

Energy-Efficient Uncertainty-Aware Biomass Composition Prediction at the Edge

Muhammad Zawish[†], Paul Albert[‡], Flavio Esposito[§], Steven Davy[¶], and Lizzy Abraham[†]

[†]Walton Institute, South East Technological University, Ireland

[‡]Center for Augmented Reasoning, Australian Institute of Machine Learning, Australia

[§]Department of Computer Science, Saint Louis University, USA

[¶]Centre for Sustainable Digital Technologies at Technological University Dublin, Ireland

{muhammad.zawish}@waltoninstitute.ie

Abstract

Clover fixates nitrogen from the atmosphere to the ground, making grass-clover mixtures highly desirable to reduce external nitrogen fertilization. Herbage containing clover additionally promotes higher food intake, resulting in higher milk production. Herbage probing however remains largely unused as it requires a time-intensive manual laboratory analysis. Without this information, farmers are unable to perform localized clover sowing or take targeted fertilization decisions. Deep learning algorithms have been proposed with the goal to estimate the dry biomass composition from images of the grass directly in the fields. The energy-intensive nature of deep learning however limits deployment to practical edge devices such as smartphones. This paper proposes to fill this gap by applying filter pruning to reduce the energy requirement of existing deep learning solutions. We report that although pruned networks are accurate on controlled, high-quality images of the grass, they struggle to generalize to real-world smartphone images that are blurry or taken from challenging angles. We address this challenge by training filter-pruned models using a variance attenuation loss so they can predict the uncertainty of their predictions. When the uncertainty exceeds a threshold, we re-infer using a more accurate unpruned model. This hybrid approach allows us to reduce energy-consumption while retaining a high accuracy. We evaluate our algorithm on two datasets: the GrassClover and the Irish clover using an NVIDIA Jetson Nano edge device. We find that we reduce energy reduction with respect to state-of-the-art solutions by 50% on average with only 4% accuracy loss.

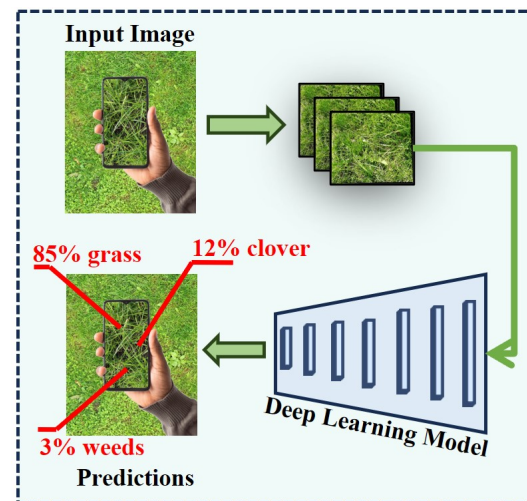


Figure 1. A typical workflow of deep learning based biomass estimation from phone / camera images.

1. Introduction

Evaluating the dry biomass composition of herbage delivers insights that can in turn be used by farmers to improve the quality of their grass. The usual biomass composition process requires cutting grass on a delimited surface in the field. The grass sample is then taken back to a laboratory where it is dried in a low-temperature oven before being manually separated in its components [12, 18]. This process is constrained by time and human labor which hinders its practical application. Computer vision techniques have the potential to reduce this laborious process by predicting dry biomass composition from in situ images. Using image processing via deep learning is also non-destructive for the soil, and a faster method to estimate the biomass can facili-

tate the adoption of fine-grained herbage probing. Figure 1 illustrates such a process when images are collected with smartphones.

Computer vision systems have been previously developed on a variety of images including high-resolution camera systems mounted on All Terrain Vehicles (ATVs) [19] or on manually operated camera tripods [12]. Other more practical approaches include using phones [3] or drone images [2]. These systems, however, were not designed to address the limited resources of the targeted edge devices, especially in terms of the energy consumption of large neural networks that can limit the granularity of biomass estimation on battery-powered devices.

Model compression, particularly structured (filter-based) pruning has proved to be a viable solution for reducing model size and computations [4]. It typically measures the saliency of filters in *conv* layers and prunes them according to a predefined threshold. The idea of mapping saliency with the loss function makes these approaches dependent on training data, hence limiting their practicability on specific computer vision tasks, such as image classification. Recent work has shown that even without measuring filter saliency, it is possible to achieve the same level of performance with random pruning [16]. Furthermore, pruning approaches in the literature mostly focus on pruning post-training to effectively cut down on inference costs, while inevitably increasing training and re-training overhead. Recently, a new method has emerged to prune models at initialization (before training) [9, 23], and show similar performances to the post-training pruning approaches. We find energy-aware pruning at *initialization* attractive for biomass estimation because although prior studies demonstrate the accuracy of deep learning models to solve the task, these large models are intended for deployment on energy-constrained edge devices like smartphones. Energy-aware pruning indeed offers the potential for large energy savings, however, end-users require reliable results to facilitate technology adoption. We aim in this paper to answer the need for low energy consumption while retaining biomass prediction accuracy.

When training state-of-the-art biomass composition estimation models [2] using initialization-pruned models, we observe a significant decrease in accuracy concentrated in a population of *harder* images. These images are clover-heavy or were taken with hand-held smartphones, leading to lower image quality and complicating the biomass estimation task. To preserve a high prediction accuracy, we propose to identify these hard images at test-time so that they can be re-predicted on using an unpruned model, inherently more robust to these challenging examples. We achieve this by training the pruned models to predict a normally distributed probability distribution over biomass values rather than predicting a single scalar biomass value. The probability distribution is characterized by a mean μ and a

variance σ . We find that σ is positively correlated with the prediction error of the pruned network on the hard images. At test-time and given the prediction of the pruned network on a target image, if we detect that σ exceeds a threshold, i.e., the prediction is under-confident, we offset such predictions using the unpruned network to increase accuracy. In the case where σ is small enough, μ is used as the energy-efficient and reliable predictions. In effect, during the deployment of our algorithm on edge devices, we can confidently save energy on simpler images and call for the more robust yet energy-hungry unpruned model for the detected challenging images.

Consequently, the proposed approach conserves better accuracy than the pruned model while yielding higher energy efficiency than the unpruned model.

In particular, in this paper, we present the following key contributions:

- We perform stochastic filter pruning at initialization as in [23], iteratively removing filters until the pruning quota is met, with high-energy *conv* layers being more likely to be pruned. We consequently present an evaluation of the detrimental effect of pruning rate on prediction accuracy.
- We introduce a uncertainty-guided approach to accommodate performance loss in pruned models. A high variance σ of the biomass probability distribution predicted by the pruned models informs the algorithm on the subset of harder images that need to be more accurately predicted by the unpruned model.
- We compare with uncompressed deep learning baselines from the state-of-the-art [2] using two network architectures [13, 17] on two grass biomass estimation datasets [12, 19] where we show we can maintain a comparable accuracy while decreasing energy usage.
- We measure the real-world energy efficiency of our algorithm on a NVIDIA Jetson Nano, a resource-constrained edge device where we realise a 40% to 60% energy consumption decrease depending on the network and dataset. We hope that this large decrease in energy consumption further motivates the deployment of biomass estimation algorithms to edge devices without compromising much on accuracy.

2. Related Work

This section presents into existing literature on dry biomass estimation from images and model compression. While significant research exists on model pruning and accurate biomass estimation using deep learning, we notice a notable gap in integrating both for biomass estimation at the edge.

2.1. Grass and clover dry biomass estimation

We first detail datasets and research relevant to grass and clover dry biomass estimation using computer vision. The GrassClover dataset [18, 19] proposes to predict the dry

biomass composition of grass from images taken from a high resolution camera mounted on an ATV. 179 images of grass-clover mixtures are collected, oven-dried and manually labeled to obtain the ground-truth. The baseline biomass estimation proposed by the authors is a two stage approach where a large semantic segmentation model trained on synthetic data is used to predict the species of each pixel in the grass canopy. Once the canopy composition is estimated, a linear regression infers the dry biomass composition. In a further iteration [20], the authors utilize a style transfer generative network [22] to generate data augmentations on the collected images to simulate different weather conditions. This effectively increases the variety of the training dataset and enables more robust predictions. Because the GrassClover dataset contains a set of partially annotated images, researchers have proposed to use a one stage regression from image to dry biomass using a mean-imputation approach using a VGG16 [17] neural network [12]. Other solutions utilize the semantic segmentation into linear separation baseline to approximately label unlabeled images so they can be introduced in the training set [1], adapt the MixMatch semi-supervised classification algorithm [6] to the regression problem [2] or to use them for self-supervised pre-training [3]. In addition to the GrassClover dataset, the IrishClover [12] dataset consists of 424 training images captured using cameras mounted on tripods and proposes a similar challenge of dry biomass estimation with added difficulty of unbalanced species representation and the presence of handheld phone images in the test set.

2.2. Lightweight neural networks for edge AI in agriculture

Generally, majority of the deep learning-based agricultural applications discussed above are cloud-centric. This means that during the test-time, the incoming data from mobile edge devices would be offloaded to a remote cloud server to perform the prediction process. For example, authors in [5] proposed a cloud-based AI solution for predicting biomass yield maps from apple orchards using images captured through UAV. Cloud has remained a viable solution for these tasks, since deep learning models are computationally-intensive and consume significant energy and storage capacity. However, it requires continuous and stable internet connection to approach cloud for deep learning tasks, which is hard to realize in remote farming environments. Recently, researchers have proposed to process the data directly on the edge devices to achieve fast responses while ensuring data privacy. For example, authors in [8] propose to directly classify the crop types using a deep learning model deployed on an embedded board over a UAV. Similarly, Camargo et al. [7] customized a lightweight version of ResNet model to fit on an NVIDIA Jetson Xavier board for automatic weed mapping. How-

ever, these customized lightweight approaches do not provide flexible and scalable models, since these models are only developed for specific tasks. For these edge applications, model compression approaches have recently emerged, particularly pruning to achieve compression from coarse (low compression) to fine (high compression) scales [4, 9, 10, 23]. In this paper, we implement filter pruning as a type of model compression to realize lightweight models for biomass estimation on different scales. We also show how to achieve energy efficiency on a NVIDIA Jetson Nano edge device while preserving accuracy using the best of both pruned and unpruned models.

3. Methodology

In this section, we first motivate the idea applying model compression to biomass composition estimation algorithms. We then show how to circumvent the performance loss of pruned models using our variance-guided approach.

3.1. Motivation

Biomass estimation using deep learning is generally carried out on a high-end remote GPU / CPU cloud servers [1]. This is because SoTA deep learning models are resource-intensive in nature, requiring significant amounts of storage and power consumption. Using cloud-based inference implies that input images captured through mobile phones need to be offloaded – via a wireless network – to the cloud. This, however, is not considered reliable for making real-time predictions, since continuous and stable internet connectivity is not guaranteed in remote agricultural farms. Moreover, offloading raw images from farms to cloud poses privacy concerns, which can be avoided by leveraging on-device processing capabilities of mobile edge devices. Numerous research efforts have been in place to accommodate deep models directly on mobile phones. Model compression, particularly filter pruning, has been widely used to enable on-device AI for several computer vision applications but at the cost of slight reduction in accuracy [10]. Moreover, most of the compression approaches are evaluated on image classification which raises questions about their functionality for other tasks. This paper proposes to apply model pruning at initialization on a vision-based regression task for biomass estimation. In the subsequent section, we further discuss the pruning approach.

3.2. Pruning at Initialization

Pruning in neural networks is typically of two types; *i) structured* and *ii) unstructured*. Early works on pruning favoured unstructured approach due to its capacity to provide high compression ratios at a finer granularity such as individual weight-level [10]. However, recent works have practically proven that unstructured sparsity merely achieves realistic acceleration due to (weight) index storage

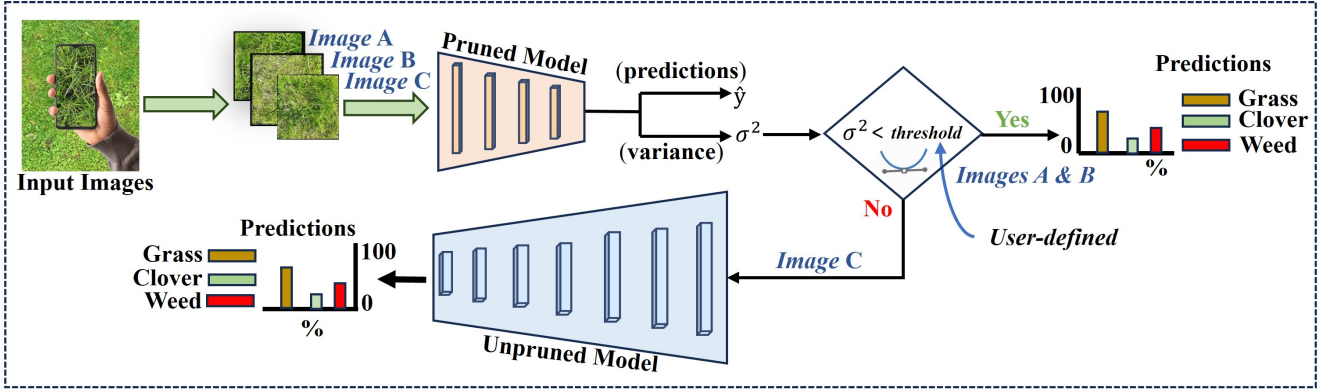


Figure 2. Our proposed framework based on variance-guided energy efficient deep learning for biomass estimation on edge.

overhead, and requires specialized libraries for matrix operation [15]. In contrast, removing whole filters/kernels from *convolution* layers is a structured form of pruning since it preserves the structure of the network. Therefore, filter pruning has gained attention in reducing the model size and computations while being implementation-friendly [4]. For this reason, we opt for filter pruning to induce structured sparsity in the baseline (unpruned) models.

Traditionally, filter pruning observes a three-stage procedure: – *train, prune and re-train* [4]. Although this setting is widely adopted, it introduces computational overhead characterized by multiple rounds of training and re-training. An alternative paradigm has however emerged namely the – Lottery Ticket Pruning, where pruning is conducted prior to training [9]. In the case of lottery ticket pruning, sparse models are created at initialization from a dense (unpruned) model without incurring a pre-training overhead and while realising similar accuracy guarantees [23]. Additionally, during the pruning phase, the saliency of filters is assessed using predefined criteria such as the l_1 -norm [14], enabling the removal of filters with saliency scores below a certain threshold. In contrast, “randomly” pruning filters has recently been proposed to also match the performance of saliency-based approaches, without introducing unnecessary computations of measuring saliency [16]. Nevertheless, it is important to note that the success of random pruning lies in careful selection of candidate *convolution* layers for filter pruning [16, 23].

3.3. Energy-driven layer selection

Typically, in the majority of pruning scenarios, networks are pruned *uniformly* where each *convolution* layer end up with the same proportion of removed filters. This not only results in a larger loss of performance but also leads to an unfair pruning regime, since every layer exhibits different level of underlying complexity. Inspired from [23], we select the candidate layer for random filter-pruning in an energy-

driven manner, until the desired compression scale (ϵ) is met. In particular, during each pruning iteration, the probability P_k of *conv* layer k to be selected is determined by its relative energy consumption as shown in Eq. 1:

$$P_k = \frac{E_k}{\sum_{k \in K} E_k}, \quad (1)$$

where (E_k) is total energy requirement of a layer k , which is sum of computational energy (E_k^{FLOPs}) and energy for memory access (E_k^{Access}). In this way, layers are penalized (pruned) in order of their contribution to the overall energy consumption of the model. This approach was shown in [23] to not only restricts the layer collapse but also keeps sufficient capacity within the model to capture important patterns from the data. Below we elaborate further on the individual terms E_k^{FLOPs} and E_k^{Access} .

Energy for computations. The computational energy (E_k^{FLOPs}) requirement within a *convolution* layer emerges from its underlying matrix multiplications and additions. These arithmetic operations are represented theoretically as Floating-point operations (FLOPs). For a *convolution* layer k , the number of FLOPs can be captured using Eq. 2:

$$k_{flops} = [C_{in} \times (\Omega)^2 \times C_{out} \times S_{out}], \quad (2)$$

then, the energy consumption for performing these computations can be modelled using Eq. 3:

$$E_k^{FLOPs} = k_{flops} \times A, \quad (3)$$

where C_{in} and C_{out} represent the input and output channels or filters, respectively. Ω^2 indicates the filter dimension, S_{out} is the feature size of the output layer, and A denotes the energy required for a single FLOP, i.e., 2.3pJ [11].

Energy for memory access. Accessing data stored in the memory also contributes to significant amount of energy consumption. For a *convolution* layer k , the amount of




Figure 3. Examples of predicted high and low variance images from the Irish dataset, ResNet18 80% pruned.

memory required can be formulated using Eq. 4:

$$k_{mem} = [(k_{params} \times 3) + \{(S_{out} \times C_{out}^k) \times 2\}] \times 4, \quad (4)$$

then, the energy consumption for accessing this amount of data can be modelled using Eq. 5:

$$E_k^{Access} = k_{mem} \times B, \quad (5)$$

where k_{params} are the number of parameters, and B denotes the energy required to access 1MB of data from DRAM i.e., 640pJ [11]. In Eq. 4, the constant 3 accounts for gradients and momentum variables apart from parameters, and 2 accounts for storing the activations and their errors during the backpropagation. Finally, 4 refers to number of bytes required to store in a 32-bit system.

3.4. Uncertainty and biomass estimation

Although model compression reduces the energy consumption of neural networks, this cannot be attained without sacrificing on accuracy. In the case of biomass composition estimation from phone images, we specifically observe that the loss in accuracy is constrained to specific input images that become harder to predict for the less generic energy-constrained models. Therefore, we propose to identify unreliable predictions of the pruned model on the fly and to refine them with accurate predictions from the unpruned model at test time. This enables high energy efficiency by using the energy-pruned model on most images while retaining high accuracy using the unpruned model on the most challenging ones.

Variance attenuation loss. Since Figure 5(b) and Figure 6(b) evidence that the errors made by the pruned models are concentrated in a few images, we aim to detect these harder samples. Although prediction confidence is often




Figure 4. Examples of predicted high and low variance images from the GrassClover dataset, ResNet18 80% pruned.

linked to prediction accuracy, estimating confidence in the context of regression is non evident. We propose to train our network to expect the ground-truth biomass target y for an input image x to not be a scalar biomass value but to be an observation from a normal distribution, centered around a mean $y(x)$ and corrupted by some input dependant noise $\sigma(x)$: $y \sim \mathcal{N}(\mu(x), \sigma(x))$. $\mu(x)$ here is the most likely value for y (mean of the distribution) predicted by the network from the input image x .

We train our network to maximize the following likelihood function over the dataset $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$

$$\prod_{i=1}^N \frac{1}{\sigma(x_i) \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{y_i - \mu(x_i)}{\sigma(x_i)} \right)^2} \quad (6)$$

equivalent to

$$\frac{1}{(\sqrt{2\pi})^2} e^{-\frac{1}{2} \sum_{i=1}^N 2 \ln \sigma(x_i) + \left(\frac{y_i - \mu(x_i)}{\sigma(x_i)} \right)^2} \quad (7)$$

which amounts to minimizing the exponent

$$\sum_{i=1}^N 2 \ln(\sigma(x_i)) + \left(\frac{(y_i - \mu(x_i))}{\sigma(x_i)} \right)^2 \quad (8)$$

Note that when considering σ is constant, minimizing L_{uncert} amounts to minimizing $(y - \mu(x))^2$, the MSE.

In practice, when training the neural network using a batch size $B < N$ and over multiple biomass categories C , we minimize

$$L_{uncert} = \frac{1}{B} \sum_{i=1}^B \frac{1}{C} \sum_{c=1}^C 2 \ln(\sigma_c(x_i)) + \left(\frac{(y_{i,c} - \mu_c(x_i))}{\sigma_c(x_i)} \right)^2 \quad (9)$$




Figure 5. (a) showing the correlation between RMSE and Variance, (b) showing the RMSE distribution over test set of Irish clover dataset for VGG-16.




Figure 6. (a) showing the correlation between RMSE and Variance, (b) showing the RMSE distribution over test set of Irish clover dataset for ResNet-18.

also known as the variance attenuation loss [21].

Intuitively, if $(y_{i,c} - \mu_c(x_i))$ is close to zero, L_{uncert} encourages the network to predict a small value for $\sigma_c(x_i)$ effectively minimizing the $2 \ln(\sigma_c(x_i))$ term. However, if the prediction for $\mu_c(x_i)$ is inaccurate, a large $\sigma_c(x_i)$ is encouraged so that the $\left(\frac{y_{i,c} - \mu_c(x_i)}{\sigma_c(x_i)}\right)^2$ fraction can be minimized while the $2 \ln \sigma_c(x_i)$ balances naive high $\sigma_c(x_i)$ solutions.

In practice, we predict $\mu(x_i)$ and $\ln \sigma(x_i)$ (for stability purposes) using two linear prediction heads sharing the same feature extractor, see Figure 2. At test time, we can evaluate the confidence of a given prediction by averaged spread of the predicted distributions, indicated by $\sigma(x_i) = \sum_{c=1}^C \sigma_c(x_i)$.

Visualizing high variance images. Figure 3 displays examples of low and high variance phone images predicted at test-time on the Irish clover dataset. While the low variance images in the top row of Figure 3 appears to be composed mostly of tall grass, higher variance images contain more clover. It is also interesting to notice that the detected high error images in the bottom row contain non uniform herbage height (left), many clover and clover flowers (all images) or foreign objects (piece of red rope in the rightmost image). As for the GrassClover dataset in Figure 4, low variance images in the top row appear to contain mostly grass with low amounts of visible soil. High loss images in

the bottom row however contain more clover and present a denser herbage with no visible soil.

For both datasets, clover-heavy images appear harder to correctly predict, especially for a pruned model. This problem is expectedly exacerbated for the GrassClover dataset that expects the algorithm to further differentiate between red and white clover species. It is also possible that the clover leaves in the herbage canopy grow above and occult a part of the grass underneath (middle and rightmost image in Figure 4’s bottom row). In the rest of this paper, we refer to images where the pruned model outputs high error predictions as the hard images.

Correlation between variance and RMSE. Figure 5(a) and Figure 6(a) shows the correlation between the test RMSE and variance (σ) over different held-out phone images from the Irish dataset. We observe that the test time RMSE and variance are positively correlated with each other, even on images unseen during training. For example, VGG-16 in Figure 5(a) exhibits a significant and positive pearson correlation of 0.71. Given this information, we can re-infer high variance predictions using the unpruned model which yield a lower error. For example, as shown in Figure 2, initially all incoming images in a batch are inferred through pruned (energy-efficient) model. Then, for images where model’s output variance does not fall under given threshold, they are re-inferred using original (unpruned) model. This approach circumvent the direct usage of the energy-intensive unpruned model for all images, effectively reducing energy consumption.

3.5. Hybridizing Models for Energy-efficient and Accurate Biomass Prediction

To utilize the strengths of pruned and unpruned models jointly, we propose to leverage both settings by creating a hybrid system. Given a test image, we first predict biomass composition and variance from the pruned (energy-efficient) model, then if the variance exceeds a hyper-parameter threshold we re-infer using the unpruned (energy-intensive) model to obtain the final prediction. In essence, the threshold on the log-variance effectively acts as a control knob to balance the energy-accuracy trade-off in real time. Our hybrid approach is illustrated in Figure 2.

4. Evaluation

We conduct here experiments to quantify the effect of various pruning ratios ϵ on the regression accuracy of a state-of-the-art prediction algorithm [2]. We run experiments on the IrishClover and the GrassClover datasets using ResNet18 [13] and VGG16 [17] architectures where although energy consumption diminishes, we observe an accuracy drop after aggressive pruning ($\epsilon > 50$). We then evaluate the performance of our pruned/unpruned hybrid algorithm on an edge device where we demonstrate that we

Table 1. Results for pruned models, averaged over 5 runs \pm std Irish dataset . Unpruned results ($\epsilon = 0$) are obtained using publicly available code from [2]

Network	ϵ	RMSE Camera	RMSE Phone
ResNet18	0	3.11 ± 0.09	4.81 ± 0.68
	20	4.38 ± 0.19	5.55 ± 1.32
	40	4.44 ± 0.05	5.32 ± 0.16
	60	4.45 ± 0.16	5.75 ± 1.02
	80	4.44 ± 0.10	5.40 ± 0.44
VGG16	0	3.84 ± 0.12	5.28 ± 0.37
	20	3.75 ± 0.10	6.82 ± 1.57
	60	3.84 ± 0.22	6.11 ± 0.74
	90	3.66 ± 0.20	7.50 ± 2.96

Table 2. Results for pruned models, averaged over 5 runs \pm std. GrassClover dataset . Unpruned results ($\epsilon = 0$) are obtained using publicly available code from [2]

Network	ϵ	RMSE Camera
ResNet18	0	10.06 ± 0.28
	20	10.91 ± 0.42
	40	11.25 ± 0.37
	60	11.28 ± 0.35
	80	11.21 ± 0.10
VGG16	0	11.83 ± 0.32
	20	13.16 ± 0.42
	60	13.67 ± 1.10
	90	13.52 ± 0.80

can maintain accuracy while decreasing energy consumption from 40% to 60%.

4.1. Evaluation Settings

We conduct biomass composition estimation experiments on the GrassClover [19] and Irish clover [12] datasets. We propose to train two different network architectures: ResNet18 [13] and VGG-16 [17]. We run experiments for 100 epochs and adopt a cosine learning rate decay. Since we do not need to estimate uncertainty for the unpruned models, we use the work of [2] as the uncompressed baseline without making use of any unlabeled images. The uncompressed baseline is trained using a Root Mean Square Error (RMSE) from a learning rate of 0.1 for ResNet18 and 0.03 for VGG16. For the pruned models, we train with L_{uncert} and an initial learning rate of 0.0003. For all experiments we train at a resolution of 224×224 and utilize data augmentations that preserve all the information in the image: random horizontal and vertical flipping as well as random grayscaling with a probability of 0.2. These training schedules follow previous implementations [1–3].

4.2. Compression rate vs prediction accuracy

We first study the effect of the compression rate ϵ on the prediction error of the network. To do so, we generate five independent pruned models for each compression percentage ϵ . We then train the pruned networks using L_{uncert} from equation 9 while the unpruned models are trained using a RMSE loss. For the Irish dataset, we train the network on the high quality camera images and evaluate on a held out set of camera or phone images to simulate edge conditions. For the GrassClover dataset, we report our results on the validation set which only contains high quality images. We repeat the experiments for VGG16 and ResNet18. Results are available in Tables 1 and 2. We observe in the case of the Irish dataset that the pruned models manage to maintain good accuracy results on the higher quality camera images despite the compression but that the performance drops on the more challenging camera images. We observe a similar increase in RMSE when training pruned networks on the GrassClover dataset in Table 2 even if all the images we test on are of high quality.

Interestingly, increasing ϵ from 20% to 90% does not significantly increase the test error for ResNet18 or VGG16 on either dataset. These results may hint toward dry biomass composition being a task not requiring large amounts of parameters to solve, at least regarding high quality camera images. The accuracy difference between unpruned and pruned models, especially in the case of the phone images does however motivate a hybrid unpruned-pruned approach.

4.3. Energy efficiency and RMSE on edge

We evaluate the energy efficiency and RMSE on a real edge device to compare the proposed approach with pruned-only and unpruned-only.

Edge Device. We performed experiments for inference on a resource-constrained edge device i.e., NVIDIA Jetson Nano. This is a single-board device, similar to Raspberry Pi, but with a GPU to handle edge AI workloads. This edge device comprises a GPU with 128 NVIDIA CUDA® cores, a Quadcore ARM Cortex-A57 CPU, 4GB of RAM, and 16GB of microSD storage.

Proposed vs pruned / Unpruned. Figures 8 and 7 show how our approach can reduce energy consumption by 50% on average with respect to the original unpruned model, while only losing an average of 4% in accuracy. These experiments are repeated over 5 runs, and the mean results are presented with 95% confidence interval.

In Figures 8(a) and 8(b), we show a performance evaluation over different threshold values for ResNet-18 and VGG-16 models respectively on the GrassClover dataset [19]. Similarly, in Figures 7(a) and 7(b), different threshold values are evaluated for VGG16 and ResNet18 models respectively on the Irish clover [12] dataset. In par-

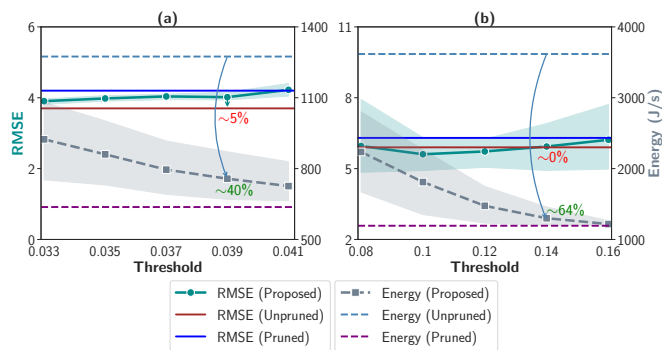


Figure 7. Comparison of proposed approach with pruned-only and unpruned-only methods for energy efficiency and RMSE on Irish clover dataset, and (a) VGG-16, and (b) ResNet-18 models.

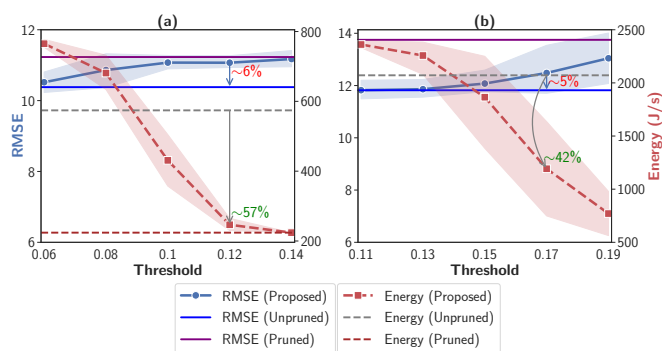


Figure 8. Comparison of proposed approach with pruned-only and unpruned-only methods for energy efficiency and RMSE on Grass-Clover dataset, and (a) ResNet-18, and (b) VGG-16 models.

ticular, we indicate the optimal threshold value for each model on both datasets to achieve high energy efficiency and low RMSE we achieve using our approach.

Number of re-inferred images. Table 3 shows the proportion of camera and phone images of the Irish dataset being refined by unpruned model according to different variance thresholds. We observe that phone images are re-inferred more often than the high resolution camera images. This indicates that the pruned algorithm struggles to generate to these harder images, unseen during training. Although the unpruned model was not seen phone images during training, it manages to generalize better than the pruned one, resulting in smaller prediction errors.

5. Conclusion

Evaluating the composition of dry biomass in grass provides valuable information for farmers. This assessment helps farmers enhance the quality of their grass and make informed decisions about local fertility improvement. Although computer vision algorithms based on deep learning have shown great potential in predicting accurate com-

Table 3. Number of Camera and Phone images being re-inferred by unpruned model on Irish dataset. Each validation image has a phone and a camera view for a total of 214 images (107 in each view)

Network	Threshold	# of Camera	# of Phone	Total
ResNet18	0.033	16	32	48
	0.035	9	18	27
	0.037	9	17	26
	0.039	6	13	19
	0.041	4	11	15
VGG16	0.08	28	32	60
	0.10	20	15	35
	0.12	8	9	17
	0.14	2	5	7
	0.16	0	2	2

position of dry biomass content, these algorithms are energy-intensive hence hampering their usage on resource-constrained edge devices. We first propose to apply filter pruning via energy-driven layer selection to reduce energy consumption. We however find that pruned model can suffer to generalize to challenging samples containing large amounts of clover as well as imperfect real world phone images. To retain accuracy while reducing energy usage, we propose to evaluate the uncertainty of the regression prediction using a variance attenuation loss. At test-time, only harder images whose predicted variance exceeds a threshold are refined by the more accurate unpruned network. We observe that this hybrid approach only slightly increases the error rate of previous algorithms while largely reducing the energy requirements of deep learning biomass estimation algorithms. Finally, generating a variance metric for the network’s predictions, which correlate with the actual biomass prediction error, is crucial for encouraging acceptance among human experts. In practical scenarios, identifying a high predicted variance would prompt end users to either take a clearer picture of the grass or to approach the network’s predictions with due caution.

Acknowledgment

This work was supported by the Science Foundation Ireland (SFI), by the Department of Agriculture, Food, and Marine on behalf of the Government of Ireland VistaMilk research centre under award 16/RC/3835, and in part by SFI under Grant 21/FFP-A/9174, by NGI Enrichers program under Horizon Europe grant: 101070125, and by the National Science Foundation (NSF) under the CPS award # 2133407.

References

- [1] 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 syn-

- thetic images. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 3, 7
- [2] Paul Albert, Mohamed Saadeldin, Badri Narayanan, Brian Mac Namee, Deirdre Hennessy, Noel E O'Connor, and Kevin McGuinness. Unsupervised domain adaptation and super resolution on drone images for autonomous dry herbage biomass estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2022. 2, 3, 6, 7
- [3] Paul Albert, Mohamed Saadeldin, Badri Narayanan, Brian Mac Namee, Noel E. O'Connor, Deirdre Hennessy, Aisling H. O'Connor, and Kevin McGuinness. Utilizing unsupervised learning to improve sward content prediction and herbage mass estimation. In *European Grassland Federation Symposium*, 2022. 2, 3, 7
- [4] Sajid Anwar, Kyuyeon Hwang, and Wonyong Sung. Structured pruning of deep convolutional neural networks. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 13(3):1–18, 2017. 2, 3, 4
- [5] Orly Enrique Apolo-Apolo, Manuel Pérez-Ruiz, Jorge Martínez-Guanter, and Joao Valente. A cloud-based environment for generating yield estimation maps from apple orchards using uav imagery and a deep learning technique. *Frontiers in plant science*, 11:1086, 2020. 3
- [6] 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
- [7] Tibor de Camargo, Michael Schirrmann, Niels Landwehr, Karl-Heinz Dammer, and Michael Pflanz. Optimized deep learning model as a basis for fast uav mapping of weed species in winter wheat crops. *Remote Sensing*, 13(9):1704, 2021. 3
- [8] Ming Der Yang, Hsin Hung Tseng, Yu Chun Hsu, and Wei Chen Tseng. Real-time crop classification using edge computing and deep learning. In *2020 IEEE 17th Annual Consumer Communications & Networking Conference (CCNC)*, pages 1–4. IEEE, 2020. 3
- [9] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2018. 2, 3, 4
- [10] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*, 2015. 3
- [11] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A Horowitz, and William J Dally. Eie: Efficient inference engine on compressed deep neural network. *ACM SIGARCH Computer Architecture News*, 44(3):243–254, 2016. 4, 5
- [12] 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. *Grassland Science in Europe, Vol. 26 – Sensing – New Insights into Grassland Science and Practice*, pages 132–134, 2021. 1, 2, 3, 7
- [13] 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. 2, 6, 7
- [14] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *International Conference on Learning Representations*, 2016. 4
- [15] Xiaolong Ma, Sheng Lin, Shaokai Ye, Zhezhi He, Linfeng Zhang, Geng Yuan, Sia Huat Tan, Zhengang Li, Deliang Fan, Xuehai Qian, et al. Non-structured dnn weight pruning—is it beneficial in any platform? *IEEE transactions on neural networks and learning systems*, 33(9):4930–4944, 2021. 4
- [16] Deepak Mittal, Shweta Bhardwaj, Mitesh M Khapra, and Balaraman Ravindran. Studying the plasticity in deep convolutional neural networks using random pruning. *Machine Vision and Applications*, 30(2):203–216, 2019. 2, 4
- [17] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *IEEE International Conference on Learning Representations (ICLR)*, 2014. 2, 3, 6, 7
- [18] 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. 1, 2
- [19] 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, 7
- [20] 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
- [21] Matias Valdenegro-Toro and Daniel Saromo Mori. A deeper look into aleatoric and epistemic uncertainty disentanglement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2022. 6
- [22] Zhizhong Wang, Lei Zhao, Sihuan Lin, Qihang Mo, Huiming Zhang, Wei Xing, and Dongming Lu. GLStyleNet: exquisite style transfer combining global and local pyramid features. *IET Computer Vision*, 2020. 3
- [23] Muhammad Zawish, Steven Davy, and Lizy Abraham. Complexity-driven model compression for resource-constrained deep learning on edge. *IEEE Transactions on Artificial Intelligence*, pages 1–15, 2024. 2, 3, 4