

Domain Adaptation and Active Learning for Fine-Grained Recognition in the Field of Biodiversity

Bernd Gruner^{1,2}, Matthias Körschens¹, Björn Barz¹ and Joachim Denzler^{1,2}

¹Computer Vision Group, Friedrich Schiller University Jena, Jena, Germany

{matthias.koerschens,bjoern.barz,joachim.denzler}@uni-jena.de

²Institute of Data Science, German Aerospace Center (DLR), Jena, Germany

bernd.gruner@dlr.de

Abstract

Deep-learning methods offer unsurpassed recognition performance in a wide range of domains, including fine-grained recognition tasks. However, in most problem areas there are insufficient annotated training samples. Therefore, the topic of transfer learning respectively domain adaptation is particularly important. In this work, we investigate to what extent unsupervised domain adaptation can be used for fine-grained recognition in a biodiversity context to learn a real-world classifier based on idealized training data, e.g. preserved butterflies and plants. Moreover, we investigate the influence of different normalization layers, such as Group Normalization in combination with Weight Standardization, on the classifier. We discovered that domain adaptation works very well for fine-grained recognition and that the normalization methods have a great influence on the results. Using domain adaptation and Transferable Normalization, the accuracy of the classifier could be increased by up to 12.35 % compared to the baseline. Furthermore, the domain adaptation system is combined with an active learning component to improve the results. We compare different active learning strategies with each other. Surprisingly, we found that more sophisticated strategies provide better results than the random selection baseline for only one of the two datasets. In this case, the distance and diversity strategy performed best. Finally, we present a problem analysis of the datasets.

1. Introduction

Deep learning methods are versatile and offer outstanding recognition performance in a wide range of tasks. One example for that is the visual recognition of classes that differ only in subtle optical details, which is also known as fine-grained classification. In various competitions,

like Fine-Grained Visual Categorization ¹, these methods achieved the best results. This work deals with fine-grained image classification of butterflies and plants in the wild, which we also refer to as the target domain. The problem in this area is that we do not have annotated data for the target domain, but labeled preserved butterflies and plants, which is our source domain. Unfortunately, the classifier learned on the source domain does not give sufficient results on the data of the target domain. The data from the two domains are usually acquired under different conditions and are only a snapshot of the reality. The resulting difference is called domain shift and limits the recognition performance on the target domain.

In section 2 we give an overview of several methods to bridge the domain shift. The focus of this work is on the Unsupervised Adversarial Domain Adaptation [10] framework, as it still provides state-of-the-art performance. It reduces the domain shift by including the unlabeled data into the training process. The aim of our work is to investigate how well this method works in the area of fine-grained classification.

Furthermore, normalization layers have proven to be an essential part of the network architecture, as they allow the training to converge faster and reduce the sensitivity to the initialization of parameters and hyperparameters [23]. However, the popular method Batch Normalization [15] has a disadvantage in the area of domain adaptation, which affects the results negatively [30]. Therefore, we test and compare Batch Normalization against the alternatives, which are Domain Agnostic Normalization [30], Transferable Normalization [45] and Group Normalization [47] in combination with Weight Standardization [28]. The latter is already successfully applied in transfer learning [22], but to the best of our knowledge not yet tested in a domain adaptation scenario.

¹<https://sites.google.com/view/fgvc7/home>

To further improve the results of the classifier, we added an active learning component, which helps us to select target samples for annotation in a structured way. In our experimental setup, we only simulate the annotation process as we already have all the labels for the target dataset in order to enable an evaluation. This can be regarded in contrast to the real application scenario where the annotation is carried out by a domain expert and is much more time-consuming.

We investigate how many target samples have to be annotated randomly in order to achieve a result comparable to fully supervised learning on the target domain. Further, this baseline experiment is expanded by comparing it to various active learning strategies: Expected Model Output Change [20], Distance and Diversity [4], Importance Weight [38] and Certainty sampling. To the best of our knowledge, Expected Model Output Change has not yet been used in the field of domain adaptation.

2. Related Work

There are various ways to deal with a domain shift, but it is particularly difficult when there are no or only a few annotated data points from the target domain. Therefore, the simplest solution is to collect more labeled target data, which is in most cases very expensive and time-consuming. Other possibilities are the use of improved methods for feature extraction, further data augmentation [42, 6], Mix Match [3, 33], semi-supervised methods [36, 27], domain adaptation [1, 50], etc. We will focus on the latter in this work.

2.1. Domain Adaptation

The domain adaptation approach tries to adjust the domains so that the learned classifier can generalize over different domains. This can be done either in the pixel space by aligning the data points before the actual processing or in the representation space by aligning the representations during feature extraction. In the following, we present a selection of approaches based on different principles and rest on the surveys [46] and [44].

Divergence-based approaches try to minimize some divergence criterion between the distributions of the two datasets. For this reason, it leads to a domain invariant feature representation and the classifier’s ability to generalize well over the domains. Criteria for measuring the distance for example are Maximum Mean Discrepancy (MMD) [32], Contrastive Domain Discrepancy (CCD) [39], Correlation Alignment (CORAL) [21] or the Wasserstein metric [8].

Other approaches are adversarial-based, which refers to any method that utilizes an adversary or an adversarial process during training. Some of the solutions here are based on Generative Adversarial Networks (GANs) [13], as Coupled Generative Adversarial Networks (CoGAN) [26]. This

approach generates synthetic target data for training. A further procedure is presented in [49], which uses a source-target converter network architecture. The Domain Adversarial Neural Network (DANN) [10], which is built on the Unsupervised Adversarial Domain Adaptation (UADA) framework, uses a discriminator for the domain classification and tries to optimize the feature extractor by a confusion loss so that the resulting representation is domain invariant.

In the reconstruction-based approaches, shared representations for the domains are created by an auxiliary reconstruction task. This is applied in the Deep Reconstruction Classification Network (DRCN) [11], which classifies the source data and reconstructs the unlabeled target data. The network learns to transform source data into data that is similar to the target dataset, while creating a common representation. Another possibility to use domain adaptation is in Cycle GANs [50]. They can learn to convert data from one domain to another, which is achieved by training two GANs simultaneously. Furthermore, Conditional GANs [16] can be applied to transfer from one domain to another by using a simple encoder-decoder architecture.

For the domain adaptation in this work the UADA approach is used, as it provides still state-of-the-art performance and can be used for unsupervised and semi-supervised domain adaptation.

2.2. Normalization

Normalization layers have become an important part of the network architecture because of their advantages during the training process (see subsection 3.2). However, Batch Normalization [15] brings some problems by introducing dependencies between the different domains when using multi-domain batches [30]. Some further developments are addressing this weakness. Therefore, Group Normalization [47] in combination with Weight Standardization [28], no longer uses batch statistics for normalization. The Domain Agnostic Normalization [30] tries to counteract the weakness by using only the source domain’s statistics. The Domain-Specific Batch Normalization [5] solves the dependencies by using the statistics of the respective domain for normalization. Another possibility is to use Adaptive Batch Normalization [25], which modulates the statistics from the source domain to the target domain. Furthermore, there is the Transferable Normalization [45], which includes a transferability factor besides the independent statistics.

Group Normalization in combination with Weight Standardization is the only method mentioned that has not been specially developed for domain adaptation and is not evaluated in this field. We test and compare Group Normalization in combination with Weight Standardization, Transferable Normalization and Domain Agnostic Normalization, because they all outperform Batch Normalization in their

application, but they are not compared with each other in their publication.

2.3. Active Learning

Active learning is a structured approach to select samples from unlabeled data using a selection strategy. Afterward, these samples are annotated. There are different strategies, for example criteria based on uncertainty [48, 24], reducing expected error [31, 51], diversity [17] or maximizing expected label changes [9, 19]. Further combinations of these criteria have been developed such as Importance Weight [38] or Distance and Diversity [4].

In previous research, the combination of domain adaptation and active learning methods was investigated. This is accomplished in Active Learning Domain Adapted (ALDA) [29, 35] for sentiment and landmine classification and in Active Adversarial Domain Adaptation (AADA) [38] for vision tasks like object detection.

In this work, we also combine domain adaptation with active learning but with a focus on fine-grained classification in the field of biodiversity. Hence, we test and evaluate different selection strategies, including the Expected Model Output Change strategy, which has not been evaluated in this field before.

3. Methods

In the following, we will briefly review the domain adaptation, normalization, and active learning methods, whose feasibility for a fine-grained recognition scenario we are going to evaluate. The core of the system is the framework Unsupervised Adversarial Domain Adaptation [10], which we present in the first section of this chapter. Then we explain the different normalization methods and finally describe the active learning approach by which the system has been extended.

3.1. Unsupervised Adversarial Domain Adaptation

The Unsupervised Adversarial Domain Adaptation (UADA) framework deals with the scenario where no labels are available in the target domain. It is based on Ben-David’s theory about domain adaptation [2, 1], which addresses the problem that the source and target data come from two different datasets, thus from two different distributions. Hence, the classification error on the target data may not be close to that on the source data. A solution developed in the paper of Ben-David et al. [1] is to align the representations of the source and target domain so that an algorithm can no longer distinguish between them. In the UADA framework a Convolutional Neural Network (CNN) is used as a feature extractor G_F , which in our case is a ResNet-50 [14] pretrained on ImageNet [34]. These features are used in the label predictor G_L for classification. This classical structure is extended by a so-called domain

discriminator G_D , which uses the features to predict the domain of the representation. For the end-to-end training, a cross-entropy classification loss is used to train the feature extractor and the label predictor. Furthermore, the domain discriminator is trained with a binary cross-entropy loss. Additionally, a confusion loss L_{conf} is used in the feature extractor, which is defined in [41] as follows:

$$L_{conf} = - \sum_{x \in T} \log(1 - G_D(G_F(x))), \quad (1)$$

where T stands for all target data points and $G_D(G_F(x))$ for the output of the domain discriminator when the feature representation of x is passed in. It should result in the representations being more aligned so that the domain discriminator can no longer distinguish between the domains.

3.2. Normalization Methods

In deep neural networks with many hidden layers, the use of normalization layers is standard. It has the positive effect that the training converges faster and reduces the sensitivity to the initialization of parameters and hyperparameters [7, 23]. A major disadvantage concerning domain adaptation is that the popular Batch Normalization (BN) [15] is based on statistics that depend on the complete batch. Since it is common for domain adaptation approaches that a batch contains data from a source and target domain, the normalization of an exclusive representation also depends on the representations of the other domain. This is a problem concerning the unlabeled target domain because these representations get no supervision from a prediction only from other images via the calculated statistics, which results in havoc on the target representation [30]. Some further developments address and resolve these weaknesses. Domain Agnostic Normalization (DAN) [30] is a very recent technique and has achieved excellent results in image segmentation. It has outperformed Batch Normalization in the domain adaptation task in the mentioned publication by 8.6 % accuracy. Therefore, this method is investigated for our fine-grained recognition task. It overcomes the problem of Batch Normalization by resolving the dependency between the different domains using only the statistics of the source domain for normalization.

A similar approach is followed by Transferable Normalization (TransNorm) [45], which is also a promising new approach. It eliminates the dependency between the domains and weights the channels according to transferability. In the paper it also improved the accuracy by up to 10 % compared to Batch Normalization depending on the dataset used.

Group Normalization [47] in combination with Weight Standardization (GN + WS) [28] is another alternative used, which, unlike the other two, was not specifically designed for domain adaptation and was not applied in this area be-

fore. It performs a batch-independent normalization of a group of channels. From the papers, we know that it has outperformed Batch Normalization by 0.79 % accuracy in the ImageNet classification task. Furthermore, this good performance is confirmed for transfer learning in [22].

3.3. Active Learning

In the case of unsupervised domain adaptation, we have a labeled source and an unlabeled target dataset. It is impractical to label the complete target dataset for a small project as it would be too time-consuming and costly. However, it is possible to annotate a part of it, and the active learning approach helps us to select informative samples. In our experimental setup, we have the labels for the target domain to be able to evaluate the results of the system, in contrast to a real-world application.

Batch-wise active learning with a batch size of k is used, because otherwise, the network would need to be trained completely from scratch again for every newly annotated example, which would be too resource-intensive. For the selection of a suitable sample for annotation, a criterion is needed. There are many criteria, and they all try to find the most informative data point based on various principles. Random selection is used as a baseline in our study. The advantage of this method is that it consumes very few resources and is independent of the size of the unlabeled dataset.

A simple and resource-saving method is Certainty (Cer) sampling. It uses entropy to select the sample where the network is most certain. The aim is to avoid errors in the prediction in case of high scores.

A more sophisticated strategy is Distance and Diversity (DivDis) [4]. It is a combination of the distance to the decision boundary, which is a kind of uncertainty measure of the classifier, and the diversity. Therefore, we use the cosine similarity. This method, first developed in 2003, is a combination of criteria and has achieved excellent results in the field of Support Vector Machines in the mentioned paper. The approach is then transferred to Deep Neural Network-based classifiers.

Importance Weight (IWERM) [38] combines the importance of an example with the uncertainty. This recent strategy was developed for the adversarial domain adaptation approach and uses the discriminator to calculate the importance.

The Expected Model Output Change (EMOC) [20] is a more sophisticated strategy and uses a higher amount of computational resources. It tries to find the sample that causes the greatest possible change in the network output after the retraining. EMOC has not yet been used in the field of domain adaptation.



Figure 1. These are two samples of the butterfly *Algaia Urticae* - *Small Tortoiseshell*. In the left column a specimen from the source domain can be seen and on the right side, there is the corresponding image from the target domain.

4. Datasets

In this section, we introduce our self-created datasets used in this study. They consist of already existing datasets, which were enriched with samples from the Global Biodiversity Information Facility (GBIF)². We downloaded images from the GBIF database and manually rechecked them and sorted out images without the expected object. The two datasets are created for the fine-grained recognition of plants and butterflies and each of them has two sub-datasets. The first subset used is the source dataset which contains images from the source domain, the second subset is the target dataset which contains images from the target domain. Both subsets consist of the same classes hence there is no category shift. These two datasets are split into 72 % training, 8 % validation, and 20 % test data. The following two subsections explain the specifics of each dataset.

4.1. Butterfly Dataset

The butterfly dataset includes 15 different species of butterflies. Each of these classes contains between 150 to 440 images. While the target domain consists of observations of wild butterflies, the source domain is preserved specimens of butterflies. Some examples are illustrated in figure 1. In the source domain, the butterflies are always photographed from above at almost the same angle. Furthermore, they are completely visible in the pictures and the conservation process has not changed the colors. In the target domain, the butterflies are photographed from different angles and distances. The background varies considerably in the target domain compared to the source domain. In summary, a certain shift between the domains can be observed.

The dataset is a combination of the UCL Butterflies [18] and Leeds Butterflies [43] datasets enriched with samples from the GBIF database.

4.2. Plant Dataset

The plant dataset includes 11 different plant species with 37 images per class. This dataset is based on the Herbarium Challenge dataset [40] with 680 species and over 46000

²<https://www.gbif.org>



Figure 2. These are two samples of the plant *Conostegia xalapensis* (Bonpl.) D.Don ex DC. In the left column a specimen from the source domain can be seen and on the right side, there is the corresponding image from the target domain.

images, but they are not uniformly distributed. There are only a few classes with over 40 images. The target subset is downloaded from the GBIF database and contains only a few examples for each plant. We have selected the overlapping classes, which have at least 37 images so that we have a balanced source and target dataset.

In figure 2, some examples are shown. Preserved specimens are the constituents of the source domain. The classes of plants range from flowers to bushes and trees. Hence, the sizes are also completely different. Often, only the leaves can be seen in the images. The conservation process changes the shape of the flowers and their color is lost. In the target domain, which shows the plants in the wild, key parts of the plants, such as flowers, are photographed and in some cases, even the whole plant can be seen on the image. Usually, the leaves of the plant are not in focus. It can be summarized that the domains are very different and it is difficult to find common features.

5. Experimental Results and Evaluation

In the beginning, we introduce the baseline experiments that are used to rank the performance of the domain adaptation system. This is followed by the presentation of the domain adaptation experiments with different normalization methods. We perform a hyperparameter optimization of learning rate and L2-regularization using grid search. The results of the grid search for all setups can be found in the supplementary materials A.1 and A.2. In general, we ran all experiments three times for 100 epochs with a batch size of 32 and report the averaged mean accuracy per class. In preprocessing, random flip and scale augmentation are used as data augmentation [14] and the images are scaled to an input size of 224 x 224.

In a further experiment, the combination of active learning and domain adaptation is examined. Therefore, a baseline for the active learning approach is created in the first

Evaluation / Training	Source Domain	Target Domain
Source Domain	100 %	89.04 %
Target Domain	67.52 %	94.59 %

Table 1. In this table, the accuracy of the classifier (source-only, baseline) trained on the source butterfly dataset and the classifier (target-only, upper bound) trained on the target butterfly dataset are shown.

Evaluation / Training	Source Domain	Target Domain
Source Domain	84.13 %	18.88 %
Target Domain	15.04 %	81.23 %

Table 2. In this table, the accuracy of the classifier (source-only, baseline) trained on the source plant dataset and the classifier (target-only, upper bound) trained on the target plant dataset are shown.

part of the experiment. The second part investigates and compares different active learning strategies. Finally, we analyze existing problems with the datasets.

5.1. Baseline: CNN-Classifier

In order to determine whether the unsupervised domain adaptation approach works in the field of fine-grained recognition and how well it performs, we create a baseline. This is done by using a CNN-classifier in two different setups and thus determining an upper bound and a baseline. Therefore, only the feature extractor and the label predictor are used from the domain adaptation approach, which corresponds to a CNN-classifier architecture. For a baseline (source-only), it is trained with the annotated source data and evaluated on the target test data. The upper bound (target-only) is created by training and testing on the annotated target data.

Table 1 shows an overview of the results when training on the butterfly dataset. The baseline achieves an accuracy of 67.52 % when evaluated on the target test set. Comparing this with the result of the target-only classifier, it is evident that there is a gap of 27.07 % and this is likely caused by the domain shift. In general, the result of the target-only classifier confirms that the fine-grained classification task can be completed with an accuracy of 94.59 %.

Furthermore, in table 2, we present the results of the source-only classifier on the plant dataset. There is a clear difference to the results of the training on the butterfly dataset. The few images per class likely cause that the classifier cannot generalize well enough to overcome the large shift between domains. The best result achieved here is 15.04 %, which is, nevertheless, slightly better than a random guess. In contrast, the target-only classifier achieves a result of 81.23 % accuracy. There is a domain shift in both

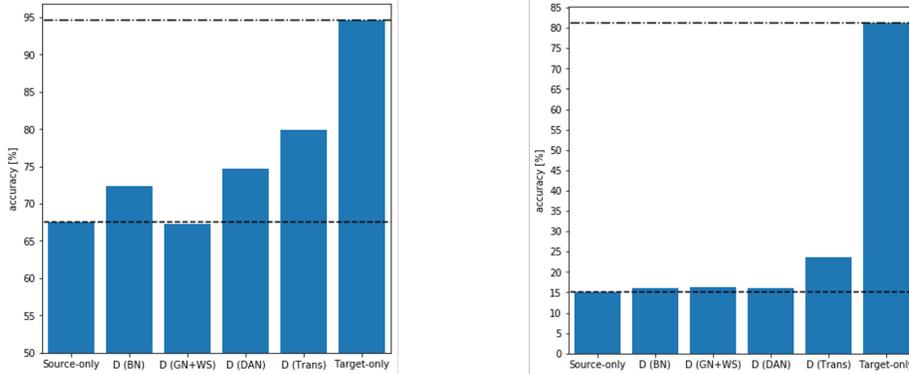


Figure 3. A comparison between the results (source-only and target-only) from the experiments in section 5.1 and the domain adaptation approach is shown. The diagram on the left shows results for the butterfly dataset and the one on the right for the plant dataset. The domain adaptation setup is divided into the different normalization methods used: Batch Normalization - D (BN), Group Normalization in combination with Weight Standardization - D (GN+WS), Domain Agnostic Normalization - D (DAN) and Transferable Normalization - D (Trans).

datasets and great potential when comparing the results of the source- and target-only classifier, which should be used by the domain adaptation approach in the following experiment.

5.2. Domain Adaptation Approach

This experiment aims to determine whether and how well the unsupervised domain adaptation approach works for our fine-grained recognition tasks. We test four setups of the domain adaptation architecture with different normalization methods. For the training, annotated source and non-annotated target data are used. Afterward, the evaluation is performed on the target test data.

In figure 3 on the left side, we present the results of the domain adaptation approach for the butterfly dataset. The upper bound and baseline from the experiments in section 5.1 are used for comparison. In between, there are the domain adaptation architectures with the different normalization methods. It has to be noted that the GN+WS [47, 28] normalization performs worse than BN [15] and does not even offer any improvement to the source-only classifier. The other three normalization methods show an improvement compared to the baseline. TransNorm [45] performs best with 79.87 % accuracy. This confirms for the butterfly dataset that the domain adaptation for the fine-grained recognition task works. However, between the results of the domain adaptation and the upper bound, there is still a gap of 14.72 % accuracy. In the upcoming experiment, we try to reduce this gap by using active learning methods.

On the right side in figure 3 an overview of the experiments on the plant dataset is provided. The results from section 5.1 are also used here for comparison. It becomes clear that this time only the TransNorm shows an advantage over the baseline. BN, GN+WS and DAN [30] do not provide an improvement. The reason could be the small dataset and

the difficult classification task with the large domain shift. We discuss the specific difficulties with the dataset in more detail in section 5.4. TransNorm can achieve a substantially better result of 23.54 % accuracy by weighting the different channels according to their transferability. However, the gap between this result and the one of the target-only classifier is with 57.69 % accuracy huge, compared to the gap in the corresponding experiment with the butterfly dataset.

In summary, it can be said that the normalization layers have a substantial influence on the classification result of the domain adaptation classifier. When using the TransNorm, the domain adaptation approach works well in the field of fine-grained recognition of butterflies and plants. This classifier outperformed the source-only training up to 12.35 % accuracy only by incorporating the unlabeled target data into the training process. Despite the aforementioned weakness, BN offered better results than the baseline. In contrast, the classifier with GN+WS only achieved results comparable to the source-only classifier. To the best of our knowledge, the latter method was applied for the first time in this area. However, it is not recommended for the datasets used in this work.

5.3. Domain Adaptation combined with Active Learning

To further improve the results of the domain adaptation classifier, parts of the target dataset can be annotated using active learning methods. Firstly, we determine the number of target samples that have to be annotated to obtain a comparable result to the upper bound from the experiment in section 5.1. Therefore, random selection is used as a baseline and the determined hyperparameters from section 5.2 are applied. For both datasets, ten images per round are annotated and then the network is trained from scratch. Within 30 rounds, 300 images are annotated for

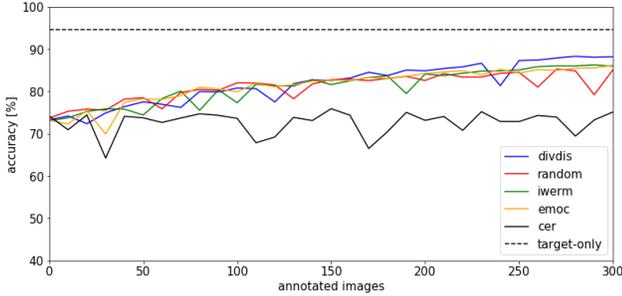


Figure 4. In the diagram, the five different active learning strategies, which are color-coded, are compared. For this purpose, ten images per round have been annotated from the butterfly target dataset over 30 rounds.

both datasets. In the case of the plant dataset, this is already the complete training set. When using the butterfly dataset, we set $k = 500$ after the first 30 rounds. Thus, the active learning approach can be applied with an acceptable amount of resources.

In the second part of the experiment, we use the same setup to test and compare more sophisticated selection strategies as a key component of the active learning approach. The strategies EMOC [20], DivDis [4], Cer and IWERM [38] are compared with the random selection tested in the first part. Furthermore, a second setup is used, where we increase the number of runs to 17 and reduce the number of rounds to ten. Thus, a total of 100 images are annotated. This setup was only carried out for the plant dataset in order to identify possible differences between the strategies.

Figure 4 shows a plot that illustrates the results of the training on the butterfly dataset for the first 30 rounds. The random selection baseline, which is shown in red, increased by 10 % after 300 annotations. A clear upward trend could be seen in all curves as well, except that of the Cer strategy. They are also not strictly monotonously increasing. After some rounds, they are even decreasing. There are some outliers in the curves, which are not averaged out due to the low number of repetitions. In general, the plot clearly shows that all results of the strategies are very close to each other except for Cer. The Cer strategy drops far below the others after the first 50 annotated images and the result after 300 annotated images is only about one percent better. It can also be seen that the DivDis strategy gives the best results from about 200 annotated images onwards, whereas the random strategy lies below DivDis, IWERM and EMOC.

As the accuracies are more than 10 % below the upper bound (dashed line), we have increased the number of images to be annotated per round to 500. The results of this setup are shown in figure 5. The baseline (red curve) needs about 4000 annotations to achieve a comparable

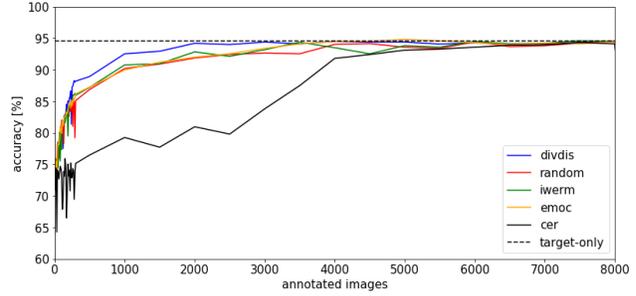


Figure 5. This plot presents the results of active learning with five different active learning strategies on the butterfly dataset compared to the result of the target-only training (black, dashed) from the experiment in section 5.1. There are 16 rounds, and 500 images are annotated in each round. The first 30 data points are taken from figure 4.

result to the upper bound, which corresponds to half of the target training set. The Cer strategy performs much worse with about 5000 images to annotate. The curves of the IWERM and EMOC strategies are slightly above the baseline and require about 3500 annotations. The best strategy is DivDis, which requires the fewest annotations with about 2000.

When using the plant dataset, ten images are annotated step by step until the complete target training set is labeled. The curve rises steepest for the first hundred annotated images that were added and then flattens further and further, as seen at the red line in figure 6. The accuracy is increased by 64.5 % over all 30 rounds but does not reach the results of target-only training even after the completed annotation of the target training set. This problem can be explained by the fact that the large domain shift means that the features have to be aligned more, so that the classifier can generalize better. However, this does not lead to a high performance in a specific domain. This can be seen when comparing the results on the other domains. The target-only classifier gives very poor results of 18 % accuracy on the source test data compared to the domain adaptation classifier on the source data which gives very good results of 83 % accuracy.

Looking at the other selection strategies, all curves show a clear upward trend with some outliers. Up to the first 50 annotated target samples, all strategies have similar results. Then the Cer strategy drops and only after 250 annotations, a comparable result to the others is reached again. In general, there is no strategy that can be regarded as being definitely better. In most cases the IWERM and random selection curves show the best results.

In order to obtain a more stable mean value, the experiment was performed 14 additional times on the plant dataset. Due to resource restrictions, the number of rounds

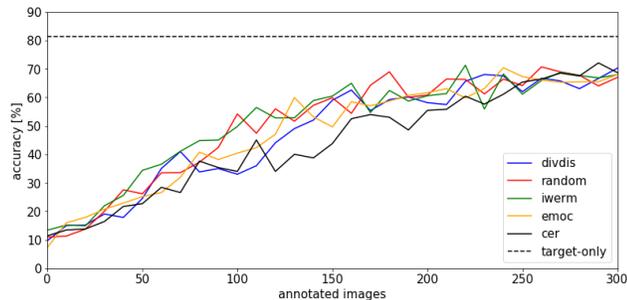


Figure 6. In the plot, the five different active learning strategies, which are color-coded, are compared. For this purpose, ten images per round have been annotated from the plant target dataset over 30 rounds.

was reduced to ten. We still observe that the strategies barely differ. However, the 17 rounds make it possible to perform a meaningful significance test for the results. Therefore, we tested the results of every combination of two strategies for significance, using student’s t-test [37]. Two results differ significantly if the p-value is smaller than the threshold of 0.05. The test shows that none of the strategies differ significantly from each other over all rounds. Between the strategies IWERM & Cer and random & Cer, the p-values of three results after different rounds show that Cer is significantly worse than IWERM and the random selection. Furthermore, the random selection also performs significantly better than the DivDis strategy in three results. In all other cases, only individual values are classified as significantly different. An overview of all p-values is presented in the supplementary materials C.

In summary, when training on the butterfly dataset, only the Cer strategy was worse than the baseline. The EMOC and IWERM strategies performed slightly better than baseline and the DivDis strategy works best, requiring only half the annotations of the baseline. Using the plant dataset, the results could be increased by up to 55 %, but even after the complete training data annotation, no comparable result to the target-only training was reached. Even the use of more sophisticated selection strategies did not lead to significantly better results. This is probably due to the massive domain shift and the small dataset. For the plant dataset we recommend the random selection due to its resource-friendliness.

5.4. Discussion of Dataset Problems

In this section, we discuss possible difficulties that occur with the datasets. The results of the classifiers trained on the butterfly dataset are very good, despite the differences between the domains. Figure 1 shows one example from the source domain and one from the target domain. It is immediately noticeable that the preserved specimen is huge and the patterns on the wings are clearly visible, when contrast-

ing them to the images from the target dataset, where the butterflies are tiny. This can be challenging for the CNN, as it generally does not produce scale-invariant features [12].

Furthermore, some butterflies are captured in flight and are therefore only visible out of focus. In the images it can be seen that butterflies are also photographed from different angles and sometimes the pattern on the wing is not visible. Additionally, the undersides of the wings look different from the topsides. This can also lead to recognition problems. However, the dataset contains enough images to overcome the obstacles most of the time. A collection of more examples from the dataset can be found in the supplementary materials B.1.

The domain adaptation classifier trained on the plant dataset provides only minor improvements over the baseline. On the left-hand side in figure 2, an image from the source domain is shown and on the right-hand side, there is an example from the target domain. In general, it can be seen that the preserved specimens mostly show the leaves and only occasionally the flowers. Besides, the flowers have a different shape due to the conservation and the colors are also different. The images of plants in the wild often show the flowers. However, there are also examples, where the complete plant is pictured. Usually the leaves are not in focus in these images and often, they can only be found very small in the background. Another point is that the source images sometimes contain only fragments of the plant. Some more examples of difficult images can be found in the supplementary materials B.2.

These reasons clarify the enormous difference between the domains and show how little information can only be transferred and used for classification. The low number of training images makes it even more difficult.

6. Conclusion

In this work, we have shown that Unsupervised Adversarial Domain Adaptation works for fine-grained recognition in the field of biodiversity. To do so, we created two datasets of butterflies and plants. Our experiments demonstrated that the domain adaptation worked for our fine-grained recognition task and that the normalization method has a great influence on the result of the classifier. Among the methods that have been investigated, we found that GN+WS is similar to the baseline, DAN and BN are in the middle and TransNorm provides the best results. Moreover, we extended the classifier with an active learning component and compared different selection strategies with each other. In general, the Cer strategy performed worst on both datasets. For the butterfly dataset we recommend using DivDis as it gives the best results. When training on the plant dataset, none of the tested strategies could significantly outperform the baseline, therefore the random selection is advised.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Vaughan. A theory of learning from different domains. *Machine Learning*, 79:151–175, 2010. 2, 3
- [2] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In B. Schoelkopf, J. C. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 137–144. MIT Press, 2007. 3
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning, 2019. 2
- [4] Klaus Brinker. Incorporating diversity in active learning with support vector machines. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML03, pages 59–66. AAAI Press, 2003. 2, 3, 4, 7
- [5] Woong-Gi Chang, Tackgeun You, Seonguk Seo, Suha Kwak, and Bohyung Han. Domain-specific batch normalization for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [6] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-ensembling with gan-based data augmentation for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [7] Koby Crammer, Michael Kearns, and Jennifer Wortman. Learning from data of variable quality. In Y. Weiss, B. Schoelkopf, and J. C. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 219–226. MIT Press, 2006. 3
- [8] Bharath Bhushan Damodaran, Benjamin Kellenberger, Remi Flamary, Devis Tuia, and Nicolas Courty. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 2
- [9] Alexander Freytag, Erik Rodner, and Joachim Denzler. Selecting influential examples: Active learning with expected model output changes. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 562–577, Cham, 2014. Springer International Publishing. 3
- [10] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 1, 2, 3
- [11] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In *European Conference on Computer Vision*, pages 597–613. Springer, 2016. 2
- [12] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. 8
- [13] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 5
- [15] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 1, 2, 3, 6
- [16] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks, 2016. 2
- [17] S. D. Jain and K. Grauman. Active image segmentation propagation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2864–2873, 2016. 3
- [18] Edward Johns, Oisín Mac Aodha, and Gabriel J Brostow. Becoming the expert-interactive multi-class machine teaching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2616–2624, 2015. 4
- [19] C. Kaeding, A. Freytag, E. Rodner, P. Bodesheim, and J. Denzler. Active learning and discovery of object categories in the presence of unnameable instances. In *2015 IEEE Conference on Computer Vision and Pattern Recognition CVPR*, pages 4343–4352, 2015. 3
- [20] Christoph Kaeding, Erik Rodner, Alexander Freytag, and Joachim Denzler. Active and continuous exploration with deep neural networks and expected model output changes, 2016. 2, 4, 7
- [21] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4893–4902, 2019. 2
- [22] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 491–507, Cham, 2020. Springer International Publishing. 1, 4
- [23] Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012. 1, 3
- [24] David D. Lewis and Jason Catlett. Heterogeneous uncertainty sampling for supervised learning. In William W. Cohen and Haym Hirsh, editors, *Machine Learning Proceedings 1994*, pages 148 – 156. Morgan Kaufmann, San Francisco (CA), 1994. 3
- [25] Yanghao Li, Naiyan Wang, Jianping Shi, Xiaodi Hou, and Jiaying Liu. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition*, 80:109 – 117, 2018. 2
- [26] Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. In D. Lee, M. Sugiyama, U. Luxburg, I.

- Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. [2](#)
- [27] Samarth Mishra, Kate Saenko, and Venkatesh Saligrama. Surprisingly simple semi-supervised domain adaptation with pretraining and consistency, 2021. [2](#)
- [28] Siyuan Qiao, Huiyu Wang, Chenxi Liu, Wei Shen, and Alan Yuille. Weight standardization, 2019. [1](#), [2](#), [3](#), [6](#)
- [29] Piyush Rai, Avishek Saha, Hal Daume, and Suresh Venkatasubramanian. Domain adaptation meets active learning. In *Proceedings of the NAACL HLT 2010 Workshop on Active Learning for Natural Language Processing*, ALNLP 10, pages 27–32, USA, 2010. Association for Computational Linguistics. [3](#)
- [30] Rob Romijnders, Panagiotis Meletis, and Gijs Dubbelman. A domain agnostic normalization layer for unsupervised adversarial domain adaptation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1866–1875. IEEE, 2019. [1](#), [2](#), [3](#), [6](#)
- [31] Nicholas Roy and Andrew McCallum. Toward optimal active learning through sampling estimation of error reduction. In *Proceedings of the Eighteenth International Conference on Machine Learning*, ICML 01, pages 441–448, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. [3](#)
- [32] Artem Rozantsev, Mathieu Salzmann, and Pascal Fua. Beyond sharing weights for deep domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(4):801–814, Apr 2019. [2](#)
- [33] Danila Rukhovich and Danil Galeev. Mixmatch domain adaptation: Prize-winning solution for both tracks of visda 2019 challenge, 2019. [2](#)
- [34] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. [3](#)
- [35] Avishek Saha, Piyush Rai, Hal Daumé, Suresh Venkatasubramanian, and Scott L. DuVall. Active supervised domain adaptation. In *Proceedings of the 2011 European Conference on Machine Learning and Knowledge Discovery in Databases - Volume Part III*, ECML PKDD 11, pages 97–112, Berlin, Heidelberg, 2011. Springer-Verlag. [3](#)
- [36] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8050–8058, 2019. [2](#)
- [37] Student. The probable error of a mean. *Biometrika*, 6(1):1–25, 1908. [8](#)
- [38] Jong-Chyi Su, Yi-Hsuan Tsai, Kihyuk Sohn, Buyu Liu, Subhransu Maji, and Manmohan Chandraker. Active adversarial domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 739–748, 2020. [2](#), [3](#), [4](#), [7](#)
- [39] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, pages 443–450. Springer, 2016. [2](#)
- [40] Kiat Chuan Tan, Yulong Liu, Barbara Ambrose, Melissa Tulig, and Serge Belongie. The herbarium challenge 2019 dataset, 2019. [4](#)
- [41] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE international conference on computer vision*, pages 4068–4076, 2015. [3](#)
- [42] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, page 5339–5349, Red Hook, NY, USA, 2018. Curran Associates Inc. [2](#)
- [43] Josiah Wang, Katja Markert, and Mark Everingham. Learning models for object recognition from natural language descriptions. In *Proceedings of the British Machine Vision Conference*, 2009. [4](#)
- [44] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018. [2](#)
- [45] Ximei Wang, Ying Jin, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Transferable normalization: Towards improving transferability of deep neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alch e Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 1953–1963. Curran Associates, Inc., 2019. [1](#), [2](#), [3](#), [6](#)
- [46] Garrett Wilson and Diane J Cook. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5):1–46, 2020. [2](#)
- [47] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. [1](#), [2](#), [3](#), [6](#)
- [48] Yazhou Yang and Marco Loog. Active learning using uncertainty information. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 2646–2651. IEEE, 2016. [3](#)
- [49] Donggeun Yoo, Namil Kim, Sunggyun Park, Anthony S Paek, and In So Kweon. Pixel-level domain transfer. In *European Conference on Computer Vision*, pages 517–532. Springer, 2016. [2](#)
- [50] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. [2](#)
- [51] Xiaojin Zhu, John Lafferty, and Zoubin Ghahramani. Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In *ICML 2003 workshop on The Continuum from Labeled to Unlabeled Data in Machine Learning and Data Mining*, pages 58–65, 2003. [3](#)

Supplementary Materials

A Experimental Setups

A.1 Butterfly Classification

Hyperparameter/Classifier	Source-only	Domain Adaptation	Target-only
L2-Regularization	0.01	0.05	0.01
Learning Rate	0.001	0.0001	0.001

Table 1: This table shows the results of the grid search for the different setups.

A.2 Plant Classification

Hyperparameter/Classifier	Source-only	Domain Adaptation	Target-only
L2-Regularization	0.001	0.1	0.15
Learning Rate	0.001	0.0001	0.0001

Table 2: This table shows the results of the grid search for the different setups.

B Further Examples from the Datasets

B.1 Butterfly Dataset



Figure S1: Here four images from the butterfly dataset are shown, which should illustrate possible difficulties in classification. On the left side images from the source domain and on the right side the corresponding images from the target domain can be seen.

B.2 Plant Dataset



Figure S2: Here six images from the plant dataset are shown, which should illustrate possible difficulties in classification. On the left side images from the source domain and on the right side the corresponding images from the target domain are shown.

C Results of the Significance Test

Round	0	1	2	3	4	5	6	7	8	9	10
rnd - emoc	0.66	0.56	0.85	0.03	0.03	0.1	0.13	0.12	0.5	0.27	0.07
rnd - iworm	0.05	0.91	0.18	0.28	0.25	0.71	0.37	0.3	0.61	0.48	0.14
rnd - divdis	0.25	0.9	0.31	0.7	0.59	0.76	0.01	0.03	0.28	0.15	0.02
rnd - cer	0.48	0.98	0.07	0.42	0.14	0.28	0.43	0.02	0.11	0.002	0.001
iworm - emoc	0.07	0.52	0.1	0.13	0.21	0.23	0.29	0.38	0.29	0.46	0.3
iworm - divdis	0.54	0.82	0.76	0.44	0.5	0.87	0.03	0.31	0.07	0.32	0.22
iworm - cer	0.19	0.9	0.43	0.92	0.73	0.56	0.94	0.2	0.01	0.004	0.01
divdis - emoc	0.21	0.6	0.2	0.05	0.07	0.12	0.81	0.75	0.95	0.94	0.77
divdis - cer	0.57	0.94	0.33	0.61	0.31	0.34	0.04	0.6	0.56	0.17	0.27
cer - emoc	0.35	0.59	0.04	0.15	0.33	0.43	0.28	0.99	0.67	0.43	0.7

Table 3: This table presents the results of the student's t-test of the different active learning strategies on the plant dataset. Every combination of two strategies has been examined. If the calculated p-value is less than 0.05 the results are significantly different.