



## A Comparison of Deep Learning Architectures for Optical Galaxy Morphology Classification

---

Ezra Fielding, Clement Nyirenda and Mattia Vaccari

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

November 6, 2021

# A Comparison of Deep Learning Architectures for Optical Galaxy Morphology Classification

Ezra Fielding

*Department of Computer Science  
University of the Western Cape  
Cape Town, South Africa  
0000-0002-7936-0222*

Clement N. Nyirenda

*Department of Computer Science  
University of the Western Cape  
Cape Town, South Africa  
0000-0002-4181-0478*

Mattia Vaccari

*Department of Physics & Astronomy  
University of the Western Cape  
Cape Town, South Africa  
0000-0002-6748-0577*

**Abstract**—The classification of galaxy morphology plays a crucial role in understanding galaxy formation and evolution. Traditionally, this process is done manually. The emergence of deep learning techniques has given room for the automation of this process. As such, this paper offers a comparison of deep learning architectures to determine which is best suited for optical galaxy morphology classification. Adapting the model training method proposed by Walmsley *et al* in 2021, the *Zoobot* Python library is used to train models to predict Galaxy Zoo DECaLS decision tree responses, made by volunteers, using EfficientNet B0, DenseNet121 and ResNet50 as core model architectures. The predicted results are then used to generate accuracy metrics per decision tree question to determine architecture performance. DenseNet121 was found to produce the best results, in terms of accuracy, with a reasonable training time. In future, further testing with more deep learning architectures could prove beneficial.

**Index Terms**—Deep Learning, Optical Galaxy Morphology, Classification, Astronomy

## I. INTRODUCTION

Morphology plays a crucial role in the investigation of galaxy formation and evolution [1]. As such, accurate and detailed morphological classifications are required to advance research in this area. This process of classification has historically been performed visually by scientists. However, modern astronomical surveys have been producing more detailed images than scientists can visually classify [1]. Projects, such as Galaxy Zoo, have attempted to solve this problem by asking so called ‘citizen scientists’, volunteer members of the public, to provide classifications through a web interface [2]. While these efforts have been largely successful and have greatly benefited the scientific community, it is only a matter of time until the output of modern astronomical surveys exceeds efforts which rely solely on human intelligence [3], [4]. In other words, projects which rely on humans to visually classify morphology, will not be able to keep up with the sheer volume of data produced by next generation surveys, such as the Square Kilometer Array (SKA), resulting in missed discoveries and insights [5], [6].

Different approaches have been used to solve the problem of classifying galaxies according to their morphology. In

the original Galaxy Zoo study, approximately one hundred thousand participants collectively provided more than  $4 \times 10^7$  individual classifications for nearly one million galaxies from the Sloan Digital Sky Survey (SDSS) [2]. Volunteers submitted classifications on the Galaxy Zoo website, where they were shown an image cutout of an area of sky centered on a galaxy. These raw classifications were then aggregated and used to form a final catalogue. These Galaxy Zoo results were consistent with those for a subset of SDSS classified by professional astronomers [2]. Subsequent projects such as Radio Galaxy Zoo and Galaxy Zoo DECaLS have also found success in using ‘citizen scientists’, to visually classify and inspect galaxies [1], [7]. The results of these Galaxy Zoo projects have been used as the basis for a large number of studies on galaxy morphology [3].

However, despite these successes, as the amount of survey data in need of classification continues to increase, approaches relying mainly on human classification and inspection will require too much time and effort to be viable [4]. For example, the recent Galaxy Zoo DECaLS project would have taken approximately 8 years to collect the standard 40 classifications for each of the 311488 suitable galaxies with the current number of volunteers available [1]. By that time, new surveys would have begun. This highlights the need for classification methods which do not rely heavily on human input.

A fair amount of work has been done in implementing deep learning algorithms for the purpose of classifying galaxy morphology [6]. A large portion of this work makes use of artificial neural networks to predict galaxy morphology based on images. These networks are trained using a labeled subset of existing surveys. Convolutional neural networks (CNNs) have been shown to be particularly useful and accurate in the classification of astronomical objects, provided the sufficient availability of labeled datasets [8]–[13]. Of particular interest is the work of Becker *et al*, which comparatively investigates the performance of different CNN architectures applied to radio galaxy classification. In their work, MCRGNET, RADIO GALAXY ZOO, and CONVXPRESS (a novel classifier) were the architectures which had the best balance of computational

requirements and recognition performance [14]. Walmsley *et al* used Galaxy Zoo DECaLS response data to train an EfficientNet B0 architecture to predict the responses to the GDZ-5 decision tree used by the Galaxy Zoo project to classify galaxies [1].

Most existing works have focused on comparing state-of-the-art deep learning architectures on radio-spectrum images. In terms of optical-spectrum images, the work in [1], which is the leading one, only uses the EfficientNet B0 and does not present a comparison with other prominent deep learning architectures. Therefore, this paper aims to assess the performance of three different, popular deep learning architectures in predicting volunteer vote fractions for optical-spectrum galaxy images from the recent Galaxy Zoo DECaLS project. EfficientNet B0 [15], DenseNet121 [16] and ResNet50 [17] will be compared in this work. Performance criteria will include training time and accuracy metrics such as precision, recall and F1-score.

The rest of this paper is organised as follows. Section II presents a brief overview of the deep learning architectures used in this work. Section III presents the materials and methods used in this work, while Section IV presents the results and their discussion. Section V draws the conclusion.

## II. DEEP LEARNING ARCHITECTURES USED IN THIS WORK

Three popular deep learning architectures were selected for this comparison. DenseNet121 and ResNet50 were selected alongside EfficientNet B0, which is the default architecture included with Walmsley’s *Zoobot* library [18]. These architectures were chosen based on their size, popularity, and ease of integration into the existing *Zoobot* library code base. These architectures are briefly described as follows.

*a) EfficientNet B0:* EfficientNet is a CNN-based architecture that uses a scaling method that uniformly scales all dimensions of depth, width, and resolution using a compound coefficient. The EfficientNet family of models are, therefore, optimised for both accuracy and efficiency, with the aim of performing better than traditional CNNs [15]. Apart from the classification of galaxies, EfficientNet has been applied to many other areas such as Covid-19 diagnosis [19], automated fruit recognition [20] and multi-label classification of fundus images [21].

*b) DenseNet121:* The DenseNet family of models differentiates itself from traditional CNN’s by connecting each layer to every other layer in a feed-forward fashion [16]. This works to reduce the number of parameters and alleviate the vanishing gradient problem. DenseNets achieve high performance while requiring less memory and computation when compared to other CNNs. DenseNet has also been applied to many areas such as remote sensing image classification [22], malware detection [23], and classification of Alzheimer’s disease [24].

*c) ResNet50:* The ResNet Family of models reformulate the model layers as learning residual functions with reference to the layer inputs. This results in models which are easier to optimise and which show a gain in accuracy to increased

model depth [17]. Some of the recent applications include crime prediction [25], [26] and skin cancer classification [27].

Table I presents the generic model that illustrates how the architectures were implemented as layers in the *Zoobot* library’s default model. The top fully connected layer for each architecture was removed, and replaced with by the three augmentation layers which feed directly in to the network.

TABLE I  
GENERIC MODEL ARCHITECTURE

Layer (type)	Output Shape
RandomRotation	(None, 300, 300, 1)
RandomFlip	(None, 300, 300, 1)
RandomCrop	(None, 224, 224, 1)
Core Architecture <sup>a</sup>	(None, 7, 7, )
GlobalAveragePooling2D	(None, )
Dense	(None, 34)

<sup>a</sup>This layer will be swapped out to test each architecture.

With the top fully connected layer removed, EfficientNet B0 resulted in approximately 4.0 million trainable parameters, while DenseNet121 and ResNet50 resulted in approximately 7.0 million and 23.6 million trainable parameters respectively. The final dense layer of the core model architecture was modified to give predictions smoothly, from 1 to 100 using softmax activation. This was done in order to provide plausible Dirichlet posteriors for galaxy morphology [1].

## III. MATERIALS AND METHODS

The method used for this architecture comparison, follows closely with that presented by Walmsley *et al*, 2021 for predicting Galaxy Zoo GZD-5 decision tree volunteer responses [1]. However, instead of training three models per architecture to form an ensemble network, only one model per architecture was trained. The training was also conducted at a decreased batch size of 64. A modified version of the *Zoobot* library was used for data pre-processing, model creation, model training, and making predictions [18].

### A. Description of the Dataset

The Galaxy Zoo DECaLS dataset includes 253286 424x424 pixel color images (3 channels in the g, r, and z bands) [1]. Volunteer response data is provided for each galaxy. This response data includes the total votes and vote fractions for each question in the GZD-5 decision tree. Any galaxy with less than 3 volunteer classifications was removed resulting in a 1.5% reduction in the dataset size. This left a total of 249581 images for training and testing. A standard random 80/20 train-test split was applied to the dataset. This resulted in a training dataset containing 199664 galaxies and a test dataset containing 49917 galaxies. The test dataset will serve as the validation dataset.

The GZD-5 decision tree used for this dataset includes ten questions which have a different number of options depending on the question. A chart explaining the GZD-5 decision tree can be found in Fig. 1.

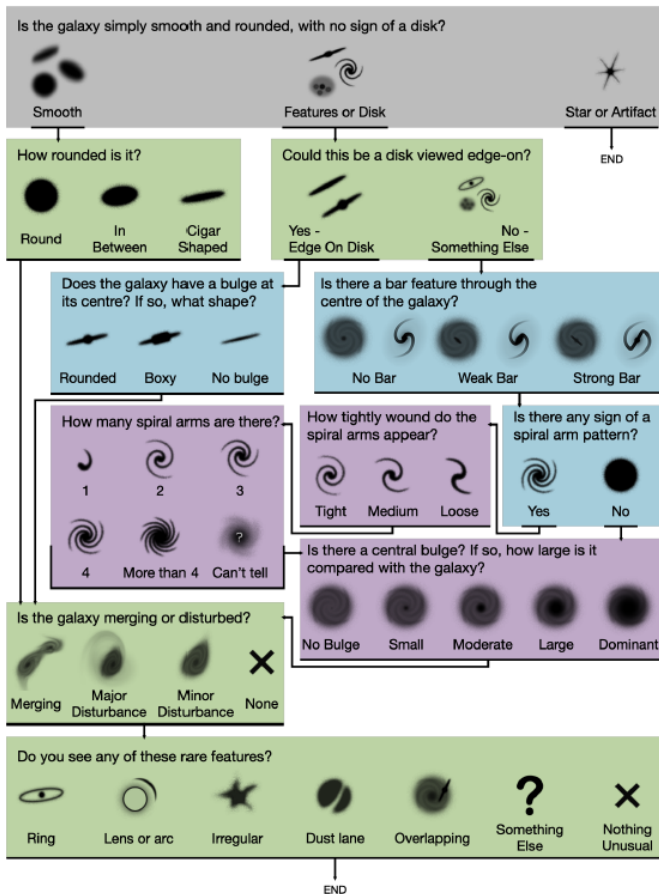


Fig. 1. The GZD-5 Decision Tree [1].

## B. Data Processing

The color information of each image was removed by taking an average over its 3 channels to avoid biasing the morphology predictions [28]. The images were then resized and saved as 300x300x1 matrices. Further augmentations were done when images were loaded into the model. Random horizontal and vertical flips were first applied, followed by a random rotation in the range  $(0, \pi)$ . The image was finally cropped to a size of 224x224 pixels about a random centroid. These augmentations were applied as the images were fed into the model.

## C. Training and Testing

Tensorflow and Python 3.8 were used for model training. The three models were identically trained using the Adam optimizer and loss was calculated by taking the negative log likelihood of observing volunteer answers to all questions from Dirichlet-Multinomial distributions [1]. Training was done with a batch size of 64 and ended after the model loss failed to improve after 10 consecutive epochs. A batch size of 64 was chosen due to the memory limitations of the training hardware. The test dataset was repeated once to reduce variance from random dropouts and augmentations. Training was completed on the South African Center for High Performance Computing’s (CHPC) Legau Cluster on two

NVIDIA V100 GPUs. Training prototyping and debugging, along with model prediction and evaluation was performed on the ilifu cloud computing facility.

To test and evaluate model performance, a new 80/20 train-test split was applied to the dataset. The new test dataset which resulted from this split was used for evaluation. A galaxy was only used to evaluate a question if 50% or more of the volunteers shown that galaxy was asked that specific question. Five forward passes per image were made with each model to account for random dropout configurations and augmentations. This resulted in five sets of predicted Dirichlet-Multinomial posteriors per image per model. The mean posterior of the total votes for each model were then calculated and recorded, to be used for calculating the predicted vote fraction for each question option.

The highest predicted vote fraction for a question was rounded to 1, while the remaining fractions were rounded to 0. This was done to produce a set of discrete classifications for each of the GZD-5 decision tree questions. The discrete classifications for each question were then used to generate accuracy metrics, namely precision, recall and F1-score.

a) *Precision*: This is defined as the ratio of images that was correctly defined as belonging to the positive class, to the total number of images that were classified as belonging to the positive class.

b) *Recall*: This refers to the ratio of the number of images that were correctly classified as belonging to the positive class, to the total number of images in the positive class.

c) *F1-Score*: This is the weighted average of recall and precision.

The weighted average metrics are considered for each question. The weighted average metrics are calculated by calculating the metrics for each question option, and finding their average weighted by support. This allows the accuracy metrics to account for label imbalance. Table II shows the support for the metrics of each decision tree question.

TABLE II  
QUESTION SUPPORT

Question	Support
smooth-or-featured	49917
disk-edge-on	15445
has-spiral-arms	11380
bar	11380
bulge-size	11380
how-rounded	32526
edge-on-bulge	2475
spiral-winding	7499
spiral-arm-count	7499
merging	49247

## IV. RESULTS AND DISCUSSION

The following results were produced after running the experiment. Table III shows the total number of epochs run during training and the total training time.

TABLE III  
STATISTICS FROM TRAINING

	<i>EfficientNet B0</i>	<i>DenseNet121</i>	<i>ResNet50</i>
Total Epochs	52	37	37
Total Time (hours)	19.795	11.723	9.642

DenseNet121 and ResNet50 both ran for 37 epochs, while EfficientNet B0 ran for 52 epochs. Despite being the largest model by number of parameters, ResNet50 trained for a total of 9.642 hours, while DenseNet121 completed training after 11.723 hours. EfficientNet B0 spent the longest time training, at 19.795 hours.

The weighted average metrics for each question for each of the three models were placed into tables for comparison. Table IV presents a comparison of the precision metric, Table V presents a comparison of the recall metric, and Table VI presents a comparison of the F1-score metric. The weighted average for all questions is displayed for each model at the bottom of each comparison table.

TABLE IV  
PRECISION COMPARISON

Question Name	Core Model Architecture		
	<i>EfficientNet B0</i>	<i>DenseNet121</i>	<i>ResNet50</i>
smooth-or-featured	0.877	<b>0.880</b>	<u>0.869</u>
disk-edge-on	<u>0.954</u>	<b>0.957</b>	0.955
has-spiral-arms	0.891	<b>0.893</b>	<u>0.868</u>
bar	0.697	<b>0.698</b>	<u>0.673</u>
bulge-size	0.684	<b>0.690</b>	<u>0.675</u>
how-rounded	0.870	<b>0.876</b>	<u>0.867</u>
edge-on-bulge	0.794	<b>0.826</b>	<u>0.785</u>
spiral-winding	0.685	<b>0.695</b>	<u>0.677</u>
spiral-arm-count	0.664	<b>0.695</b>	<u>0.652</u>
merging	0.849	<b>0.851</b>	<u>0.845</u>
<b>Weighted Average</b>	0.838	<b>0.843</b>	<u>0.831</u>

Best model result highlighted in bold. Worst model result underlined.

TABLE V  
RECALL COMPARISON

Question Name	Core Model Architecture		
	<i>EfficientNet B0</i>	<i>DenseNet121</i>	<i>ResNet50</i>
smooth-or-featured	0.876	<b>0.881</b>	<u>0.863</u>
disk-edge-on	<u>0.955</u>	<b>0.957</b>	0.955
has-spiral-arms	0.887	<b>0.888</b>	<u>0.875</u>
bar	<b>0.715</b>	0.710	<u>0.696</u>
bulge-size	0.694	<b>0.696</b>	<u>0.691</u>
how-rounded	0.867	<b>0.875</b>	<u>0.864</u>
edge-on-bulge	0.820	<b>0.825</b>	<u>0.816</u>
spiral-winding	0.698	<b>0.707</b>	<u>0.690</u>
spiral-arm-count	<b>0.676</b>	0.674	<u>0.664</u>
merging	0.881	<b>0.882</b>	<u>0.880</u>
<b>Weighted Average</b>	0.848	<b>0.851</b>	<u>0.841</u>

Best model result highlighted in bold. Worst model result underlined.

From the model metric comparison tables, as far as accuracy is concerned, it is clear that DenseNet121 produces the best predictions for the given dataset. ResNet50 generally performs the worst out of the three, and EfficientNet B0 lies in the middle between the two in terms of accuracy.

TABLE VI  
F1-SCORE COMPARISON

Question Name	Core Model Architecture		
	<i>EfficientNet B0</i>	<i>DenseNet121</i>	<i>ResNet50</i>
smooth-or-featured	0.876	<b>0.881</b>	<u>0.865</u>
disk-edge-on	<u>0.955</u>	<b>0.956</b>	0.955
has-spiral-arms	0.889	<b>0.890</b>	<u>0.869</u>
bar	<b>0.693</b>	0.677	<u>0.662</u>
bulge-size	<u>0.669</u>	<b>0.674</b>	0.672
how-rounded	<u>0.867</u>	<b>0.875</b>	<u>0.865</u>
edge-on-bulge	0.787	<b>0.794</b>	<u>0.784</u>
spiral-winding	0.687	<b>0.696</b>	<u>0.678</u>
spiral-arm-count	0.654	<b>0.668</b>	<u>0.642</u>
merging	0.851	<b>0.855</b>	<u>0.848</u>
<b>Weighted Average</b>	0.836	<b>0.840</b>	<u>0.829</u>

Best model result highlighted in bold. Worst model result underlined.

Taking into account accuracy and training time, DenseNet121 appears to be the most optimal architecture for this dataset. While ResNet50 had the fastest training time, it generally performed the worst for each question in each metric. EfficientNet B0, on the other hand, performed second best in terms of accuracy. However, it took significantly longer to train compared to the other two architectures. DenseNet121 provides the best accuracy performance at a reasonable model training time.

## V. CONCLUSION

In conclusion, three popular deep learning architectures (EfficientNet B0, DenseNet121, and ResNet50) were evaluated and compared in terms of accuracy and training performance with regards to classifying optical galaxy morphology. Performance criteria included training time and accuracy metrics such as precision, recall and F1-score. The method for training outlined by Walmsley *et al* in [1] was followed with some minor changes. In terms of accuracy, DenseNet121 was the best performing architecture with the weighted averages of 0.843, 0.851 and 0.840 for precision, recall and F1-score respectively. On the other hand, ResNet50 exhibited the lowest performance with the weighted average scores of 0.831, 0.841, and 0.829 for precision, recall and F1-score respectively. The training time for DenseNet121 was 11.723 hours, which was approximately 2 hours more than ResNet50. On the other hand, EfficientNet B0 took 19.795 hours. Taking accuracy, training time and total parameters into account, DenseNet121 is the best model to use for this use case.

In future, further testing with a larger selection of deep learning architectures could prove useful and beneficial for the morphology classification of optical galaxies. Applying these architectures to other Galaxy Zoo projects could also prove insightful.

## ACKNOWLEDGMENT

The authors would like to thank the Galaxy Zoo volunteers, whose efforts made this dataset possible. Their efforts are individually acknowledged at <https://authors.galaxyzoo.org/>. Special thanks to Mike Walmsley for granting access to the Galaxy Zoo DECaLS dataset and making the *Zoobot* library

publicly available. The authors acknowledge the Centre for High Performance Computing (CHPC), South Africa, and the ilifu cloud computing facility – [www.ilifu.ac.za](http://www.ilifu.ac.za) - for providing computational resources to this research project.

## REFERENCES

- [1] M. Walmsley, C. Lintott, T. Geron, S. Kruk, C. Krawczyk, K. W. Willett, S. Bamford, W. Keel, L. S. Kelvin, L. Fortson *et al.*, “Galaxy zoo decals: Detailed visual morphology measurements from volunteers and deep learning for 314,000 galaxies,” *arXiv preprint arXiv:2102.08414*, 2021.
- [2] C. J. Lintott, K. Schawinski, A. Slosar, K. Land, S. Bamford, D. Thomas, M. J. Raddick, R. C. Nichol, A. Szalay, D. Andreescu *et al.*, “Galaxy zoo: morphologies derived from visual inspection of galaxies from the sloan digital sky survey,” *Monthly Notices of the Royal Astronomical Society*, vol. 389, no. 3, pp. 1179–1189, 2008.
- [3] K. L. Masters and G. Z. Team, “Twelve years of galaxy zoo,” *Proceedings of the International Astronomical Union*, vol. 14, no. S353, pp. 205–212, 2019.
- [4] T. J. Galvin, M. T. Huynh, R. P. Norris, X. Wang, E. Hopkins, K. Polsterer, N. O. Ralph, A. N. O’Brien, and G. H. Heald, “Cataloguing the radio-sky with unsupervised machine learning: a new approach for the ska era,” *Monthly Notices of the Royal Astronomical Society*, vol. 497, no. 3, pp. 2730–2758, 2020.
- [5] P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. Lazio, “The square kilometre array,” *Proceedings of the IEEE*, vol. 97, no. 8, pp. 1482–1496, 2009.
- [6] C. J. Fluke and C. Jacobs, “Surveying the reach and maturity of machine learning and artificial intelligence in astronomy,” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 10, no. 2, p. e1349, 2020.
- [7] J. K. Banfield, O. Wong, K. W. Willett, R. P. Norris, L. Rudnick, S. S. Shabala, B. D. Simmons, C. Snyder, A. Garon, N. Seymour *et al.*, “Radio galaxy zoo: host galaxies and radio morphologies derived from visual inspection,” *Monthly Notices of the Royal Astronomical Society*, vol. 453, no. 3, pp. 2326–2340, 2015.
- [8] S. Dieleman, K. W. Willett, and J. Dambre, “Rotation-invariant convolutional neural networks for galaxy morphology prediction,” *Monthly notices of the royal astronomical society*, vol. 450, no. 2, pp. 1441–1459, 2015.
- [9] W. Alhassan, A. Taylor, and M. Vaccari, “The first classifier: compact and extended radio galaxy classification using deep convolutional neural networks,” *Monthly Notices of the Royal Astronomical Society*, vol. 480, no. 2, pp. 2085–2093, 2018.
- [10] H. Domínguez Sánchez, M. Huertas-Company, M. Bernardi, D. Tuccillo, and J. Fischer, “Improving galaxy morphologies for sdss with deep learning,” *Monthly Notices of the Royal Astronomical Society*, vol. 476, no. 3, pp. 3661–3676, 2018.
- [11] C. J. Burke, P. D. Aleo, Y.-C. Chen, X. Liu, J. R. Peterson, G. H. Sembroski, and J. Y.-Y. Lin, “Deblending and classifying astronomical sources with mask r-cnn deep learning,” *Monthly Notices of the Royal Astronomical Society*, vol. 490, no. 3, pp. 3952–3965, 2019.
- [12] V. Lukic, M. Brüggen, J. K. Banfield, O. Wong, L. Rudnick, R. P. Norris, and B. Simmons, “Radio galaxy zoo: compact and extended radio source classification with deep learning,” *Monthly Notices of the Royal Astronomical Society*, vol. 476, no. 1, pp. 246–260, 2018.
- [13] Z. Ma, H. Xu, J. Zhu, D. Hu, W. Li, C. Shan, Z. Zhu, L. Gu, J. Li, C. Liu *et al.*, “A machine learning based morphological classification of 14,245 radio agns selected from the best-heckman sample,” *The Astrophysical Journal Supplement Series*, vol. 240, no. 2, p. 34, 2019.
- [14] B. Becker, M. Vaccari, M. Prescott, and T. Grobler, “Cnn architecture comparison for radio galaxy classification,” *Monthly Notices of the Royal Astronomical Society*, vol. 503, no. 2, pp. 1828–1846, 2021.
- [15] M. Tan and Q. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 6105–6114.
- [16] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.
- [17] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [18] M. Walmsley, “Zoobot,” <https://github.com/mwalmsley/zoobot>, 2021.
- [19] G. Marques, D. Agarwal, and I. de la Torre Díez, “Automated medical diagnosis of covid-19 through efficientnet convolutional neural network,” *Applied soft computing*, vol. 96, p. 106691, 2020.
- [20] L. T. Duong, P. T. Nguyen, C. Di Sipio, and D. Di Ruscio, “Automated fruit recognition using efficientnet and mixnet,” *Computers and Electronics in Agriculture*, vol. 171, p. 105326, 2020.
- [21] J. Wang, L. Yang, Z. Huo, W. He, and J. Luo, “Multi-label classification of fundus images with efficientnet,” *IEEE Access*, vol. 8, pp. 212499–212508, 2020.
- [22] W. Tong, W. Chen, W. Han, X. Li, and L. Wang, “Channel-attention-based densenet network for remote sensing image scene classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 13, pp. 4121–4132, 2020.
- [23] J. Hemalatha, S. A. Roseline, S. Geetha, S. Kadry, and R. Damaševičius, “An efficient densenet-based deep learning model for malware detection,” *Entropy*, vol. 23, no. 3, p. 344, 2021.
- [24] J. Ruiz, M. Mahmud, M. Modasshir, M. S. Kaiser, f. t. Alzheimer’s Disease Neuroimaging Initiative *et al.*, “3d densenet ensemble in 4-way classification of alzheimer’s disease,” in *International Conference on Brain Informatics*. Springer, 2020, pp. 85–96.
- [25] T. Materoke, C. Nyirenda, and M. Ghaziasgar, “A comparative evaluation of spatio temporal deep learning techniques for crime prediction,” in *Proceedings of the IEEE Africon Conference*. IEEE, 2021.
- [26] B. Wang, P. Yin, A. L. Bertozzi, P. J. Brantingham, S. J. Osher, and J. Xin, “Deep learning for real-time crime forecasting and its ternarization,” *Chinese Annals of Mathematics, Series B*, vol. 40, no. 6, pp. 949–966, 2019.
- [27] N. Gouda and J. Amudha, “Skin cancer classification using resnet,” in *2020 IEEE 5th International Conference on Computing Communication and Automation (ICCCA)*. IEEE, 2020, pp. 536–541.
- [28] M. Walmsley, L. Smith, C. Lintott, Y. Gal, S. Bamford, H. Dickinson, L. Fortson, S. Kruk, K. Masters, C. Scarlata *et al.*, “Galaxy zoo: probabilistic morphology through bayesian cnns and active learning,” *Monthly Notices of the Royal Astronomical Society*, vol. 491, no. 2, pp. 1554–1574, 2020.