

DeepAstroUDA: Semi-Supervised Universal Domain Adaptation for Cross-Survey Galaxy Morphology Classification and Anomaly Detection

A. Ćiprijanović^{1,2,‡}, A. Lewis^{1,‡}, K. Pedro¹, S. Madireddy³
B. Nord^{1,2,4,5}, G. N. Perdue¹, S. M. Wild^{3,6}

¹Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

²Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA

³Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, USA

⁴Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

⁵Laboratory for Nuclear Science, MIT, Cambridge MA, 02139, USA

⁶Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

E-mail: aleksand@fnal.gov

Abstract. Artificial intelligence methods show great promise in increasing the quality and speed of work with large astronomical datasets, but the high complexity of these methods leads to the extraction of dataset-specific, non-robust features. Therefore, such methods do not generalize well across multiple datasets. We present a universal domain adaptation method, *DeepAstroUDA*, as an approach to overcome this challenge. This algorithm performs semi-supervised domain adaptation and can be applied to datasets with different data distributions and class overlaps. Non-overlapping classes can be present in any of the two datasets (the labeled source domain, or the unlabeled target domain), and the method can even be used in the presence of unknown classes. We apply our method to three examples of galaxy morphology classification tasks of different complexities (3-class and 10-class problems), with anomaly detection: 1) datasets created after different numbers of observing years from a single survey (LSST mock data of 1 and 10 years of observations); 2) data from different surveys (SDSS and DECaLS); and 3) data from observing fields with different depths within one survey (wide field and Stripe 82 deep field of SDSS). For the first time, we demonstrate the successful use of domain adaptation between very discrepant observational datasets. *DeepAstroUDA* is capable of bridging the gap between two astronomical surveys, increasing classification accuracy in both domains (up to 40% on the unlabeled data), and making model performance consistent across datasets. Furthermore, our method also performs well as an anomaly detection algorithm and successfully clusters unknown class samples even in the unlabeled target dataset.

Keywords: domain adaptation, convolutional neural networks, deep learning, model robustness, galaxy morphological classification, sky surveys

‡ Equal contribution

1. Introduction

With big datasets from current and next generation astronomical surveys like the Dark Energy Survey (DES; Dark Energy Survey Collaboration et al., 2016), the Hyper Suprime-Cam Subaru Strategic Program (HSC-SSP; Aihara et al., 2018), the Vera Rubin Legacy Survey of Space and Time (LSST; Ivezić et al., 2019), and the Nancy Grace Roman Space Telescope[§], development of artificial intelligence (AI) algorithms capable of combining knowledge from different telescopes will open doors to many new insights. For a review of the impact of AI on analysis of galaxy surveys, see Huertas-Company and Lanusse (2022). Furthermore, many astrophysics and cosmology studies often start from simulated data that are well suited for learning the connection between the physical parameters and the observables. Simulations are the best (and sometimes the only) resource for creating labeled datasets needed for training supervised learning algorithms, which are ultimately intended to be used on real observational data. Unfortunately, standard deep learning algorithms are not well suited for working with multiple datasets. Very small, even pixel-level, differences between datasets can cause an algorithm trained on one dataset to experience a substantial drop in performance or even not work at all on another dataset (Gide et al., 2016; Dodge and Karam, 2016, 2017; Ford et al., 2019; Ćiprijanović et al., 2022). In astronomy, differences between datasets are often substantially larger. Simulated data can be different from observations due to computational constraints, approximations, missing or unknown physics, or imperfect addition of observational effects; while differences between observational datasets come from different telescope characteristics, observing times, and observing conditions.

Domain adaptation (DA) research includes the development of methods designed to bridge the gap between datasets and enable the creation of deep learning models that perform well on multiple datasets at the same time. This is done by guiding the model to learn and use only features present in both datasets (Csurka, 2017; Wang and Deng, 2018; Wilson and Cook, 2020), that is, domain-invariant features. These methods can be divided into 1) distance-based methods such as Maximum Mean Discrepancy (MMD; Gretton et al., 2007; Gretton et al., 2012), Deep Correlation Alignment (CORAL; Sun and Saenko, 2016), Central Moment Discrepancy (CMD; Zellinger et al., 2019); and 2) adversarial-based methods such as Domain Adversarial Neural Networks (DANNs; Ganin et al., 2016) and Conditional Domain Adversarial Networks (CDAN; Long et al., 2017).

In astronomy, DA was first studied and used by Vi-

lalta et al. (2019) for two different tasks: classification of Supernovae Ia and the identification of Mars landforms. Then, in our previous work Ćiprijanović et al. (2021), DA was used for the problem of identifying merging galaxies in simulated and real data. In this paper, we have used two DA techniques, MMD (Gretton et al., 2012) and DANNs (Ganin et al., 2016), and have successfully trained a model that works on two Illustris-1 (Vogelsberger et al., 2014) simulated data sets of distant merging galaxies, as well as models that work on Illustris-1 simulated data of nearby merging galaxies and observed data from the Sloan Digital Sky Survey (SDSS; York et al., 2000). In Ćiprijanović et al. (2022), we showed that even smaller differences between datasets, such as inadvertent data perturbations that can naturally occur in astronomical data pipelines, can also cause the model trained on clean data to make catastrophic errors. Here, we developed deep learning models to classify galaxy morphology and include DA to increase model robustness and allow the model to perform well even in the presence of data perturbations.

Unfortunately, DA methods can be hard to fine tune and the choice of good hyperparameters can vary substantially for different datasets and types of studies. They also often include the addition of multiple loss functions to the training procedure. This leads to harder optimization tasks when searching for the true total loss function minima that will result in the best model performance. Finally, in real astrophysical applications, researchers will not always be able to work with nicely curated datasets. For example, when training models for classification tasks, we might encounter situations where two datasets contain only a portion of overlapping classes, while there might be other dataset-specific classes, not present in both datasets. Furthermore, datasets can also include unknown object classes, anomalies, or simply unlabeled portions of the data. Most standard DA methods, such as MMD, try to align entire data distributions. The presence of any kind of non-overlapping classes makes the two data distributions very different, and so the DA method cannot be applied successfully.

In this work, we develop a more universal DA method that does not require exact overlap between classes in the two datasets and can handle dataset-specific and non-overlapping classes in any of the two datasets. Furthermore, it works even in the presence of unknown classes and can even be used for anomaly detection. As with all DA methods, model training requires two datasets: the source data domain (with labeled images) and the target data domain (that can be unlabeled, as labels are not used during model training). Our aim is to develop a universal DA method that can be applied to a plethora of astronomical

[§] <https://roman.gsfc.nasa.gov>

tasks and that can successfully perform DA on both simulated and astronomical survey data.

Here we focus on galaxy morphology classification, generally into spiral, elliptical, and merging galaxies, and more granularly by leveraging sub-classes that more closely describe galaxy shapes (such as ellipticity, bulge prominence, presence of a bar in spiral galaxies, etc.). Understanding galaxy morphology is an important stepping stone for a full understanding of the formation and evolution of galaxies, mass assembly, and structure formation. As galaxies evolve and interact with each other, complex structures are formed, such as bulges, spiral arms, bars, tidal tails, and clumps. Morphology is also related to other physical properties of galaxies, such as their color, gas mass, stellar mass, and star formation rate (Kauffmann et al., 2003; Lianou et al., 2019). The problem of galaxy morphology classification has been approached in several ways: visual inspection by experts (Hubble, 1926; van den Bergh, 1960); visual inspection by volunteers and crowd sourcing like in the Galaxy Zoo project (Lintott et al., 2008); developing multiple sets of morphology parameters, such as the Sérsic index (Sérsic, 1963), or concentration, asymmetry, and clumpiness (CAS) (Conselice et al., 2003); using simple machine learning algorithms on the morphology parameters (Snyder et al., 2019); and even using more complex deep learning algorithms on the galaxy images themselves (Walmsley et al., 2020; Cavanagh et al., 2021; Cheng et al., 2021).

In this work, we develop a robust deep learning algorithm, capable of handling multi-dataset galaxy morphology problems. We used simulated LSST mock images made from IllustrisTNG (Nelson et al., 2019), as well as multiple observational datasets available in the Galaxy Zoo project (Lintott et al., 2008, 2011; Willett et al., 2013). This way, we create a multi-dataset test bed well suited for the development of a universal DA method for galaxy morphology classification across different simulated and real datasets. Additionally, our method is equally applicable to other classification, regression, and anomaly detection tasks that use multiple datasets.

In Section 2, we describe the DA method we develop and how it is implemented, and in Section 3, we describe the deep learning model we train to perform the galaxy morphology classification. In Section 4, we describe all of the datasets we use in our studies. We present the results on different types of cross-dataset tasks in Section 5, with a discussion and conclusion in Section 6. All the code used in this work can be found on our GitHub page^{||}, and we also make all of the datasets available on Zenodo[¶].

^{||} <https://github.com/deepskies/DeepAstroUDA>
[¶] <https://doi.org/10.5281/zenodo.7473597>

2. Methods

DA methods are designed to help deep learning models extract and use only those features that are present in all datasets the model is intended to handle. Internally, using DA in model training helps align latent data distributions, allowing a model to find a common decision boundary between the classes for more consistent performance on multiple datasets (Csurka, 2017; Wang and Deng, 2018; Wilson and Cook, 2020). Most distance-based DA methods (such as MMD or CORAL) and adversarial DA methods (such as DANNs) will only work efficiently under the assumption that the two latent data distributions are similar. These methods treat the entire data distribution as a whole; in other words, they are not class-aware. Therefore, correct alignment is only possible if both distributions contain the same classes of objects with similar properties. This condition is not always satisfied in scientific applications, so more flexible DA methods are needed.

We denote the source domain dataset as $\mathcal{D}_s = (X_s, Y_s)$ and the target domain dataset as $\mathcal{D}_t = (X_t, Y_t)$, where X_s and X_t are image sets from the source and target dataset, respectively, and Y_s and Y_t are the corresponding label sets. We can then divide class-aware DA methods into closed DA approaches that assume that the two datasets include the same classes, i.e. $Y_s = Y_t$ (Gretton et al., 2012; Ganin et al., 2016), and methods that assume that one of the domains contains more classes, such as open DA for $Y_s \subset Y_t$ (Busto and Gall, 2017; Fu et al., 2019; Liu et al., 2019; Yuan et al., 2022), partial DA for $Y_t \subset Y_s$ (Zhang et al., 2018; Cao et al., 2019; Xu et al., 2021), or DA for problems that are a mix of open and partial (You et al., 2019). For a graphical depiction of the different types of DA problems, see Figure 1. Since the inclusion of DA methods is needed when the target domain is unlabeled (otherwise one could simply train a regular supervised learning model on the target dataset directly), we may not know which of these situations will occur in our experiments. Hence, the development of more general methods, capable of handling all types of dataset overlaps (closed, open, partial, or mixed), is needed for full implementation of these methods in the sciences.

Here, we aim to develop such a Universal Domain Adaptation (UDA) method, capable of performing domain alignment even in the presence of extra classes (known or unknown). These additional classes can be present in both the labeled source or unlabeled target domain. The general principle is that extra classes, present in only one of the domains, should not be aligned with the other domain that does not contain them. This means that the unlabeled target domain can also contain completely unknown objects




Figure 1: Types of DA problems. The source domain is represented with a solid line ellipse, while the target domain uses a dashed line ellipse. Source domain classes are represented by filled shapes and target domain by empty shapes.

or anomalies that the method needs to recognize, group as a new class, and exclude from the domain alignment procedure.

One such method, called *Domain Adaptive Neighborhood Clustering via Entropy optimization* (DANCE), was introduced in Saito et al. (2020). Their method includes two novel ideas: 1) a neighborhood clustering technique to learn the structure of the target domain in a self-supervised way and to cluster neighboring source and target examples, and 2) using entropy to align the known or to reject unknown target classes. The authors show that DANCE can handle arbitrary domain shifts and they apply it to several benchmarking datasets, such as Office (3 domains and 31 classes, Saenko et al., 2010), OfficeHome (4 domains and 65 classes, Venkateswara et al., 2017), and VisDA (2 domains and 12 classes, Peng et al., 2017).

Furthermore, the idea of contrastive self-supervised learning has become a key component for learning semantically meaningful representations of the data in situation where no labels are available (Wu et al., 2018; Tian et al., 2019; Slijepcevic et al., 2022). Representations are learned by comparing and contrasting positive (drawn by pairing the augmentations of the same image) and negative pairs of samples (different images) in an unsupervised setting without access to any labels. However, contrastive learning in the context of domain adaptation remains underexplored. In Thota and Leontidis (2021), authors propose to extend contrastive learning approach to a situation where

none of the data domains contain any labeled data. In our work we combine ideas of contrastive learning on unlabeled samples and supervised learning using the labeled source domain data to build our UDA method for astronomy.

2.1. DeepAstroUDA Method

In our work we introduce *DeepAstroUDA*, which performs domain alignment and clustering of similar objects into classes via two loss functions: adaptive clustering and entropy separation. Clustering of target samples is performed using adaptive clustering ideas introduced in Li et al. (2021). To further improve the clustering of the unlabeled target domain, we use the entropy separation loss used in Saito et al. (2020), but we improve upon this method by adding active hyper-parameter tuning for this loss function⁺. In the following sections, we take a closer look at the different components of our *DeepAstroUDA* method (see Figure 2) and how to fine-tune their hyperparameters for the best performance on multiple cross-dataset astrophysics applications.

2.2. Adaptive Clustering Loss

To build our Adaptive Clustering (AC) loss, we follow ideas from Li et al. (2021) and Saito et al. (2020). The main idea of this type of semi-supervised clustering is to group target domain samples into clusters by computing pairwise similarities among features of unlabeled samples in the target domain. The loss minimization then forces the classifier to predict consistent class labels for samples with high pairwise feature similarities. This is achieved by training the model with a binary cross-entropy loss, where binary pairwise feature similarities are used as ground truth labels.

To facilitate semi-supervised clustering and to increase the number of similar samples we perform a data augmentation step. For each image in the source and target batches, we create two transformed versions (90° rotation and scaling by zooming to 300 pixels than cropping back to 256) of that image. Then, for any pair of unlabeled target samples x_1 and x_2 , we predict a pairwise similarity label by using the classifier output prediction vectors $\mathbf{p}_1$ and $\mathbf{p}_2$ and rank ordering their elements (Han et al., 2020). We then require the top k elements to have the same ordering ($k = 3$ for 3-class problem and $k = 7$ for 10-class problem, for computational efficiency) to decide that the paired

⁺ We have also tested using regular neighborhood clustering loss from DANCE, but we have found it is hard to fine tune and has worse performance than the adaptive clustering we describe below. We include both losses in the code available on our GitHub page, so an interested user can test both on their problem.




Figure 2: DeepAstroUDA method and the effects of different loss functions. The cross-entropy loss (filled red arrows) clusters the labeled source domain data (filled circles and triangles). The adaptive clustering loss (empty red arrows) pushes unlabeled target domain data (empty circles and triangles), towards data it shares the most similar features with (both source and target data). Finally, the entropy loss (empty violet arrows) uses entropy to push unknown classes away from the known ones. In this open DA example, the unknown class is present in the target domain (empty squares).

samples belong to the same class, which is denoted by a similarity label $s_{12} = 1$; otherwise, $s_{12} = 0$. For the labeled source domain images, we use their class labels to generate similarity labels. We also calculate a similarity score between samples as $\hat{s}_{12} = \mathbf{p}_1^\top \mathbf{p}_2$. Finally, we can write the AC loss as a binary cross-entropy loss, where similarity labels are used as ground truth labels:

$$\mathcal{L}_{AC} = - \sum_{i \in B} \sum_{j \in b_t} s_{ij} \log(\hat{s}_{ij}) + (1 - s_{ij}) \log(1 - \hat{s}_{ij}), \quad (1)$$

where B is the bank that contains samples from all previous source and target batches, and b_t is the current target batch (Saito et al., 2020). By comparing similarities between unlabeled target samples from the current target batch to all elements stored in the bank, current target samples are pushed towards the source and target samples with which they share the most similar features.

2.3. Entropy Separation Loss

Entropy Separation (ES) has an explicit objective to encourage the alignment (known classes, present in both domains) or rejection (unknown classes present

only in one of the domains) of target samples (Saito et al., 2020). This is possible because unknown samples often do not share features with known samples, which leads to larger entropy of the classifier output for unknown samples compared to entropy of the output for known classes (You et al., 2019). Therefore, the entropy can be used to decide the boundary between known and unknown samples. If we denote the mean entropy of the classifier output $\mathbf{p}_i$ of sample i from the target batch b_t as $H(\mathbf{p}_i)$, we can define a boundary ρ around the entropy value so that:

$$\mathcal{L}_{ES}(\mathbf{p}_i) = \begin{cases} -|H(\mathbf{p}_i) - \rho| & |H(\mathbf{p}_i) - \rho| > m, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Here m is a confidence threshold around the boundary ρ , which is used to decide if we are confident about whether a particular sample belongs in a known or unknown class. Only those samples that are far enough from the entropy boundary ρ will be moved towards known classes or pushed away as an unknown class. The boundary ρ and confidence threshold m start from preset values (determined from experiments), but are actively fine-tuned during training. Section 2.5 describes the active fine-tuning of parameters in more detail. Finally, the total ES loss is:

$$\mathcal{L}_{ES} = \frac{1}{|b_t|} \sum_{i \in b_t} \mathcal{L}_{ES}(\mathbf{p}_i). \quad (3)$$

2.4. Total Loss and Model Training

While the ES loss is applied only to the target domain data, the AC loss is used to cluster the target data by aligning the samples to both the source and target data via the bank that stores all previously seen samples from both source and target batches. Finally, the main classification loss, applied only to the labeled source domain data, is the standard weighted Cross-Entropy (CE) loss:

$$\mathcal{L}_{CE} = \frac{- \sum_{k=1}^K w_k y_k \log \hat{y}_k}{\sum_{k=1}^K w_k}, \quad (4)$$

where the class weight (distinct from the network weight parameters) for each class is calculated as $w_k = N_s / (K n_k)$, where n_k is the number of images in class k , K is the total number of source classes, and N_s is the total number of images in the training source dataset. The true and predicted labels are y_k and $\hat{y}_k$, respectively.

The final objective of the model training is:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda(\mathcal{L}_{AC} + \mathcal{L}_{ES}). \quad (5)$$

We utilize domain-specific batch normalization, which eliminates domain style information leakage, which can be viewed as a form of weak domain alignment. Finally, the importance of the clustering losses that perform domain adaptation is governed by the DA weight parameter λ . We tested values in the range of 0.0001 – 100 and find that $\lambda = 0.005$ is a good value that achieves the best model performance. Example of the behaviour of each of the three losses during training is shown on Figure 5 in Section 5.1

2.5. Hyperparameter Tuner

Fine-tuning network hyperparameters can be a challenging task. Most DA methods include multiple hyperparameters inside different loss functions, as well as loss weights that control the contribution of each loss term to the total loss function used in the training. Often, the final model performance is very dependent on the hyperparameters. In most cases, the hyperparameter values are hardcoded for a specific dataset and use case scenario, which is not useful when applying the same algorithm to a different dataset. Furthermore, to determine the optimal hyperparameter values, we often rely on simple random or grid search, or Bayesian optimization (Bergstra et al., 2011; Shahriari et al., 2016) over the parameter space, which can be slow or require strong computational resources. More importantly, as DA training progresses and the source and target latent data representations change, the choice of optimal hyperparameters might also change, which is not taken into account in most DA approaches.

In this work we develop and utilize a hyperparameter tuner that actively changes and fine-tunes the hyperparameters of the ES loss to help successfully cluster unlabeled target samples. It focuses on two parameters: the boundary ρ around the sample entropy value and the confidence interval m (see Section 2.3). During training, our hyperparameter tuner iteratively checks the performance of the model and changes hyperparameter values to improve it. It first changes the boundary ρ and then adjusts the confidence m . Once the loss stops improving, the tuner repeats this cycle. We found that active hyperparameter tuning of the ES loss is crucial for good performance, since both the boundary and confidence interval values will certainly change as the training progresses, the target samples become more clustered, and the distinction between classes becomes clearer.

The initial guess for the boundary ρ is calculated as $\rho = \log(K)/2$, where K is the number of known source classes. The initial value for m and the step value used by the tuner for both ρ and m are found through experimentation. We run experiments for problems between 3 and 10 classes to set the initial parameter values for the tuner. For example, we found

that for problems with less than 5 classes initial values for m are usually in the range of 0.2 – 0.8, while for more classes the range is 1.1 – 1.6. Once the training is performed on any new dataset, the tuner will use the closest values for the specific class-size problem. Since the preset initial value is not necessarily the best choice for all datasets, using active tuning is encouraged. Different initial tuner values and step sizes for both parameters to further improve the performance can be specified manually, and other hyperparameters can be added to the tuner.

2.6. Latent Space Visualization

Visualization of the latent data distribution is important to understand the model behavior, performance, and trustworthiness, and ultimately to conduct further model refinement. It is particularly useful in DA tasks, where the model is trained to align data distributions from two different data domains. Furthermore, our method not only tries to align classes present in both data domains, but also pushes away unknown samples or classes that are present in only one of the domains and therefore should not be used in the domain alignment. This makes latent space visualizations even more important.

To visualize the latent data distributions, we utilize manifold learning methods because of their capability to recognize non-linear data patterns. We use a combination of the isomap (Tenenbaum et al., 2000) and the t-distributed stochastic neighbor embedding (tSNE; van der Maaten and Hinton, 2008).

An isomap is a lower-dimensional embedding of the latent space, such that geodesic distances in the original higher-dimensional space are also respected in the lower-dimensional space. It is constructed in three stages. First, a weighted neighborhood graph G over all data points is constructed. Then, the edge weight values of the graph G are assigned the distances between neighboring points and the geodesic distances between all pairs of points are estimated as their shortest-path distances in the graph G . Finally, the lower d -dimensional (often $d = 2$ or $d = 3$ for visualization purposes) embedding that best preserves the manifold’s estimated intrinsic geometry is produced by applying classical Metric Multidimensional Scaling (MDS; Borg and Groenen, 2005) to the matrix of the shortest-graph distances. We use the `scikit` implementation of isomaps (Pedregosa et al., 2011). Isomaps are a great tool for realistic latent data representation, but they do not always show the cleanest and most visually distinct clumps for different classes, which can make interpretation harder.

The t-SNE method calculates the probability distribution over data point pairs in the latent space of the model. It assigns a higher probability to

similar objects and a lower probability to dissimilar pairs. It then constructs a 2D or 3D representation of the data, in which data pairs share the same probabilities. By minimizing the Kullback–Leibler (KL) divergence (Kullback and Leibler, 1951) between the two distributions, the t-SNE method ensures the similarity between the actual distribution and the low-dimensional projection. The t-SNE algorithm on its own can be slow, but it produces well separated clumps for different classes. Furthermore, since it adapts to data by performing different transformations in different regions, it is difficult to compare the relative sizes of clusters in t-SNE plots. In addition, the final appearance of the t-SNE plots is highly dependent on several user-defined parameters Wattenberg et al. (2016).

In this work, we combine isomaps and t-SNE plots and utilize the best properties of both methods. First, we create a 2D isomap representation of the latent space of the model, ensuring that the intrinsic geometry of the original high-dimensional space is preserved. We then use the isomap as an input to the t-SNE algorithm, to create a more visually distinct classes that are easier to interpret. Since t-SNE in this case does not need to further reduce the dimensionality of the space, we do not expect dramatic changes in the positions of the class clumps. In the following text, we refer to our visualizations as t-SNE plots, but we emphasize here that the presented plots actually show t-SNE plots of the isomaps of the latent space of the model.

3. Neural Network Model and Training

In all of our experiments, we use the *ResNet50* (He et al., 2016) network (with random initialization of model weights). We train it with early stopping, which monitors the classification accuracy and stops the training when there is no improvement after 12 consecutive epochs. The model is trained using stochastic gradient descent with Nesterov momentum (Sutskever et al., 2013) and an initial learning rate of 0.001. The learning rate is tuned using an inverse learning rate scheduler, whereby the learning rate is decayed by a factor of 0.1 every set number of epochs, determined by the specific initial dataset configuration (every 10 epochs in all of our 10-class experiments and every 7 epochs for our 3-class experiment). We train our models on 4 NVIDIA RTX A6000 GPUs (available from Google Colab and LambdaLabs), and on average, the training converges in ≈ 5 hours (dependent on the dataset size and the complexity of the experiment).

4. Data

4.1. Simulations - LSST Mock Data

IllustrisTNG (Marinacci et al., 2018; Naiman et al., 2018; Springel et al., 2018; Naiman et al., 2018; Pillepich et al., 2018; Nelson et al., 2019) is a state-of-the-art cosmological magneto-hydrodynamical simulation that includes gas, stars, dark matter, supermassive black holes, and magnetic fields. It uses a galaxy formation model built on the cosmological simulation code AREPO (Springel, 2010), which solves the coupled equations of ideal magneto-hydrodynamics and self-gravity. IllustrisTNG builds on the successes of the older simulation, Illustris-1 (Vogelsberger et al., 2014), but extends the mass range of the simulated galaxies and halos, uses improved numerical and astrophysical modeling, and addresses some of the identified shortcomings and tensions with observations (see Nelson et al. (2015)).

The Illustris project is ideally suited for studying galaxy formation and morphology, merging galaxies, gas accretion, formation of galaxy clusters, and large-scale structures. It is also a great resource for the creation of labeled mock datasets that can be made to mimic different astronomical surveys and used to train different kinds of supervised learning algorithms.

We use simulated LSST mocks from Ćiprijanović et al. (2022), which are made from the IllustrisTNG100 (Nelson et al., 2019) simulation. Mock images include three filters (g, r, i), made from the two simulation snapshots: 95 (redshift $z = 0.05$) and 99 (redshift $z = 0$). All images were converted to an effective redshift of $z = 0.05$, to create a larger single-redshift dataset. This dataset includes three galaxy morphology classes—spiral, elliptical and merging galaxies—made by following Lotz et al. (2004) and Snyder et al. (2015), who use the $G-M_{20}$ “bulge statistic” to distinguish between galaxy morphology classes. It also includes augmentation of the least numerous class (merging galaxies) to create the final balanced dataset of $\approx 35,000$ galaxy images (with dimensions 100×100 pixels), which is the approach we follow in this work.

The creation of LSST mock observations from IllustrisTNG100 was done using the *GalSim* package (Rowe et al., 2015). We use both datasets created in Ćiprijanović et al. (2022)—a high-noise one-year survey (“Y1”) and a low-noise ten-year survey (“Y10”)—made by applying an exposure time corresponding to one year or ten years of observations directly to the raw images (552s per year for the r and i filters and 240s for the g filter). The images also include both atmospheric and optical point spread function (PSF) blurring, mimicking LSST observations. Finally, the images include arcsinh stretching to make fainter ob-

jects more apparent while preserving the original color ratios in each pixel. For more details on the mock dataset creation, see Ćiprijanović et al. (2022).

4.2. Observations

The Galaxy Zoo project (GZ; Lintott et al., 2008) was the first to provide morphological classifications of nearly one million galaxies from the Sloan Digital Sky Survey (SDSS; York et al., 2000), performed by $\approx 10^5$ volunteers that classified images using a web-based interface. While the first project used a simplified classification into spiral, elliptical and merging galaxies, its successor, Galaxy Zoo 2 (GZ2; Willett et al., 2013), used a more complex classification system that considered the presence of bars and bulges, the shapes of edge-on disks, and quantification of the relative strengths of galactic bulges and spiral arms. It provided more than 300,000 reliable morphological classifications of galaxies in the full Data Release 7 (DR7) of SDSS as well as the deeper Stripe 82. Following the success of these volunteer-based projects, the Galaxy Zoo expanded to include several other classification-based projects that use data from other telescopes and astronomical surveys*. The newer projects include Galaxy Zoo 3: Hubble (Willett et al., 2017); Galaxy Zoo 4: CANDELS (Simmons et al., 2017), DECaLS (Walmsley et al., 2022), UKIDSS (Galloway, 2017), Ferengi (Galloway, 2017), GAMA-KiDS (Holwerda et al., 2019); and Galaxy Zoo Illustris (Dickinson et al., 2018), using simulated data.

In this work, we apply our model to several GZ datasets to test its performance on different observational cross-dataset scenarios and use cases.

4.2.1. Galaxy Zoo: SDSS and DECaLS

We use two GZ datasets (Lintott et al., 2008; Willett et al., 2013)‡: the source domain dataset is from GZ2 SDSS (Lintott et al., 2008; Lintott et al., 2010; Willett et al., 2013), and the target domain dataset is from GZ3 DECaLS (Walmsley et al., 2022), which uses DR7 of the the DECam Legacy Survey (part of the Dark Energy Spectroscopic Instrument (DESI) Legacy Imaging Surveys, Dey et al., 2019).

Specifically, for the target domain, we use an $\approx 18k$ subset of the DECaLS dataset, with images that passed more rigorous vote filtering and better class separation, named Galaxy10 DECaLS††. We perform a 10-class experiment, in which we use 9 classes from

this dataset (disturbed, merging, round smooth, cigar-shaped smooth, barred spiral, unbarred tight spiral, unbarred loose spiral, edge-on without bulge, edge-on with bulge). We add one more class from the full GZ3 DECaLS dataset, gravitationally lensed galaxies, which we will treat as an unknown class present in the target domain. We chose gravitational lenses as our unknown class because they are rare and difficult to find, and creating automated AI methods to search for these objects in new observations is crucial for inferring cosmological parameters.

We use the same labels for our source domain GZ2 SDSS dataset. For each class, we use galaxy IDs to find these objects in the GZ2 SDSS data and extract a similar number of examples, as was done in the DECaLS data. Finally, in both domains all classes have between 1–2.6 thousand images, except a much smaller cigar-shaped smooth class, with 334 images, which we purposefully do not augment, to test the performance of the model in the presence of a class imbalance. Finally, the source and target datasets each contain $\approx 20,000$ images (GZ2 has 21784, and DECaLS 19840).

Both the SDSS and DECaLS data include three filter images (i , r , g). The SDSS DR7 data has a pixel scale of $0.396''$, while the DECaLS DR7 data has a scale of $0.262''$. Furthermore, the SDSS data includes galaxies with Petrosian half-light magnitude in the r -band $m_r < 17.0$ (after the galactic extinction correction was applied). Deeper DECaLS images (with the lowest r -band magnitude of $m_r = 23.6$, versus $m_r = 22.2$ from SDSS) reveal spiral arms, weak bars, and tidal features not previously visible in SDSS imaging. The GZ3 DECaLS dataset was derived from SDSS DR8 imaging (Aihara et al., 2011), so it only includes galaxies that are within both the DECaLS and SDSS DR8 footprint and have a slightly fainter magnitude limit of $m_r < 17.77$.

4.2.2. Galaxy Zoo: SDSS Wide and Deep Field

Additionally, we use DeepAstroUDA on a problem that includes different types of observational data from the same astronomical survey. The source domain includes wide-field SDSS data from GZ2, and the target domain comes from the deeper Stripe 82 region of SDSS, also available in GZ2 (Lintott et al., 2008; Lintott et al., 2010; Willett et al., 2013). Stripe 82 is a multiply-imaged section along the celestial equator in the Southern Galactic Cap, which results in better visibility of fainter objects.

For the source domain data, we use the same 9-class wide-field dataset described in Section 4.2.1. For our target data from Stripe 82, we use images from the co-added depth (set 2) †. These images

* For decision trees used in different GZ projects, see https://data.galaxyzoo.org/gz_trees/gz_trees.html

‡ Current publicly available GZ datasets can be found at <https://data.galaxyzoo.org>

†† Galaxy10 data is available at <https://astronn.readthedocs.io/en/latest/galaxy10.html>

† Stripe 82 data is available at <https://data.galaxyzoo.org>

are made from combining between 47–55 individual exposures, resulting in better detection of fainter features and improved seeing, compared to the source domain data. This deeper dataset includes objects with magnitude limit $m_r < 17.77$. The images also include a modest color desaturation to deemphasize background noise in the co-added data. In order to make both datasets as unique as possible, we allow a maximum of 5% of extracted object overlap between the two datasets (objects found in a region covered by both wide and deep field observations). We create the target domain dataset by extracting $\approx 2,000$ images for each of the known 9 classes and the unknown lens class, balancing the classes in the target dataset with a final size of $\approx 20,000$ images. These images did not pass any curation, so unlike the source domain, they also include examples that are harder to distinguish between classes. We chose to do this to observe how *DeepAstroUDA* performs in more realistic scenarios, where new data did not yet pass any human-level filtering. The fact that this dataset includes harder to distinguish examples will also likely lead to more confusion in the crowd-sourced true labels, with possibly more incorrect true labels.

4.3. Data Pre-Processing

For all of our observational data, we use SExtractor (Bertin and Arnouts, 1996) to determine the center and radius of objects in the downloaded images, and then crop images to 256×256 pixels, using the extracted object properties to ensure no pertinent parts of the image have been inappropriately cut off.

All datasets (simulated and observational) are normalized to have pixel values in $[0, 1]$. Finally, the datasets are divided into training, validation, and test sets in proportions 60% : 20% : 20%.

5. Results

In this section, we present the results of the *DeepAstroUDA* model on three different types of DA problems that include observational data: 1) DA between different data releases of the same survey (LSST); 2) DA between two different astronomical surveys (SDSS and DECaLS); and 3) DA between two observing fields of different depths within the same survey (wide and deep fields of SDSS). In all three examples, we focus not only on the proper domain alignment and overlap of known classes, but also on adding an unknown class to all target domain datasets. This makes all of our examples harder open DA problems and also allows us to test anomaly detection (by discovering and clustering the unknown class samples). We show that our model is good at both aligning the known classes and discovering and

clustering the unknown anomaly class, in all types of observational cross-dataset problems.

5.1. Domain Adaptation Within a Survey: Different Data Releases of LSST

We first apply the *DeepAstroUDA* model to the simplest example: two data releases from the same astronomical survey. In this case, the difference between the domains arises due to the different noise levels in the data (with subsequent data releases having lower noise levels due to the longer observation time). For this test, we used simulated mock LSST data (one and ten years of observations), described in detail in Section 4.1. The dataset includes three galaxy morphology classes: spiral, elliptical, and merging galaxies. As was shown already in Ćiprijanović et al. (2022), without DA, a model trained on one of these datasets cannot be used on the other one. In that paper, the model was trained on the low-noise Y10 data and applied to the noisier Y1 data.

In contrast to Ćiprijanović et al. (2022), here we focus mainly on a more realistic scenario, in which the model is trained on the first noisier data release (Y1), with the intention of later being applied to all other lower-noise data releases (for example Y10). In Figure 3, we show randomly selected example images from the source Y1 and target Y10 test data sets, denoted as High→Low training. We additionally train another model on the reverse case (Low→High), where the Y10 data is the source domain, as in Ćiprijanović et al. (2022). This facilitates a closer comparison to that paper. In this work, we consider a harder open DA problem in which spiral and elliptical galaxies are known classes present in both domains and merging galaxies are an additional unknown class present only in the target domain (while in Ćiprijanović et al. (2022), all three classes are present in both domains). In order to show how much the inclusion of DA helps with the performance in the target domain, we will compare the results of two models: 1) regular training on the source domain without any DA, and 2) training on the source domain with the inclusion of DA for the unlabeled target domain data.

In Table 1, we show common performance metrics on the source and target test sets (accuracy, precision, recall, and F1 score, averaged over all classes present in each of the domains), with regular training, i.e., training with just $\mathcal{L}_{CE}$ on the source domain (top row), or training with DA (bottom row). All models in this paper were retrained 5 times, with different random seeds used to initialize the weights, and we report the values of all performance metrics averaged across these runs.

We can see that in both cases training without DA makes the model not work at all on the target




Figure 3: Example images from the source domain high-noise Y1 (left) and target domain low-noise Y10 (right) test set. The source domain contains two classes: spiral and elliptical galaxies, while the target domain also contain an unknown merging galaxies class.

domain data. Using DA, on the other hand, increases the accuracy in the target domain up to 74%. Adding DA leads to a substantial increase in the mean accuracy even in the source domain (by 9 – 12%), which ultimately reaches similarly high levels as the target domain accuracy (up to 76%). Furthermore, even though we are working on a harder problem, where mergers are an unknown target domain class, employing *DeepAstroUDA* leads to 8% better target domain accuracies compared to the results in Čiprijanović et al. (2022). In addition, our methods produce more balanced performance across the source and target domains, considering accuracy, precision, and recall. By allowing the model to use more robust features, DA acts as a regularizer, making overfitting harder. Finally, by including active tuning of the DA loss parameters, the model can reach higher accuracies in fewer epochs (this example was impacted the most, with training stopping 15 epoch earlier in the

Table 1: Mean performance metrics for *ResNet50* when trained with regular training (top row) and with DA (bottom row). The top table shows results when low-noise data is used as the source domain and high-noise data as the target domain, while the bottom table shows results for the reversed case. The inclusion of DA increases the accuracy and other metrics for both source and target data, in both experiments.

Low → High			
Training	Metric	Source	Target
Reg.	Accuracy	0.64 ± 0.22	0.36 ± 0.11
	Precision	0.59 ± 0.34	0.35 ± 0.31
	Recall	0.65 ± 0.31	0.41 ± 0.29
	F1 Score	0.62 ± 0.23	0.38 ± 0.22
DA	Accuracy	0.76 ± 0.12	0.74 ± 0.09
	Precision	0.74 ± 0.13	0.71 ± 0.19
	Recall	0.75 ± 0.08	0.74 ± 0.21
	F1 Score	0.74 ± 0.08	0.72 ± 0.32

High → Low			
Training	Metric	Source	Target
Reg.	Accuracy	0.61 ± 0.32	0.33 ± 0.48
	Precision	0.65 ± 0.34	0.31 ± 0.46
	Recall	0.54 ± 0.28	0.33 ± 0.39
	F1 Score	0.59 ± 0.22	0.32 ± 0.31
DA	Accuracy	0.70 ± 0.20	0.73 ± 0.18
	Precision	0.64 ± 0.30	0.72 ± 0.17
	Recall	0.66 ± 0.24	0.74 ± 0.20
	F1 Score	0.65 ± 0.19	0.73 ± 0.13

Low→High case).

To better illustrate the effects of DA training, in Figure 4 we show how the accuracies for all three classes change during the more realistic High→ Low noise training case. Vertical dashed lines mark epochs where our hyperparameter tuner changes the ES loss parameters to try to improve clustering. We also show how all three loss functions change during training in Figure 5.

Finally, to better understand how the inclusion of DA influences the way the model represents the data in its latent space, we show the t-SNE plot (of the isomap of the latent space representation of the data) in Figure 6. We can see that without DA (top), the domains do not overlap very well. In this example, the domains are quite similar (both datasets include the same galaxy images, and the only difference is the noise levels), so the model trained without DA is able to place some target-domain images correctly, but the unknown merger class is completely overlapping with the known classes. The inclusion of DA (bottom) makes the known classes overlap much better, while the unknown target class is pushed away from the known data.




Figure 4: The source domain Y1 (top) and target domain Y10 (bottom) accuracies during model training. Elliptical (violet) and spiral galaxies (navy) are present in both domains, while merging galaxies (yellow) represent the unknown anomaly class, present only in the target domain. Vertical gray dashed lines show epochs in which the tuner changes the ES loss hyperparameters, which results in improved clustering and better performance on the anomaly class.

5.2. Domain adaptation across surveys: SDSS to DECaLS

As mentioned above, DA in astronomy has only been tried between simulated mock data mimicking a particular telescope and real observations from the same telescope Čiprijanović et al. (2021). The success of the methods used so far strongly depended on the minimization of the dataset shift, by making simulated mock data as similar to real data as possible, before any model training was performed. The difference between data from two different astronomical surveys can, on the other hand, be much larger: from different noise levels, PSF blurring, pixel scale, survey depth, etc. Here, we present the first successful cross-survey DA result in astronomy. We will focus on a




Figure 5: The loss function values during the High $\rightarrow$ Low training: CE loss (top, yellow), AC loss (middle, green), and ES loss (bottom, dark green).

more complicated 10-class galaxy morphology problem that includes data from two different astronomical surveys: SDSS and DECaLS. These more complex dataset shifts require more flexible DA methods. If two data distributions look very different, or do not share all of the same classes, simpler methods that explicitly minimize a distance metric between the latent data distributions and are not class-aware might not be powerful enough to properly align the two data distributions. For that reason, in this work we develop a much more flexible semi-supervised method that is able to align similar samples for any kind of data distribution shapes and possible overlaps.

In Table 2, we again report the mean accuracy, precision, recall, and F1 score for the source and target test sets: regular training (top row), and training with DA (bottom row). Again, the use of DA improves performance in both data domains, from a 5% increase in accuracy in the source domain to a 36% increase in accuracy in the target domain. Furthermore, the performance and behavior of the model (reflected in all performance metrics) is similar across the domains after the inclusion of DA. Also, the model performance reaches levels similar to our first, much simpler example in Section 5.1, even though this problem has a much larger distribution shift and more classes.

To further illustrate the performance of the model on each of the classes, in Figure 7 we show the individual target domain class accuracies and how they change during training. For better readability, we separate classes into lower performing (top figure) and higher performing (middle figure) classes. Additionally, we present the confusion matrix (bottom figure), to further illustrate the confusion between similar classes when the final trained model




Figure 6: The t-SNE plots when training with Y1 data as the source domain and Y10 data as the target domain (High $\rightarrow$ Low). The top plot shows the latent space of the model trained without DA, while the bottom plot shows the model trained with DA. Using DA aligns known classes much better, but also allows the unknown class to be pushed to the outskirts.

is applied to our target data test set of images. We can clearly see that the model confuses merging and disturbed galaxies, which can both exhibit asymmetric disturbed morphology and are both associated with galaxy interaction processes. Also, barred spiral and unbarred tight spiral are often incorrectly classified as the class edge-on without bulge, which also makes sense, given that all three classes actually include spiral galaxies with different orientations. Similar conclusions can be made for unbarred tight spiral and unbarred loose spiral, which are often confused with the class edge-on with bulge. We have also noticed some occurrence of errors in which a spiral galaxy

Table 2: Mean performance metrics for *ResNet50* on SDSS (source) and DECaLS (target) test data for regular training without DA (top row) and training with DA (bottom row). The inclusion of DA increases accuracy and other metrics for both the source and target data.

SDSS $\rightarrow$ DECaLS			
Training	Metric	Source	Target
Reg.	Accuracy	0.77 ± 0.20	0.43 ± 0.32
	Precision	0.52 ± 0.23	0.42 ± 0.43
	Recall	0.70 ± 0.25	0.33 ± 0.30
	F1 Score	0.60 ± 0.18	0.37 ± 0.25
DA	Accuracy	0.82 ± 0.09	0.79 ± 0.10
	Precision	0.79 ± 0.06	0.73 ± 0.21
	Recall	0.84 ± 0.13	0.81 ± 0.18
	F1 Score	0.81 ± 0.07	0.77 ± 0.14

that is very small with hard-to-distinguish features gets classified as one of the smooth classes.

Figure 8 shows t-SNE plots from the model with regular training and training with DA. For better visibility, we choose to plot only three out of ten classes: round smooth, barred spiral, and the unknown gravitational lens class. In this more complicated and larger domain shift problem, the model trained using regular training on the source domain data tends to flip classes from the two domain, overlapping the barred spiral class from the source domain with the round smooth class from the target and vice versa. The unknown lens class is also completely overlapping with known classes. The inclusion of DA during training corrects this behavior and correctly aligns the known classes, while pushing the unknown class to the outskirts.

In Figure 9, we show example images from the source SDSS (top) and target DECaLS (bottom) test data sets. For each of the domains, the top two rows show examples of images most often correctly classified, while the bottom two rows show examples of most often incorrectly classified (considering the 5 models trained with different random seed initializations). All images show the true class in the top left corner and predicted class in the top right corner, for one of our trained models. In Figure 10 (left), we also plot examples of correctly classified lenses in DECaLS data. Each image also includes the number of models (out of 5) that were able to discover the lens.

In Figure 11 (left plot), we show how the accuracies of the target domain change during training. We present the mean accuracy of the known 9 classes and the unknown class (gravitational lens) accuracy separately. We can see that the known classes reach quite high accuracy, around 80%, which is




Figure 7: Individual class accuracies in the target domain during training. The top figure shows lower performing classes, while the middle figure shows higher performing classes. Finally, the bottom figure shows the confusion matrix of the final trained model, when applied to the target domain test set (disturbed (0), merging (1), round smooth (2), cigar-shaped smooth (3), barred spiral (4), unbarred tight spiral (5), unbarred loose spiral (6), edge-on without bulge (7), edge-on with bulge (8), lenses (9)). The performance on all classes is very good, with slight confusion between several classes that indeed look visually similar, for example disturbed and merging galaxies.




Figure 8: The t-SNE plots when training on 10-class SDSS data (source) and testing on DECaLS data (target). We choose to plot 3 classes for better visibility: round smooth (class 2), barred spiral (class 4), and lens (class 9) anomaly class in the target domain. Again, training with DA (bottom plot) helps correctly overlap the known classes and push away the anomaly class.

very good given the fact that this dataset includes several morphologically or visually similar classes (for example, disturbed and merging galaxies, or cigar-shaped smooth and edge-on without bulge). Furthermore, the model is even better at finding and classifying the unknown anomaly class, with the lens class reaching an accuracy of almost 90%, proving that this DA method can be used successfully not only to bridge the gap between different observational datasets, but also to search for unknown objects of interest, such as gravitational lenses.

When examining images from our datasets, we noticed examples in which true labels from crowd sourcing do not seem to be correct. In some of these cases, we find the predicted labels to be more appropriate, which makes us believe that the model had enough correctly labeled samples to properly understand the galaxy morphology and correctly classify even those examples that include questionable true labels. For example, in Figure 9 the last image in the bottom row has a true lens label, while the network predicts class edge-on with bulge, which upon visual inspection seem like a more appropriate class. More evidence for this conclusion can be seen in the confusion matrix, where the biggest confusion between classes occurs for truly morphologically similar classes.

5.3. Domain adaptation within a survey: wide and deep fields of SDSS

Finally, DA can help bridge the gap between data from wide and deep fields within the same survey. Here, we focus on the wide field (source data) and Stripe 82 deep field (target data) of SDSS. We use the same 10-class problem as in the previous cross-survey example. As we describe in Section 4.2.2, we use the same source dataset as in the cross-survey example, and we make the target domain using randomly chosen examples from Stripe 82, which also contains crowd-sourced GZ2 labels. Compared to the previous cross-survey example, where we use curated datasets that contain cleaner examples, the target domain Stripe 82 data in this test is not curated and can include galaxies much harder to correctly classify into one of the ten classes. This makes classification using AI harder, but will also make human labeling more difficult, probably containing an even larger percentage of incorrect true labels.

In Table 3, we again present the performance metrics for the models trained with and without DA, when applied to the source and target test datasets. As before, training without DA makes the model useless on target data, whereas the inclusion of DA allows the model to perform well on both wide and deep fields of the same survey, with similar performance metrics across domains. The mean accuracy with DA

Table 3: Performance metrics for *ResNet50* on SDSS wide field (source) and SDSS Stripe 82 deep field (target) test data for regular training without DA (top row) and training with DA (bottom row). The inclusion of DA increases the accuracy for both source and target data.

		Wide → Deep	
Training	Metric	Source	Target
Reg.	Accuracy	0.76 ± 0.22	0.43 ± 0.37
	Precision	0.68 ± 0.12	0.42 ± 0.22
	Recall	0.76 ± 0.21	0.44 ± 0.32
	F1 Score	0.72 ± 0.12	0.43 ± 0.19
DA	Accuracy	0.84 ± 0.08	0.81 ± 0.12
	Precision	0.83 ± 0.15	0.80 ± 0.07
	Recall	0.79 ± 0.12	0.81 ± 0.07
	F1 Score	0.81 ± 0.10	0.80 ± 0.05

reaches 81% in the target domain and 84% in the source domain, which is even slightly higher than in our previous cross-survey example.

We omit plotting the confusion matrix for this example, since it looks very similar to the cross-survey example shown in Figure 7. Again, the biggest confusions occur between disturbed and merging galaxies (classes 0 and 1) and between different subclasses of spiral galaxies (classes 4, 5 and 6 incorrectly classified as 7 or 8). The most notable difference, in this example, is the presence of the somewhat unexpected confusion between the class edge-on with bulge (class 8) and the unknown gravitational lens class (class 9). This probably leads to lower performance on the unknown class, which we can see in Figure 11 (right), showing the mean accuracy for all known classes and the accuracy of the unknown class in the target domain during model training. As in the cross-survey example, the mean known class accuracy reaches levels around 80%, but the unknown class accuracy is in this case somewhat lower. After examining lens images from the SDSS deep field, we concluded that the lenses indeed look less prominent compared to the DECaLS data, which we believe has led to both more incorrect true labels from crowd-sourcing, and more confusion in the predictions by the *DeepAstroUDA* model.

In Figure 12, we again show examples most often correctly classified and incorrectly classified in both domains, the wide field (top) and the target Stripe 82 deep field of SDSS data (bottom), as well as their true and predicted classes. In Figure 10 (right), we show examples of discovered lenses, and give the number of times the models correctly classified each of the lenses (out of 5 models trained with different random seed initializations).




Figure 9: Example images from source SDSS data (top) and target DECaLS data (bottom). Each image contains the true label in the top left corner and the predicted label in the top right corner. For each domain, we give most often correctly classified examples in the top two rows and most often incorrectly classified examples in the bottom two rows. Classes are: disturbed (0), merging (1), round smooth (2), cigar-shaped smooth (3), barred spiral (4), unbarred tight spiral (5), unbarred loose spiral (6), edge-on without bulge (7), edge-on with bulge (8), lenses (9).




Figure 10: Randomly selected example images of correctly classified lenses from DECaLS data (left) and SDSS Stripe 82 deep field (right). Each image contains the number of models (out of 5, trained with different random seed initializations) that correctly classified the example.




Figure 11: The target domain mean known (blue) and unknown class (red) accuracies during training. The mean accuracy for all classes is plotted with a dashed black line. The left plot shows results for the cross-survey problem (SDSS to DECaLS), where the target domain includes well curated examples, with clearer class distinctions. The right plot shows DA from wide to deep field of SDSS, where the target domain also contains harder examples. This leads to a lower increase in performance on the unknown lens class, which often gets confused with the edge-on with bulge class. Still, the increase in performance during model training is evident and the final performance on the unknown class is much better than for the model trained without any DA (since that model does not work at all on both known classes and the unknown class in the target domain).




Figure 12: Example images from the source SDSS wide field (top) and the target SDSS Stripe 82 deep field (bottom). Each image contains the true label in the top left corner and the predicted label in the top right corner. For each domain, we give most often correctly classified examples in the top row and most often incorrectly classified examples in the bottom row. Classes are again: disturbed (0), merging (1), round smooth (2), cigar-shaped smooth (3), barred spiral (4), unbarred tight spiral (5), unbarred loose spiral (6), edge-on without bulge (7), edge-on with bulge (8), lenses (9).

6. Discussion and Conclusion

DeepAstroUDA is a flexible semi-supervised DA algorithm that can handle any kind of cross-dataset problem. It can be used for classification, regression, and anomaly detection, and it works well even in the presence of non-overlapping (or unknown) classes in any of the two data domains. This means that it can be used for open, partial, and mixes of open and partial DA problems. This is a crucial feature for successful implementation of DA algorithms in real

scientific applications, where we often will not have well curated datasets.

The inclusion of DA in astronomy and cosmology is a necessity, as most types of studies include the use of multiple simulated and observational datasets. In this work, we address open DA problems related to galaxy morphology classification. We focus on observational data and demonstrate the use of the method in three different scenarios: 1) data collected after different numbers of observing years, i.e., different data releases of the same survey; 2) data from two

different surveys; and 3) data from wide and deep fields of the same survey. These three scenarios will have different combinations of factors that make the datasets different: pixel scale, noise levels, PSF blurring, the depth of the survey, and the magnitude limit of the objects that can be observed. All the results presented in this paper show that *DeepAstroUDA*:

- (i) can successfully be used on difficult domain shift problems that include multiple observational survey datasets;
- (ii) allows the trained model to perform well even on the unlabeled target domain data, with an observed increase in accuracy up to 40%;
- (iii) handles any type of domain overlap and performs even in the presence of unknown classes, which can be used for anomaly detection tasks, like searching for merging galaxies or gravitational lenses;
- (iv) increases accuracy on both source and target data, and makes performance of the model consistent across both domains.

In this work, we use Galaxy Zoo labels, made by thousands of volunteers. This process naturally makes true labels prone to errors since it is often very hard to visually distinguish between different morphological classes. Still, most true labels should be correct, since all of the models we train seemed to be able to correctly learn the distinction between galaxy classes. For example, upon closer visual inspection of examples from our test set, we noticed images for which we actually agree more with the (incorrect) predicted class given by the model than with the true crowd-sourced label. Furthermore, the biggest confusion between classes occurs for truly hard to distinguish objects, like disturbed and merging galaxies, where both classes contain very asymmetric objects, with tidal tails and other interesting features, which are a product of galaxy interactions.

An important feature of *DeepAstroUDA* is its ability to be used as an anomaly detection algorithm in the unlabeled target domain. We focus here on merging and gravitationally lensed galaxies. Merging galaxies are crucial for understanding the process of galaxy evolution, formation of different galaxy morphologies, and star formation. Unfortunately, they are very difficult to find in real photometric observations. Without spectroscopic observations, we often cannot be sure if a pair of galaxies is truly merging or if it is just visually overlapping. With future surveys like LSST, where we expect that over 60% of objects will be overlapping, discovering true mergers will be even more difficult. Methods such as *DeepAstroUDA* will be a necessity to allow us to use the information available in other simulated and

real datasets to search for merging galaxies in new observations.

Strong gravitational lensing of galaxies (as well as supernovae and quasars) is an important cosmological probe. Future surveys like LSST will discover unprecedented numbers of lensed objects, and efficient discovery and creation of complete catalogs will allow us to put tighter constraints on cosmological parameters and better understand both dark energy and dark matter. Regular AI models trained on simulated samples of strong gravitation lenses or on old observations will not be able to perform well on new observations, but the inclusion of DA will open doors for much easier discovery of new strong gravitational lenses.

Finally, with the increase in size and complexity of next-generation astronomical datasets, true comparison and understanding of these datasets will become increasingly difficult. The power of AI algorithms is in their ability to work and extract information from very complex multidimensional datasets. Our future work will focus on further development and refinement of DA algorithms. With their ability to extract similarities and differences between datasets, they can potentially be used to help us improve our simulations, and even understand the underlining physics.

Acknowledgments

This manuscript has been supported by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy (DOE), Office of Science, Office of High Energy Physics. This research has been partially supported by the High Velocity Artificial Intelligence grant as part of the DOE High Energy Physics Computational HEP program. This research has been partially supported by the DOE Office of Science, Office of Advanced Scientific Computing Research, applied mathematics and SciDAC programs under Contract No. DE-AC02-06CH11357. This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is a user facility supported by the DOE Office of Science.

The authors of this paper have committed themselves to performing this work in an equitable, inclusive, and just environment, and we hold ourselves accountable, believing that the best science is contingent on a good research environment. We acknowledge the Deep Skies Lab as a community of multi-domain experts and collaborators who have facilitated an environment of open discussion, idea-generation, and collaboration. This community was important for the development of this project.

Author Contributions

A. Ćiprijanović: *Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Resources, Supervision, Visualization, Writing of original draft*; A. Lewis: *Formal analysis, Investigation, Methodology, Resources, Software, Visualization, Writing of original draft*; K. Pedro: *Conceptualization, Methodology, Project administration, Resources, Supervision, Writing (review & editing)*; S. Madireddy: *Conceptualization, Methodology, Resources, Supervision, Writing (review & editing)*; B. Nord: *Conceptualization, Methodology, Supervision, Writing (review & editing)*; G. N. Perdue: *Conceptualization, Methodology, Project administration, Resources, Supervision, Writing (review & editing)*; S. M. Wild: *Conceptualization, Methodology, Resources, Supervision, Writing (review & editing)*.

References

- Aihara, H., Arimoto, N., Armstrong, R., Arnouts, S., Bahcall, N.A., Bickerton, S., Bosch, J., Bundy, K., et al., 2018. The Hyper Suprime-Cam SSP Survey: Overview and survey design. *PASJ* 70, S4. doi:10.1093/pasj/psx066, arXiv:1704.05858.
- Aihara, H., Prieto, C.A., An, D., Anderson, S.F., Aubourg, É., Balbinot, E., Beers, T.C., Berlind, A.A., Bickerton, S.J., Bizyaev, D., et al., 2011. ERRATUM: “THE EIGHTH DATA RELEASE OF THE SLOAN DIGITAL SKY SURVEY: FIRST DATA FROM SDSS-III” (2011, *ApJS*, 193, 29). *The Astrophysical Journal Supplement Series* 195, 26. URL: <https://doi.org/10.1088/0067-0049/195/2/26>, doi:10.1088/0067-0049/195/2/26.
- Bergstra, J., Bardenet, R., Bengio, Y., Kégl, B., 2011. Algorithms for hyper-parameter optimization, in: Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc. URL: <https://proceedings.neurips.cc/paper/2011/file/86e8f7ab32cfd12577bc2619bc635690-Paper.pdf>.
- Bertin, E., Arnouts, S., 1996. SExtractor: Software for source extraction. *A&AS* 117, 393–404. doi:10.1051/aas:1996164.
- Borg, I., Groenen, P., 2005. *Modern Multidimensional Scaling: Theory and Applications* (Springer Series in Statistics). doi:10.1007/978-1-4757-2711-1.
- Busto, P.P., Gall, J., 2017. Open set domain adaptation, in: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 754–763. doi:10.1109/ICCV.2017.88.
- Cao, Z., You, K., Long, M., Wang, J., Yang, Q., 2019. Learning to Transfer Examples for Partial Domain Adaptation. arXiv e-prints, arXiv:1903.12230 arXiv:1903.12230.
- Cavanagh, M.K., Bekki, K., Groves, B.A., 2021. Morphological classification of galaxies with deep learning: comparing 3-way and 4-way CNNs. *MNRAS* 506, 659–676. doi:10.1093/mnras/stab1552, arXiv:2106.01571.
- Cheng, T.Y., Conselice, C.J., Aragón-Salamanca, A., Agüena, M., Allam, S., Andrade-Oliveira, F., Annis, J., Bluck, A.F.L., Brooks, D., Burke, D.L., et al., 2021. Galaxy morphological classification catalogue of the Dark Energy Survey Year 3 data with convolutional neural networks. *MNRAS* 507, 4425–4444. doi:10.1093/mnras/stab2142, arXiv:2107.10210.
- Ćiprijanović, A., Kafkes, D., Downey, K., Jenkins, S., Perdue, G.N., Madireddy, S., Johnston, T., Snyder, G.F., Nord, B., 2021. DeepMerge - II. Building robust deep learning algorithms for merging galaxy identification across domains. *MNRAS* 506, 677–691. doi:10.1093/mnras/stab1677, arXiv:2103.01373.
- Ćiprijanović, A., Kafkes, D., Snyder, G., Sánchez, F.J., Perdue, G.N., Pedro, K., Nord, B., Madireddy, S., Wild, S.M., 2022. DeepAdversaries: examining the robustness of deep learning models for galaxy morphology classification. *Machine Learning: Science and Technology* 3, 035007. doi:10.1088/2632-2153/ac7f1a, arXiv:2112.14299.
- Conselice, C.J., Bershad, M.A., Dickinson, M., Papovich, C., 2003. A Direct Measurement of Major Galaxy Mergers at z less than ~ 3 . *Astron. J.* 126, 1183–1207. doi:10.1086/377318, arXiv:astro-ph/0306106.
- Csurka, G., 2017. A comprehensive survey on domain adaptation for visual applications, in: *Domain Adaptation in Computer Vision Applications*. Springer International Publishing, Cham, pp. 1–35. doi:10.1007/978-3-319-58347-1_1.
- Dark Energy Survey Collaboration, Abbott, T., Abdalla, F.B., Aleksić, J., Allam, S., Amara, A., Bacon, D., Balbinot, E., Banerji, M., et al., 2016. The Dark Energy Survey: more than dark energy - an overview. *MNRAS* 460, 1270–1299. doi:10.1093/MNRAS/stw641, arXiv:1601.00329.
- Dey, A., Schlegel, D.J., Lang, D., Blum, R., Burleigh, K., Fan, X., Findlay, J.R., Finkbeiner, D., Herrera, D., Juneau, S., et al., 2019. Overview of the DESI Legacy Imaging Surveys. *Astron. J.* 157, 168. doi:10.3847/1538-3881/ab089d, arXiv:1804.08657.
- Dickinson, H., Fortson, L., Lintott, C., Scarlata, C., Willett, K., Bamford, S., Beck, M., Cardamone, C., Galloway, M., Simmons, B., Keel, W., Kruk, S., Masters, K., Vogelsberger, M., Torrey, P.,

- Snyder, G.F., 2018. Galaxy Zoo: Morphological Classification of Galaxy Images from the Illustris Simulation. *Astrophys. J.* 853, 194. doi:10.3847/1538-4357/aaa250, arXiv:1801.08541.
- Dodge, S., Karam, L., 2016. Understanding How Image Quality Affects Deep Neural Networks. arXiv e-prints arXiv:1604.04004.
- Dodge, S., Karam, L., 2017. A study and comparison of human and deep learning recognition performance under visual distortions. arXiv e-prints arXiv:1705.02498.
- Ford, N., Gilmer, J., Carlini, N., Cubuk, D., 2019. Adversarial examples are a natural consequence of test error in noise. arXiv e-prints arXiv:1901.10513.
- Fu, J., Wu, X., Zhang, S., Yan, J., 2019. Improved open set domain adaptation with backpropagation, in: 2019 IEEE International Conference on Image Processing (ICIP), pp. 2506–2510. doi:10.1109/ICIP.2019.8803287.
- Galloway, M.A., 2017. Morphology Is a Link to the Past: Examining Formative and Secular Galactic Evolution through Morphology. Ph.D. thesis. University of Minnesota, Twin Cities.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., Lempitsky, V., 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research* 17, 1–35. URL: <http://jmlr.org/papers/v17/15-239.html>.
- Gide, M.S., Dodge, S.F., Karam, L.J., 2016. The effect of distortions on the prediction of visual attention. arXiv e-prints arXiv:1604.03882.
- Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., Smola, A., 2007. A kernel method for the two-sample-problem, in: *Advances in Neural Information Processing Systems 19*, MIT Press. pp. 513–520.
- Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A., 2012. A kernel two-sample test. *Journal of Machine Learning Research* 13, 723–773. URL: <http://jmlr.org/papers/v13/gretton12a.html>.
- Han, K., Rebuffi, S.A., Ehrhardt, S., Vedaldi, A., Zisserman, A., 2020. Automatically Discovering and Learning New Visual Categories with Ranking Statistics. arXiv e-prints , arXiv:2002.05714 arXiv:2002.05714.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778. doi:10.1109/CVPR.2016.90.
- Holwerda, B.W., Kelvin, L., Baldry, I., Lintott, C., Alpaslan, M., Pimblet, K.A., Liske, J., Kitching, T., Bamford, S., de Jong, J., et al., 2019. The Frequency of Dust Lanes in Edge-on Spiral Galaxies Identified by Galaxy Zoo in KiDS Imaging of GAMA Targets. *Astron. J.* 158, 103. doi:10.3847/1538-3881/ab2886, arXiv:1909.07461.
- Hubble, E.P., 1926. Extragalactic nebulae. *Astrophys. J.* 64, 321–369. doi:10.1086/143018.
- Huertas-Company, M., Lanusse, F., 2022. The DAWES review 10: The impact of deep learning for the analysis of galaxy surveys. arXiv e-prints , arXiv:2210.01813 arXiv:2210.01813.
- Ivezić, Ž., Kahn, S.M., Tyson, J.A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., et al., 2019. LSST: From science drivers to reference design and anticipated data products. *Astrophys. J.* 873, 111. doi:10.3847/1538-4357/ab042c, arXiv:0805.2366.
- Kauffmann, G., Heckman, T.M., White, S.D.M., Charlot, S., Tremonti, C., Peng, E.W., Seibert, M., Brinkmann, J., Nichol, R.C., SubbaRao, M., York, D., 2003. The dependence of star formation history and internal structure on stellar mass for 10^5 low-redshift galaxies. *MNRAS* 341, 54–69. doi:10.1046/j.1365-8711.2003.06292.x, arXiv:astro-ph/0205070.
- Kullback, S., Leibler, R.A., 1951. On information and sufficiency. *Ann. Math. Statist.* 22, 79–86. URL: <https://doi.org/10.1214/aoms/1177729694>, doi:10.1214/aoms/1177729694.
- Li, J., Li, G., Shi, Y., Yu, Y., 2021. Cross-domain adaptive clustering for semi-supervised domain adaptation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2505–2514.
- Lianou, S., Barmby, P., Mosenkov, A.A., Lehnert, M., Karczewski, O., 2019. Dust properties and star formation of approximately a thousand local galaxies. *Astron. Astrophys.* 631, A38. doi:10.1051/0004-6361/201834553, arXiv:1906.02712.
- Lintott, C., Schawinski, K., Bamford, S., Slosar, A., Land, K., Thomas, D., Edmondson, E., Masters, K., Nichol, R.C., Raddick, M.J., Szalay, A., Andreescu, D., Murray, P., Vandenberg, J., 2010. Galaxy Zoo 1: data release of morphological classifications for nearly 900 000 galaxies*. *MNRAS* 410, 166–178. URL: <https://doi.org/10.1111/j.1365-2966.2010.17432.x>, doi:10.1111/j.1365-2966.2010.17432.x, arXiv:<https://academic.oup.com/mnras/article-pdf/410/1/166/18442057/mnras0410-0166.pdf>.
- Lintott, C., Schawinski, K., Bamford, S., Slosar, A., Land, K., Thomas, D., Edmondson, E., Masters, K., Nichol, R.C., Raddick, M.J., Szalay, A., Andreescu,

- D., Murray, P., Vandenberg, J., 2011. Galaxy Zoo 1: data release of morphological classifications for nearly 900 000 galaxies. *MNRAS* 410, 166–178. doi:10.1111/j.1365-2966.2010.17432.x, arXiv:1007.3265.
- Lintott, C.J., Schawinski, K., Slosar, A., Land, K., Bamford, S., Thomas, D., Raddick, M.J., Nichol, R.C., Szalay, A., Andreescu, D., Murray, P., Vandenberg, J., 2008. Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey. *MNRAS* 389, 1179–1189. doi:10.1111/j.1365-2966.2008.13689.x, arXiv:0804.4483.
- Liu, H., Cao, Z., Long, M., Wang, J., Yang, Q., 2019. Separate to adapt: Open set domain adaptation via progressive separation, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2922–2931. doi:10.1109/CVPR.2019.00304.
- Long, M., Cao, Z., Wang, J., Jordan, M.I., 2017. Conditional adversarial domain adaptation. arXiv e-prints arXiv:1705.10667.
- Lotz, J.M., Primack, J., Madau, P., 2004. A new nonparametric approach to galaxy morphological classification. *Astron. J.* 128, 163–182. doi:10.1086/421849, arXiv:astro-ph/0311352.
- Marinacci, F., Vogelsberger, M., Pakmor, R., Torrey, P., Springel, V., Hernquist, L., Nelson, D., Weinberger, R., et al., 2018. First results from the IllustrisTNG simulations: radio haloes and magnetic fields. *MNRAS* 480, 5113–5139. doi:10.1093/MNRAS/sty2206, arXiv:1707.03396.
- Naiman, J.P., Pillepich, A., Springel, V., Ramirez-Ruiz, E., Torrey, P., Vogelsberger, M., Pakmor, R., Nelson, D., et al., 2018. First results from the IllustrisTNG simulations: a tale of two elements - chemical evolution of magnesium and europium. *MNRAS* 477, 1206–1224. doi:10.1093/MNRAS/sty618, arXiv:1707.03401.
- Nelson, D., Pillepich, A., Genel, S., Vogelsberger, M., Springel, V., Torrey, P., Rodriguez-Gomez, V., Sijacki, D., Snyder, G.F., Griffen, B., Marinacci, F., Blecha, L., Sales, L., Xu, D., Hernquist, L., 2015. The illustris simulation: Public data release. *Astronomy and Computing* 13, 12–37. doi:10.1016/j.ascom.2015.09.003, arXiv:1504.00362.
- Nelson, D., Springel, V., Pillepich, A., Rodriguez-Gomez, V., Torrey, P., Genel, S., Vogelsberger, M., Pakmor, R., Marinacci, F., Weinberger, R., Kelley, L., Lovell, M., Diemer, B., Hernquist, L., 2019. The IllustrisTNG simulations: public data release. *Computational Astrophysics and Cosmology* 6, 2. doi:10.1186/s40668-019-0028-x, arXiv:1812.05609.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., et al., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., Saenko, K., 2017. VisDA: The Visual Domain Adaptation Challenge. arXiv e-prints, arXiv:1710.06924 arXiv:1710.06924.
- Pillepich, A., Nelson, D., Hernquist, L., Springel, V., Pakmor, R., Torrey, P., Weinberger, R., Genel, S., et al., 2018. First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies. *MNRAS* 475, 648–675. doi:10.1093/MNRAS/stx3112, arXiv:1707.03406.
- Rowe, B.T.P., Jarvis, M., Mandelbaum, R., Bernstein, G.M., Bosch, J., Simet, M., Meyers, J.E., Kacprzak, T., et al., 2015. GALSIM: The modular galaxy image simulation toolkit. *Astronomy and Computing* 10, 121–150. doi:10.1016/j.ascom.2015.02.002, arXiv:1407.7676.
- Saenko, K., Kulis, B., Fritz, M., Darrell, T., 2010. Adapting visual category models to new domains, in: Proceedings of the 11th European Conference on Computer Vision: Part IV, Springer-Verlag, Berlin, Heidelberg. p. 213–226.
- Saito, K., Kim, D., Sclaroff, S., Saenko, K., 2020. Universal Domain Adaptation through Self Supervision. arXiv e-prints, arXiv:2002.07953 arXiv:2002.07953.
- Sérsic, J.L., 1963. Photometry of southern galaxies IX: NGC 1313. *Boletín de la Asociacion Argentina de Astronomia La Plata Argentina* 6, 99–99.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., de Freitas, N., 2016. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE* 104, 148–175. doi:10.1109/JPROC.2015.2494218.
- Simmons, B.D., Lintott, C., Willett, K.W., Masters, K.L., Kartaltepe, J.S., Häußler, B., Kaviraj, S., Krawczyk, C., Kruk, S.J., McIntosh, D.H., et al., 2017. Galaxy Zoo: quantitative visual morphological classifications for 48 000 galaxies from CANDELS. *MNRAS* 464, 4420–4447. doi:10.1093/mnras/stw2587, arXiv:1610.03070.
- Slijepcevic, I.V., Scaife, A.M.M., Walmsley, M., Bowles, M., 2022. Learning useful representations for radio astronomy “in the wild” with contrastive learning. arXiv e-prints, arXiv:2207.08666doi:10.48550/arXiv.2207.08666, arXiv:2207.08666.
- Snyder, G.F., Rodriguez-Gomez, V., Lotz, J.M., Torrey, P., Quirk, A.C.N., Hernquist, L., Vogelsberger, M., Freeman, P.E., 2019. Automated dis-

- tant galaxy merger classifications from Space Telescope images using the Illustris simulation. *MNRAS* 486, 3702–3720. doi:10.1093/mnras/stz1059, arXiv:1809.02136.
- Snyder, G.F., Torrey, P., Lotz, J.M., Genel, S., McBride, C.K., Vogelsberger, M., Pillepich, A., Nelson, D., et al., 2015. Galaxy morphology and star formation in the Illustris Simulation at $z = 0$. *MNRAS* 454, 1886–1908. doi:10.1093/MNRAS/stv2078.
- Springel, V., 2010. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *MNRAS* 401, 791–851. doi:10.1111/j.1365-2966.2009.15715.x, arXiv:0901.4107.
- Springel, V., Pakmor, R., Pillepich, A., Weinberger, R., Nelson, D., Hernquist, L., Vogelsberger, M., Genel, S., et al., 2018. First results from the IllustrisTNG simulations: matter and galaxy clustering. *MNRAS* 475, 676–698. doi:10.1093/MNRAS/stx3304, arXiv:1707.03397.
- Sun, B., Saenko, K., 2016. Deep CORAL: Correlation alignment for deep domain adaptation, in: *ECCV Workshops*.
- Sutskever, I., Martens, J., Dahl, G., Hinton, G., 2013. On the importance of initialization and momentum in deep learning, in: Dasgupta, S., McAllester, D. (Eds.), *Proceedings of the 30th International Conference on Machine Learning*, PMLR, Atlanta, Georgia, USA. pp. 1139–1147. URL: <https://proceedings.mlr.press/v28/sutskever13.html>.
- Tenenbaum, J.B., de Silva, V., Langford, J.C., 2000. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science* 290, 2319–2323. doi:10.1126/science.290.5500.2319.
- Thota, M., Leontidis, G., 2021. Contrastive Domain Adaptation. arXiv e-prints, arXiv:2103.15566doi:10.48550/arXiv.2103.15566, arXiv:2103.15566.
- Tian, Y., Krishnan, D., Isola, P., 2019. Contrastive Multiview Coding. arXiv e-prints, arXiv:1906.05849doi:10.48550/arXiv.1906.05849, arXiv:1906.05849.
- van den Bergh, S., 1960. A Preliminary Luminosity Classification for Galaxies of Type Sb. *Astrophys. J.* 131, 558. doi:10.1086/146869.
- van der Maaten, L., Hinton, G., 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9, 2579–2605. URL: <http://jmlr.org/papers/v9/vandermaaten08a.html>.
- Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S., 2017. Deep hashing network for unsupervised domain adaptation. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 5385–5394.
- Vilalta, R., Dhar Gupta, K., Boubner, D., Meskhi, M.M., 2019. A General Approach to Domain Adaptation with Applications in Astronomy. *PASP* 131, 108008. doi:10.1088/1538-3873/aaf1fc, arXiv:1812.08839.
- Vogelsberger, M., Genel, S., Springel, V., Torrey, P., Sijacki, D., Xu, D., Snyder, G., Nelson, D., Hernquist, L., 2014. Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe. *MNRAS* 444, 1518–1547. doi:10.1093/mnras/stu1536, arXiv:1405.2921.
- Walmsley, M., Lintott, C., Geron, T., Kruk, S., Krawczyk, C., Willett, K.W., Bamford, S., Kelvin, L.S., Fortson, L., Gal, Y., Keel, W., Masters, K.L., Mehta, V., Simmons, B.D., Smethurst, R., Smith, L., Baeten, E.M., Macmillan, C., 2022. Galaxy Zoo DECaLS: Detailed visual morphology measurements from volunteers and deep learning for 314 000 galaxies. *MNRAS* 509, 3966–3988. doi:10.1093/mnras/stab2093, arXiv:2102.08414.
- Walmsley, M., Smith, L., Lintott, C., Gal, Y., Bamford, S., Dickinson, H., Fortson, L., Kruk, S., Masters, K., Scarlata, C., Simmons, B., Smethurst, R., Wright, D., 2020. Galaxy Zoo: probabilistic morphology through Bayesian CNNs and active learning. *MNRAS* 491, 1554–1574. doi:10.1093/mnras/stz2816, arXiv:1905.07424.
- Wang, M., Deng, W., 2018. Deep visual domain adaptation: A survey. *Neurocomputing* 312, 135–153. doi:10.1016/j.neucom.2018.05.083.
- Wattenberg, M., Viégas, F., Johnson, I., 2016. How to use t-sne effectively. Distill URL: <http://distill.pub/2016/misread-tsne>, doi:10.23915/distill.00002.
- Willett, K.W., Galloway, M.A., Bamford, S.P., Lintott, C.J., Masters, K.L., Scarlata, C., Simmons, B.D., Beck, M., Cardamone, C.N., Cheung, E., et al., 2017. Galaxy Zoo: morphological classifications for 120 000 galaxies in HST legacy imaging. *MNRAS* 464, 4176–4203. doi:10.1093/mnras/stw2568, arXiv:1610.03068.
- Willett, K.W., Lintott, C.J., Bamford, S.P., Masters, K.L., Simmons, B.D., Casteels, K.R.V., Edmondson, E.M., Fortson, L.F., Kaviraj, S., Keel, W.C., et al., 2013. Galaxy Zoo 2: detailed morphological classifications for 304 122 galaxies from the Sloan Digital Sky Survey. *MNRAS* 435, 2835–2860. doi:10.1093/mnras/stt1458, arXiv:1308.3496.
- Wilson, G., Cook, D.J., 2020. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology* 11. doi:10.1145/3400066.

- Wu, Z., Xiong, Y., Yu, S., Lin, D., 2018. Unsupervised Feature Learning via Non-Parametric Instance-level Discrimination. arXiv e-prints , arXiv:1805.01978doi:10.48550/arXiv.1805.01978, arXiv:1805.01978.
- Xu, Y., Yang, J., Cao, H., Li, Q., Mao, K., Chen, Z., 2021. Partial video domain adaptation with partial adversarial temporal attentive network. 2021 IEEE/CVF International Conference on Computer Vision (ICCV) , 9312–9321.
- York, D.G., Adelman, J., Anderson, John E., J., Anderson, S.F., Annis, J., Bahcall, N.A., Bakken, J.A., Barkhouser, R., Bastian, S., Berman, E., et al., SDSS Collaboration, 2000. The Sloan Digital Sky Survey: Technical Summary. Astron. J. 120, 1579–1587. doi:10.1086/301513, arXiv:astro-ph/0006396.
- You, K., Long, M., Cao, Z., Wang, J., Jordan, M.I., 2019. Universal domain adaptation, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2715–2724. doi:10.1109/CVPR.2019.00283.
- Yuan, Y., He, X., Jiang, Z., 2022. Adaptive open domain recognition by coarse-to-fine prototype-based network. Pattern Recognition 128, 108657. doi:10.1016/j.patcog.2022.108657.
- Zellinger, W., Moser, B.A., Grubinger, T., Lughofer, E., Natschläger, T., Saminger-Platz, S., 2019. Robust unsupervised domain adaptation for neural networks via moment alignment. Information Sciences 483, 174–191. doi:10.1016/j.ins.2019.01.025.
- Zhang, J., Ding, Z., Li, W., Ogunbona, P., 2018. Importance weighted adversarial nets for partial domain adaptation. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition , 8156–8164.