

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Unsupervised domain adaptation and super resolution on drone images for autonomous dry herbage biomass estimation

Paul Albert^{1,3,5}, Mohamed Saadeldin^{2,3,5}, Badri Narayanan^{2,3,5}, Brian Mac Namee^{2,3,5}, Deirdre Hennessy^{4,5}, Noel E. O'Connor^{1,3,5}, Kevin McGuinness^{1,3,5}

¹School of Electronic Engineering, Dublin City University ²School of Computer Science, University College Dublin ³Insight Centre for Data Analytics ⁴Teagasc, ⁵VistaMilk paul.albert@insight-centre.org

Abstract

Herbage mass yield and composition estimation is an important tool for dairy farmers to ensure an adequate supply of high quality herbage for grazing and subsequently milk production. By accurately estimating herbage mass and composition, targeted nitrogen fertiliser application strategies can be deployed to improve localised regions in a herbage field, effectively reducing the negative impacts of over-fertilization on biodiversity and the environment. In this context, deep learning algorithms offer a tempting alternative to the usual means of sward composition estimation, which involves the destructive process of cutting a sample from the herbage field and sorting by hand all plant species in the herbage. The process is labour intensive and time consuming and so not utilised by farmers. Deep learning has been successfully applied in this context on images collected by high-resolution cameras on the ground. Moving the deep learning solution to drone imaging, however, has the potential to further improve the herbage mass yield and composition estimation task by extending the ground-level estimation to the large surfaces occupied by fields/paddocks. Drone images come at the cost of lower resolution views of the fields taken from a high altitude and requires further herbage ground-truth collection from the large surfaces covered by drone images. This paper proposes to transfer knowledge learned on groundlevel images to raw drone images in an unsupervised manner. To do so, we use unpaired image style translation to enhance the resolution of drone images by a factor of eight and modify them to appear closer to their groundlevel counterparts. We then use the enhanced drone images to train a semi-supervised algorithm that uses groundtruthed, ground-level images as the labelled data together



8 fold upsampling, deblurring and visual domain transfer

Figure 1. Up-sampling drone images by a factor of 8. Images at the top are 64×64 crops from drone images. Images at the bottom are up-sampled to 512×512 , deblurred and transferred to the ground-level visual domain in an unpaired fashion. We use the transformed images in a semi-supervised regression objective.

with a large amount of unlabeled drone images. We validate our results on a small held-out drone image test set to show the validity of our approach, which opens the way for automated dry herbage biomass monitoring www.github. com/PaulAlbert31/Clover_SSL.

1. Introduction

Nitrogen fertilization has proven to be efficient in enhancing herbage quantity and quality, yet over-fertilization has detrimental effects on biodiversity and on the environment in general [4, 26, 35]. In this context, clover proves to be an important ally to the farmer for two reasons. First, clover naturally captures widely available nitrogen from the atmosphere and renders it available in the soil for the grass to use [40, 49]. Second, having proper amounts of clover in the feed has been shown to increase cow appetite, which in turn translates to higher milk production [15, 34]. Monitoring clover content in the herbage then becomes an important aspect of milk production and regular herbage biomass probing is performed by humans to ensure a proper grass to clover balance. The herbage probing process involves cutting a sample from the field, drying it in lab before manually separating each component of the herbage by hand [15]. Knowing the herbage composition and dry mass is valuable for the farmer but because existing probing processes are destructive and time consuming it is never probed at the farm level. In this context, deep learning has the capacity to provide a simpler, non-destructive alternative to dry herbage phenotyping and mass estimation from images alone. The feasibility of the method has been shown in previous works, often relying on partially labeled or unlabeled images to reduce the strain of the lengthy data collection process [2, 36, 47]. These works, however, only studied the application of deep learning algorithms to ground-level images using handheld devices and tripods [19] or all terrain vehicles (ATV) [46]. In this paper, we propose to extend the dry biomass and herbage mass estimation problem to drone images, which are more suitable for covering large herbage fields. Because drones operate at higher altitudes, large land areas that can span from tens to hundreds of square meters depending on the altitude are captured in every drone image, rendering the fine ground-truthing of data very challenging. To mitigate this issue, we propose to transfer knowledge learned from few high-resolution ground-level images to drone images in an unsupervised manner. To do so, we apply an unpaired domain transfer algorithm [42] to the drone images to enhance their resolution to 2048×2048 and to reduce the visual domain gap with the ground-level images (see Figure 1). We then train a semi-supervised neural network for regression on a small number of labeled groundlevel images together with unlabeled drone images to effectively transfer knowledge between the two domains. To evaluate the quality of our regression algorithm, we test it on a small data set of ground-truthed drone images collected in Ireland and evaluate the benefit the large quantities of unlabeled images drone imagery provides to improve the ground-level predictions. Our contributions are:

- 1. 328 drone images of herbage fields in Ireland;
- 2. An unpaired image transfer pipeline, increasing the resolution of drone images 8 fold and transferring them to the ground-level camera visual domain;

3. A semi-supervised regression that learns to estimate dry herbage biomass from a small set of annotated ground-level images and unlabeled drone images.

2. Related work

2.1. Computer vision for agriculture

Computer vision offers possibilities to revolutionize smart agriculture by providing farmers with automated solutions to address deficiencies in their fields and also to drive efficiencies in their daily practice. Weed detection, for example, is a topic that received significant attention where undesirable weeds are automatically detected in a field using image analysis. The weed detection process encompasses simple edge detection or color filtering approaches [41, 44, 50]; random forest classifiers trained on color features [21]; or more recently semantic segmentation neural networks [29]. Other popular phenotyping tasks include fruit detection and counting in trees [1, 9, 45] or wheat head identification [13]. Generative adversarial networks (GANs) are of special interest in the plant domain. We separate here GANs architectures between the conditional architectures that are trained with pairs of input and outputs in the two different visual domains [22] and unpaired architectures, i.e. CycleGAN [55] or Contrastive Unpaired Translation [42] where images from both visual domains are not semantically linked. Conditional GANs have been successfully applied to generate RGB images from semantic segmentation masks [56], to predict cabbage growth [14], or plant super-resolution to improve feature detection [12]. Unpaired GANs have been used to estimate disease spreading on leaves [31, 37, 38] or to improve the realism of synthetic images [6, 18]. Finally, drone (UAV) imaging holds important potential for automating farm tasks since drones can easily cover large areas of uneven terrain [53]. Deep learning has been successfully applied to derive growth rate from nitrogen fertilization on drone images [20], estimate the emergence rate of seeds in the field [33], wheat density [25], weed detection [16], and land classification [11]. The main drawback when applying deep learning on drone images for plant phenotyping remains the difficulty of ground-truthing the images because of the large areas covered [53].

2.2. Biomass composition prediction from canopy images

Herbage biomass composition from images gained traction after the publication of the GrassClover image dataset for semantic and hierarchical species understanding in agriculture [46,47]. To solve the biomass composition problem, Skovsen *et al.* [47] propose to create artificial images where grass/clover/weeds elements are manually cropped from the raw images and pasted in a random fashion on a soil background image to create a synthetic but fully segmented image. A semantic segmentation network is then trained on the synthetic data and predicts species pixel percentages from the real RGB images and a least-squares regression algorithm predicts the dry biomass from the pixel percentages. The Irish grass clover dataset [19] proposes, additionally to the dry biomass percentages, to predict the herbage height pre-grasing (cm) and the dry matter per hectare (kg DM/ha) from the canopy images. Although both datasets provide an additional large amount of unlabeled images, the respectively baseline are purely supervised and do not make use of the raw images. Subsequent algorithms were published and tested on both datasets to attempt to use the raw data to improve the biomass prediction. Narayannan *et al.* [36] proposed to use mean imputation to infer labels for the partially labeled samples before training a convolutional neural network (CNN) on the larger dataset. Albert et al. [2] generate synthetic semantic segmentation images in a similar fashion to Skovsen et al. [47] but instead of using the linear regression algorithm to predict at test time, the regressor is used to automatically label raw images. The automatically labeled data is then used together with the few groundtruthed images to train a regression neural network robust to label noise. Another work by Albert *et al.* [3] instead uses an unsupervised learning algorithm [30] on the unlabeled data to learn better initial representations that allow for better accuracy numbers with limited amounts of labels. Finally, Skovsen et al. [48] published an updated version of their segmentation algorithm from synthetic data using style transfer GANs to simulate different weather conditions and multi-resolution prediction.

2.3. Semi-supervised regression from images

Semi-supervised regression (SSR) solves a regression task on a dataset where the labeled data is limited but the unlabeled data is plentiful. Although semi-supervised classification received many important contributions in the last years, the attention given to SSR has been limited. Timilisina et al. [51] construct a fully connected graph from the feature representations of every sample before performing a bounded heat diffusion process to annotate the unlabeled data. Jean et al. [23] adopt a Bayesian approach by fitting the labeled representations with gaussian processes and training an auxiliary regularization objective to minimize the predictive variance with regards to the unlabeled points. Bzdok et al. [10] apply an autoencoder on top of medical images of brain voxels to solve a action regression task. The autoencoder is used to compress the input vectors and to ensure that the features extracted from labeled and unlabeled images will be compatible with the end logistic regressor. Li et al. [32] propose a process to aggregate the predictions from multiple regression predictions into a safe pseudo label for the unlabeled samples by means of solving a convex linear combination of each regressor output. Zhou *et al.* [54] co-train two KNN regressors with different distance metrics that predict pseudo labels to be used by the other regressor on the unlabeled data, effectively reducing confirmation bias. Note that semi-supervised classification algorithms such as consistency regularization approaches [8, 52] or pseudo-labeling [5] should translate to the regression setting.

3. Unsupervised domain adaptation and super resolution on drone images

We aim to solve the biomass prediction task jointly from a small set of ground-level images \mathcal{X}_l with biomass labels \mathcal{Y}_l (ground-level images) together with a large set of unlabeled (raw) images \mathcal{X}_u from a different visual domain (drone images) in an unsupervised fashion. To do so, we use two neural networks: Ψ performing super resolution and visual shift from the domain of \mathcal{X}_u to \mathcal{X}_l and Φ , a regression network we use to learn jointly from \mathcal{X}_l and \mathcal{X}_u by optimizing a semi-supervised objective.

3.1. Dataset presentation

We consider here three different herbage biomass estimation datasets. The first one is the publicly available GrassClover dataset [46]. This dataset is composed of 157 annotated images (to be divided between training set and validation set) and 31.600 unlabeled images. The image acquisition was carried out in Danish fields between 2017 and 2018 using for the most part an ATV mounted camera. The ground-truth collected is composed of the dry biomass percentages for the grass, white clover, red clover, total clover and weeds. The second dataset is the Irish clover dataset [19], which is composed of 424 training images, 104 held out test images, and 594 unlabeled images. The images were captured in the south of Ireland in the Summer of 2020 using a camera mounted on a tripod. The ground-truth collected is composed of the dry biomass percentages for grass, total clover and weeds (%), the herbage height (cm), and the herbage dry matter per ha (kg DM/ha). Finally, we propose in this paper an extension of the Irish dataset where we collect drone images in the same 23 herbage paddocks originally studied in Ireland in late Autumn of 2021. We collect between 36 and 7 drone images per paddock at an altitude between 6 and 12 meters. The drone we use is the DJI Mavic 2 Pro¹ with its default camera, taking pictures at a resolution of 5472×3648 . Although our drone is not capable of capturing its altitude relative to the land below, we subtract the above sea level GPS altitude of the drone from the land altitude at the associated GPS coordinates to obtain an approximate relative altitude using an open source

¹https://www.dji.com/ie/mavic-2



Figure 2. Drone image cropping process at different altitudes. Given that resolution of the image is fixed, we increase or reduce the cropped area. All crops are then bicubicly upscaled to 2048×2048 before deblurring.

API². We obtain 328 drone images in total with their associated altitude. Because of the huge areas covered by drone images, the ground-truth we collect is limited to the dry herbage mass at the paddock level and we omit the grass height and biomass percentage information. The resulting 80 labeled drone images are only used as a means to test the knowledge transfer from the ground-level to the drone images and not used for training. We propose two ground-truth estimations for the drone images: the first is a visual estimation performed on site at the time of the image collection by two human experts, very familiar with the site and that visually estimate the herbage on site every week. The second is obtained following the protocol of Egan et al. [15]: we cut two 1.2×8 meters strips in the paddocks 4 cm above ground level (typical cow grazing height) using an Etesia lawn mower (Etesia UK. Ltd., Warwick, UK). A 100 grams sample is collected from the cut material and dried at 95°C for 16 hours to obtain the dry herbage mass. We compare our algorithm against the human estimation and the exact ground-truth.

3.2. Contrastive Unpaired Translation (CUT)

The first step of our algorithm is to increase the resolution of drone images and to modify them to appear visually closer to the few ground-truthed images captured using high resolution cameras on the ground. To do so, we use Contrastive Unpaired Translation (CUT) [42]. CUT trains an adversarial network (GAN) to perform unpaired image style transfer using three principal components. G is the generator part of the network, competing to fool D the discriminator in an alternative adversarial optimization and Fthe projection head is used to optimize the contrastive part of the algorithm, which promotes semantic similarities between the same image before and after the visual transformation. CUT minimizes a combination of three losses to learn the parameters for G, D, and F. First the adversarial loss [17]

$$\mathbf{L}_{adv}(G, D, \mathcal{X}_l, \mathcal{X}_u) = \mathbb{E}_{x_l \sim \mathcal{X}_l} \log D(x_l) \\ + \mathbb{E}_{x_u \sim \mathcal{X}_u} (1 - \log D(G(x_u))), \quad (1)$$

promotes the generator G to transform images from \mathcal{X}_u (the drone images) so that they become indistinguishable by the discriminator D from the high resolution ground-level images in \mathcal{X}_l . Second, once the image has been transformed by the generator, a patch contrastive regularization objective is applied where patches at the same location in the image before and after the transformation are encouraged to have similar features after projection through F while being dissimilar to any other random patch from the image. This results in a constrastive patch objective

$$\mathbf{L}_{patch}(G, F, X) = -\frac{1}{P} \sum_{i=1}^{P} \log\left(\frac{\exp\left(ip(p_i, p'_i)/\tau\right)}{\sum_{k=1}^{P} \exp\left(ip(p_k, p'_i)/\tau\right)}\right)$$
(2)

where P = 64 random patches are cropped out from the input image, their feature representations encoded through G (stopping half way), projected through F and L2 normalized. The process is repeated for the transformed version of the image to form P pairs of random patches $\{(p_i, p'_i)\}_{i=1}^{P}$ for a given image $x \in \mathcal{X}_u$ where p_i is the representation before the domain shift and p'_i after. The dot product between the representations of corresponding pairs is encouraged to be close to one and close to zero for different patches. L_{patch} can also be applied to images in \mathcal{X}_l to enforce that Gwill perform the identity operation, *i.e.* $\forall x \in \mathcal{X}_l, G(x) = x$.

²opentopodata.org



Figure 3. Overview of our up-sampling and knowledge transfer algorithm. We use an up-sampling and visual domain transfer network Ψ and a semi-supervised network Φ that we use to learn jointly from few labeled ground-level examples (N = 54) and unlabeled drone images.

 $\tau=0.07$ is the contrastive temperature parameter. The final objective minimized by CUT where $\lambda_1=\lambda_2=0.5$ is

$$\mathbf{L} = \mathbf{L}_{adv}(G, D, \mathcal{X}_l, \mathcal{X}_u) + \lambda_1 \mathbf{L}_{patch}(G, F, \mathcal{X}_u) + \lambda_2 \mathbf{L}_{patch}(G, F, \mathcal{X}_l).$$
(3)

3.3. CUT for super resolution and style transfer

Cropping the drone images. We propose to crop squared areas from the drone images to obtain similar amounts of elements per image as ground level images. Because the drone images were not all captured at the same altitude, we adjust the area cropped out from the drone images depending on the altitude at which the image was captured. We observe visually that at an altitude of 8 meters, a 256×256 pixel crop of the drone data yields similar numbers of grass elements and of similar size to the ground-level images. Given the altitude of the drone at the time the picture was taken, we multiply the edge of the crop by the ratio between the altitude and the standard value of 6 meters *i.e.* for an altitude of 12 meters, the edge of the square crop will be $6/12 \times 256 = 128$. Figure 2 illustrates the height adjusted cropping process. This process allows us to capture

the same area of land independently of the height of the drone.

Deblurring the crops Although CUT is originally designed to transfer styles between two unpaired visual domains, we propose here to task the algorithm with improving the resolution of the drone images while at the same time transferring their visual style to ground-level images. Note that the super-resolution task is usually performed by conditional GANs (e.g. [39]) but we propose here to use an unpaired algorithm. For each image $x \in \mathcal{X}_u$, we upscale the image from the original resolution to 2048×2048 . Ψ is then trained to transfer the visual style of the ground-level high resolution images to the up-sampled crops, effectively deblurring them to appear closer to the higher resolution images (see Figure 1).

3.4. Semi-supervised regression on drone data

By up-sampling and visually transforming drone images to appear closer to the ground-level visual domain, we are now able to learn jointly from X_l and X_u . Since it is only practical to obtain labels for ground-level camera images, we propose to optimize a semi-supervised regression objective using \mathcal{X}_l as the labeled set and \mathcal{X}_u as the unlabeled data. After an initial pretraining of Φ on \mathcal{X}_l , we start guessing biomass labels for X_u using a consistency regularization approach [8]. Using two data augmented views x'_u and x''_u (vertical and horizontal random flipping), we use an exponential moving average (EMA) on the weights of Φ to guess two approximate biomass labels y'_u and y''_u for x_u . Rather than averaging the two approximate labels with equal importance like consistency regularization algorithms for image classification [7, 8], we draw a random mixing parameter λ from a uniform distribution to improve the regularization of the predictions and avoid confirmation bias [5]. We obtain an approximate label $\tilde{y} = \lambda y'_u + (1 - \lambda)y''_u$ for every unlabeled images. We enforce the distribution of the predictions on unlabeled samples to match the observed ground-truth distribution on the labeled data (distribution alignment) by multiplying the label prediction by the ratio between a sliding window average (50 mini-batches in practice) of \tilde{y} and the observed distribution on the labeled data. EMAs and distribution alignment are common principles of consistency regularization algorithms for semi-supervised learning [7, 24]. Finally, we normalize the biomass composition to sum to 1 in the approximate label \tilde{y} . Figure 3 presents an overview of the proposed semi-supervised training algorithm

3.5. Regression from images

We predict biomass labels from images in \mathcal{X}_l and \mathcal{X}_u using Φ to extract visual features. For the Irish dataset [19], we use three different linear heads, separating the predictions of the herbage mass, herbage height, and biomass composition. We normalize the herbage mass and herbage height values between 0 and 1 using fixed normalization values (4000 kg DM/ha for the herbage mass and 20cm for the height) and offset them by +0.2 to improve the prediction for low values originally too close to 0. To obtain values between zero and one for each target prediction and ensure that the sum of the biomass percentages equals one, we apply a softmax function on the three outputs from the biomass head and a sigmoid function for each of the other two values. For the GrassClover dataset [47], we use a single linear head, predicting the biomass percentages (grass, white clover, red clover, weeds) and sum the predictions for white and red clover to obtain the total clover content. These configurations follow the work of Albert et al. [2]. We use the root mean squared error (RMSE) as the training objective.

4. Experiments

4.1. Experimental setup

We conduct experiments on two biomass prediction datasets from canopy images. For the GrassClover dataset,



Figure 4. Overview of the deblurring effect of Ψ on drone data. A high variance indicates a sharper image.

we use 100 labeled and 1,000 unlabeled images for training and 57 images for validation. We report the test accuracy results on the evaluation server for the GrassClover dataset. For the Irish dataset, we use 52 labeled and 595 unlabeled images for training and 104 images for validation. For the drone images, we extract 50 random crops from each of the 328 images to create a dataset of 16, 400 unlabeled images. We train using stochastic gradient descend at a resolution of 512×512 with a batch size of 32 and a fixed learning rate of 0.03. We update the EMA with a multiplication parameter of 0.99 at every mini-batch. The training augmentations are resize, random crop, random horizontal and vertical flipping, and normalization. When we perform semisupervised learning, we create each mini-batch by aggregating 4 labeled samples with unlabeled images as in Albert et al. [2]. For the neural networks, we use a ResNet18 [27] pretrained on ImageNet [28] for the regression network Φ and the 9 ResNet blocks version of the CUT model Ψ .

4.2. Drone image deblurring and style transfer

We evaluate the deblurring capacity of Ψ by computing the variance of the Laplacian on the grayscale view of an image. Computing the Laplacian of the image allows us to extract edges in the image and the variance of the resulting value quantifies the sharpness of the edges: sharpness = $var(\nabla^2(grayscale(image)))$ [43] with ∇^2 the Laplacian operator. A higher variance indicates a bet-

	HRMSE				RMSE						
	Total	Grass	Clover	Weeds	Avg.	HRE	Grass	Clover	Weeds	Avg.	HE
Albert <i>et al.</i> [2] Albert <i>et al.</i> [3]	230.10 229.12	220.84 218.02	34.86 37.65	27.13 29.21	94.28 94.96	1.14 1.09	4.81 4.58	4.75 4.22	3.42 3.44	4.33 4.08	2.15 2.03
Labeled only Semi-sup + distribution alignment + EMA	229.23 234.50 220.79 217.28	268.90 224.10 215.88 208.96	107.39 43.03 40.00 34.16	39.82 26.74 26.78 26.50	138.71 97.96 94.22 89.88	1.08 1.04 1.08 1.08	17.15 5.85 5.53 4.86	14.08 5.51 5.51 4.73	4.74 3.19 3.25 3.26	11.99 4.85 4.76 4.28	2.28 2.24 2.09 2.09
+ Self-sup [30] drone unlabeled	217.20 211.57 209.69	203.90 202.18 199.61	28.93 33.59	26.80 27.13	85.97 86.78	1.00 1.09 1.02	4.54 4.74	4.50 4.65	3.21 3.33	4.08 4.24	2.09 2.18

Table 1. Ablation study and comparison against state-of-the-art algorithms on the Irish dataset. The last row denotes replacing the ground-level unlabeled camera images with deblurred drone images. The best results are in bold.

	HRMSE	HRE				
Against harvested ground-truth						
Labeled only	1094.18	0.43				
Semi-sup.	566.68	0.81				
Semi-sup. drone	219.15	0.97				
Human expert	170.03	1.04				

Table 2. Results on drone images. Errors are computed against the absolute harvested ground-truth.



Figure 5. Visualization of the HRE on the validation set of the Irish dataset with 95% confidence intervals.

ter defined (sharper) image. The sharpness estimation process is illustrated in Figure 4. We observe that the variance averaged over all the crops changes from 5.32 when cropping directly from the drone images to 1261.05 for the same images deblurred by Ψ .

4.3. Semi-supervised biomass, herbage height and herbage mass prediction

We evaluate the capacity of our algorithm to predict the biomass composition of herbage (%) together with an estimation of the dry herbage mass (kg DM/ha) and the grass

height (cm) in a semi-supervised manner. We run experiments on the Irish dataset with the original unlabeled images but also study replacing them with an equal number (N = 596) of deblurred drone images. This is to evaluate the capacity for drones to capture unlabeled data that can be used to improve the prediction at the groundlevel. For evaluation purposes, we compute the Herbage Root Mean Square Error (HRMSE) which is the RMSE between predicted and ground-truth herbage mass for grass, clover, weeds and for the total mass and the Herbage Relative Error (HRE) which is the ratio between the groundtruth value of the total herbage mass and the prediction: $HRE = \frac{pred_{herbage}}{gt_{herbage}}$. The HRE measure is typically used to compare human visual estimation against the collected ground-truth [15]. We additionally compute the RMSE over the predicted percentages of grass, clover, weeds in the herbage and the Height Error (HE) which is the RMSE between the predicted herbage height and ground-truth height. We compare against state-of-the-art results on the Irish dataset where the validation images are ground-level images in Table 1. We study the importance of the different elements of the semi-supervised algorithm on the validation error, including enforcing distribution alignment for the label guesses and using an exponential moving average (EMA) on the weights of the semi-supervised network. We also report results when initializing the weights of the network using an unsupervised representation learning algorithm [30] on the unlabeled data as in Albert *et al.* [3]. We finally point out that using an equal number of deblurred drone images produces comparable results to using the original unlabeled images (last two rows in Table 1). This result motivates the use of drone images to easily capture large amounts of unlabeled images. Figure 5 shows a line plot of the HRE compared against the ground-truth herbage mass where we observe that the algorithm struggles the most on high or low herbage mass outliers (< 500 to > 2000 kg DM/ha). This is most likely due to the low amount of high or low herbage mass examples seen during training.

HRMSE							
1	5	20	50	100	all		
252.71 ± 51.24	227.18 ± 20.15	222.97 ± 12.56	219.15 ± 6.93	220.09 ± 4.14	219.53		

Table 3. RMSE errors on drone image for varying amounts of random crops per image for the best performing model. Averaged over 5 random sets of crops.

	Grass	Total	White	Red	Weeds	Avg.
Skovsen et al. [47]	9.05	9.91	9.51	6.68	6.50	8.33
Naranayan et al. [36]	8.64	8.73	8.16	10.11	6.95	8.52
Albert et al. [2]	8.78	8.35	7.72	7.35	7.17	7.87
Labeled only	9.81	8.49	7.99	8.58	7.25	8.57
Semi-sup.	6.68	7.76	8.08	8.66	6.72	7.58

Table 4. Results on the GrassClover test set (RMSE). Lowest errors are in bold.

4.4. Prediction on drone images

Table 2 reports the Herbage Root Mean Square Error (HRMSE) when predicting on drone data without the need to gather additional labels. First, we evaluate the accuracy of a CNN model learned on ground-level data only to predict on deblurred drone patches using Ψ . We then evaluate the accuracy benefits of training on deblurred drone patches in a semi-supervised manner for the herbage mass prediction and compare the error rate of our algorithm against human experts. We observe a significant reduction in error rates when the drone images are used in the semi-supervised objective (semi-sup. drone row) over using only the ground-level data (labeled only or semi-sup. with the unlabeled ground-level images). The performance proposed by our low supervision algorithm is close to be on par with human experts at the paddock level.

Table 3 reports on how augmenting the number of random crops improves the prediction of the best performing model. We report the average herbage mass error and standard deviation over 5 random sets of crops. "All" denotes cropping the image in a checkerboard fashion and using all crops (from 250 to 1,000 crops per image depending on the altitude). Although 1 random crop per image yield interestingly good results, we validate our choice of 50 crops per images for more stable predictions.

4.5. Semi-supervised biomass prediction on the GrassClover dataset

We compare our semi-supervised approach against stateof-the-art algorithms on the publicly available GrassClover dataset in Table 4 where we report RMSE errors for the biomass percentage prediction on the held out test set³ where we perform on par with existing approaches.

5. Conclusion

This paper investigates how to extend the biomass estimation and herbage mass prediction problem from groundlevel to drone images. By its nature, the herbage biomass information of drone images is hard to annotate finely because of the huge areas covered. Ground-level data, however, has the advantage of providing easier to acquire, finely annotated, and high resolution images of the herbage but would not be a good solution to generalize to targeted fertilization on entire herbage fields. To successfully transfer knowledge from the ground-truth images to the drone data, we propose to train an unpaired style transfer algorithm to deblur height adjusted crops of drone images, increase resolution by a factor of 8 and to transfer the visual style of ground-level images, captured using different cameras, to look more similar to their ground-level counterparts. The large set of transformed unlabeled drone images is used together with the finely annotated ground-level images to learn unsupervised initialization weights and to train a semi-supervised regression algorithm. The neural network trained on the partially labeled set largely improves the regression accuracy on the ground-level data and the herbage mass prediction on drone images. We significantly reduce the prediction gap compared to a human expert, achieving error rates close to experienced technicians familiar with the land with little extra annotation cost than few ground-level images. This challenge prompts further research in the field of low supervision computer vision for herbage biomass prediction from drone images. Future work would involve using the approximate predictions made by the human experts to improve the results and a smart algorithm capable of sampling areas of interest in the drone images to be predicting upon by the herbage biomass algorithm.

Acknowledgements

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under grant number SFI/15/SIRG/3283, SFI/16/RC/3835 and SFI/12/RC/2289_P2 and the computing support of the Irish Centre for High End Computing (ICHEC).

³https://competitions.codalab.org/competitions/21122

References

- [1] Manya Afonso, Hubert Fonteijn, Felipe Schadeck Fiorentin, Dick Lensink, Marcel Mooij, Nanne Faber, Gerrit Polder, and Ron Wehrens. Tomato Fruit Detection and Counting in Greenhouses Using Deep Learning. *Frontiers in plant science*, 2020. 2
- [2] Paul Albert, Mohamed Saadeldin, Badri Narayanan, Brian Mac Namee, Deirdre Hennessy, Aisling H. O'Connor, Noel E. O'Connor, and Kevin McGuinness. Semi-supervised dry herbage mass estimation using automatic data and synthetic images. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2, 3, 6, 7, 8
- [3] Paul Albert, Mohamed Saadeldin, Badri Narayanan, Brian Mac Namee, Noel E. O'Connor, Deirdre Hennessy, Aisling H. O'Connor, and Kevin McGuinness. Using image analysis and machine learning to estimate sward clover content. In *European Grassland Federation Symposium*, 2022. 3, 7
- [4] Francisco Albornoz. Crop responses to nitrogen overfertilization: A review. *Scientia horticulturae*, 2016. 2
- [5] E. Arazo, D. Ortego, P. Albert, N.E. O'Connor, and K. McGuinness. Pseudo-Labeling and Confirmation Bias in Deep Semi-Supervised Learning. In *International Joint Conference on Neural Networks (IJCNN)*, 2020. 3, 6
- [6] Ruud Barth, Jochen Hemming, and Eldert J Van Henten. Optimising realism of synthetic images using cycle generative adversarial networks for improved part segmentation. *Computers and Electronics in Agriculture*, 2020. 2
- [7] D. Berthelot, N. Carlini, E. Cubuk, A. Kurakin, K. Sohn, H. Zhang, and C. Raffel. ReMixMatch: Semi-Supervised Learning with Distribution Matching and Augmentation Anchoring. In *International Conference on Learning Representations (ICLR)*, 2020. 6
- [8] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 3, 6
- [9] Kushtrim Bresilla, Giulio Demetrio Perulli, Alexandra Boini, Brunella Morandi, Luca Corelli Grappadelli, and Luigi Manfrini. Single-shot convolution neural networks for real-time fruit detection within the tree. *Frontiers in plant science*, 2019. 2
- [10] Danilo Bzdok, Michael Eickenberg, Olivier Grisel, Bertrand Thirion, and Gaël Varoquaux. Semi-supervised factored logistic regression for high-dimensional neuroimaging data. Advances in neural information processing systems (NeurIPS), 2015. 3
- [11] Mang Tik Chiu, Xingqian Xu, Yunchao Wei, Zilong Huang, Alexander G. Schwing, Robert Brunner, Hrant Khachatrian, Hovnatan Karapetyan, Ivan Dozier, Greg Rose, David Wilson, Adrian Tudor, Naira Hovakimyan, Thomas S. Huang, and Honghui Shi. Agriculture-Vision: A Large Aerial Image Database for Agricultural Pattern Analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2

- [12] Qiang Dai, Xi Cheng, Yan Qiao, and Youhua Zhang. Crop leaf disease image super-resolution and identification with dual attention and topology fusion generative adversarial network. *IEEE Access*, 2020. 2
- [13] Etienne David, Simon Madec, Pouria Sadeghi-Tehran, Helge Aasen, Bangyou Zheng, Shouyang Liu, Norbert Kirchgessner, Goro Ishikawa, Koichi Nagasawa, and Minhajul A Badhon. Global Wheat Head Detection (GWHD) dataset: a large and diverse dataset of high-resolution RGB-labelled images to develop and benchmark wheat head detection methods. *Plant Phenomics*, 2020. 2
- [14] Lukas Drees, Laura Verena Junker-Frohn, Jana Kierdorf, and Ribana Roscher. Temporal prediction and evaluation of Brassica growth in the field using conditional generative adversarial networks. *Computers and Electronics in Agriculture*, 2021. 2
- [15] Michael Egan, Norann Galvin, and Deirdre Hennessy. Incorporating white clover (Trifolium repens L.) into perennial ryegrass (Lolium perenne L.) swards receiving varying levels of nitrogen fertilizer: Effects on milk and herbage production. *Journal of Dairy Science*, 101(4):3412–3427, 2018. 2, 4, 7
- [16] Aaron Etienne and Dharmendra Saraswat. Machine learning approaches to automate weed detection by UAV based sensors. In Autonomous Air and Ground Sensing Systems for Agricultural Optimization and Phenotyping, 2019. 2
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. Advances in neural information processing systems (NeurIPS), 2014. 4
- [18] Zane KJ Hartley and Andrew P French. Domain Adaptation of Synthetic Images for Wheat Head Detection. *Plants*, 2021. 2
- [19] Deirdre Hennessy, Mohamed Saadeldin, Brian Mac Namee, Noel E. O'Connor, Kevin. McGuinness, Paul Albert, Badri Narayanan, and Aisling H. O'Connor. Using image analysis and machine learning to estimate sward clover content. In *European Grassland Federation Symposium*, 2021. 2, 3, 6
- [20] Fenner H Holman, Andrew B Riche, Adam Michalski, March Castle, Martin J Wooster, and Malcolm J Hawkesford. High throughput field phenotyping of wheat plant height and growth rate in field plot trials using UAV based remote sensing. *Remote Sensing*, 2016. 2
- [21] N Islam, MM Rashid, S Wibowo, S Wasimi, A Morshed, C Xu, and S Moore. Machine learning based approach for Weed Detection in Chilli field using RGB images. In *International Conference on Fuzzy Systems and Knowledge Discovery (FSKD)*, 2020. 2
- [22] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision* and Pattern Recognition, 2017. 2
- [23] Neal Jean, Sang Michael Xie, and Stefano Ermon. Semisupervised deep kernel learning: Regression with unlabeled data by minimizing predictive variance. Advances in Neural Information Processing Systems (NeurIPS), 2018. 3
- [24] Lu Jiang, Di Huang, Mason Liu, and Weilong Yang. Beyond Synthetic Noise: Deep Learning on Controlled Noisy Labels.

In International Conference on Machine Learning (ICML), 2020. 6

- [25] Xiuliang Jin, Shouyang Liu, Frédéric Baret, Matthieu Hemerlé, and Alexis Comar. Estimates of plant density of wheat crops at emergence from very low altitude UAV imagery. *Remote Sensing of Environment*, 2017. 2
- [26] Xiaotang Ju, Xuejun Liu, Fusuo Zhang, and Marco Roelcke. Nitrogen fertilization, soil nitrate accumulation, and policy recommendations in several agricultural regions of China. AMBIO: a Journal of the Human Environment, 2004. 2
- [27] H. Kaiming, Z. Xiangyu, R. Shaoqing, and S. Jian. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [28] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems (NeurIPS), 2012. 6
- [29] Petre Lameski, Eftim Zdravevski, Vladimir Trajkovik, and Andrea Kulakov. Weed detection dataset with RGB images taken under variable light conditions. In *International Conference on ICT Innovations*, 2017. 2
- [30] Kibok Lee, Yian Zhu, Kihyuk Sohn, Chun-Liang Li, Jinwoo Shin, and Honglak Lee. i-Mix: A Domain-Agnostic Strategy for Contrastive Representation Learning. In *International Conference on Learning Representations (ICLR)*, 2021. 3, 7
- [31] Jie Li, Junjie Jia, and Donglai Xu. Unsupervised representation learning of image-based plant disease with deep convolutional generative adversarial networks. In 2018 37th Chinese control conference (CCC), 2018. 2
- [32] Yu-Feng Li, Han-Wen Zha, and Zhi-Hua Zhou. Learning safe prediction for semi-supervised regression. In AAAI Conference on Artificial Intelligence, 2017. 3
- [33] Tao Liu, Rui Li, Xiuliang Jin, Jinfeng Ding, Xinkai Zhu, Chengming Sun, and Wenshan Guo. Evaluation of seed emergence uniformity of mechanically sown wheat with UAV RGB imagery. *Remote Sensing*, 2017. 2
- [34] Anders K Mortensen, Henrik Karstoft, Karen Søegaard, René Gislum, and Rasmus N Jørgensen. Preliminary results of clover and grass coverage and total dry matter estimation in clover-grass crops using image analysis. *Journal of Imaging*, 2017. 2
- [35] F Nájera, Y Tapia, C Baginsky, V Figueroa, R Cabeza, and O Salazar. Evaluation of soil fertility and fertilisation practices for irrigated maize (Zea mays L.) under Mediterranean conditions in central Chile. *Journal of soil science and plant nutrition*, 2015. 2
- [36] Badri Narayanan, Mohamed Saadeldin, Paul Albert, Kevin McGuinness, and Brian Mac Namee. Extracting pasture phenotype and biomass percentages using weakly supervised multi-target deep learning on a small dataset. In *Irish Machine Vision and Image Processing conference*, 2020. 2, 3, 8
- [37] Haseeb Nazki, Jaehwan Lee, Sook Yoon, and Dong Sun Park. Image-to-image translation with GAN for synthetic data augmentation in plant disease datasets. *Smart Media Journal*, 2019. 2

- [38] Haseeb Nazki, Sook Yoon, Alvaro Fuentes, and Dong Sun Park. Unsupervised image translation using adversarial networks for improved plant disease recognition. *Computers* and Electronics in Agriculture, 2020. 2
- [39] Ben Niu, Weilei Wen, Wenqi Ren, Xiangde Zhang, Lianping Yang, Shuzhen Wang, Kaihao Zhang, Xiaochun Cao, and Haifeng Shen. Single image super-resolution via a holistic attention network. In *European conference on computer vi*sion (ECCV), 2020. 5
- [40] Daniel Nyfeler, Olivier Huguenin-Elie, Matthias Suter, Emmanuel Frossard, John Connolly, and Andreas Lüscher. Strong mixture effects among four species in fertilized agricultural grassland led to persistent and consistent transgressive overyielding. *Journal of Applied Ecology*, 2009. 2
- [41] Ajinkya Paikekari, Vrushali Ghule, Rani Meshram, and VB Raskar. Weed detection using image processing. *International Research Journal of Engineering and Technology (IR-JET)*, 2016. 2
- [42] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision* (ECCV), pages 319–345. Springer, 2020. 2, 4
- [43] José Luis Pech-Pacheco, Gabriel Cristóbal, Jesús Chamorro-Martinez, and Joaquín Fernández-Valdivia. Diatom autofocusing in brightfield microscopy: a comparative study. In *International Conference on Pattern Recognition (ICPR)*, 2000. 6
- [44] Camilo Andra Pulido-Rojas, Manuel Alejandro Molina-Villa, and Leonardo Enrique Solaque-GuzmÃin. Machine vision system for weed detection using image filtering in vegetables crops. *Revista Facultad de IngenierÃa Universidad de Antioquia*, 2016. 2
- [45] Inkyu Sa, Zongyuan Ge, Feras Dayoub, Ben Upcroft, Tristan Perez, and Chris McCool. Deepfruits: A fruit detection system using deep neural networks. *Sensors*, 2016. 2
- [46] Søren Skovsen, Mads Dyrmann, Jørgen Eriksen, René Gislum, Henrik Karstoft, and Rasmus Nyholm Jørgensen. Predicting dry matter composition of grass clover leys using data simulation and camera-based segmentation of field canopies into white clover, red clover, grass and weeds. In *International Conference on Precision Agriculture*, 2018. 2, 3
- [47] Soren Skovsen, Mads Dyrmann, Anders K Mortensen, Morten S Laursen, René Gislum, Jorgen Eriksen, Sadaf Farkhani, Henrik Karstoft, and Rasmus N Jorgensen. The GrassClover image dataset for semantic and hierarchical species understanding in agriculture. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019. 2, 3, 6, 8
- [48] Søren Kelstrup Skovsen, Morten Stigaard Laursen, Rebekka Kjeldgaard Kristensen, Jim Rasmussen, Mads Dyrmann, Jørgen Eriksen, René Gislum, Rasmus Nyholm Jørgensen, and Henrik Karstoft. Robust species distribution mapping of crop mixtures using color images and convolutional neural networks. *Sensors*, 2021. 3
- [49] Karen Søegaard. Nitrogen fertilization of grass/clover swards under cutting or grazing by dairy cows. Acta Agriculturae Scandinavica Section B–Soil and Plant Science, 2009.
 2

- [50] Jing-Lei Tang, Xiao-Qian Chen, Rong-Hui Miao, and Dong Wang. Weed detection using image processing under different illumination for site-specific areas spraying. *Computers* and Electronics in Agriculture, 2016. 2
- [51] Mohan Timilsina, Alejandro Figueroa, Mathieu d'Aquin, and Haixuan Yang. Semi-supervised regression using diffusion on graphs. *Applied Soft Computing*, 104:107188, 2021.
 3
- [52] Vikas Verma, Kenji Kawaguchi, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. In *International Joint Conference on Artifical Intelligence (IJCAI)*, 2019. 3
- [53] Michelle Watt, Fabio Fiorani, Björn Usadel, Uwe Rascher, Onno Muller, and Ulrich Schurr. Phenotyping: new windows into the plant for breeders. *Annual review of plant biology*, 2020. 2
- [54] Zhi-Hua Zhou, Ming Li, et al. Semi-supervised regression with co-training. In *International Joint Conference on Artifical Intelligence (IJCAI)*, 2005. 3
- [55] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycleconsistent adversarial networks. In *IEEE international conference on computer vision (ICCV)*, 2017. 2
- [56] Yezi Zhu, Marc Aoun, Marcel Krijn, Joaquin Vanschoren, and High Tech Campus. Data Augmentation using Conditional Generative Adversarial Networks for Leaf Counting in Arabidopsis Plants. In *British Machine Vision Conference* (*BMVC*), 2018. 2