RADIUS

Each injected galaxy is generated as the sum of two Sérsic components: a disk and a halo, by calling GalSim’s *Sérsic* twice, once for the disk with the Sérsic index between 0.5 and 1.5 and once for the halo with the Sérsic index fixed at 0.5. The half-light radius of the disk is in range {3, 5, 7, 9} arcsec, to cover a wide range of possible values, the halo’s half-light radius is set to twice that value. The fluxes of the two components are set by choosing a random value between {0.0, 0.3}, which will represent the halo’s contribution to the total flux, while the disk receives the remainder.

5.2.2 CENTRAL SURFACE BRIGHTNESS, COLOR AND ELLIPTICITY

Central surface brightness takes values in the range {23.0, 24.0, 25.0, 26.0} mag arcsec⁻². We read the r-band’s zero point from the image header (defaulting to 22.5 if it is missing) and use it to convert central surface brightness to total model flux. Small color offsets are then applied in the g and i bands (having a fixed offset of -0.3 in the i-band and +0.3 in g-band) and a Gaussian scatter of $\sigma = 0.05$ mag. Each galaxy is given an ellipticity from {0.1, 0.3, 0.5, 0.7} and a random position angle from $[0^\circ, 180^\circ]$.

5.2.3 BACKGROUND NORMALIZATION, INJECTION AND PSF CONVOLUTION

The background needs to be normalized before injecting the low-surface-brightness galaxy in order to standardize the background to the same zero offset. The normalization of each SDSS sky is done by subtracting the global median sky level. Concretely, for each band we compute $\text{sky}_{\text{med}} = \text{median}\{I_{ij}\}$ where I_{ij} are the raw pixel values in the FITS (Flexible Image Transport

System) array. We then replace every pixel by $I'_{ij} = I_{ij} - \text{sky}_{\text{med}}$. Then we choose the injection coordinates at a random position, within a safe margin of four times the disk’s effective radius from the edges. For each unique combination of background density, central surface brightness, effective radius, ellipticity, we inject at five distinct positions, resulting in a total of $6 \cdot 4 \cdot 4 \cdot 4 \cdot 5 = 1920$ different images. To reproduce the blurring effects of atmospheric turbulence and telescope optics, we convolve the injected galaxy with a Moffat Point Spread Function (PSF) [15] with $\beta = 4$ [16] and FWHM (full width at half maximum) of $1.2''$, which is below the $1.5''$ quality cutoff [22].

5.2.4 PREPARING DATA FOR LSBGNET AND MTO

After injecting each synthetic LSBG into all three bands, we generate two distinct data products of the same model, to match the requirements of the two detection pipelines. To ensure both models receive the same photon count, the first step is to convert each band into a gray-scale image. Before this conversion, we apply a percentile clipping (1st to 99th) to reduce the influence or extreme outliers. For MTO, which expects a single-channel FITS file, we sum the three gray-scale bands into a single image and save it as a FITS file. For LSBGnet, which expects three-channels JPG inputs, we use the same gray-scale images as three color bands, forming a RGB image. All injection metadata is saved to a CSV file that is later used for the assessment of the detection performance of both pipelines as functions of each parameter.


5.3 EXAMPLE PRODUCTS




(a) LSBGnet input before



(b) LSBGnet input after



(c) MTO input before (displayed in SAOImageDS9)



(d) MTO input after (displayed in SAOImageDS9)

Figure 8: Comparison of data products for LSBGnet (top row) and MTO (bottom row), before and after synthetic LSBG injection.

Figure 8 presents the before/after comparison for LSBGnet’s RGB inputs and MTO’s FITS for a synthetic galaxy injected in a background with low density.

5.4 DETECTION PIPELINES

5.4.1 LSBGNET

The evaluation of the LSBGnet pipeline was conducted in two phases: model training and predicting.

For training, we generated a dedicated dataset by applying our synthetic-galaxy injection procedure to a disjoint set of SDSS background fields, to ensure no overlap with the test images and an unbiased evaluation. Each JPG was paired with a Pascal VOC-format XML annotation, encoding a square bounding box of side length $6R_e$ (effective radius) around the injection center [9]. These pairs were organized into the VOC2007 directory structure (JPEGImages, Annotations, ImageSets) and partitioned into a 90% training and 10% validation split, yielding the files

for LSBGnet’s training.

Training ran for 200 epochs: during the first 50 epochs, the CSPDarkNet backbone remained frozen and only the detection heads were updated; in the following 150 epochs all layers were unfrozen. The best performing model checkpoint in the training was used for prediction.

In the prediction phase, the trained LSBGnet model was applied to the full set of test JPEGs (both before and after images), producing predicted bounding boxes and associated confidence scores for each detection. For each injected LSBG, we checked whether the predicted boxes in each ”before” and ”after” image contained the known injection center. The prediction was considered successful if the LSBG was found at that location in the ”after” image, but not in the ”before” image.

5.4.2 VISUALIZATION OF LSBGNET OUTPUTS




Figure 9: True positive example 1: RGB preview before (left) and after (right) injection ($\mu_0 = 24.0$, $R_e = 5.0$, $e = 0.7$).




Figure 10: True positive example 2: RGB preview before (left) and after (right) injection ($\mu_0 = 26.0$, $R_e = 9.0$, $e = 0.1$).




Figure 11: True positive example 3: RGB preview before (left) and after (right) injection ($\mu_0 = 23.0$, $R_e = 5.0$, $e = 0.5$).



Figure 12: False negative example 1: RGB preview before (left) and after (right) injection ($\mu_0 = 25.0$, $R_e = 3.0$, $e = 0.7$).



Figure 13: False negative example 2: RGB preview before (left) and after (right) injection ($\mu_0 = 26.0$, $R_e = 3.0$, $e = 0.9$).




Figure 14: False negative example 3: RGB preview before (left) and after (right) injection ($\mu_0 = 26.0$, $R_e = 5.0$, $e = 0.7$).

5.4.3 MTO

We evaluated MTO by applying it - with its default parameters to the FITS files of the same testing set as LSBGnet. For each “before”/“after” pair, the algorithm produces segmentation maps. We declare a synthetic LSBG successfully detected if the injection coordinate lies outside all segments in the “before” map but falls inside at least one segment in the “after” map.

5.4.4 VISUALIZATION OF MTO OUTPUTS




Figure 15: True positive example 1: preview before (left) and after (right) injection ($\mu_0 = 23.0$, $R_e = 3.0$, $e = 0.7$).





Figure 16: True positive example 2: preview before (left) and after (right) injection ($\mu_0 = 26.0$, $R_e = 3.0$, $e = 0.5$).





Figure 17: True positive example 3: preview before (left) and after (right) injection ($\mu_0 = 24.0$, $R_e = 3.0$, $e = 0.3$).





Figure 18: False negative example 1: preview before (left) and after (right) injection ($\mu_0 = 26.0$, $R_e = 5.0$, $e = 0.7$).





Figure 19: False negative example 2: preview before (left) and after (right) injection ($\mu_0 = 26.0$, $R_e = 3.0$, $e = 0.7$).




Figure 20: False negative example 3: preview before (left) and after (right) injection ($\mu_0 = 26.0$, $R_e = 3.0$, $e = 0.3$).

6 RESULTS

We now present the performance of LSBGnet and MTO on the 1920 synthetic injections.

6.1 DETECTION OUTCOMES

We first quantify the ability of both detection pipelines to find the low-surface-brightness galaxies. Figure 21 shows the fraction of successful detections, misses, and uncertain cases for both MTO and LSBGnet, while Figure 22 details the distribution of "before" and "after" detections. To verify our false-positive rate, we additionally applied both algorithms to 500 pure-noise cutouts and observed zero spurious detections.




Figure 21: Overall detection rates of MTO (left) and LSBGnet (right).




Figure 22: Confusion Matrices for MTO (left) and LSBGnet (right).

6.1.1 STATISTICAL COMPARISON OF MTO AND LSBGNET

To quantify whether the LSBGnet's increased sensitivity to low-surface-brightness galaxy detection is statistically significant, we performed a McNemar test. Out of all 1920 injections, LSBGnet alone recovered 309 that MTO missed, and MTO alone recovered 13 that LSBGnet missed.

$$\chi^2 = \frac{(309 - 13)^2}{309 + 13} = 272.09, \quad p < 10^{-4},$$

demonstrating that LSBGnet's higher detection rate is statistically significant.

6.2 DETECTION RATES VS. INJECTION PARAMETERS

To understand how the detection performance depends on the galaxy and background properties, we plot the injection fraction as a function of each parameter.




Figure 23: Detection rates vs. μ_0 of MTO (left) and LSBGnet (right).

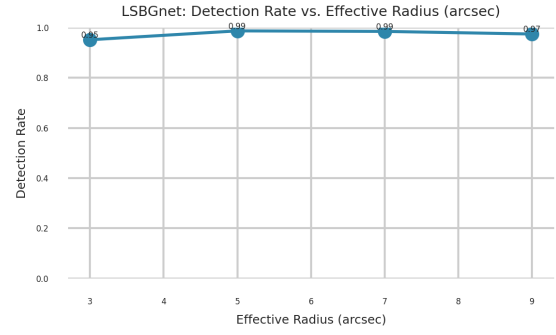
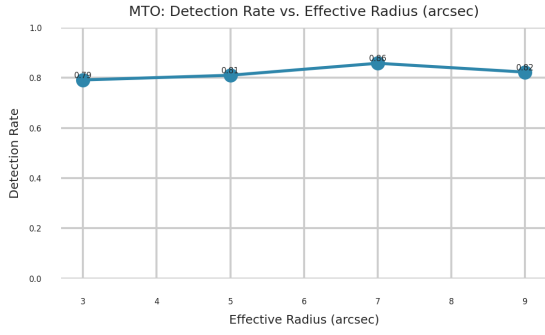


Figure 24: Detection rates vs. R_e of MTO (left) and LSBGnet (right).

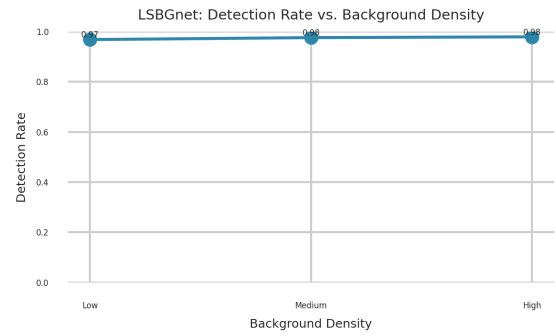



Figure 25: Detection rates vs. background density of MTO (left) and LSBGnet (right).

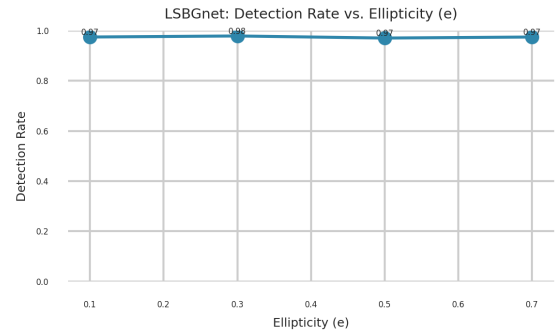
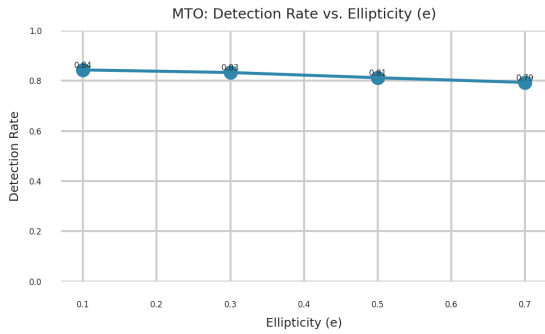


Figure 26: Detection rates vs. ellipticity of MTO (left) and LSBGnet (right).





Figure 27: Detection rates vs. Sérsic index of MTO (left) and LSBGnet (right).





Figure 28: Detection rates vs. halo flux fraction of MTO (left) and LSBGnet (right).

6.3 DISTRIBUTION HISTOGRAMS

Histograms of injection parameters distribution for detected, missed, and uncertain cases illustrates further where each detection pipeline succeeds or fails.





Figure 29: Histogram of μ_0 for detected and missed injections, for MTO (left) and LSBGnet (right).





Figure 30: Histogram of R_e for detected and missed injections, for MTO (left) and LSBGnet (right).




Figure 31: Histogram of background density for detected and missed injections, for MTO (left) and LSBGnet (right).




Figure 32: Histogram of ellipticity for detected and missed injections, for MTO (left) and LSBGnet (right).




Figure 33: Histogram of Sérsic index for detected and missed injections, for MTO (left) and LSBGnet (right).





Figure 34: Histogram of halo flux fraction for detected and missed injections, for MTO (left) and LSBGnet (right).

6.4 STACKED BARS

Stacked bar charts present the normalized fraction of successes, misses and uncertain cases in each parameter bin.

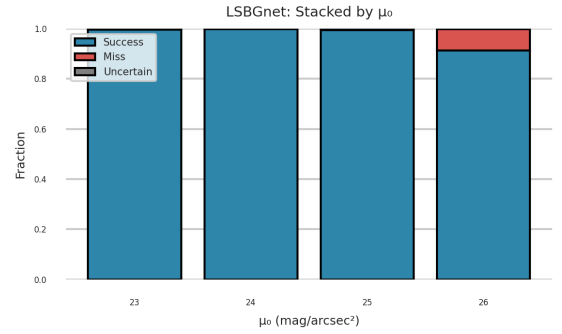
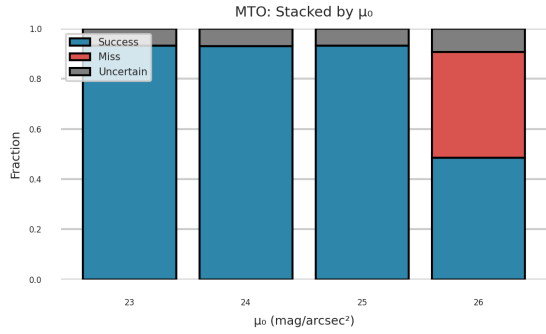


Figure 35: Stacked bar of μ_0 , for MTO (left) and LSBGnet (right).




Figure 36: Stacked bar of R_e , for MTO (left) and LSBGnet (right).




Figure 37: Stacked bar of ellipticity, for MTO (left) and LSBGnet (right).

6.5 COMPARATIVE METHOD PERFORMANCE

Comparing LSBGnet and MTO on the same plot highlights their strengths and weaknesses (Figure 38). The heatmaps in Figure 39 show the detection performance as a function of central surface brightness and effective radius and Figure 39 also plots which method succeeds without the other.




Figure 38: Detection fraction vs. μ_0 (left) and Detection fraction vs. R_e (right).




Figure 39: Heatmap of μ_0 vs. R_e for MTO (left), LSBGnet (right), and Detection overlap by μ_0 .

6.6 CONFIDENCE AND DETECTION PROBABILITY MODELING

We further explore the relation between LSBGnet's confidence and detection success and model the dependence of detection probability on the galaxy's parameters through logistic regression.

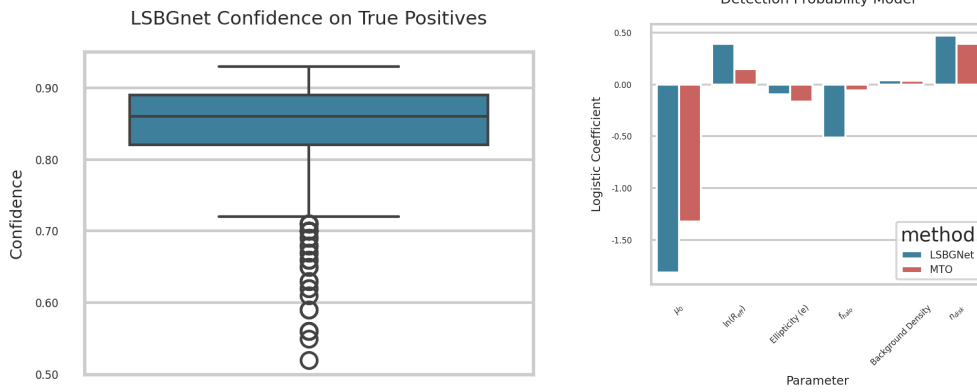


Figure 40: Confidence on True Positives (left) and Logistic Regression Model

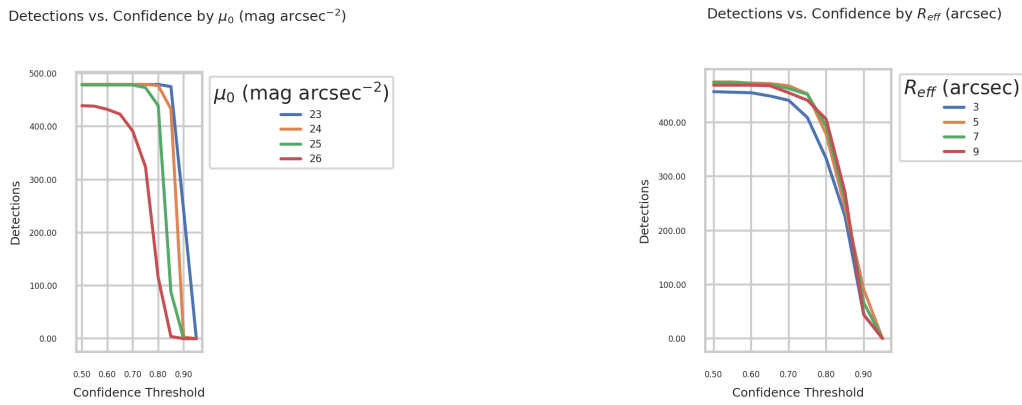


Figure 41: Confidence vs. μ_0 (left) and Confidence vs. R_e (right).

6.7 RECOVERY ACCURACY

Another key requirement of a good detection method for LSBGs is a faithful object reconstruction, therefore, Figure 42 showcases each method's spatial recovery by measuring the fraction of injected flux captured (Recovery Fraction) and Figure 43 demonstrates how the recovery fraction changes with the Sérsic index.




Figure 42: Recovery Fraction of MTO (left) and LSBGnet (right).




Figure 43: Recovery Fraction vs. Sérsic Index of MTO (left) and LSBGnet (right).

7 DISCUSSION AND CONCLUSION

7.1 KEY FINDINGS

The experiment, conducted on 1920 synthetic LSBG injections, revealed that LSBGnet consistently outperforms MTO in recovering low-surface-brightness galaxies across all parameter values. In particular, LSBGnet achieves an overall detection rate of 97% vs. 82% for MTO, with the largest performance gap occurring at $\mu_0 = 26.0 \text{ mag arcsec}^{-2}$, where LSBGnet has a success rate of 91% vs. MTO's, which is 49%. The effective radius has little influence on detection success for either LSBGnet or MTO, both methods have the lowest success rate at $R_e = 3 \text{ arcsec}$, remaining over 95% for LSBGnet and 79% for MTO at all effective radii. Both detection pipelines show little sensitivity to the density of the background and the ellipticity of the galaxy: success fractions vary by at most 3% between density bins and by 5% between ellipticity bins. Disk Sérsic index affects MTO more than LSBGnet: the detection success varying by 17% for the Max-Tree method, and only by 6% for LSBGnet. MTO is also more sensitive to the halo light fraction (the success rate varying by 14%) than LSBGnet (the success rate varying by 6%).

LSBGnet also assigns each detection a confidence score (Figure 40). True positives show a median confidence of around 0.86, indicating not only that LSBGnet detects fainter galaxies than MTO but does so at a high level of certainty.

The logistic regression model quantifies each parameter's importance: the central surface brightness (μ_0) having by far the biggest effect of both detection pipelines. Recovery fractions indicate that MTO has a higher chance to recover the whole object, detecting 100% of the low-surface-brightness galaxies in around 1200 cases, while LSBGnet recovers the whole galaxy in only around 500 cases.

7.2 METHOD COMPARISON

LSBGnet excels because of its deep learning backbone, trained on diverse synthetic LSBGs, so it can adapt better to faint galaxies, while MTO relies on intensity-threshold segmentation, which misses diffuse outskirts. However, MTO is faster on CPU and recovers the galaxy better, offering advantages when computer resources are limited.

7.3 LIMITATIONS AND FUTURE WORK

The two detection pipelines are tested on synthetically injected galaxies, testing on real survey data would lead to additional complications.

Applying this framework to the multi-band MTO on the g, r, i bands and comparing it to LSBGnet is a natural extension. Validation on real-sky surveys with known low-surface-brightness galaxies would further test real-world performance. Finally, exploring prefiltering methods could reduce MTO's failure at faint central surface brightnesses.

7.4 CONCLUSION

In summary, LSBGnet demonstrates superior sensitivity to detecting the injected synthetic low-surface-brightness galaxies, showing the most increased sensitivity compared to MTO at $\mu_0 = 26 \text{ mag arcsec}^{-2}$. MTO, while less sensitive, offers simplicity and better recovery. Choosing between them depends on the application: for surveys requiring maximal completeness at the cost of training overhead, LSBGnet is preferable, for faster scans, MTO remains a viable option.

REFERENCES

- [1] M. Baes and G. Gentile. “Analytical expressions for the deprojected Sérsic model”. In: *Astronomy & Astrophysics* 525 (Dec. 2010), A136. ISSN: 1432-0746. DOI: 10.1051/0004-6361/201015716. URL: <http://dx.doi.org/10.1051/0004-6361/201015716>.
- [2] Nils Bergvall, Erik Zackrisson, and Brady Caldwell. “The red haloes of SDSS low surface brightness disc galaxies”. In: 405.4 (July 2010), pp. 2697–2716. DOI: 10.1111/j.1365-2966.2010.16650.x. arXiv: 0909.4296 [astro-ph.CO].
- [3] Willem Jan Geert de Blok. “The Properties and Evolution of Low Surface Brightness Galaxies”. PhD thesis. University of Groningen, 1997. URL: <https://research.rug.nl/en/publications/the-properties-and-evolution-of-low-surface-brightness-galaxies>.
- [4] Claire Burke. “The Assembly of Stellar Mass in Galaxy Cluster Cores over Cosmic Time”. PhD thesis. Apr. 2013.
- [5] Enhui Chai et al. “An Efficient Asymmetric Nonlinear Activation Function for Deep Neural Networks”. In: *Symmetry* 14.5 (2022). ISSN: 2073-8994. DOI: 10.3390/sym14051027. URL: <https://www.mdpi.com/2073-8994/14/5/1027>.
- [6] Ran Chen et al. “PANetW: PANet with wider receptive fields for object detection”. In: *Multimedia Tools and Applications* 83 (Jan. 2024), pp. 1–22. DOI: 10.1007/s11042-024-18219-7.
- [7] Mohammad Faezi, Reynier Peletier, and Michael Wilkinson. “Multi-Spectral Source-Segmentation Using Semantically-Informed Max-Trees”. In: *IEEE Access* PP (Jan. 2024), pp. 1–1. DOI: 10.1109/ACCESS.2024.3403309.
- [8] K. C. Freeman. “On the Disks of Spiral and S0 Galaxies”. In: *Astrophysical Journal* 160 (June 1970), p. 811. DOI: 10.1086/150474.
- [9] Alister W. Graham and Simon P. Driver. “A Concise Reference to (Projected) Sérsic R1/n Quantities, Including Concentration, Profile Slopes, Petrosian Indices, and Kron Magnitudes”. In: *Publications of the Astronomical Society of Australia* 22.2 (2005), pp. 118–127. ISSN: 1448-6083. DOI: 10.1071/as05001. URL: <http://dx.doi.org/10.1071/AS05001>.
- [10] Johnny P. Greco et al. “Illuminating low surface brightness galaxies with the hyper supprime-cam survey”. In: *The Astrophysical Journal* 857.2 (Apr. 2018), p. 104. DOI: 10.3847/1538-4357/aab842.
- [11] Caroline Haigh et al. “Optimising and comparing source-extraction tools using objective segmentation quality criteria”. In: *Astronomy & Astrophysics* 645 (Jan. 2021), A107. ISSN: 1432-0746. DOI: 10.1051/0004-6361/201936561. URL: <http://dx.doi.org/10.1051/0004-6361/201936561>.
- [12] Guangshun Huang and Rinaldo Baldini Ferroli. “Probing the internal structure of baryons”. In: *National Science Review* 8.11 (Oct. 2021). DOI: 10.1093/nsr/nwab187.
- [13] Erin Kado-Fong et al. “The intrinsic shapes of low surface brightness galaxies (LSBGs): A discriminant of LSBG Galaxy Formation Mechanisms”. In: *The Astrophysical Journal* 920.2 (Oct. 2021), p. 72. DOI: 10.3847/1538-4357/ac15f0.
- [14] Diederik Kingma and Jimmy Ba. “Adam: A Method for Stochastic Optimization”. In: *International Conference on Learning Representations* (Dec. 2014). DOI: <https://doi.org/10.48550/arXiv.1412.6980>.
- [15] A. F. J. Moffat. “A Theoretical Investigation of Focal Stellar Images in the Photographic Emulsion and Application to Photographic Photometry”. In: *Astronomy & Astrophysics* 3 (Dec. 1969), p. 455. URL: <https://ui.adsabs.harvard.edu/abs/1969A&A....3..455M>.

- [16] Rene Racine. “The Telescope Point Spread Function”. In: *Publications of the Astronomical Society of the Pacific* 108 (Aug. 1996), p. 699. DOI: 10.1086/133788.
- [17] Javier Román, Aida Castilla, and Javier Pascual-Granado. “Discovery and analysis of low-surface-brightness galaxies in the environment of NGC 1052”. In: *Astronomy & Astrophysics* 656 (Dec. 2021). DOI: 10.1051/0004-6361/202142161.
- [18] Hao Su et al. “LSBGnet: An Improved Detection Model for Low-Surface Brightness Galaxies”. In: *Monthly Notices of the Royal Astronomical Society* 528.1 (Feb. 2024), pp. 873–882. DOI: 10.1093/mnras/stae001. URL: <https://doi.org/10.1093/mnras/stae001>.
- [19] Paul Teeninga et al. “Statistical attribute filtering to detect faint extended astronomical sources”. In: *Mathematical Morphology - Theory and Applications* 1 (Mar. 2016). DOI: 10.1515/mathm-2016-0006.
- [20] Aku Venhola et al. “The Fornax Deep Survey (FDS) with VST XII: Low surface brightness dwarf galaxies in the Fornax cluster”. In: *arXiv e-prints*, arXiv:2111.01855 (Nov. 2021), arXiv:2111.01855. DOI: 10.48550/arXiv.2111.01855. arXiv: 2111.01855 [astro-ph.GA].
- [21] Chien-Yao Wang et al. “ CSPNet: A New Backbone that can Enhance Learning Capability of CNN ”. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. Los Alamitos, CA, USA: IEEE Computer Society, June 2020, pp. 1571–1580. DOI: 10.1109/CVPRW50498.2020.00203. URL: <https://doi.ieeeecomputersociety.org/10.1109/CVPRW50498.2020.00203>.
- [22] Donald G. York et al. “The Sloan Digital Sky Survey: Technical Summary”. In: *Astrophysical Journal* 120.3 (Sept. 2000), pp. 1579–1587. DOI: 10.1086/301513. arXiv: astro-ph/0006396 [astro-ph].
- [23] G. H. Zhong et al. “The star formation histories of red and Blue Low Surface Brightness Disk Galaxies”. In: *Astronomy and Astrophysics* 520 (Sept. 2010). DOI: 10.1051/0004-6361/201014195.