

Automatic Plant Cover Estimation with Convolutional Neural Networks

Matthias Körschens^{1,2}, Paul Bodesheim³, Christine Römermann⁴, Solveig Franziska Bucher⁵, Mirco Migliavacca⁶, Josephine Ulrich⁷, Joachim Denzler⁸

Abstract: Monitoring the responses of plants to environmental changes is essential for plant biodiversity research. This, however, is currently still being done manually by botanists in the field. This work is very laborious, and the data obtained is, though following a standardized method to estimate plant coverage, usually subjective and has a coarse temporal resolution. To remedy these caveats, we investigate approaches using convolutional neural networks (CNNs) to automatically extract the relevant data from images, focusing on plant community composition and species coverages of 9 herbaceous plant species. To this end, we investigate several standard CNN architectures and different pretraining methods. We find that we outperform our previous approach at higher image resolutions using a custom CNN with a mean absolute error of 5.16%. In addition to these investigations, we also conduct an error analysis based on the temporal aspect of the plant cover images. This analysis gives insight into where problems for automatic approaches lie, like occlusion and likely misclassifications caused by temporal changes.

Keywords: Computer Vision; Biodiversity; Deep Learning; Plant Cover

1 Introduction

Environmental changes can have a large impact not only on human life, but also on animals and plants, and can have drastic effects on biodiversity. This gives reason to steadily monitor nature for such changes and their impacts. Plants, for example, are strong indicators of environmental changes like climate change. This can be seen in phenological responses [Ro07, Me06, Bu18, MRP08, C112, FF02], and also in changes in plant community compositions [Ro07, Li18, Li09]. However, even more environmental aspects can be

¹ Computer Vision Group, Friedrich Schiller University Jena, 07737 Jena, Germany; matthias.koerschens@uni-jena.de

² Plant Biodiversity Group, Friedrich Schiller University Jena, 07737 Jena, Germany

³ Computer Vision Group, Friedrich Schiller University Jena, 07737 Jena, Germany; paul.bodesheim@uni-jena.de

⁴ Plant Biodiversity Group, Friedrich Schiller University Jena, 07737 Jena, Germany; christine.roemermann@uni-jena.de

⁵ Plant Biodiversity Group, Friedrich Schiller University Jena, 07737 Jena, Germany; solveig.franziska.bucher@uni-jena.de

⁶ Department Biogeochemical Integration, Max Planck Institute for Biogeochemistry, 07701 Jena, Germany; mmigliavacca@bgc-jena.mpg.de

⁷ Plant Biodiversity Group, Friedrich Schiller University Jena, 07737 Jena, Germany; josephine.ulrich@uni-jena.de

⁸ Computer Vision Group, Friedrich Schiller University Jena, 07737 Jena, Germany; joachim.denzler@uni-jena.de



monitored through plants, like land use [Ge14, AC12] and insect abundance [SZS16]. Therefore, plant community compositions are the focus of a large number of experiments [Li18, Ge14, SZS16, Br18], where the composition is usually recorded at regular time intervals. This is usually done by estimating the plant cover of each species in the predefined vegetation plots. The cover of a plant species describes the percentage of soil area covered by the respective species. In many experiments, this value is estimated at intervals of at least about a week, usually much longer, due to it being a very laborious task, and the estimation itself is usually done by a single biologist. This introduces two problems: a low temporal resolution of recordings of the plant community composition and noisy data caused by human error and subjectivity of the estimations. These problems can be mitigated by automatic image analysis methods. However, while in current times there are many possibilities to collect images and videos of vegetation plots in an automated manner (e.g., [Br16]), performing more complicated analyses automatically like the cover estimation mentioned above is still not possible as we are still missing the methods to do so. Convolutional neural networks (CNNs) have proven to be a very effective means for image analysis in general and hence could also be utilized as a basis for such methods. However, automatic plant cover prediction is a rather complicated task, as it usually includes large images containing many plant species with only subtle differences, which strongly occlude each other. In addition to this, the annotations for plant cover are relative values (percentages), which is a mostly unexplored kind of annotation in the area of computer vision. In our previous work [Kö20] we laid the groundwork to solve the task of plant cover prediction. In this work, we will investigate that approach, i.e., the cover prediction via pixel-wise classification using the previously proposed network head, in more detail regarding multiple aspects. Firstly, we aim to investigate how different standard CNN architectures with differing characteristics perform on our task in comparison to the network from [Kö20]. The architectures we compare are a ResNet50 [He16], InceptionV3 [Sz16] and DenseNet121 [Ia14]. Secondly, we investigate whether the addition of information from earlier network layers of the abovementioned networks and an additional increase of the output resolution by using a Feature Pyramid Network [Li17a] can improve the performance in comparison to the networks' standalone performance. Thirdly, as changing the image resolution during training and evaluation can strongly affect the network performance, for example, due to the limitations of the networks' receptive field size, we evaluate the performance of the abovementioned CNN architectures on different image resolutions. Fourthly, as pretraining in a similar domain to the target task usually results in better network performances than out-of-domain pretraining [Cu18], we aim to find out how large the benefits of this method are for plant cover prediction. To this end, we compare the previously mentioned configurations in two different pretraining settings. For our domain-similar ("in-domain") pretraining setting, we constructed a dedicated plant image dataset using images from the GBIF website⁹, here referred to as the GBIF dataset. In our out-of-domain setting, we utilize freely available off-the-shelf Imagenet [Ru15] pretraining weights. Lastly, to explain the errors made by the networks during the plant cover prediction, we conduct

⁹ Global Biodiversity Information Facility, <https://www.gbif.org/>

a temporal analysis of the prediction error of a selected network based on the week of recording of the respective images. In the following, we will elaborate on related work in the area of automatic plant analysis (section 2), followed by an introduction of the datasets and methods we use in section 3 and section 4, respectively. After that, we present the results of our experiments (section 5) and end the paper with a conclusion and an outlook on possible future work.

2 Related Work

2.1 Plant Identification

In the last years, there have been several CNN-based approaches developed for the identification of plant species from images [YR16, Ba17, Le15, GYA17]. Barré et al. [Ba17] developed a rather simple CNN architecture called LeafNet, whose aim is to determine the plant species from leaf images. Yalcin et al. [YR16] utilize a pretrained CNN with 11 layers to classify images of agricultural plants. However, simple identification of plant species via images is not the only area where Deep Learning approaches like CNNs are applied. Other tasks include the detection of fruits via segmentations [Ga19], counting of fruits of agricultural plants via regression networks [Lu17, Xi19] and plant disease prediction [Ch20]. A more comprehensive list of approaches can be found, for example, in the survey by Kamilaris et al. [KPB18].

2.2 Plant Cover Prediction

While the number of approaches for problems concerning plant identification is rather abundant, only two approaches try to solve the problem of cover prediction via images. The first one is by Kattenborn et al. [Ka20], who tried to determine the cover percentages of woody species from UAV-based remote sensing data. For their analysis, they utilized a simple custom CNN with eight layers to calculate the cover percentages for several herb, shrub and woody species. In contrast to the data used in this paper, there was not only no occlusion to be taken into account, but the annotations were delineations in the images. In addition to that, the woody plant species investigated mostly had heterogeneous appearances, making automatic discrimination between them relatively easy compared to the much more similar herbaceous species found in the InsectArmageddon dataset [Ul10, Kö20] we use. The second approach in the area of plant cover prediction is our recent work presented in [Kö20], which we will build on in this paper. In [Kö20] we utilized a novel custom network architecture with 12 layers and a block similar to those of the Inception network [Sz16] for feature aggregation. The respective experiments were also done on the InsectArmageddon dataset from Ulrich et al. [Ul10]. In contrast to our previous work, we will not utilize a custom network here but evaluate the performance of several standard CNN architectures in



Fig. 1: Example images from our GBIF dataset. The plant species shown are from left to right: *Achillea millefolium* (Yarrow), *Centaurea jacea* (brown knapweed), *Plantago lanceolata* (ribwort plantain), *Trifolium pratense* (red clover), *Scorzonerooides autumnalis* (autumn hawkbit) and *Grasses*, which are not differentiated into different species.

conjunction with Feature Pyramid Networks (FPNs) [Li17a]. This this end, we utilize the network head proposed in [Kö20] and hence adopt the pixel-wise classification approach. This way, we aim to determine how well these architectures are suited for this task and if the shortcut connections and higher output resolution added by the FPN can be beneficial for plant cover prediction. In addition to this, we also investigate different image resolutions as well as different types of pretraining, as they can strongly affect the performance of CNNs.

3 Datasets

In our experiments, we utilize two datasets. The first one is used for pretraining and will be referred to as GBIF dataset, as it contains plant images taken from the Global Biodiversity Information Facility (GBIF)¹⁰. The second dataset is the so-called InsectArmageddon dataset introduced in [UI20] and [Kö20]. These datasets will be explained in subsection 3.1 and subsection 3.2 respectively.

3.1 The GBIF Dataset

The GBIF dataset is a classification dataset we designed specifically for pretraining models in preparation for using them for cover prediction on the InsectArmageddon dataset. Hence, it contains mostly the same classes as the latter, namely the plant species *Achillea millefolium* (Ach_mil), *Centaurea jacea* (Cen_jac), *Lotus corniculatus* (Lot_cor), *Medicago lupulina* (Med_lup), *Plantago lanceolata* (Pla_lan), *Scorzonerooides autumnalis* (Sco_aut) and *Trifolium pratense* (Tri_pra), as well as a collective class encompassing a group of grasses in which the individuals are not further identified to the species-level.

This dataset contains 7,200 images with 750 training and 150 validation images per class and is therefore also balanced. A selection of example images from the dataset is shown in Figure 1.

¹⁰ <https://www.gbif.org/>



Fig. 2: Example images from the InsectArmageddon dataset. We can see temporal changes in plant coverages in the experimental units of the Ecotron, where different plant species have been grown from seed.

3.2 The InsectArmageddon Dataset

The InsectArmageddon dataset was generated during the eponymous project¹¹ in 2018, in which the effect of invertebrate density on plant composition and growth were investigated [UI20]. In this experiment, 24 so-called EcoUnits, i.e., enclosed boxes with a base area of about $1.5\text{ m} \times 1.5\text{ m}$ containing small, closed ecosystems with different experimental treatments, were equipped with two cameras each, collecting images from a height of about 2 m over a duration of 18 weeks. During this time, plants of 9 different species grew in the EcoUnits, resulting in large variations of the plants in the images as seen in Figure 2. The plant species in this dataset are the same as in section 3.1 with the addition of dead biomass, referred to as *Dead litter*, which was introduced due to lack of visual distinguishability in the images. For more information on the InsectArmageddon experiment itself or the image collection process, we would like to refer to [UI20, ET18, Tü17] and [Kö20], respectively.

The images themselves contain the plant species in largely different distributions, partially caused by the different treatments. Therefore, the contents of the images can differ strongly. In addition to this, the dataset is heavily imbalanced. On the one hand, multiple plant species are either very abundant or have large leaf sizes, each leading to large cover percentages. On the other hand, there are also plant species with tiny leaves, leading to excessively small cover percentages. This distribution is visualized in Figure 3, which shows that the most prevalent plant is *Trifolium pratense*, representing about a third of the complete dataset. In contrast, *Achillea millefolium* can be found on the other end of the distribution with only about 3% cover.

¹¹ https://www.idiv.de/en/research/platforms_and_networks/idiv_ecotron/experiments/insect_armageddon.html

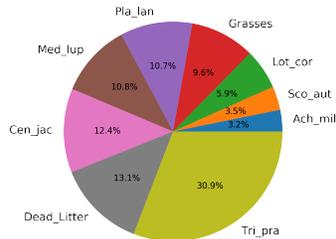


Fig. 3: The distribution of plant species in the InsectArmageddon dataset, displayed as fractions of the total sum of cover percentages over the complete dataset. *Trifolium pratense* is the most abundant plant in the dataset, while *Achillea millefolium* represents the least abundant species.

The dataset contains a large number of difficulties for automatic image analysis methods. Due to technical reasons, the zoom levels and image quality are inconsistent within and across the EcoUnits. However, one of the biggest challenges in the dataset is the occlusion caused by massive mutual overlaps of the plants, which is inherent to plant communities, irrespective of the number of co-occurring species. This makes a correct plant cover prediction very difficult.

While images from the EcoUnits have been collected daily, due to the necessity of annotating them, we can only utilize one image per week per camera for our investigations, leading to only 682 images with annotations.

3.2.1 Plant Cover Annotations

The plant cover is defined as the area of ground covered by each plant species. This value is typically estimated by a biologist directly in the field and can thus be highly subjective. Therefore, the annotations of the InsectArmageddon dataset, which were provided by Ulrich et al. [Ul20], can be considered noisy. In addition to this, as an accurate percentage estimation is practically infeasible for humans, the estimated plant cover is usually done with a certain quantization, which in the case of this dataset is the so-called Schmidt-scale [Pf97]. The latter includes the percentages 0, 0.5, 1, 3, 5, 8, 10, 15, 20, 25, 30, 40, 50, 60, 70, 75, 80, 90 and 100 percent. The quantization process introduces additional label noise.

4 Method

We apply the same approach as in [Kö20]: we view the plant cover prediction problem as a pixel-wise classification problem, where these classifications can then be aggregated into a final cover prediction. One additional advantage of this approach is that the pixel-wise classifications can be viewed as segmentation maps, serving as a way for the user to verify the correctness of the result. Taking this approach requires us to identify the plant species

at each location in the images before aggregating them into a single cover value for each species. Figure 2 shows that the plants are, in part, very similar to each other, and, due to the relatively large area captured in the single images, rather small compared to the whole image. In order to still be able to identify the plants correctly and be able to take into account their small and essential features, we need to process the images in a relatively high resolution. Often, when high-resolution images are used during CNN training, the images are divided into smaller parts, which are then used for training (e.g., [He17]). This approach, however, is not possible in our case, as the numerical plant cover annotations are only valid for the complete image. Therefore, we are required to train on high-resolution images directly.

While this was done already in our previous work [Kö20] with a resolution of 672x336 px, here we test the effect of training on even higher resolutions. Hence, we evaluate our method on images with resolutions of 768x384 and 1536x768 px, respectively. To receive a better insight on the effect of different network architectures on the task of plant cover prediction, we also measure the performance of standard classification network architectures in conjunction with a Feature Pyramid Network (FPN) [Li17a]. The networks we will compare are ResNet50 [He16], DenseNet121 [Ia14], and InceptionV3 [Sz16]. ResNet50 was selected as it is used for comparison and as backbone in almost all possible settings and tasks, from simple classification [He16] over object detection [Li17b] to instance segmentation [He17]. InceptionV3, in turn, is commonly used in fine-grained classification tasks (e.g., [Cu18]), and DenseNet121 was chosen due to its particular network structure, which incorporates information from earlier layers directly in the decision process in later layers via long-range skip connections. As these networks only have output feature maps of small sizes, we also apply the FPN, which is usually utilized in tasks like object detection (e.g., [Re15, Li17b]) and semantic segmentation (e.g., [He17]), and could therefore be a sensible choice due to the similarity of the underlying task. The FPN increases the size of the output feature maps by incrementally fusing feature maps from the end of the network with higher resolution ones generated in earlier layers of the network. Hence, this structure also enables us to generate high-resolution output feature maps when joined with the abovementioned standard networks. This also allows us to evaluate whether generating higher resolution class outputs results in a better cover prediction, and to receive information on the usefulness of features from earlier layers.

As mentioned above, we also want to quantify the effect of in-domain and out-of-domain transfer learning. To this end, we use, for the in-domain pretraining, the previously introduced GBIF dataset, as due to the specifically selected plant species it is much more similar to our target task than, for example, the widely used ImageNet [Ru15] dataset, which we use as out-of-domain pretraining for comparison.

4.1 Basic Network Structure

Our network consists of two main components: a feature extractor backbone and a network head. The backbone consists of one of the abovementioned standard classification networks

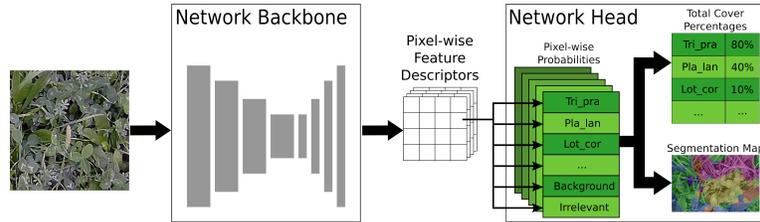


Fig. 4: The basic structure of the network. The network is input in the network backbone, which extracts pixel-wise feature descriptors. These are then used by the network head to calculate the pixel-wise probabilities for each plant, which are afterwards aggregated into the total cover percentages. The generated pixel-wise probabilities can also be interpreted as a segmentation map, so that the predictions of the network can also be interpreted by the user.

potentially in conjunction with a Feature Pyramid Network to increase the output resolution. It hence generates a pixel-wise feature map, which is then used as input for the network head. The latter comprises a calculation model, explained in detail in section 4.2, which initially determines the pixel-wise probabilities of each plant species and afterward aggregates the predicted values into the total cover percentage of each species. As already mentioned, the generated class or probability map can also be used as a segmentation map for the user to confirm the correctness of the results. An overview of this network structure is shown in Figure 4.

4.2 Calculation Model

The calculation model represents the computations taking place in the network head to lastly generate the cover predictions. This calculation is separated into two parts: pixel-wise class predictions and aggregations.

Class Prediction. As the pixel-wise predictions are largely dependent on the actual cover calculation, we will first introduce the intuitive formula to calculate the cover for a single plant species:

$$cover^{plant} = \frac{A^{plant}}{A^{plant} + A^{uncovered}}, \quad (1)$$

where A^{plant} is the area covered by the single plant species and $A^{uncovered}$ is the area not covered by it. Under this premise, there are several things to take into account for the pixel-wise class prediction. First, as occlusion is usually ignored for calculating the plant cover, it has to be possible to predict multiple different plant species in the same location. Second, the areas in which no plants exist also have to be taken into account in the calculation. Third, as seen in the example images in Figure 3, there are areas in the

image, which are irrelevant for cover calculation (e.g., the walls of the EcoUnits) and should thus be ignored. While the existence of multiple plants in a single pixel does not have to be mutually exclusive, plants, soil and irrelevant pixels cannot be predicted in the same location and therefore have to be mutually exclusive. To be able to calculate the plant cover accurately, we also include these premises in our calculation model. As these factors are already included in the calculation model from our previous work [Kö20], we also utilize this model here with slight modifications.

We define $P_{x,y}^p$ as the probability of a plant p existing at location x, y . We also define $P_{x,y}^{bio}$, $P_{x,y}^{bg}$, and $P_{x,y}^{irr}$ as the probability of any plant, background (uncovered soil), and irrelevant data being at location x, y , respectively. Note that, as mentioned above, the single plant probabilities $P_{x,y}^p$ do not necessarily add up to one, as they are not mutually exclusive, while the three probabilities $P_{x,y}^{bio}$, $P_{x,y}^{bg}$ and $P_{x,y}^{irr}$ do.

As $P_{x,y}^{bio}$ is naturally dependent on the probabilities $P_{x,y}^p$, we can set them in relation, as in [Kö20]:

$$P_{x,y}^{bio} = 1 - (P_{x,y}^{bg} + P_{x,y}^{irr}) = 1 - \frac{\kappa}{\kappa + \sum_p P_{x,y}^p} = \frac{\sum_p P_{x,y}^p}{\kappa + \sum_p P_{x,y}^p}, \quad (2)$$

where κ is a threshold value, which determines if a pixel is considered a plant pixel or not. While κ is a tunable hyperparameter in [Kö20], we use it as a learnable parameter in our experiments.

Aggregation. Now, to aggregate the predicted probabilities into the plant-wise cover percentages, we can calculate the areas

$$(A^{bg}, A^{irr}, A^{bio}) = \left(\sum_{x,y} P_{x,y}^{bg}, \sum_{x,y} P_{x,y}^{irr}, \sum_{x,y} P_{x,y}^{bio} \right), \quad (3)$$

$$A^{total} = A^{bio} + A^{bg} + A^{irr}. \quad (4)$$

Thus, we can finally calculate the cover values for each plant p :

$$cover^p = \frac{\sum_{x,y} P_{x,y}^p}{A^{total} - A^{irr}} = \frac{\sum_{x,y} P_{x,y}^p}{A^{bio} + A^{bg}} \quad (5)$$

While in [Kö20] we also applied a vegetation percentage, i.e., the percentage of ground covered by plants in general, as an auxiliary value, we found that this is not necessary when using pretrained networks, which in turn removes the need for additional annotations.

5 Experiments

In our experiments, we compare the performance of the previously introduced network architectures with each other using two different resolutions. This comparison is done first using in-domain pretraining for each of the networks and then repeated using ImageNet pretraining. At the end of this section, we provide an error analysis based on the experiment duration of the InsectArmageddon dataset. The metrics we will use in our experiments for performance evaluation are the mean absolute error (MAE) and the mean scaled absolute error (MSAE) as defined in [Kö20]. As the MAE is heavily influenced by the average cover percentages of each plant species, we introduced the MSAE, which is the mean absolute error divided by the species-wise plant cover mean, to allow for a fairer comparison between the different species. For the exact values used to scale the errors, we would like to refer the reader to [Kö20].

5.1 Comparison of Different Networks

In this experiment different network architectures are utilized as backbone and their performance is compared. Three different architectures are used: ResNet50 [He16], InceptionV3 [Sz16] and DenseNet121 [Ia14]. We apply these networks in two different versions: with and without an FPN. In the latter case, we utilize an FPN [Li17a] with depth 512 to increase spatial resolution and hence to be able to better predict the classes pixel-wise. We use the P2 layer of the FPN, i.e., the second-to-last layer of the FPN, resulting in a final resolution of $\frac{1}{4}$ th of the original input image resolution.

Experimental Setup. In the following experiments, pretraining on the GBIF dataset (see section 3.1) was utilized, whereafter the main training on the InsectArmageddon (section 3.2) dataset took place.

The pretraining on GBIF was done on the base network and the FPN, if used. All networks were initialized with off-the-shelf ImageNet weights provided by the Keras framework [Ch15]. During the pretraining we used an image resolution of 448×448 px for ResNet50 and DenseNet121, and 427×427 px for InceptionV3 due to its different network structure. Each network was trained using a batch size of 12 for 20 epochs with an initial learning rate of 0.01, decaying by a factor of 0.1 at epoch 10 and 15. To improve the adaptability of the learned weights on the InsectArmageddon dataset, we also utilized several data augmentations: random rotations, random crops, and random horizontal flipping. The optimizer during training is SGD with a momentum of 0.9, and the loss is the standard categorical cross-entropy.

During training on the InsectArmageddon dataset, we exchange the classification head of the network from the pretraining step with the plant cover prediction head explained in section 4.2, which we then train in conjunction with the FPN, if used. The rest of the network

Tab. 1: The numerical results of the experiment using several standard networks with pretraining on the GBIF dataset. Top results are marked in bold font. MAE and MSAE represent the *mean absolute error* and the *mean scaled absolute error*, respectively.

Image Resolution	768 × 384		1536 × 768	
Metric	MAE	MSAE	MAE	MSAE
ResNet50 [He16]	5.49%	0.528	5.22%	0.503
ResNet50 [He16] + FPN [Li17a]	5.35%	0.512	5.18%	0.496
InceptionV3 [Sz16]	5.51%	0.531	5.30%	0.506
InceptionV3 [Sz16] + FPN [Li17a]	5.51%	0.528	5.27%	0.503
DenseNet121 [Ia14]	5.37%	0.522	5.18%	0.495
DenseNet121 [Ia14] + FPN [Li17a]	5.30%	0.515	5.16%	0.494

is frozen during training, as in previous experiments not shown here we found that fully fine-tuning the network does either not change the results or is even detrimental to them. We also found that for this training, a batch size of 1 is sufficient, and larger batch sizes also do not change the results. We assume this is the case due to the calculation performed in the network head. The latter can be interpreted as performing many single classifications with strongly changing and weakly defined target values. This, in turn, massively regularizes the network to a point at which bigger batches and additional regularization do not improve the results in any way. The networks are trained using the Adam [KB15] optimizer for faster convergence for 40 epochs starting with a learning rate of 0.001 (0.00001 in case of the network containing an FPN), decaying with a factor of 0.1 at epoch 20 and 30. We apply the mean absolute error as loss during the training process. Only horizontal flips are applied as data augmentation, as we found that the effect of most augmentations was either negligible or detrimental to the results. All experiments are performed in a special 12-fold cross-validation, i.e., during each fold, we select two EcoUnits, whose images are used for testing, while the images of the remaining units are used for training. For generating the experimental results below, the results of these folds are averaged.

Results. The results of the GBIF-pretraining experiments can be seen in Table 1. It is visible that the best results concerning the mean absolute error (MAE) and mean scaled absolute error (MSAE) were achieved with a DenseNet121 with FPN. We can also note that the DenseNet outperforms the two other networks in almost every single setting, which can likely be attributed to the structure of the DenseNet, which directly incorporates information from earlier layers in the network for the final prediction. Similarly, the ResNet50 outperforms InceptionV3 in all settings, which leads to the assumption that the structure of the Inception modules might be detrimental in this task. From these results, we can also observe that using a higher resolution is strictly better than using the lower one, which means that, even though the 768 × 384 px resolution is already quite high, there is still more relevant information contained in the images, which can only be utilized when using an even higher image resolution. In the experimental data, we can also see that the results improve upon using the FPN for some networks more than for others. The network profiting most

from the FPN is the ResNet50, which seems to produce useful information for identification at multiple levels of the FPN and therefore receives a substantial performance boost. The DenseNet, as mentioned before, already includes low-level features by default in its network structure, which is why the FPN does not add much novel information and therefore only improves results slightly. While the InceptionV3 does not have long-range skip connections, the improvements by addition of the FPN are also only minor. While the reason for this is not entirely apparent, this could be caused by the differing network structure as compared to the ResNet50 and DenseNet121, which might not benefit as much from the FPN. The top result of the DenseNet121 outperforms the network proposed in [Kö20] in terms of mean absolute error (5.16% vs. 5.30%) and mean scaled absolute error (0.494 vs. 0.500). However, it should also be noted that in this work, we completely abstain from using the auxiliary vegetation-percent annotations, which were used in [Kö20] in addition to the plant cover annotations, and the amount of training necessary here is also much less, while still outperforming the network from [Kö20].

5.2 Comparison with ImageNet Pretraining

To evaluate the effect of the additional in-domain pretraining using the GBIF dataset in section 5.1 we conduct an ablation study and compare it with a standard ImageNet pretraining. It should be noted that the ImageNet pretraining has only been done on the base classification networks without FPN, while the pretraining on the GBIF dataset in section 5.1 has been performed on the base network with FPN.

Experimental Setup. In these experiments, the setup is very similar to the one in the experiments before. However, instead of pretraining the network on GBIF, we merely load off-the-shelf ImageNet weights for the base networks, which are provided by the Keras framework [Ch15]. During training, in case of not using an FPN, we again only train the last layer using the same scheme as described above with an initial learning rate of 0.001. In the cases where an FPN is used, as its weights are not initialized, we train the FPN and the cover prediction head jointly with a learning rate of 0.00001.

Results. The results of our experiments with off-the-shelf ImageNet weights can be seen in Table 2. We can see that the top results are generated again by the DenseNet121, with an MAE of 5.17% and an MSAE of 0.491. These results differ only slightly from the previous experiments. This difference likely results from the structure of the DenseNet, which directly incorporates features from earlier layers into top-level predictions, which are beneficial for high-resolution plant predictions. Here, we can also see that the FPN appears to be beneficial in all cases for the lower image resolution. For the higher resolutions, we notice that the performances of the networks with FPN are very similar to the ones of the standalone base networks in all cases. This leads to the assumption that the learned lower-level features

Tab. 2: The numerical results of the experiment using several standard networks with only ImageNet pretraining. Top results are marked in bold font. MAE and MSAE represent the *mean absolute error* and the *mean scaled absolute error*, respectively.

Image Resolution	768 × 384		1536 × 768	
Metric	MAE	MSAE	MAE	MSAE
ResNet50 [He16]	5.55%	0.537	5.32%	0.510
ResNet50 [He16] + FPN [Li17a]	5.37%	0.521	5.30%	0.510
InceptionV3 [Sz16]	5.74%	0.546	5.34%	0.505
InceptionV3 [Sz16] + FPN [Li17a]	5.62%	0.535	5.33%	0.506
DenseNet121 [Ia14]	5.39%	0.515	5.17%	0.491
DenseNet121 [Ia14] + FPN [Li17a]	5.31%	0.510	5.20%	0.494

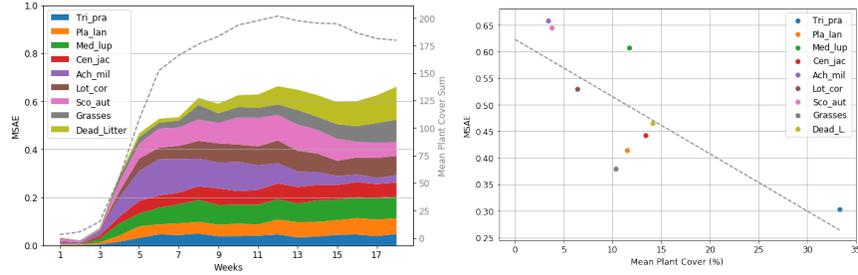


Fig. 5: Left: The mean scaled absolute errors (MSAE) averaged over all plant species in the dataset in dependence on the week of the image recording. For reference the plant cover sum averaged over all weeks is shown as dashed line. Right: The relationship between the species-wise mean plant cover and the mean scaled absolute error (MSAE). We can see that the values are not or even negatively correlated, i.e., as opposed to the MAE the MSAE does not depend strongly on the plant abundances (correlation of $p = -0.774$; $R^2 = 0.600$).

incorporated by the FPN are not very relevant on this resolution anymore, and higher-resolution ones become more crucial. Compared to the GBIF pretraining experiments, we can see that if we only consider the base networks without FPN, the GBIF pretraining is beneficial in almost all cases, with only a slight deviation from this pattern for DenseNet121.

5.3 Error Analysis

To finally also analyze the output and errors of our method, we take a look at the error of the DenseNet121 pretrained on GBIF with FPN for an input image resolution of 1536×784 px from section 5.1 plotted over the weeks recorded in the dataset. As mentioned above, the images contain time series of images collected weekly over a range of 18 weeks, with the plants growing progressively over time while dying towards the end of the experiment. As these changing conditions can hugely affect the outcome of the automatic analysis, we observe the error of our method over the single weeks. Here, we will only focus on the

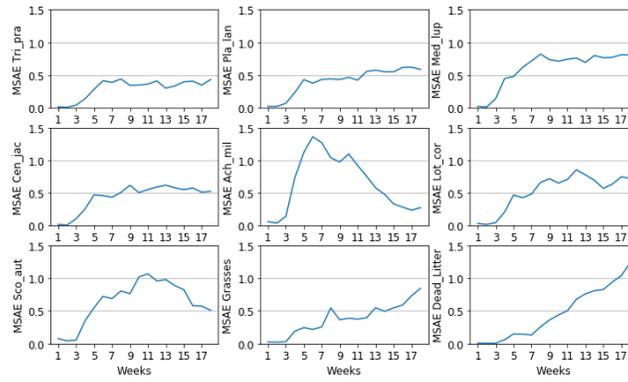


Fig. 6: The mean scaled absolute error in dependence on the week of the image recording.

MSAE. An overview of the MSAE changing over time can be seen in Figure 5, in which the change of the average plant cover percentage summed over all plants is also shown for comparison. We see that in the first few weeks, where not many plants have grown, the error is small but starts rising strongly around week 4, which coincides with the plant growth. Following this sudden rise, the error stays mostly consistent, also coinciding with the cover percentages. This high error can probably be attributed to the large amount of occlusion taking place in the images. It is visible in the graph that the latter peaks at 200%, meaning about half of the plant area used for estimating the ground truth values is occluded, making a correct estimation for the network difficult. Towards the end of the experiment, we can see that the error does not coincide with the changes in cover percentages anymore, as the former is rising while the latter are dropping. This can be explained when taking a look at the plant-wise errors in Figure 6. There we can see that the error for most plants is decreasing, as they are dying towards the end of the experiment. This process, i.e., the death of plants by either age (senescence) or due to the experimental setup, leads to plants being counted not as their original class anymore, but as the *Dead litter* class, where we see a rise in the error. This creates the assumption that the network cannot really identify the *Dead litter* correctly and possibly still assigns the respective pixels to its original class. From this, we can assume that the network could not work out a suitable discrimination between *Dead litter* and the other classes. This is understandable, as the point in time at which a plant starts to be considered *Dead litter* can be quite arbitrary and subjective. While its label changes over the duration of one week, the plants might still bear significant resemblances to their previous label. In summary, we can therefore say that most of the errors are likely caused by the amount of occlusion in the images and the problem of discriminating between the regular plant classes and *Dead litter*. This, however, could also be caused by the non-existence of the *Dead litter* class in the pretraining dataset.

6 Conclusions and Future Work

In this work, we found that utilizing standard CNNs can yield similar or slightly better results than using, for example, our previously proposed network from [Kö20]. We also found that the usage of FPNs is in general beneficial, while the ResNet50 benefits most from them. It should also be noted that FPNs increase the output resolution of the resulting segmentation map, which, in turn, leads to better interpretability of the results compared to their non-FPN counterparts. We also found that using higher image resolutions is always beneficial in our setting. Therefore, collecting images with high resolutions should be prioritized to improve prediction accuracy. Pretraining the network on a more related dataset rather than using off-the-shelf ImageNet weights proved beneficial for all the networks with only few exceptions when using DenseNet121. Hence an in-domain pretraining should, in most cases, be preferred to an out-of-domain one, if available. Also, with pretraining, we can achieve similar or better results to [Kö20] with only training a part of the network on the final task, saving time and hardware capacities for training on the plant cover prediction task. From our week-based error analysis, we were able to conclude that some parts of the errors are most likely caused by the massive occlusions in the images, as well as the network's inability to differentiate between *Dead litter* pixels and its preceding plant species before its death.

In future work, we aim to tackle the abovementioned problems. For occlusion, some approaches in the area of amodal segmentation [LM16] might be suitable. For a better differentiation between plants, other pretraining datasets like iNaturalist [Va18] or even pretraining the network on segmentation datasets like MSCOCO [Li14] to prime it for pixel-wise classification might be a possibility.

Acknowledgement

Matthias Körschens thanks the Carl Zeiss Foundation for the financial support. We would also like to thank Alban Gebler and the iDiv for providing the data for our investigations.

Bibliography

- [AC12] Aggemyr, Elsa; Cousins, Sara AO: Landscape structure and land use history influence changes in island plant composition after 100 years. *Journal of Biogeography*, 39(9):1645–1656, 2012.
- [Ba17] Barré, Pierre; Stöver, Ben C; Müller, Kai F; Steinhage, Volker: LeafNet: A computer vision system for automatic plant species identification. *Ecological Informatics*, 40:50–56, 2017.
- [Br16] Brown, Tim B; Hultine, Kevin R; Steltzer, Heidi; Denny, Ellen G; Denslow, Michael W; Granados, Joel; Henderson, Sandra; Moore, David; Nagai, Shin; SanClements, Michael et al.: Using phenocams to monitor our changing Earth: toward a global phenocam network. *Frontiers in Ecology and the Environment*, 14(2):84–93, 2016.

-
- [Br18] Bruelheide, Helge; Dengler, Jürgen; Purschke, Oliver; Lenoir, Jonathan; Jiménez-Alfaro, Borja; Hennekens, Stephan M; Botta-Dukát, Zoltán; Chytrý, Milan; Field, Richard; Jansen, Florian et al.: Global trait–environment relationships of plant communities. *Nature Ecology & Evolution*, 2(12):1906–1917, 2018.
- [Bu18] Bucher, Solveig Franziska; König, Patrizia; Menzel, Annette; Migliavacca, Mirco; Ewald, Jörg; Römermann, Christine: Traits and climate are associated with first flowering day in herbaceous species along elevational gradients. *Ecology and Evolution*, 8(2):1147–1158, 2018.
- [Ch15] Chollet, François et al.: , Keras. <https://keras.io>, 2015.
- [Ch20] Chen, Junde; Chen, Jinxiu; Zhang, Defu; Sun, Yuandong; Nanehkaran, YA: Using deep transfer learning for image-based plant disease identification. *Computers and Electronics in Agriculture*, 173:105393, 2020.
- [Cl12] Cleland, Elsa E; Allen, Jenica M; Crimmins, Theresa M; Dunne, Jennifer A; Pau, Stephanie; Travers, Steven E; Zavaleta, Erika S; Wolkovich, Elizabeth M: Phenological tracking enables positive species responses to climate change. *Ecology*, 93(8):1765–1771, 2012.
- [Cu18] Cui, Yin; Song, Yang; Sun, Chen; Howard, Andrew; Belongie, Serge: Large scale fine-grained categorization and domain-specific transfer learning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4109–4118, 2018.
- [ET18] Eisenhauer, Nico; Türke, Manfred: From climate chambers to biodiversity chambers. *Frontiers in Ecology and the Environment*, 16(3):136–137, 2018.
- [FF02] Fitter, AH; Fitter, RSR: Rapid changes in flowering time in British plants. *Science*, 296(5573):1689–1691, 2002.
- [Ga19] Ganesh, P; Volle, K; Burks, TF; Mehta, SS: Deep Orange: Mask R-CNN based Orange Detection and Segmentation. *IFAC-PapersOnLine*, 52(30):70–75, 2019.
- [Ge14] Gerstner, Katharina; Dormann, Carsten F; Stein, Anke; Manceur, Amour M; Seppelt, Ralf: EDITOR’S CHOICE: REVIEW: Effects of land use on plant diversity—A global meta-analysis. *Journal of Applied Ecology*, 51(6):1690–1700, 2014.
- [GYA17] Ghazi, Mostafa Mehdipour; Yanikoglu, Berrin; Aptoula, Erchan: Plant identification using deep neural networks via optimization of transfer learning parameters. *Neurocomputing*, 235:228–235, 2017.
- [He16] He, Kaiming; Zhang, Xiangyu; Ren, Shaoqing; Sun, Jian: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, pp. 770–778, 2016.
- [He17] He, Kaiming; Gkioxari, Georgia; Dollár, Piotr; Girshick, Ross: Mask r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. IEEE, pp. 2961–2969, 2017.
- [Ia14] Iandola, Forrest; Moskewicz, Matt; Karayev, Sergey; Girshick, Ross; Darrell, Trevor; Keutzer, Kurt: Densenet: Implementing efficient convnet descriptor pyramids. *arXiv preprint arXiv:1404.1869*, 2014.
- [Ka20] Kattenborn, Teja; Eichel, Jana; Wisser, Susan; Burrows, Larry; Fassnacht, Fabian E; Schmidlein, Sebastian: Convolutional Neural Networks accurately predict cover fractions of plant species and communities in Unmanned Aerial Vehicle imagery. *Remote Sensing in Ecology and Conservation*, 2020.

- [KB15] Kingma, Diederik P.; Ba, Jimmy: Adam: A Method for Stochastic Optimization. CoRR, abs/1412.6980, 2015.
- [Kö20] Körschens, Matthias; Bodesheim, Paul; Römermann, Christine; Bucher, Solveig Franziska; Ulrich, Josephine; Denzler, Joachim: Towards Confirmable Automated Plant Cover Determination. In: European Conference on Computer Vision. Springer, pp. 312–329, 2020.
- [KPB18] Kamilaris, Andreas; Prenafeta-Boldú, Francesc X: Deep learning in agriculture: A survey. *Computers and electronics in agriculture*, 147:70–90, 2018.
- [Le15] Lee, Sue Han; Chan, Chee Seng; Wilkin, Paul; Remagnino, Paolo: Deep-plant: Plant identification with convolutional neural networks. In: 2015 IEEE international conference on image processing (ICIP). IEEE, pp. 452–456, 2015.
- [Li14] Lin, Tsung-Yi; Maire, Michael; Belongie, Serge; Hays, James; Perona, Pietro; Ramanan, Deva; Dollár, Piotr; Zitnick, C Lawrence: Microsoft coco: Common objects in context. In: European conference on computer vision. Springer, pp. 740–755, 2014.
- [Li17a] Lin, Tsung-Yi; Dollár, Piotr; Girshick, Ross; He, Kaiming; Hariharan, Bharath; Belongie, Serge: Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2117–2125, 2017.
- [Li17b] Lin, Tsung-Yi; Goyal, Priya; Girshick, Ross; He, Kaiming; Dollár, Piotr: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988, 2017.
- [Li18] Liu, Huiying; Mi, Zhaorong; Lin, LI; Wang, Yonghui; Zhang, Zhenhua; Zhang, Fawei; Wang, Hao; Liu, Lingli; Zhu, Biao; Cao, Guangmin et al.: Shifting plant species composition in response to climate change stabilizes grassland primary production. *Proceedings of the National Academy of Sciences*, 115(16):4051–4056, 2018.
- [LI09] Lloret, Francisco; Peñuelas, Josep; Prieto, Patricia; Llorens, L; Estiarte, Marc: Plant community changes induced by experimental climate change: seedling and adult species composition. *Perspectives in Plant Ecology, Evolution and Systematics*, 11(1):53–63, 2009.
- [LM16] Li, Ke; Malik, Jitendra: Amodal instance segmentation. In: European Conference on Computer Vision. Springer, pp. 677–693, 2016.
- [Lu17] Lu, Hao; Cao, Zhiguo; Xiao, Yang; Zhuang, Bohan; Shen, Chunhua: TasselNet: counting maize tassels in the wild via local counts regression network. *Plant methods*, 13(1):79, 2017.
- [Me06] Menzel, Annette; Sparks, Tim H; Estrella, Nicole; Koch, Elisabeth; Aasa, Anto; Ahas, Rein; Alm-Kübler, Kerstin; Bissolli, Peter; Braslavská, Ol'ga; Briede, Agrita et al.: European phenological response to climate change matches the warming pattern. *Global change biology*, 12(10):1969–1976, 2006.
- [MRP08] Miller-Rushing, Abraham J; Primack, Richard B: Global warming and flowering times in Thoreau's Concord: a community perspective. *Ecology*, 89(2):332–341, 2008.
- [Pf97] Pfadenhauer, J.: Vegetationsökologie - ein Skriptum. IHW-Verlag, Eching, 2. verbesserte und erweiterte auflage edition, 1997.

-
- [Re15] Ren, Shaoqing; He, Kaiming; Girshick, Ross; Sun, Jian: Faster r-cnn: Towards real-time object detection with region proposal networks. arXiv preprint arXiv:1506.01497, 2015.
- [Ro07] Rosenzweig, Cynthia; Casassa, Gino; Karoly, David J; Imeson, Anton; Liu, Chunzhen; Menzel, Annette; Rawlins, Samuel; Root, Terry L; Seguin, Bernard; Tryjanowski, Piotr et al.: Assessment of observed changes and responses in natural and managed systems. *Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, pp. 79–131, 2007.
- [Ru15] Russakovsky, Olga; Deng, Jia; Su, Hao; Krause, Jonathan; Satheesh, Sanjeev; Ma, Sean; Huang, Zhiheng; Karpathy, Andrej; Khosla, Aditya; Bernstein, Michael et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [Sz16] Szegedy, Christian; Vanhoucke, Vincent; Ioffe, Sergey; Shlens, Jon; Wojna, Zbigniew: Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, pp. 2818–2826, 2016.
- [SZS16] Souza, Lara; Zelikova, Tamara Jane; Sanders, Nathan J: Bottom-up and top-down effects on plant communities: nutrients limit productivity, but insects determine diversity and composition. *Oikos*, 125(4):566–575, 2016.
- [Tü17] Türke, Manfred; Feldmann, Reinhart; Fürst, Bernd; Hartmann, Henrik; Herrmann, Martina; Klotz, Stefan; Mathias, Georg; Meldau, Stefan; Ottenbreit, Markus; Reth, Sascha et al.: Multitrophische Biodiversitätsmanipulation unter kontrollierten Umweltbedingungen im iDiv Ecotron. In: *Lysimetertagung*. pp. 107–114, 2017.
- [Ul20] Ulrich, Josephine; Bucher, Solveig Franziska; Eisenhauer, Nico; Schmidt, Anja; Türke, Manfred; Gebler, Alban; Barry, Kathryn; Lange, Markus; Römermann, Christine: Invertebrate decline leads to shifts in plant species abundance and phenology. *Frontiers in plant science*, 11:1410, 2020.
- [Va18] Van Horn, Grant; Mac Aodha, Oisín; Song, Yang; Cui, Yin; Sun, Chen; Shepard, Alex; Adam, Hartwig; Perona, Pietro; Belongie, Serge: The inaturalist species classification and detection dataset. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, pp. 8769–8778, 2018.
- [Xi19] Xiong, Haipeng; Cao, Zhiguo; Lu, Hao; Madec, Simon; Liu, Liang; Shen, Chunhua: TasselNetv2: in-field counting of wheat spikes with context-augmented local regression networks. *Plant Methods*, 15(1):150, 2019.
- [YR16] Yalcin, Hulya; Razavi, Salar: Plant classification using convolutional neural networks. In: *2016 Fifth International Conference on Agro-Geoinformatics (Agro-Geoinformatics)*. IEEE, pp. 1–5, 2016.