
Deep Learning Segmentation of Spiral Arms and Bars

Mike Walmsley*

Dunlap Institute for Astronomy and Astrophysics
University of Toronto
Toronto, ON M5S 3H4, Canada
m.walmsley@utoronto.ca

Ashley Spindler

Centre for Astrophysics Research, Department of Physics, Astronomy and Mathematics
University of Hertfordshire
Hatfield, Hertfordshire, AL10 9AB, UK
a.spindler@herts.ac.uk

Abstract

We present the first deep learning model for segmenting galactic spiral arms and bars. In a blinded assessment by expert astronomers, our predicted spiral arm masks are preferred over both current automated methods (99% of evaluations) and our original volunteer labels (79% of evaluations). Experts rated our spiral arm masks as ‘mostly good’ to ‘perfect’ in 89% of evaluations. Bar lengths trivially derived from our predicted bar masks are in excellent agreement with a dedicated crowdsourcing project. The pixelwise precision of our masks, previously impossible at scale, will underpin new research into how spiral arms and bars evolve.

1 Introduction

Astronomers do not know how spiral arms are made. There are two long-standing theories: density waves [1] and swing-amplification [2]. Despite decades of work, it is unclear which theory dominates [3, 4]. Spiral arms also show different characteristics in different galaxies; they vary in arrangement (e.g. grand design vs. flocculent), arm count, and arm tightness (pitch angle). Precise explanations for these characteristics remain elusive.

A closely-related galaxy structure is galactic bars; linear features in the center of galaxies, often connecting two spiral arms [5]. Barred spirals make up about 30% of disk-type galaxies in the nearby universe [6]. Bars alter the structure of a host galaxy by funnelling gas and dust from the disk into the galaxy bulge [7, 8, 9]. Bars grow over time; whether long and short bars are distinct subclasses or part of a continuum is unclear [10].

Progress in understanding both bars and spiral arms has been limited by the difficulty of measuring their shapes [4, 11, 12]. Shape measurements to-date have been made by either expert inspection [13], rule-based image processing [14, 1], or crowdsourcing [15, 12]. We develop a deep learning model to jointly identify both spiral arms and bars.

Deep learning for galaxy morphology has focused almost exclusively on classifying entire images ([16, 17, 18, 19, 20, 21, 22, 23, 24, 25], etc.) Works on segmentation are rare. Hausen and Robertson [26] trained a U-net-like architecture to predict ‘pixel level morphological classifications’ i.e. to

*Second affiliation: Jodrell Bank Centre for Astrophysics, Department of Physics & Astronomy, University of Manchester, Manchester M13 9PL, UK




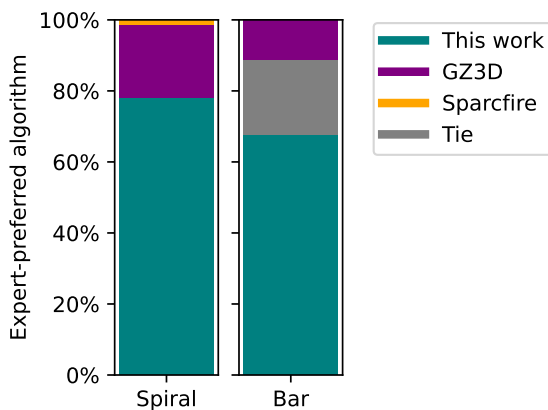
Figure 1: From left: input galaxy image, segmentation masks by volunteers (GZ3D, [11]), masks by *sparcfire* [14](spiral only), and masks by our model. Rows show randomly-selected test set galaxies, filtered to have GZ DESI [25] vote fractions above 0.5.

identify the pixels belonging to a single galaxy and then ascribe *all* those pixels to a single galaxy class (e.g. spheroid, disk, etc.) Ostdiek et al. [27] used a U-Net-like architecture to classify pixels as belonging to either lensing galaxies or lensed subhalos. Richards et al. [28] heavily extended Mask R-CNN to classify both tidal structures around galaxies and background contamination. Bhambra et al. [29] (hereafter PB+2022) trained a standard CNN (rather than U-Net or Mask R-CNN) to classify galaxy images as barred or not barred, and then used saliency mapping to identify the pixels most relevant to that classification. All of these works adopt a pixelwise classification framework.

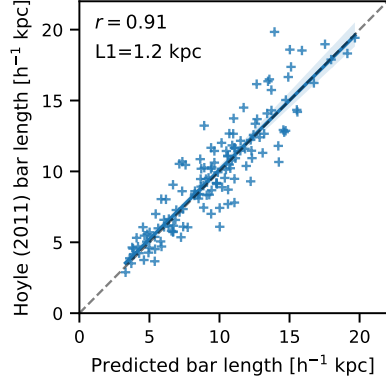
2 Methods

2.1 Data and Loss Function Design

Our segmentation labels are sourced from Galaxy Zoo: 3D, published in Masters et al. [11]. GZ3D collected crowdsourced pixel masks of the spiral arms and bars of the 29,831 galaxies listed as potential targets for the MaNGA survey [30]. While the project was successful and released a public catalogue of pixel masks, Masters et al. [11] highlighted the time-consuming and labour-intensive nature of the task — particularly in regards to labelling the spiral arms.



(a) Fraction of astronomer evaluations selecting each ‘algorithm’ as best-performing (N=367, 133). Astronomers strongly prefer our model, even over the original labels we use for training.



(b) Bar length measurements from our model vs. human measurements (Hoyle et al. [15]). Length predictions are derived by trivially processing our bar masks.

GZ3D volunteers used a polygon drawing tool to enclose areas corresponding to bars or spiral arms. Each galaxy was marked by 15 volunteers. We use the individual volunteer polygon vertices published by Masters et al. [11] to calculate how many volunteers enclosed each pixel.²

Unlike previous astronomy works, we frame our segmentation task as pixelwise regression. The fraction of volunteers enclosing each pixel reflects the confidence of the crowd. This confidence is a critical training signal when attempting to learn the often-ambiguous task of identifying spiral arms. We choose the loss function:

$$\mathcal{L}(X, V, w) \propto \sum_k^C \left| \sum_{i,j} \sum_l^V \delta(X_{ij} \in V_{l,k}) - \phi_w(X)_{ijk} \right| \quad (1)$$

where, for spatial pixel index (i, j) in image X , we check if each pixel X_{ij} is enclosed by the vertices marked by volunteer V_l for class k (spiral or bar) and compare that to the model output for the same pixel in channel k , $\phi_w(X)_{ijk}$. In short, we calculate the mean absolute error when predicting, for each pixel, the fraction of volunteers who included that pixel in their mask, and sum over each output class/channel.

We implement this loss by encoding the pixelwise volunteer vote fractions as greyscale JPEG images (where e.g. a value of 255 corresponds to all 15 volunteers marking a pixel) and then loading and augmenting these masks as if they were conventional images.³ Our final convolutional layer uses ReLU activation [31] to ensure non-negative predictions. The same network jointly learns to predict masks for both spiral arms and bars using a two-channel output.

Our images are sourced from the DESI Legacy Surveys [32]. DESI-LS images are comparable to the original GZ3D images for identifying spiral arms and bars in nearby galaxies, but using them for training allows for future work to make direct measurements on the full DESI-LS footprint. We use fixed train/validation/test sets with 70/10/20% division. We only include galaxies with GZ3D masks for spiral arms. We do not penalise the network for predicting bars (or, in principle, spiral arms) on galaxies that volunteers chose not to mark as such, as we judge the lack of marks to be an unreliable indicator for the lack of a feature.

²Unlike Masters et al. [11], we do not discount self-intersecting areas. Many volunteer marks are only self-intersecting where closing line segments slightly intersect, and so discarding these areas removes a significant amount of otherwise-useful annotations.

³We use only flip and interpolated rotation augmentations, thus preserving the pixel (vote fraction) values.

2.2 Model Architecture

Our model uses a U-Net architecture which consists of a downsampling encoder and an upsampling decoder. For both submodels, a single step is comprised of two residual blocks, followed by a strided convolution for the downsampling or a transpose convolution for the upsampling. Dropout is used between the residual blocks and convolutions in the encoder to control overfitting. Skip connections between each level of the encoder pass the outputs of the downsampling step to the inputs of the corresponding upsampling step. Following Smith et al. [33] we use Mish activation functions [34] in the residual blocks. We perform a hyperparameter sweep (on the validation set) to select the batch size, dropout rates, and down/upsampling dimensions. Our code is publicly available at <https://github.com/mwalmsley/zoobot-3d>.

2.3 Results

Figure 1 presents a random selection of predicted spiral and bar masks. Our predicted masks typically capture spiral and bar features well. Further, they are typically smoother and more precise than the original volunteer labels. The polygon drawing tool used by GZ3D volunteers creates a ‘blocky’ mask which fails to capture the soft edges and ambiguous locations of spiral arms. Small number statistics (each galaxy is marked by 1-15 volunteers) then lead to unphysical step functions where masks from individual volunteers overlap. Our model learns from these observed masks to predict the expected mask at each pixel. Unlike the labels, the expected mask is smoothly varying.

Comparing to previous work is difficult because, to our knowledge, there are no existing automated methods for segmenting spiral arms. The closest comparable code is `sparcfire` [14], which segments as an intermediate step towards mathematically describing the arm segments. Our model is notably more accurate than `sparcfire` on the random examples shown. We also compared with the recent general purpose ‘foundation’ model `SegmentAnything` [35]; we find that `SegmentAnything` typically identifies galaxy vs. background pixels but cannot segment structures within galaxies.

To quantitatively assess performance, we recruited 20 astronomers for a blinded comparison between methods. Participants were shown a grid of images (similar to 1) and asked to rate their preferred and second-preferred mask. For each preferred algorithm, they then rated each mask by quality on a four-point scale: ‘perfect’, ‘mostly perfect (some minor errors)’, ‘mostly poor (some major errors)’, and ‘totally failed’. Participants made a total of 500 evaluations on 100 galaxies. Participants were not involved in this research and did not know which algorithm generated each mask.

We find that our model is strongly preferred by astronomers over both `sparcfire` and our original volunteer labels. Fig. 2a shows the fraction of evaluations where astronomers rated each ‘algorithm’ as the best performing; for spiral arms, our model was preferred by an individual expert in 78% of spiral evaluations (vs. 20% for the multi-volunteer GZ3D labels and 1% for `sparcfire`) and 68% of bar evaluations (vs. 11% for the GZ3D labels, with the remainder tied). Our model ranked higher than `sparcfire` in 99% of spiral evaluations and higher than GZ3D in 79% of evaluations⁴. Astronomers rated our spiral arm masks as ‘mostly good’ to ‘perfect’ in 89% of evaluations, vs. 77% for the GZ3D labels.

To quantitatively demonstrate that our masks are useful for astronomy, we process our bar masks to calculate bar lengths. We calculate the length of the confidently-barred⁵ pixel mask and convert this length to physical distances. We find (Fig. 2b excellent agreement with the measurements gathered by Hoyle et al. [15] in a crowdsourcing project dedicated to finding bars.

3 Conclusion and Discussion

We have presented the first deep learning model for segmenting spiral arms and bars. Our model outperforms both the closest existing automated approach and the original crowdsourced labels on which it was trained, as judged by independent astronomers in a blinded evaluation.

⁴Formally, the one-tailed Binomial odds of an ensemble of experts preferring our model in k out of N galaxies (excluding ties) under the null hypothesis of no true preference between our model and the original labels is $p \ll 0.01$ for spirals ($N = 91$, $k = 74$) and $p = 0.04$ for bars ($N = 26$, $k = 17$)

⁵Defined as having predicted bar vote fractions greater than the median-confidence pixel

Accurate segmentation of spiral arms and bars is especially relevant for Integral Field Spectroscopy (IFS) surveys, which measure spectra at every pixel. This was the original motivation for GZ3D. There are now multiple large IFS surveys (e.g. MaNGA, SAMI) and IFS is available on observational platforms including NIRSpec and MIRI on JWST and ESO-MUSE on the VLT. Accurate automated segmentation masks could enhance research with any of these platforms.

Accurate segmentation is also relevant for the very largest surveys. Our model is trained on DESI-LS images, enabling spiral and bar shape measurements for galaxies with upcoming DESI spectra. Euclid [36] will shortly begin resolving the detailed morphology of higher-redshift ($z \approx 1$) galaxies, allowing us to measure bars and spirals as they form and grow to their present-day appearance.

We chose to use a U-Net architecture because it is appropriate for smaller labelled datasets (here, 5054 training galaxies) and to focus this work on introducing a first solution to this new task. More recent architectures may improve performance.

In closing, we would like to note a negative result. We found that, by adding a Dirichlet regression head [37] to the latent space of our U-Net, we could also predict Galaxy Zoo volunteer vote fractions competitively with a dedicated classification model. However, adding this joint prediction task did not improve the quality of our segmentation masks. This complicates the ‘astronomy foundation model’ view presented by Walmsley et al. [38] (see also [39, 40, 41]). General purpose models may not always outperform dedicated ones.

Acknowledgments and Disclosure of Funding

We are grateful to the Zooniverse volunteers and GZ:3D authors for creating the labels we rely on here.

We would like to thank the astronomers who contributed blinded assessments, including: Dominic Adams, Elisabeth M Baeten, Hugh Dickinson, James Tropp Garland, Karen L Masters, James Pearson, Jonathon C S Pierce, Shravya Shenoy, Brooke Simmons, Josh Speagle, Grant Stevens, Laura Trouille, and Klaas Wiersema.

This work was supported by a grant from Meta. MW is a Dunlap Fellow. The Dunlap Institute is funded through an endowment established by the David Dunlap family and the University of Toronto. This work has made use of the University of Hertfordshire high-performance computing facility (<https://uhhpc.herts.ac.uk/>).

References

- [1] C. C. LIN and FRANK H. SHU. ON THE SPIRAL STRUCTURE OF DISK GALAXIES. *Selected Papers of C C Lin with Commentary*, pages 561–570, 8 1987. doi: 10.1142/9789814415651{_}0033.
- [2] J. A. Sellwood, L. S. Sparke, J. A. Sellwood, and L. S. Sparke. Pattern speeds in barred spiral galaxies. *MNRAS*, 231(1):25P–31, 3 1988. ISSN 0035-8711. doi: 10.1093/MNRAS/231.1.25P. URL <https://ui.adsabs.harvard.edu/abs/1988MNRAS.231P..25S/abstract>.
- [3] Ross E. Hart, Steven P. Bamford, William C. Keel, Sandor J. Kruk, Karen L. Masters, Brooke D. Simmons, and Rebecca J. Smethurst. Galaxy Zoo: constraining the origin of spiral arms. *Monthly Notices of the Royal Astronomical Society*, 478(1):932–949, 7 2018. ISSN 0035-8711. doi: 10.1093/MNRAS/STY1201. URL <https://researchportal.port.ac.uk/en/publications/galaxy-zoo-constraining-the-origin-of-spiral-arms>.
- [4] Si-Yue Yu and Luis C. Ho. The Statistical Properties of Spiral Arms in Nearby Disk Galaxies. *The Astrophysical Journal*, 900(2):150, 9 2020. ISSN 0004-637X. doi: 10.3847/1538-4357/ABAC5B. URL <https://iopscience.iop.org/article/10.3847/1538-4357/abac5bhttps://iopscience.iop.org/article/10.3847/1538-4357/abac5b/meta>.
- [5] Ronald J. Buta. Galaxy Morphology. 2 2011. URL <http://arxiv.org/abs/1102.0550>.
- [6] Karen L. Masters, Robert C. Nichol, Ben Hoyle, Chris Lintott, Steven P. Bamford, Edward M. Edmondson, Lucy Fortson, William C. Keel, Kevin Schawinski, Arfon M. Smith, and Daniel

- Thomas. Galaxy Zoo: Bars in disc galaxies. *Monthly Notices of the Royal Astronomical Society*, 411(3):2026–2034, 3 2011. ISSN 13652966. doi: 10.1111/j.1365-2966.2010.17834.x. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2010.17834.x>.
- [7] F. Combes, F. Debbasch, D. Friedli, D. Pfenniger, F. Combes, F. Debbasch, D. Friedli, and D. Pfenniger. Box and peanut shapes generated by stellar bars. *A&A*, 233:82, 1990. ISSN 0004-6361. URL <https://ui.adsabs.harvard.edu/abs/1990A&A...233...82C/abstract>.
- [8] E. Athanassoula and A. Misiriotis. Morphology, photometry and kinematics of N-body bars - I. Three models with different halo central concentrations. *Monthly Notices of the Royal Astronomical Society*, 330(1):35–52, 2 2002. ISSN 00358711. doi: 10.1046/j.1365-8711.2002.05028.x. URL <https://academic.oup.com/mnras/article-lookup/doi/10.1046/j.1365-8711.2002.05028.x>.
- [9] S. Khoperskov, M. Haywood, P. Di Matteo, M. D. Lehnert, and F. Combes. Bar quenching in gas-rich galaxies. *Astronomy and Astrophysics*, 609:A60, 1 2018. ISSN 14320746. doi: 10.1051/0004-6361/201731211. URL <https://ui.adsabs.harvard.edu/abs/2018A&A...609A..60K/abstract>.
- [10] Tobias Geron, Rebecca J Smethurst, Chris Lintott, Sandor Kruk, Karen L Masters, Brooke Simmons, Kameswara Bharadwaj Mantha, Mike Walmsley, L Garma-Oehmichen, Niv Drory, and Richard R Lane. Galaxy Zoo: kinematics of strongly and weakly barred galaxies. *Monthly Notices of the Royal Astronomical Society*, 521(2):1775–1793, 3 2023. ISSN 0035-8711. doi: 10.1093/MNRAS/STAD501. URL <https://academic.oup.com/mnras/article/521/2/1775/7043473>.
- [11] Karen L. Masters, Coleman Krawczyk, Shoaib Shamsi, Alexander Todd, Daniel Finnegan, Matthew Bershady, Kevin Bundy, Brian Cherinka, Amelia Fraser-Mckelvie, Dhanesh Krishnarao, Sandor Kruk, Richard R. Lane, David Law, Chris Lintott, Michael Merrifield, Brooke Simmons, Anne Marie Weijmans, and Renbin Yan. Galaxy Zoo: 3D – crowdsourced bar, spiral, and foreground star masks for MaNGA target galaxies. *Monthly Notices of the Royal Astronomical Society*, 507(3):3923–3935, 9 2021. ISSN 0035-8711. doi: 10.1093/MNRAS/STAB2282. URL <https://dx.doi.org/10.1093/mnras/stab2282>.
- [12] Ross E. Hart, Steven P. Bamford, Kyle W. Willett, Karen L. Masters, Carolin Cardamone, Chris J. Lintott, Robert J. Mackay, Robert C. Nichol, Christopher K. Rosslowe, Brooke D. Simmons, and Rebecca J. Smethurst. Galaxy Zoo: Comparing the demographics of spiral arm number and a new method for correcting redshift bias. *Monthly Notices of the Royal Astronomical Society*, 461(4):3663–3682, 2016. ISSN 13652966. doi: 10.1093/mnras/stw1588.
- [13] Robert Kennicutt and Paul Hodge. A COMPARISON OF MEASURED SPIRAL ARM PROPERTIES WITH MODEL PREDICTIONS. *The Astrophysical Journal*, 101:101–107, 1982.
- [14] Darren R. Davis and Wayne B. Hayes. SpArcFiRe: Scalable automated detection of spiral galaxy arm segments. *Astrophysical Journal*, 790(2):87, 7 2014. ISSN 15384357. doi: 10.1088/0004-637X/790/2/87. URL <http://stacks.iop.org/0004-637X/790/i=2/a=87?key=crossref.73c5cbd1359d0e439cd559ab23a013d0>.
- [15] Ben Hoyle, Karen L. Masters, Robert C. Nichol, Edward M. Edmondson, Arfon M. Smith, Chris Lintott, Ryan Scranton, Steven Bamford, Kevin Schawinski, and Daniel Thomas. Galaxy Zoo: bar lengths in local disc galaxies. *Monthly Notices of the Royal Astronomical Society*, 415(4):3627–3640, 8 2011. ISSN 1365-2966. doi: 10.1111/J.1365-2966.2011.18979.X. URL <https://onlinelibrary.wiley.com/doi/full/10.1111/j.1365-2966.2011.18979.x>
<https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2966.2011.18979.x>
<https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2966.2011.18979.x>.
- [16] S. Dieleman, K. W. Willett, and J. Dambre. Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Monthly Notices of the Royal Astronomical Society*, 450(2):1441–1459, 2015. ISSN 0035-8711. doi: 10.1093/mnras/stv632. URL <http://arxiv.org/abs/1503.07077>.

- [17] M. Huertas-Company, R. Gravet, G. Cabrera-Vives, P. G. Pérez-González, J. S. Kartaltepe, G. Barro, M. Bernardi, S. Mei, F. Shankar, P. Dimauro, E. F. Bell, D. Kocevski, D. C. Koo, S. M. Faber, and D. H. McIntosh. A catalog of visual-like morphologies in the 5 candels fields using deep learning. *Astrophysical Journal, Supplement Series*, 221(1):8, 2015. ISSN 00670049. doi: 10.1088/0067-0049/221/1/8. URL <http://dx.doi.org/10.1088/0067-0049/221/1/8>.
- [18] Asad Khan, E. A. Huerta, Sibow Wang, Robert Gruendl, Elise Jennings, and Huihuo Zheng. Deep learning at scale for the construction of galaxy catalogs in the Dark Energy Survey. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, 795:248–258, 2019. ISSN 03702693. doi: 10.1016/j.physletb.2019.06.009. URL <https://www.sciencedirect.com/science/article/pii/S0370269319303879?via%3Dihub>.
- [19] M. Huertas-Company, J. R. Primack, A. Dekel, D. C. Koo, S. Lapiner, D. Ceverino, R. C. Simons, G. F. Snyder, M. Bernardi, Z. Chen, H. Domínguez-Sánchez, Z. Chen, C. T. Lee, B. Margalef-Bentabol, and D. Tuccillo. Deep Learning Identifies High-z Galaxies in a Central Blue Nugget Phase in a Characteristic Mass Range. *The Astrophysical Journal*, 858(2):114, 5 2018. ISSN 1538-4357. doi: 10.3847/1538-4357/aabfed. URL <http://stacks.iop.org/0004-637X/858/i=2/a=114?key=crossref.4e09299072a147355484d972ba56e818>.
- [20] H. Domínguez Sanchez, M. Huertas-Company, M. Bernardi, S. Kaviraj, J. L. Fischer, T. M.C. Abbott, F. B. Abdalla, J. Annis, S. Avila, D. Brooks, E. Buckley-Geer, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, C. E. Cunha, C. B. D’Andrea, L. N. Da Costa, C. Davis, J. De Vicente, P. Doel, A. E. Evrard, P. Fosalba, J. Frieman, J. García-Bellido, E. Gaztanaga, D. W. Gerdes, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, W. G. Hartley, D. L. Hollowood, K. Honscheid, B. Hoyle, D. J. James, K. Kuehn, N. Kuropatkin, O. Lahav, M. A.G. Maia, M. March, P. Melchior, F. Menanteau, R. Miquel, B. Nord, A. A. Plazas, E. Sanchez, V. Scarpine, R. Schindler, M. Schubnell, M. Smith, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, M. E.C. Swanson, G. Tarle, D. Thomas, A. R. Walker, and J. Zuntz. Transfer learning for galaxy morphology from one survey to another. *Monthly Notices of the Royal Astronomical Society*, 484(1):93–100, 3 2019. ISSN 13652966. doi: 10.1093/mnras/sty3497. URL <https://academic.oup.com/mnras/article/484/1/93/5266389>.
- [21] Mike Walmsley, Lewis Smith, Chris Lintott, Yarin Gal, Steven Bamford, Hugh Dickinson, Lucy Fortson, Sandor Kruk, Karen Masters, Claudia Scarlata, Brooke Simmons, Rebecca Smethurst, and Darryl Wright. Galaxy Zoo: probabilistic morphology through Bayesian CNNs and active learning. *Monthly Notices of the Royal Astronomical Society*, 491(2):1554–1574, 1 2020. ISSN 0035-8711. doi: 10.1093/mnras/stz2816. URL <https://academic.oup.com/mnras/article/491/2/1554/5583078>.
- [22] Mike Walmsley, Chris Lintott, Tobias Géron, Sandor Kruk, Coleman Krawczyk, Kyle W. Willett, Steven Bamford, Lee S. Kelvin, Lucy Fortson, Yarin Gal, William Keel, Karen L. Masters, Vihang Mehta, Brooke D. Simmons, Rebecca Smethurst, Lewis Smith, Elisabeth M. Baeten, and Christine MacMillan. Galaxy Zoo DECaLS: Detailed visual morphology measurements from volunteers and deep learning for 314 000 galaxies. *Monthly Notices of the Royal Astronomical Society*, 509(3):3966–3988, 12 2022. ISSN 13652966. doi: 10.1093/mnras/stab2093. URL <https://doi.org/10.1093/mnras/stab2093>.
- [23] Ting-Yun Cheng, H Domínguez Sánchez, J Vega-Ferrero, C J Conselice, M Siudek, A Aragón-Salamanca, M Bernardi, R Cooke, L Ferreira, M Huertas-Company, J Krywult, A Palmese, A Pieres, A A Plazas Malagón, A Carnero Rosell, D Gruen, D Thomas, D Bacon, D Brooks, D J James, D L Hollowood, D Friedel, E Suchyta, E Sanchez, F Menanteau, F Paz-Chinchón, G Gutierrez, G Tarle, I Sevilla-Noarbe, I Ferrero, J Annis, J Frieman, J García-Bellido, J Mena-Fernández, K Honscheid, K Kuehn, L N da Costa, M Gatti, M Raveri, M E S Pereira, M Rodríguez-Monroy, M Smith, M Carrasco Kind, M Aguena, M E C Swanson, M Weaverdyck, P Doel, R Miquel, R L C Ogando, R A Gruendl, S Allam, S R Hinton, S Dodelson, S Bocquet, S Desai, S Everett, and V Scarpine. Lessons learned from the two largest Galaxy morphological classification catalogues built by convolutional neural networks. *Monthly Notices of the Royal Astronomical Society*, 518(2):2794–2809, 11 2022. ISSN 0035-8711. doi: 10.1093/MNRAS/STAC3228. URL <https://academic.oup.com/mnras/article/518/2/2794/6823716>.
- [24] J. Vega-Ferrero, H. Domínguez Sánchez, M. Bernardi, M. Huertas-Company, R. Morgan, B. Margalef, M. Aguena, S. Allam, J. Annis, S. Avila, D. Bacon, E. Bertin, D. Brooks,

- A. Carnero Rosell, M. Carrasco Kind, J. Carretero, A. Choi, C. Conselice, M. Costanzi, L. N. Da Costa, M. E.S. Pereira, J. De Vicente, S. Desai, I. Ferrero, P. Fosalba, J. Frieman, J. García-Bellido, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, W. G. Hartley, S. R. Hinton, D. L. Hollowood, K. Honscheid, B. Hoyle, M. Jarvis, A. G. Kim, K. Kuehn, N. Kuropatkin, M. Lima, M. A.G. Maia, F. Menanteau, R. Miquel, R. L.C. Ogando, A. Palmese, F. Paz-Chinchón, A. A. Plazas, A. K. Romer, E. Sanchez, V. Scarpine, M. Schubnell, S. Serrano, I. Sevilla-Noarbe, M. Smith, E. Suchyta, M. E.C. Swanson, G. Tarle, F. Tarsitano, C. To, D. L. Tucker, T. N. Varga, and R. D. Wilkinson. Pushing automated morphological classifications to their limits with the Dark Energy Survey. *Monthly Notices of the Royal Astronomical Society*, 506(2):1927–1943, 7 2021. ISSN 0035-8711. doi: 10.1093/MNRAS/STAB594. URL <https://academic.oup.com/mnras/article/506/2/1927/6156625>.
- [25] Mike Walmsley, Tobias Géron, Sandor Kruk, Anna M M Scaife, Chris Lintott, Karen L Masters, James M Dawson, Hugh Dickinson, Lucy Fortson, Izzy L Garland, Kameswara Mantha, David O’, Ryan 10, Jürgen Popp, Brooke Simmons, Elisabeth M Baeten, and Christine Macmillan. Galaxy Zoo DESI: Detailed Morphology Measurements for 8.7M Galaxies in the DESI Legacy Imaging Surveys. *MNRAS*, 000:1–18, 9 2023. URL <https://arxiv.org/abs/2309.11425v1>.
- [26] Ryan Hausen and Brant E. Robertson. Morpheus : A Deep Learning Framework for the Pixel-level Analysis of Astronomical Image Data . *The Astrophysical Journal Supplement Series*, 248(1):20, 6 2020. ISSN 00670049. doi: 10.3847/1538-4365/ab8868. URL <http://arxiv.org/abs/1906.11248>.
- [27] Bryan Ostdiek, Ana Diaz Rivero, and Cora Dvorkin. Extracting the Subhalo Mass Function from Strong Lens Images with Image Segmentation. *The Astrophysical Journal*, 927(1):83, 3 2022. ISSN 0004-637X. doi: 10.3847/1538-4357/AC2D8D. URL <https://iopscience.iop.org/article/10.3847/1538-4357/ac2d8d><https://iopscience.iop.org/article/10.3847/1538-4357/ac2d8d/meta>.
- [28] Felix Richards, Adeline Paiement, Xianghua Xie, Elisabeth Sola, and Pierre Alain Duc. Panoptic Segmentation of Galactic Structures in LSB Images. *Proceedings of MVA 2023 - 18th International Conference on Machine Vision and Applications*, 2023. doi: 10.23919/MVA57639.2023.10216057.
- [29] Prabh Bhambra, Benjamin Joachimi, and Ofer Lahav. Explaining deep learning of galaxy morphology with saliency mapping. *Monthly Notices of the Royal Astronomical Society*, 511(4):5032–5041, 10 2022. ISSN 13652966. doi: 10.1093/mnras/stac368. URL <https://academic.oup.com/mnras/article/511/4/5032/6529251>.
- [30] Eder C.M. Gomes, Paulo R.P. Amora, Elvis M. Teixeira, Antonio G.S. Lima, Felipe T. Brito, Juliano F.C. Ciocari, and Javam C. Machado. UTTOS: A tool for testing UEFI code in OS environment. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9976 LNCS(1):218–224, 12 2016. ISSN 16113349. doi: 10.1007/978-3-319-47443-4{_}14. URL <http://stacks.iop.org/0004-637X/798/i=1/a=7?key=crossref.ce8f697114eeb3752ff6da50e3e66986>.
- [31] Abien Fred Agarap. Deep Learning using Rectified Linear Units (ReLU), 3 2018. URL <http://arxiv.org/abs/1803.08375>.
- [32] Arjun Dey, David J. Schlegel, Dustin Lang, Robert Blum, Kaylan Burleigh, Xiaohui Fan, Joseph R. Findlay, Doug Finkbeiner, David Herrera, Stéphanie Juneau, Martin Landriau, Michael Levi, Ian McGreer, Aaron Meisner, Adam D. Myers, John Moustakas, Peter Nugent, Anna Patej, Edward F. Schlafly, Alistair R. Walker, Francisco Valdes, Benjamin A. Weaver, Christophe Yèche, Hu Zou, Xu Zhou, Behzad Abareshi, T. M. C. Abbott, Bela Abolfathi, C. Aguilera, Shadab Alam, Lori Allen, A. Alvarez, James Annis, Behzad Ansarinejad, Marie Aubert, Jacqueline Beechert, Eric F. Bell, Segev Y. BenZvi, Florian Beutler, Richard M. Bielby, Adam S. Bolton, César Briceño, Elizabeth J. Buckley-Geer, Karen Butler, Annalisa Calamida, Raymond G. Carlberg, Paul Carter, Ricard Casas, Francisco J. Castander, Yumi Choi, Johan Comparat, Elena Cukanovaite, Timothée Delubac, Kaitlin DeVries, Sharmila Dey, Govinda Dhungana, Mark Dickinson, Zhejie Ding, John B. Donaldson, Yutong Duan, Christopher J. Duckworth, Sarah Eftekharzadeh, Daniel J. Eisenstein, Thomas Etourneau, Parker A. Fagrelus,

- Jay Farihi, Mike Fitzpatrick, Andreu Font-Ribera, Leah Fulmer, Boris T. Gänsicke, Enrique Gaztanaga, Koshy George, David W. Gerdes, Satya Gontcho A Gontcho, Claudio Gorgoni, Gregory Green, Julien Guy, Diane Harmer, M. Hernandez, Klaus Honscheid, Lijuan (Wendy) Huang, David J. James, Buell T. Jannuzi, Linhua Jiang, Richard Joyce, Armin Karcher, Sonia Karkar, Robert Kehoe, Jean-Paul Kneib, Andrea Kueter-Young, Ting-Wen Lan, Tod R. Lauer, Laurent Le Guillou, Auguste Le Van Suu, Jae Hyeon Lee, Michael Lesser, Laurence Perreault Levasseur, Ting S. Li, Justin L. Mann, Robert Marshall, C. E. Martínez-Vázquez, Paul Martini, Hélión du Mas des Bourboux, Sean McManus, Tobias Gabriel Meier, Brice Ménard, Nigel Metcalfe, Andrea Muñoz-Gutiérrez, Joan Najita, Kevin Napier, Gautham Narayan, Jeffrey A. Newman, Jundan Nie, Brian Nord, Dara J. Norman, Knut A. G. Olsen, Anthony Paat, Nathalie Palanque-Delabrouille, Xiyang Peng, Claire L. Poppett, Megan R. Poremba, Abhishek Prakash, David Rabinowitz, Anand Raichoor, Mehdi Rezaie, A. N. Robertson, Natalie A. Roe, Ashley J. Ross, Nicholas P. Ross, Gregory Rudnick, Sasha Safonova, Abhijit Saha, F. Javier Sánchez, Elodie Savary, Heidi Schweiker, Adam Scott, Hee-Jong Seo, Huanyuan Shan, David R. Silva, Zachary Slepian, Christian Soto, David Sprayberry, Ryan Staten, Coley M. Stillman, Robert J. Stupak, David L. Summers, Suk Sien Tie, H. Tirado, Mariana Vargas-Magaña, A. Katherina Vivas, Risa H. Wechsler, Doug Williams, Jinyi Yang, Qian Yang, Tolga Yapici, Dennis Zaritsky, A. Zenteno, Kai Zhang, Tianmeng Zhang, Rongpu Zhou, and Zhimin Zhou. Overview of the DESI Legacy Imaging Surveys. *The Astronomical Journal*, 157(5):168, 4 2019. ISSN 00046256. doi: 10.3847/1538-3881/ab089d. URL <http://arxiv.org/abs/1804.08657>.
- [33] Michael J. Smith, James E. Geach, Ryan A. Jackson, Nikhil Arora, Connor Stone, and Stéphane Courteau. Realistic galaxy image simulation via score-based generative models. *Monthly Notices of the Royal Astronomical Society*, 511(2):1808–1818, 11 2022. ISSN 13652966. doi: 10.1093/mnras/stac130. URL <https://arxiv.org/abs/2111.01713v1>.
- [34] Diganta Misra. Mish: A Self Regularized Non-Monotonic Activation Function. 8 2019. URL <https://arxiv.org/abs/1908.08681v3>.
- [35] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment Anything. 4 2023. URL <https://arxiv.org/abs/2304.02643v1>.
- [36] Nelson Christensen, Renate Meyer, Lloyd Knox, and Ben Luey. Bayesian methods for cosmological parameter estimation from cosmic microwave background measurements, 10 2001. ISSN 02649381. URL <http://arxiv.org/abs/1110.3193>.
- [37] Mike Walmsley, Campbell Allen, Ben Aussel, Micah Bowles, Kasia Gregorowicz, Inigo Val Slijepcevic, Chris J Lintott, Anna M. Scaife, Maja Jabłońska, Kosio Karchev, Denise Lanzieri, Devina Mohan, David O’Ryan, Bharath Saiguhun, Crisel Suárez, Nicolás Guerra-Varas, and Renuka Velu. Zoobot: Adaptable Deep Learning Models for Galaxy Morphology. *Journal of Open Source Software*, 8(85):5312, 2023. doi: 10.21105/joss.05312. URL <https://doi.org/10.21105/joss.05312>.
- [38] Mike Walmsley, Inigo Val Slijepcevic, Micah Bowles, and Anna M. M. Scaife. Towards Galaxy Foundation Models with Hybrid Contrastive Learning. *Machine Learning for Astrophysics Workshop at the Thirty Ninth International Conference on Machine Learning (ICML 2022)*, 6 2022. doi: 10.48550/arxiv.2206.11927.
- [39] Tomasz Różański, Yuan-Sen Ting, and Maja Jabłońska. Toward a Spectral Foundation Model: An Attention-Based Approach with Domain-Inspired Fine-Tuning and Wavelength Parameterization. 6 2023. URL <https://arxiv.org/abs/2306.15703v1>.
- [40] Henry W. Leung, Jo Bovy, Henry W. Leung, and Jo Bovy. Towards an astronomical foundation model for stars with a Transformer-based model. *arXiv*, 000:arXiv:2308.10944, 8 2023. doi: 10.48550/ARXIV.2308.10944. URL <https://ui.adsabs.harvard.edu/abs/2023arXiv230810944L/abstract>.
- [41] Inigo V. Slijepcevic, Anna M. M. Scaife, Mike Walmsley, Micah Bowles, O. Ivy Wong, Stanislav S. Shabala, and Sarah V. White. Radio Galaxy Zoo: Building a multi-purpose foundation model for radio astronomy with self-supervised learning. pages 1–13, 5 2023. URL <https://arxiv.org/abs/2305.16127v2>.