

# From fat droplets to floating forests: cross-domain transfer learning using a PatchGAN-based segmentation model

Kameswara Bharadwaj Mantha<sup>1,\*†</sup>, Ramanakumar Sankar<sup>1,\*†</sup>, Yuping Zheng<sup>1</sup>, Lucy Fortson<sup>1</sup>, Thomas Pengo<sup>2</sup>, Douglas Mashek<sup>3</sup>, Mark Sanders<sup>3</sup>, Trace Christensen<sup>4</sup>, Jeffrey Salisbury<sup>4</sup>, Laura Trouille<sup>5</sup>, Jarrett E. K. Byrnes<sup>6</sup>, Isaac Rosenthal<sup>6</sup>, Henry Houskeeper<sup>7</sup> and Kyle Cavanaugh<sup>7</sup>

<sup>1</sup>School of Physics & Astronomy, University of Minnesota, Twin Cities, 116 Church St SE, Minneapolis, MN, 55455

<sup>2</sup>University of Minnesota Informatics Institute, 2231 6th St SE, Minneapolis, MN, 55455

<sup>3</sup>Medical School, University of Minnesota, Twin Cities, 420 Delaware Street SE, Minneapolis, MN, 55455

<sup>4</sup>Mayo Clinic, 200 First Street SW, Rochester, MN, 55905

<sup>5</sup>Adler Planetarium, 1300 S DuSable Lake Shore Dr., Chicago, IL 60605

<sup>6</sup>Department of Biology, University of Massachusetts Boston 100 Morrissey Blvd; Boston, MA, 02125

<sup>7</sup>Department of Geography, University of California Los Angeles, Los Angeles, CA 90095

## Abstract

Many scientific domains gather sufficient labels to train machine algorithms through human-in-the-loop techniques provided by the Zooniverse.org citizen science platform. As the range of projects, task types and data rates increase, acceleration of model training is of paramount concern to focus volunteer effort where most needed. The application of Transfer Learning (TL) between Zooniverse projects holds promise as a solution. However, understanding the effectiveness of TL approaches that pretrain on large-scale generic image sets vs. images with similar characteristics possibly from similar tasks is an open challenge. We apply a generative segmentation model on two Zooniverse project-based data sets: (1) to identify fat droplets in liver cells (FatChecker; FC) and (2) the identification of kelp beds in satellite images (Floating Forests; FF) through transfer learning from the first project. We compare and contrast its performance with a TL model based on the COCO image set, and subsequently with baseline counterparts. We find that both the FC and COCO TL models perform better than the baseline cases when using  $> 75\%$  of the original training sample size. The COCO-based TL model generally performs better than the FC-based one, likely due to its generalized features. Our investigations provide important insights into usage of TL approaches on multi-domain data hosted across different Zooniverse projects, enabling future projects to accelerate task completion.

## Keywords

datasets, generative adversarial neural networks, UNET generator, patch-based discriminator, focal tversky loss, transfer learning

## 1. Introduction

Citizen Science has established itself as a valuable method for distributed data analysis enabling research teams from diverse domains to solve problems involving large quantities of data with complexity levels requiring human pattern recognition capabilities [1, 2]. As the largest citizen science platform, Zooniverse.org has enabled over 2.5 million volunteers to provide over half a billion annotations on hundreds of projects across the sciences and humanities. Many of these projects use the result-

ing labels to train machine learning algorithms typically training models from scratch e.g., [3, 4, 5, 6, 7, 8, 9, 10, 11]. To accelerate labeling efficiencies across the platform, the Zooniverse human-machine system should take advantage of transfer learning techniques, especially when volunteer engagement is at a premium. When applying transfer learning, a new project would require fewer labels from volunteers to achieve the same performance as training a model from scratch. Volunteer labelers would thus be able to focus on tasks more suited to humans such as anomaly detection e.g., [12].

Transfer learning (TL) is an established approach, where the feature space from a pretrained model can be transferred to another framework and fine tuned to perform analogous or different tasks. Feature extraction is typically performed using Deep Convolutional Neural Networks (CNNs) such as [13, 14]. Transfer learning generally uses models trained on data that is either “out-of-domain” (i.e., training data characteristics are different from data at hand) or “in-domain” (data that are similar

*Human-in-the-loop data curation workshop at ACM CIKM 2022, Oct 17–21, 2022, Atlanta, GA*

\*Corresponding author.

†KBM and RS contributed equally to the majority of this research. Additional authors contributed to specific aspects including initial models, data sets and project or platform development.

✉ manth145@umn.edu (K. B. Mantha); rsankar@umn.edu (R. Sankar)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

or closely relatable to the data at hand). Quantifying the gains provided by these different TL approaches is an active area of research, where studies find several factors to be at play that govern its effectiveness: Accuracy and architecture choice of the pretrained model [15], robustness of model to input adversarial noise [16], and type of task to which the TL is being applied [17]. Recent works (e.g., [12, 8]) have demonstrated that transfer learning from a model pretrained on in-domain data performs better than transfer learning from out-of-domain data. On the other hand, some studies find that TL models based on out-of-domain data (e.g., ImageNet or COCO datasets) perform on par with or better than the in-domain TL models [18, 19].

In order to leverage the Zooniverse’s large library of image-label pairs across multiple domains, there is thus a clear need to better understand the effectiveness of cross-domain transfer learning. In particular, we are interested in the application of transfer learning specifically to projects that share task similarity across a wide range of domains. For example, image segmentation tasks vary across vastly different disciplines, from cell biology to satellite imagery. Frameworks such as the U-Net [20], Recurrent Convolutional Networks such as Mask-RCNNs [21], and Generative Adversarial Networks (GANs; e.g., [22, 23]) have been used to perform such object segmentation across multiple domains and data sets. However, robust learning of such segmentation models from scratch often requires large annotated training samples that may not be available (e.g., medical imaging), which can lead to poor generalizability of the learnt features to newer data, even in related domains. While Zooniverse can provide these large annotation sets per project, this comes at the cost of volunteer effort which we seek to optimize.

In an effort to increase project completion rates, this study investigates potential machine performance gains through transfer learning across domains by leveraging the shared task similarity between Zooniverse projects. We use a PatchGAN-based [23] segmentation model<sup>1</sup> to investigate the effectiveness of segmenting kelp beds from satellite images. Particularly, we test transfer learning from the COCO dataset (i.e., out-of-domain) and microscopy imaging of lipid droplets in liver cells (pseudo-in-domain) and compare them to their corresponding “trained from scratch” counterparts.

## 2. Methods

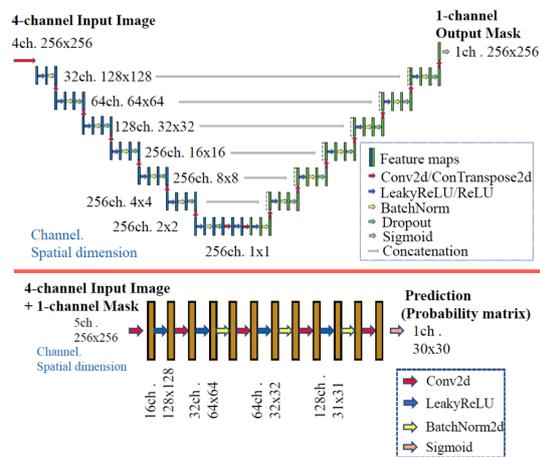
In this section, we detail our PatchGAN architecture [23], the training and testing data and its preparation, and the description of the five models analyzed in our work.

<sup>1</sup><https://github.com/ramanakumars/patchGAN/>

### 2.1. PatchGAN Framework

The implemented PatchGAN framework is inherited from the Pix2Pix GAN architecture in [23], which is a conditional GAN for realizing paired image-to-image translation. The PatchGAN architecture consists of a Generator ( $G$ ) and Discriminator ( $D$ ):

The generator is composed of a U-Net [20], a U-shaped encoder-decoder neural network, with skip connections across the bottleneck layer (Figure 1). The encoder (decoder) comprises of 6 downsampling (upsampling) blocks, each consisting of  $4 \times 4$  convolution (transposed convolution), Leaky ReLU activation, and a batch normalization layer. All the blocks in the inner layers of the network also include a dropout layer which omits 50% of the extracted features during training. The outputs of the transposed convolutions are also concatenated with the corresponding skip connection feature map from the encoder block.



**Figure 1:** U-Net Generator (top) and Discriminator (bottom) of our PatchGAN framework.

The discriminator is a patch-wise binary classifier that takes a concatenation of the input image and its corresponding ground truth or generated mask and outputs a  $30 \times 30$  probability matrix. Each unit of this matrix represents a  $70 \times 70$  patch of the input image, and provides the probability that the patch is real.

### 2.2. Data

For this study, we use three sources for our image-mask pairs: the Floating Forests dataset, Etch-a-Cell dataset and the COCO-stuff. The former two are Zooniverse projects focusing on image segmentation, while the latter represents a generic image dataset that is used in computer vision, representing an out-of-domain dataset compared to the former two. These three data sources represent a diverse feature set on which to perform our

transfer learning experiment. Figure 2 shows an example of an image-mask pair from each dataset.

### 2.2.1. Floating Forests (*FF*)

Floating Forests is an ecology-based citizen science project hosted on Zooniverse.org<sup>2</sup> to identify kelp beds in Landsat imagery. The project presents segments of Landsat data to Zooniverse volunteers, who draw outlines around the kelp beds. These annotations are aggregated using a pixel-by-pixel consensus to create masks of the kelp beds in the corresponding Landsat segments. We use 4 channels from the Landsat data (Blue, Green, Red and near Infrared) to train the patchGAN on the image-mask pairs. This FF data comprises 6,967 ( $350 \times 350$  pix) image-mask pairs. We pre-process these data such that each pair is cropped into four  $256 \times 256$  overlapping cutouts, and augment each crop 5 times (rotation and flipping). This resulted in 118,440 training and 4180 testing images.

### 2.2.2. Etch-a-Cell: Fat Checker (*FC*)

Etch-a-Cell: Fat Checker is a cell biology project hosted on Zooniverse.org<sup>3</sup> to identify lipid droplets in electron microscopy data. The Zooniverse project presents 2D slices of the data to volunteers who annotate the outline of the lipid droplet. The lipid mask is generated by aggregating the annotations by multiple volunteers based on consensus. The data set consists of 2341 image-mask pairs and each image is  $1200 \times 1200$  pix in shape, with 3 channels. We split the sample into 2,106 training and 235 testing sets. We transform these images and masks to work with our PatchGAN framework by resizing them to  $512 \times 512$  pix and generating five crops (four corners and one center crop). We further augment them by applying three rotations (90, 180, 270 deg) per image, yielding augmented training and testing samples of 42120 and 4700 images, respectively.

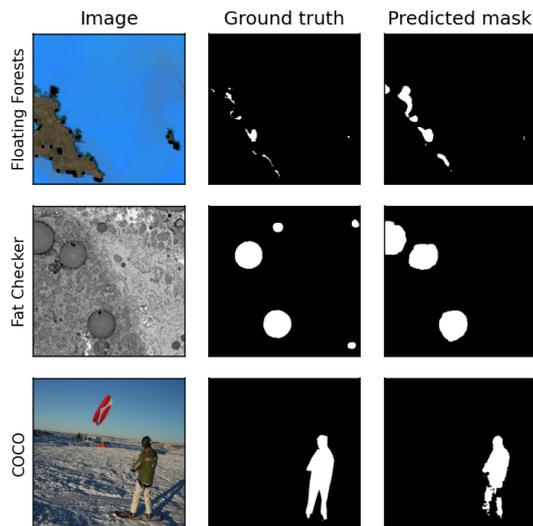
### 2.2.3. COCO-Stuff

The Common Objects in Context (COCO; [24]) is a large collection of several real-world images with objects set in various simple to complex scenes, which are annotated by outlines<sup>4</sup>. [25] further processed the COCO data set to produce dense pixel-wise annotations for them (the COCO-Stuff data set; hereafter COCO). These images and annotated masks vary widely in their shapes, and therefore, we standardize these images by resizing them to a  $256 \times 256$  pix shape. For our PatchGAN training, we limit the training and testing data to those that host

<sup>2</sup><https://www.zooniverse.org/projects/zooniverse/floating-forests/>

<sup>3</sup><https://www.zooniverse.org/projects/dwright04/etch-a-cell-fat-checker>

<sup>4</sup><https://github.com/nightrome/cocostuff>



**Figure 2:** Visualization of example input image, truth mask, and patchGAN predicted output mask.

the ‘person’ class. This amounts to 63785 training and 2673 testing image-mask pairs.

## 2.3. Experimental Design

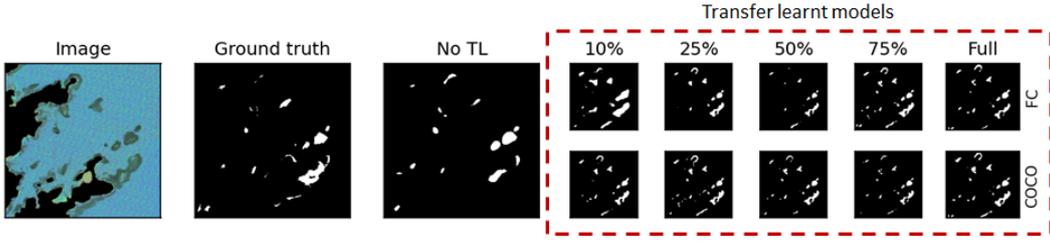
In this work, we investigate the potential of cross-domain transfer learning by training 5 models. The first 3 models are trained from scratch –  $\Lambda_{FF}$ ,  $\Lambda_{FC}$ , and  $\Lambda_{COCO}$  – using 100% of their corresponding data sets *FF*, *FC*, and *COCO*, respectively. Next, we train the  $\Lambda_{FC \rightarrow FF}$  and  $\Lambda_{COCO \rightarrow FF}$  by transferring the weights from the trained  $\Lambda_{FC}$  and  $\Lambda_{COCO}$  models to the  $\Lambda_{FF}$ . By comparing between the baseline  $\Lambda_{FF}$  to the transfer learnt models  $\Lambda_{FC \rightarrow FF}$  and  $\Lambda_{COCO \rightarrow FF}$ , we quantify the impact of performing transfer learning on the accelerated learning of the  $\Lambda_{FF}$  model from two distinct feature initializations. During this transfer learning exercise, we also vary the amount of training data used from 10%-100%.

## 3. Training & Results

In this section, we outline the training strategy and provide details of the hyper parameters. We also present the results of our training and discuss the outcomes of our transfer learning exercise.

### 3.1. Training Strategy

Our  $\Lambda_{FF}$ ,  $\Lambda_{FC}$ , and  $\Lambda_{COCO}$  models have been trained for 50 epochs. For the generator, we use the Focal Tversky Loss (FTL; [26]), which is a generalized version of the Tversky Loss (TL) defined in terms of the Tversky Index



**Figure 3:** Comparison of generated mask from different model runs on the Floating Forests data, showing different performance gains from transfer learning.

(TI) as:

$$TI = \frac{TP}{TP + \alpha FN + \beta FP} \rightarrow TL = (1 - TI) \rightarrow FTL \quad (1)$$

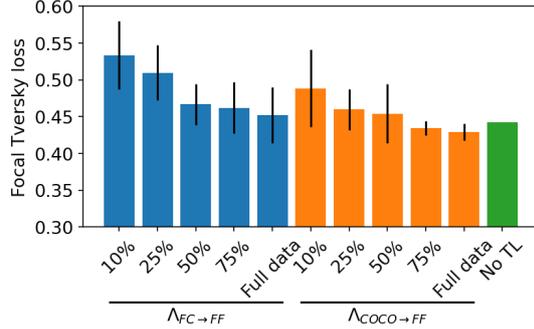
For our training, we use  $\alpha = 0.7$  and  $\beta = 0.3$ . The  $\gamma$  parameter controls the non-linearity of the TL with respect to the  $TI$ , enabling the learning to focus on easier ( $\gamma < 1$ ) vs. harder ( $\gamma > 1$ ) examples. We use  $\gamma = 0.75$  during our training. For the discriminator optimization, we use the Binary Cross-Entropy ( $BCE$ ) loss. Specifically, our total discriminator loss is the average of two components: the discriminator applied on the generated mask (i.e., against a fake label), and applied on the true mask (i.e., the real label). For both the generator and discriminator, we use the Adam optimizer with an initial learning rate  $5 \times 10^{-4}$  and  $1 \times 10^{-4}$  respectively, decayed exponentially by  $\tau = 0.95$ , applied every 5 epochs.

### 3.2. Transfer learning strategy

For our transfer learning based model training of  $\Lambda_{FC \rightarrow FF}$  and  $\Lambda_{COCO \rightarrow FF}$ , we load the weights of the  $\Lambda_{FC}$  and  $\Lambda_{COCO}$  models into the freshly initialized  $\Lambda_{FF}$  model architecture. To account for the 3 vs 4 channel mismatch between the  $\Lambda_{COCO}$ ,  $\Lambda_{FC}$  and  $\Lambda_{FF}$ , we load model layer parameters excluding the input layer. For each model, we train 5 different versions, using random subsets of 10%, 25%, 50%, 75% and 100% of the full Floating Forests data, to compare TL efficiency gains from having a smaller dataset. For these experiments, we also use only the first 6,967 un-augmented images for re-training. We train the  $\Lambda_{FC \rightarrow FF}$  and  $\Lambda_{COCO \rightarrow FF}$  models with the same hyper-parameter settings as the aforementioned “from scratch” models for 50 epochs.

### 3.3. Results and discussion

We find that our  $\Lambda_{FF}$ ,  $\Lambda_{FC}$  and  $\Lambda_{COCO}$  generally predict the annotation masks reasonably well (Figure 2), qualitatively matching with the ground truths. Figures 3 and 4 show our transfer learning results. In Figure 4, we show our average validation loss for the different



**Figure 4:** Comparison of mean final loss on Floating Forests validation data across the different models.

model training runs. As expected, larger training samples provide much better performance, but we also find that the model pretrained on the COCO dataset provides noticeably better performance on the Floating Forests data, compared to both  $\Lambda_{FC \rightarrow FF}$  and also  $\Lambda_{FF}$ . In fact, the  $\Lambda_{COCO \rightarrow FF}$  is able to match the performance of the  $\Lambda_{FF}$  model with between 50-75% of the training Floating Forests dataset.

In Figure 3, we show examples highlighting the difference between the generated masks from  $\Lambda_{FF}$  and corresponding masks from  $\Lambda_{FC \rightarrow FF}$  and  $\Lambda_{COCO \rightarrow FF}$ . The sharpness of the kelp beds is poorly reconstructed by the  $\Lambda_{FF}$  model but is well captured by the transfer learnt models (particularly when training  $\Lambda_{COCO \rightarrow FF}$  with more than 75% of the original data). The transfer learnt models are also better at capturing kelp beds not identified in the original consensus data. For example, both the ground truth and  $\Lambda_{FF}$  fail to reveal the kelp beds in the top left of the image, but these are picked up well by the transfer learnt models.

This is likely due to the large diversity of the features in the COCO dataset, making it a much more robust feature extraction network to transfer learn from. Indeed, compared to  $\Lambda_{FC \rightarrow FF}$ , the  $\Lambda_{COCO \rightarrow FF}$  model-detected kelp beds are qualitatively better visually (e.g., Figure 3), especially at lower training data sizes. This is likely compounded with the lower feature diversity in both the Floating Forest and Fat Checker data sets, given

the fewer number of samples in the training data and low variety in target classes.

### 3.3.1. Transfer learning approaches for citizen science datasets

For the Zooniverse platform, this study provides an avenue to build quick access for projects to use machine learning frameworks for simple tasks (e.g., image segmentation), by transfer learning from existing models on a small sample of volunteer annotated data sets. However, despite the results presented here, there are still several key questions which need to be answered:

**Domain dependency:** It is unclear how much of the performance gained from COCO was a ‘global truth’. That is, whether COCO (or similarly diverse datasets) are immediately applicable to out-of-domain data, for all domains, or if there are domain-specific restrictions which allow these performance gains to occur on data such as Floating Forests. This requires more experiments with increasingly different data sets on Zooniverse to investigate the range of performance gains possible.

**Task dependency:** Previous studies on transfer learning across domains show significant variations in performance across different task types. For example, image classification tasks (e.g., [12, 17]) show lower gains than image segmentation based tasks (e.g., [18]). We need to further investigate the inherent difficulty associated with different tasks on Zooniverse projects, and how effectively they can be transferred between domains. [12], for example, show that significant boosts to performance is only provided by using in-domain transfer learning.

**Target data purity:** For Zooniverse projects, data labels are generally provided by volunteers and are aggregated based on volunteer consensus. In this study, we found that transfer learning can help mitigate data purity effects, since transfer learnt feature extraction models are generally robust to mislabeled data. The extent to which transfer learning models are sensitive to data purity effects needs to be further investigated.

In conclusion, we find that transfer learning can provide a significant boost to projects that contain similar tasks on Zooniverse. However, the extent to which this can be generalized across the full Zooniverse ecosystem is a question of ongoing study.

## Acknowledgements

The authors would like to thank the Zooniverse volunteers without whom this work would not have been possible. RS, KM, LF, YZ, LT would like to acknowledge partial support from the National Science Foundation under grant numbers IIS 2006894 and OAC 1835530. Partial support by RS, KM, LF, TP, MS, TC, JS is acknowledged through Minnesota Partnership MNP IF#119.09.

We also thank Lucy Collinson and the Electron Microscopy Science Technology Platform (The Francis Crick Institute, London UK) for their input into this project. This work was supported in part by the Francis Crick Institute which receives its core funding from Cancer Research UK (FC001999), the UK Medical Research Council (FC001999), and the Wellcome Trust (FC001999). This project has been made possible in part by grant number 2020-225438 from the Chan Zuckerberg Initiative DAF, an advised fund of Silicon Valley Community Foundation (H.S.). This publication uses data generated via the Zooniverse.org platform, development of which is funded by generous support, including a Global Impact Award from Google, and by a grant from the Alfred P. Sloan Foundation.

## References

- [1] L. Trouille, C. J. Lintott, L. F. Fortson, Citizen science frontiers: Efficiency, engagement, and serendipitous discovery with human-machine systems, *Proceedings of the National Academy of Sciences* 116 (2019) 1902–1909. URL: <https://www.pnas.org/content/116/6/1902>. doi:10.1073/pnas.1807190116.
- [2] L. Fortson, D. Wright, C. Lintott, L. Trouille, Optimizing the human-machine partnership with Zooniverse, in: *CI 2018: ACM Collective Intelligence*, ACM, <http://arxiv.org/abs/1809.09738>, 2018. URL: <http://arxiv.org/abs/1809.09738>. arXiv:1809.09738.
- [3] C. N. Beaumont, A. A. Goodman, S. Kendrew, J. P. Williams, R. Simpson, The Milky Way Project: Leveraging Citizen Science and Machine Learning to Detect Interstellar Bubbles, *ApJS* 214 (2014) 3. doi:10.1088/0067-0049/214/1/3. arXiv:1406.2692.
- [4] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Rohani, S. Allen, M. Cabero, K. Crowston, A. K. Katsaggelos, S. L. Larson, et al., Gravity Spy: integrating advanced ligo detector characterization, machine learning, and citizen science, *Classical and Quantum Gravity* 34 (2017) 064003.
- [5] M. Norouzzadeh, A. Nguyen, M. Kosmala, A. Swanson, C. Packer, J. Clune, Automatically identifying wild animals in camera trap images with deep learning, arXiv preprint arXiv:1703.05830 (2017).
- [6] D. Wright, C. Lintott, S. Smartt, K. Smith, L. Fortson, L. Trouille, C. Allen, M. Beck, M. Bouslog, A. Boyer, K. Chambers, H. Flewelling, W. Granger, E. Magner, A. McMaster, G. Miller, J. O’Donnell, B. Simons, H. Spiers, J. Tonry, M. Veldhuis, R. Wainscoat, C. Waters, M. Willman, Z. Wolfenbarger, D. Young, A transient search using combined hu-

- man and machine classifications, *Monthly Notices of the Royal Astronomical Society* 472 (2017) 1315–1323. URL: <http://dx.doi.org/10.1093/mnras/stx1812>. doi:10.1093/mnras/stx1812.
- [7] H. Domínguez Sánchez, M. Huertas-Company, M. Bernardi, D. Tuccillo, J. L. Fischer, Improving galaxy morphologies for SDSS with Deep Learning, *Monthly Notices of the Royal Astronomical Society* 476 (2018) 3661–3676. URL: <https://doi.org/10.1093/mnras/sty338>. doi:10.1093/mnras/sty338.
- [8] M. Willi, R. T. Pitman, A. W. Cardoso, C. Locke, A. Swanson, A. Boyer, M. Veldthuis, L. Fortson, Identifying animal species in camera trap images using deep learning and citizen science, *Methods in Ecology and Evolution* 10 (2019) 80–91.
- [9] M. Laraia, D. Wright, H. Dickinson, A. Simenstad, K. Flanagan, S. Serjeant, L. Fortson, VERITAS Collaboration, Muon Hunter 2.0: efficient crowdsourcing of labels for IACT image analysis, in: 36th International Cosmic Ray Conference (ICRC2019), volume 36 of *International Cosmic Ray Conference*, 2019, p. 678.
- [10] O. Ranadive, S. van der Lee, V. Tang, K. Chao, Applying Machine Learning to Crowd-sourced Data from Earthquake Detective, *arXiv e-prints* (2020) arXiv:2011.04740. arXiv:2011.04740.
- [11] H. Spiers, H. Songhurst, L. Nightingale, J. de Folter, R. Hutchings, C. J. Peddie, A. Weston, A. Strange, S. Hindmarsh, C. Lintott, L. M. Collinson, M. L. Jones, Citizen science, cells and cnns – deep learning for automatic segmentation of the nuclear envelope in electron microscopy data, trained with volunteer segmentations, *bioRxiv* (2020). URL: <https://www.biorxiv.org/content/early/2020/07/29/2020.07.28.223024>. doi:10.1101/2020.07.28.223024.
- [12] M. Walmsley, A. M. M. Scaife, C. Lintott, M. Lochner, V. Etsebeth, T. Géron, H. Dickinson, L. Fortson, S. Kruk, K. L. Masters, K. B. Mantha, B. D. Simmons, Practical galaxy morphology tools from deep supervised representation learning, *Monthly Notices of the Royal Astronomical Society* 513 (2022) 1581–1599. doi:10.1093/mnras/stac525. arXiv:2110.12735.
- [13] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [14] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, *arXiv preprint arXiv:1409.1556* (2014).
- [15] S. Hosseinzadeh Kassani, P. Hosseinzadeh Kassani, M. J. Wesolowski, K. A. Schneider, R. Deters, Deep transfer learning based model for colorectal cancer histopathology segmentation: A comparative study of deep pre-trained models, *International Journal of Medical Informatics* 159 (2022) 104669. URL: <https://www.sciencedirect.com/science/article/pii/S1386505621002951>. doi:<https://doi.org/10.1016/j.ijmedinf.2021.104669>.
- [16] H. Salman, A. Ilyas, L. Engstrom, A. Kapoor, A. Madry, Do adversarially robust imagenet models transfer better?, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems*, volume 33, Curran Associates, Inc., 2020, pp. 3533–3545. URL: <https://proceedings.neurips.cc/paper/2020/file/24357dd085d2c4b1a88a7e0692e60294-Paper.pdf>.
- [17] K. Thenmozhi, U. S. Reddy, Crop pest classification based on deep convolutional neural network and transfer learning, *Computers and Electronics in Agriculture* 164 (2019) 104906.
- [18] M. Majurski, P. Manescu, S. Padi, N. Schaub, N. Hotaling, C. Simon Jr, P. Bajcsy, Cell image segmentation using generative adversarial networks, transfer learning, and augmentations, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.
- [19] J. Ma, L. Bao, Q. Lou, D. Kong, Transfer learning for automatic joint segmentation of thyroid and breast lesions from ultrasound images, *International Journal of Computer Assisted Radiology and Surgery* 17 (2022) 363–372.
- [20] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2015, pp. 234–241.
- [21] K. He, G. Gkioxari, P. Dollár, R. Girshick, Mask r-cnn, in: *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [22] Y. Huo, Z. Xu, S. Bao, A. Assad, R. G. Abramson, B. A. Landman, Adversarial synthesis learning enables segmentation without target modality ground truth, in: *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*, IEEE, 2018, pp. 1217–1220.
- [23] P. Isola, J.-Y. Zhu, T. Zhou, A. A. Efros, Image-to-image translation with conditional adversarial networks, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.
- [24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick, Microsoft coco: Common objects in context, in: *European conference on computer vision*, Springer, 2014, pp. 740–755.
- [25] H. Caesar, J. Uijlings, V. Ferrari, Coco-stuff: Thing and stuff classes in context, in: *Computer vision and pattern recognition (CVPR)*, 2018 IEEE conference

- on, IEEE, 2018.
- [26] N. Abraham, N. Mefraz Khan, A Novel Focal Tversky loss function with improved Attention U-Net for lesion segmentation, arXiv e-prints (2018) arXiv:1810.07842. arXiv: 1810.07842.