Study of DNN-based Ragweed Detection from Drones *

Martin Lechner[1,2][0000−0003−1083−0246], Lukas Steindl[1], and Axel Jantsch[1][0000−0003−2251−0004]

[1] Christian Doppler Laboratory for Embedded Machine Learning, Institute of Computer Technology, TU Wien, 1040 Vienna, Austria
[2] Corresponding author, martin.lechner@tuwien.ac.at

Abstract. Ambrosia artemisiifolia, also known as ragweed, is an invasive weed species that aggressively spreads across Europe. 4-5% of the population suffers from strong allergic reactions to its pollen. This work studies the use of aerial drones equipped with highly compressed deep neural networks (DNNs) to scan large areas of vegetation for ragweed with high speed and precision. Compared to the manual approaches with an estimated survey cost of roughly 860 EUR/km^2, the proposed technique can cut the cost of ragweed detection and monitoring by orders of magnitudes. Aerial drones are heavily limited by their battery capacity and thus require an efficient computation platform. As such, it is not possible to use standard DNN models due to memory and computational constraints. Offloading the workload into data centers introduces new issues as the drones may operate in rural areas with poor network connectivity. To overcome these challenges, we train state-of-the-art object detection and segmentation models on a ragweed dataset. The best performing segmentation models were compressed using shunt connections, fine-tuned with knowledge distillation, and further optimized with Nvidias TensorRT library for deployment on an Nvidia Jetson TX2. The highest accuracy models still achieve between 200ms and 400ms inference latency, enabling real-time ragweed survey and potentially allowing more advanced autonomous eradication use cases.

Keywords: Machine Learning · DNN · Autonomous Machines · Drones

1 Introduction

Over the last 20 years, the rise of Ambrosia Artemisiifolia extended the hayfever season for millions of sensitized people in Europe from early summer to late autumn. The plant originally native to North America came to our continent due to upcoming global trade in the 19th century. Initially not spreading too much, several factors such as industrialized agriculture and climate change speeded up the infestation in Europe [7]. Figure 1 shows the rapid distribution of this alien species across Europe.

* This work was supported in part by the Austrian Federal Ministry for Digital and Economic Affairs, in part by the National Foundation for Research, Technology and Development, and in part by the Christian Doppler Research Association.

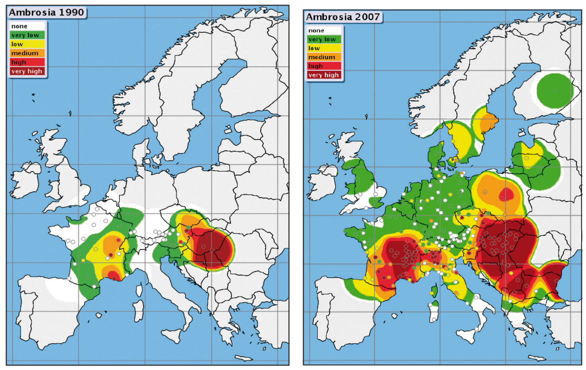

(a) (b)

Fig. 1. Ragweed pollen concentration 1990 (a) and 2007 (b). Favored by climate change Ambrosia Artemisiifolia expands roughly 25km from east to west every year. Data provided by the European Aeroallergen Network [2].

Richter et al. [16] simulated the dispersion of ragweed in Austria and Bavaria for different climate scenarios and found that taking no counteractions would raise the mean annual allergy-related cost from 133 Mio. EUR in 2005 to 422 Mio. EUR in 2050. For the same region, they calculated that investments of 15 Mio. EUR/year in traditional weed management would reduce the mean allergy-related cost by 82 Mio. EUR per year.

The effort for removing ragweed splits into the cost for detection (10%) and subsequent eradication (90%) and is estimated to be 8570 EUR per km^2 in total. Since eradication costs increase for larger plants, a survey should cover large areas in a short time. However, a person requires 25 hours to monitor a single km^2 once. One way to lower the detection cost is to crowd-source the task. People are asked to take pictures of ragweed with their mobile phones and send them to the authorities for verification and removal coordination. The disadvantage of such an unstructured approach is that large areas might not be covered. In an interdisciplinary project, a ResNet50-Image-Classifier was trained to show that deep neural networks (DNNs) can distinguish ragweed from similar-looking plants in close-up images from mobile phones.

This work extends the idea of automated detection and proposes a novel method that uses highly optimized DNNs deployed to embedded devices to scan vegetation images collected with aerial drones. Compared to the traditional manual approach the proposed automated technique has the potential to cut the survey costs (860 EUR/km^2) by 80% and the survey time by 90%.

One challenge that comes with this approach is that offloading large datasets to the cloud may not be possible due to poor network connectivity in rural areas. Another challenge is that DNN Models with millions of parameters like ResNet50 cannot be deployed efficiently to embedded systems due to the lack of available memory and computation resources. Furthermore, the DNN cannot utilize the embedded platform entirely since image pre-processing (cropping, scaling) and data post-processing need additional resources.

We identify state-of-the-art DNNs to recognize ragweeds in images collected with drones, study methods to compress and optimize the DNNs, and map them on suitable embedded hardware. An optimized system can provide a cost-effective solution to support the combat against the spread of Ambrosia Artemisiifolia and to alleviate the medical conditions experienced by millions of people. This paper makes the following main contribution:

A case study of ragweed detection from drones based on standard DNNs, and specifically,

- study and comparison of segmentation and object detection DNNs,
- optimization and customization of the used DNNs,
- a method to find the optimal operating parameters for survey drones.

The rest of the paper is organized as follows: Section 3 describes our approach for drone-based ragweed detection, with our experimental results in 4 and a conclusion drawn in Section 5.

2 Related Work

2.1 Weed detection and eradication

Plascak et al. [14] use helicopters and drones at high altitudes to collect RGB and NGB pictures. To interpret the images, they created orthographic maps and manually engineered features such as tone, color, texture, pattern, form, size, height/altitude, and location. They highlight that the spatial distribution of ragweed can be estimated well by the strong green leaf color in early summer. Budda et al. [6] suggested a combination of robotic sprayer and drone systems to eradicate different kinds of weeds automatically. They trained an object detector to recognize certain weed types in an early plant stage to apply small amounts of pesticides. They emphasize the importance of flying drones at high altitudes to improve efficiency. Increasing the distance between camera and surface reduces the number of pixels per plant and degrades recognition quality. To mitigate this problem, they propose to upsample images using generative adversarial networks (GANs) and pass the output of the GAN as input to an object detector. The authors reported a 93.8 % classification accuracy for the detected instances of the three weed classes they consider in their model. However, 50 out of 170 Ragweed examples were not detected by the Fast R-CNN Object Detector. Their solution also requires offloading the workload to the cloud, limiting the application to areas with sufficient network connectivity and bandwidth.

In 2019 the Laser Zentrum Hanover [3] proposed an eradication solution using drones and robots equipped with lasers to remove unwelcome plants. Mathiassen et al. suggest that laser treatment can indeed reduce the relative survival rate of certain weed species and demonstrated that laser treatment is an effective technique for weed control. Depending on infestation and plant development, a

Table 1. The optimal ragweed eradication method depends on the level of infestation and the size of the plant.

Plant Stage	**Late**	Manual or piloted drone (chemical / mechanical)	Tractor-supported (chemical / mechanical)
	Early	Autonomous drone (laser)	Autonomous robot (laser)
		Weak	**Strong**
		Infestation	

different removal technique might be optimal. Table 1 shows a classification of weed management strategies. In case of heavily infested vegetation with plants in the late stage, tractors and sprayers are likely the optimal eradication method. If the plants are still in an early stage, one could consider ground-based robots equipped with lasers as a viable solution [4, 1].

2.2 Deep Learning and Model Compression

Deep neural networks have been proven to outperform classical approaches in various domains like object detection [12, 5] and semantic segmentation [8], and applications like autonomous driving [9] and plant identification [13]. Object detection networks can be grouped into two approaches, two-stage detectors like Faster-RCNN [15] and single-stage detectors like YoloV4 [5] and singe-shot-detection (SSD) [12]. While two-stage detectors typically reach a higher accuracy, they are slower and require more resources than single-stage detectors. While object detectors predict bounding boxes for each found object in an image, the task of semantic segmentation is to assign a class to every pixel. DeeplabV3+ [8] is one of the latest methods for semantic segmentation, which can be combined with different sized backbones depending on the required accuracy and latency.

However, all state-of-the-art networks still need a lot of resources and require further optimization to run on embedded platforms and meet latency targets. Methods to compress deep neural networks include pruning, quantization, knowledge distillation [11], and shunt-connections [10]. Pruning is a technique to reduce the number of weights in a neural network and can be challenging in architectures containing residual blocks as found in ResNet or MobileNet variants. In knowledge distillation, one trains a smaller neural network with the output distribution of a larger DNN to achieve a comparable accuracy with a reduced number of parameters. On the other hand, Shunt-connections replace computationally expensive blocks of a DNN with much smaller ones.

3 System Architecture

3.1 Ground Sampling Distance

The ground sampling distance (GSD) describes the spatial density of image discretization. It gives the number of units of measurement on the ground that are

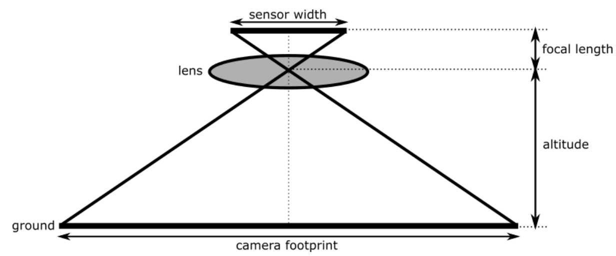

Fig. 2. The ground sampling distance is the relation between the width of the camera footprint on the ground g_w and the sensor width s_w

represented by a single pixel in the observed image. A ground sampling distance of $1cm/pixel$ means that 1 pixel in the image represents $1cm$ of ground distance in the real world. The ground-sampling distance is the relation between the width of the camera footprint on the ground g_w in cm and the sensor resolution sr_w in pixel.

$$GSD = \frac{g_w}{sr_w} \qquad PS = \frac{s_w}{sr_w} \tag{1}$$

For our considerations, we use the standard pinhole camera model to calculate the flight altitude for a given GSD and a given pixel size PS:

$$a = \frac{f \cdot g_w}{s_w} = f \cdot \frac{GSD}{PS} \tag{2}$$

where a is the altitude, f is the focal length, g_w is the width on the ground, and s_w is the sensor width.

A more expensive camera with a higher resolution sensor could outperform cheaper models that require the drone platform to fly at lower altitudes for the same GSD. We derive a model based on equation (eq 2) to decide which drone-camera-combination is the most efficient for a specific GSD. Next, we need to find the limits of the GSD necessary for detecting ragweed. When creating a training dataset, image annotation is an important task typically performed by humans. We found that annotators require at least 100x100 pixels to identify a ragweed plant in an image.

Large ragweed plants often cover $1m$ of the ground surface. To obtain 100×100 pixels the GSD should be $1cm/pixel$. Plants in an early stage cover a surface area of $\approx 10cm$. Detecting smaller objects requires a lower GSD of 0.1cm/pixel. One could argue that it is cheaper to detect ragweed when the plant is larger. Although the detection cost goes down for larger plants, the eradication cost is much higher when detected later, and laser-based removal techniques may no longer be viable.

3.2 Flight Parameter Model

A drone platform provides a maximal flight duration t and a maximal speed v. The cost of a flight is assumed to be c EUR. The mounted camera provides an

Table 2. Five final layers of the network for the original (Cityscapes) and the reduced (Ragweed) segmentation head.

Layer	Cityscapes Out Shape (Params)	Ragweed Out Shape (Params)
decoder_decoder_conv0	271, 481, 19 (2451)	271, 481, 2 (258)
decoder_feature_project0	271, 481, 19 (1691)	271, 481, 2 (178)
add_3	271, 481, 19 (0)	271, 481, 2 (0)
resize_3	2161, 3841, 19 (0)	2161, 3841, 2 (0)
activation_87	2161, 3841, 19 (0)	2161, 3841, 2 (0)

image resolution of $h \times w$ pixels. The sensor height, width and focal length is given as s_h, s_w and f. For a certain GSD in cm/pixel, the ground area covered by an image is $g_w \times g_h \mathrm{m}^2$ with

$$g_w = \frac{w \cdot GSD}{100} \text{ and } g_h = \frac{h \cdot GSD}{100}. \tag{3}$$

The number of flights f_c required to scan $1\,\mathrm{km}^2$ can be calculated as:

$$f_c = \frac{10^6}{d \times g_w} \tag{4}$$

where d is the distance, a drone can travel during one flight.

Finally, we can calculate the total cost of scanning $1km^2$ of vegetation at a given GSD by

$$\text{Total Cost per km}^2 = \frac{10^8 \cdot c}{t \cdot v \cdot w \cdot GSD} \tag{5}$$

This expression can be interpreted intuitively. The total costs decrease if a drone can fly longer (flight time t) and faster (speed v) with a higher sensor resolution. Ceteris paribus, a larger ground sampling distance allows the drone to fly at higher altitudes increasing the observed ground surface and reducing the overall data collection cost per km^2. Increasing the sensor resolution w also increases g_w resulting in the same effect.

The model also allows the calculation of the optimal flight parameters. Equation (2) gives the ideal altitude. Ignoring motion blur, the optimal speed is equal to the maximal speed of the drone. The number of pictures taken during a flight is and the required framerate to fully cover the ground are given by:

$$\text{images per flight} = \frac{d}{g_h} \qquad \text{required framerate} = \frac{t}{p} \tag{6}$$

3.3 Dataset

In September 2020, we recorded roughly 130 minutes of video data and 200 high-resolution images of ragweed-infested dense vegetation in the eastern part of Austria at heights of 2.5 and 4 meters. We used multiple configurations of camera systems and perspectives to obtain image data with the desired level of quality.

Inspired by the work of Buddha et al. [6] we annotated an object-detection dataset where we chopped the original 4k videos into eight non-overlapping video tiles. The final dataset consists of 971 image tiles with a resolution of 960x1080 pixels, of which 481 tiles show at least one region of interest. We split the images into 80% training samples and 20% test samples. We also created a segmentation dataset consisting of 139 images.

3.4 Object Detection Network

We used an SSDResNet50FPN as included in the TensorFlow Object Detection API for object detection, which is pre-trained on the MS CoCo 2017 dataset. We used a cosine decay learning rate with a base learning rate factor of 0.04 and a 2000 step warm-up phase with a warm-up learning rate of 0.013333. In total, we executed the training for 25000 steps.

3.5 Semantic Segmentation Network

Our model for semantic segmentation is based on the DeeplabV3+ architecture with a MobileNetV3 backbone and pre-trained on the Cityscapes dataset. In the next step, we trained, compressed, and fine-tuned our model using the Shunt-Connector-Framework from Haas et al. [10]. This includes modifying the segmentation head for our dataset. The cityscapes dataset is annotated on 19 output classes, whereas our ragweed dataset has only two classes (ragweed, background). However, the modifications have to be done in-place in order to load the pre-trained weights. Table 2 shows all modified layers. For training the shunt-inserted model, we used a constant learning rate of 0.05 and executed the training for 200 epochs.

4 Results

In the following sections, we first compare different drones for small and large ragweed detection in terms of cost-efficiency. Then, we describe our experimental setup and compare the results of object detection and segmentation on our ragweed dataset. Finally, we show the required optimization steps for embedded hardware deployment.

4.1 Drone Selection

As an example, we compare three off-the-shelf drones, a DJI Phantom 4 Pro 2 with a built-in camera, a DJI Matrice 300 RTK with a Zenmuse P1 (50mm) camera, and a WingtraOne with a Sony QX1. For simplicity, we assume the operating costs to be the same for each drone and only use the deprecation costs set to $1/100 \times list\_price$ as costs per flight. We consider two scenarios.

Table 3. Overview of all required UAV parameters to decide which platform-sensor combination is most efficient and to compute the optimal flight parameters.

	DJI Phantom	DJI Matrice	WingtraOne
Flight duration [s]	1800	3300	3540
Cost per flight [EUR]	30	110	45
Speed [m/s]	20	23	16
Flight distance [km]	36	75.9	56.640
Camera	Built in	Zenmuse P1	Sony QX1
Resolution [px]	5472×3648	8192×5460	5456×3532
Sensor size [mm]	13.2×8	35.9×24	23.2×15.4
Focal length [mm]	8.8	50	20

Scenario S1 aims to detect small ragweed plants that are $10cm$ in size. A GSD of $0.1cm/pixel$ will be assumed to achieve this goal. In Scenario S2, the drone should detect large ragweed ($1m$ in size). To obtain 100×100 pixels the GSD will be $1cm/pixel$.

With the parameters from Table 3, the ground dimensions mapped to a single image can be calculated using Equation (3). Considering the real-world surface dimensions of a single image, we can derive the total observed ground surface of one flight. Table 4 summarizes the results. Note that the price per km^2 decreases by a factor of 10 when the GSD increases by a factor of 10.

Although the Matrice is about three times as expensive as the Phantom, it is still more cost-efficient in collecting images with the desired GSD. The optimal flight altitude can be calculated with Equation (2).

Another metric is the minimal interval of images required to cover the entire ground surface during a flyover without overlaps. The Phantom-V4 in scenario S1 has to take $36000/3.65 = 9869$ pictures during an 1800 second long flight resulting in a framerate of 0.18 images per second. This sets a strict time limit for real-time onboard ragweed detectors. Due to the higher spatial resolution of the DJI-Matrice, the drone can fly higher and take fewer images relative to the Phantom for the same distance. However, it still takes more pictures as it can fly longer and faster. Table 4 shows the optimal flight parameters and time limits for both scenarios.

In conclusion, the relative drone efficiency does not depend on the scenario, as changing the GSD affects all drones equally. Ignoring the constant values in Equation (5) it is optimal to pick the drone-camera combination with the lowest ratio of drone *cost* to maximal_airtime $\times$ velocity $\times$ resolution.

4.2 Experimental Setup

For our experiments, we used an Nvidia Jetson TX2 as it is the sweet spot between power, performance, and weight for a drone application. All networks are converted to Nvidia's neural network inference framework Tensor RT (TRT). It can quantize models with FP16 and INT8 operations while preserving the

Table 4. Efficiency metrics and optimal flight parameters of three UAV models for small and large ragweed detection. The costs per km scales linearly with the required GSD.

Scenario	S1			S2		
Drone	DJI Phantom	DJI Matrice	Wingtra One	DJI Phantom	DJI Matrice	Wingtra One
Efficiency Metrics						
Groundwidth [m/img]	5.47	8.19	5.46	54.72	81.92	54.56
Groundlength [m/img]	3.65	5.46	3.63	36.48	54.60	36.32
Ground area [km²/flight]	0.2	0.62	0.31	1.97	6.22	3.09
Flights per km²	5.08	1.61	3.24	0.51	0.16	0.32
Time [h/km²]	**2.54**	**1.48**	**3.19**	**0.26**	**0.15**	**0.31**
Costs [EUR/km²]	**152.3**	**176.9**	**145.6**	**15.2**	**17.7**	**14.6**
Optimal Flight Parameters						
Altitude [m]	3.65	11.41	4.7	36.48	114.09	47.03
Speed [m/s]	20	23	16	20	23	16
Images per flight	9,869	13,902	15,595	987	1,391	1,560
Framerate [images/s]	5.48	4.21	4.41	0.55	0.42	0.44

Table 5. Semantic segmentation $mIoU$ results. *IoU reported on the training server.

	mIoU			
Resolution	dlv3*	dlv3s	dlv3trt	dlv3trts
961×541	0.588	0.625	0.574	0.625
1921×1081	0.670	0.695	0.634	0.695

models' accuracy and supports layer-fusing, kernel-auto-tuning, and other optimizations. This allows us to run our networks at the highest possible efficiency.

Object detection On the ragweed testset, we achieved an Average Precision (AP) of 0.525 for an Intersection over Union (IoU) of 0.5 and a mean AP of 0.229 for IoU scores in the range of 0.5 to 0.95. Figure 3 shows that while the object detector can detect individual plants, it struggles in areas with dense vegetation. Thus, object detection is a valid approach in spring or early summer when vegetation is still sparse. After running optimizations with TensorRT, we obtained a latency of 95 ms at floating point 16 (FP16) and 124ms at FP32. In FP16 mode, the AP@0.5IoU dropped slightly to 0.523, and the mean AP dropped to 0.227.

Semantic Segmentation In our experiments, we compare the latency and the IoU of the original model with TensorRT optimizations and the shunt-inserted model with and without TensorRT optimizations. However, for the original model, we could not run it on a TX2 due to memory constraints. Thus we provide the IoU achieved on our training machine. Table 5 shows the $mIoU$

(a) (b) (c)

(d) (e) (f)

Fig. 3. The model predicts reasonable (green) bounding boxes if the plant is well separated from the rest of the vegetation. However, it sometimes fails (see (b) and (f)). Especially when vegetation is dense the (red) ground truth boxes overlap with the non-ragweed background.

of the original model (dlv3), the shunt-inserted model (dlv3s), the TensorRT optimized model (dlv3trt), and the TensorRT optimized shunt-inserted model on two different resolutions (1/4 and 1/8 of the original 4k images). As a general trend, inserting the shunt connections followed by fine-tuning increases the $mIoU$. Also, the TensorRT-based optimizations have no impact on the $mIoU$ for the shunt-inserted model. In addition to the $mIoU$, Figure 4a also shows the mean latency for all models on the Jetson TX2. Platform-specific optimizations clearly boost the latency, while shunt connections do not make a huge difference in terms of latency. However, due to the increased $mIoU$, the optimized, shunt-inserted models represent the Pareto optimum. The TensorRT optimized, shunt-inserted model (dv3trts) with a resolution of 1921×1081 has a mean latency of $200ms$ which meets most of the framerate requirements from Table 4. In case of a DJI Phantom in scenario S1, only the lower resolution model with a mean latency of $45ms$ is fast enough. Figure 4b shows the impact of the batch size during training together with the achievable $mIoU$ for three different resolutions on the training server. As expected, the highest resolution model achieves the best $mIoU$ scores. However, these models do not fit our target hardware or latency constraints. Figure 5 shows the predictions of the segmentation model together with the ground-truth labels for two sample images. In both scenarios, the predictions and the labeled ground-truth overlap to a large extent.

(a) (b)

Fig. 4. Left: The Pareto-optima on the TX2 are the shunt-inserted deeplabv3+ models optimized with TensorRT (dv3trts). Shunt-optimization has little effect on a TRT-optimized models latency but significantly boosts its accuracy. The size of the points denotes the image size and the shape (dot, cross) the training batch size (BS). The test-batch size is always 1. **Right**: Higher resolution and larger training batch size improve the quality of the segmentation model.

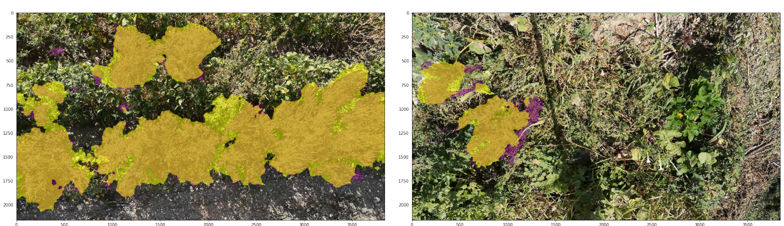

(a) (b)

Fig. 5. Predictions of the segmentation model (yellow) and ground-truth (purple).

5 Conclusion and Future Work

In this work, we presented a case study of ragweed detection from drones. The experiments show that the current state of the art in computer vision and model compression techniques is sufficient to detect ragweed with a high level of accuracy when the model runs on a resource constraint embedded system. Based on the given assumptions (860 EUR/km^2, 25h/km^2), we can conclude that using drones to detect ragweed plants (15-180 EUR/km^2, 0.15-3.2h/km^2) is between five and fifty times more cost-efficient and eight to a hundred times more time-efficient depending on the plant size.

Shunt connections and knowledge distillation allowed us to train much smaller models to a similar or even higher accuracy as the original models with a significantly reduced latency. However, the latency improvement due to TensorRT was substantially larger, highlighting the importance of hardware-specific optimizations.

The next step is to collect a larger dataset using drones at higher altitudes. This dataset should then needs careful annotations since the object detector results show that a high-quality dataset is at least as important as the model itself. Although the object detector did not work too well in the case of dense vegetation, it has benefits for sparse vegetation images. Thus, we will fuse both approaches to reliably detect ragweed plants independent of the growth state and vegetation density. Further, we plan to extend our approach to detect and locate other invasive plants.

References

1. Corteva leverages mobile robots to walk row crops, https://www.futurefarming.com/tech-in-focus/field-robots/corteva-leverages-mobile-robots-to-walk-row-crops/, accessed 2021-12-17
2. European aeroallergen network, https://ean.polleninfo.eu/Ean/, accessed 2021-12
3. Laser zentrum hannover (lzh) presents anti-weed laser system at agritechnica, https://optics.org/news/10/11/12, accessed 2021-12-17
4. Welaser project refines laser-based weed control, https://optics.org/news/12/3/52, accessed 2021-12-17
5. Bochkovskiy, A., Wang, C., Liao, H.M.: Yolov4: Optimal speed and accuracy of object detection. CoRR **abs/2004.10934** (2020)
6. Buddha, K., Nelson, H.J., Zermas, D., Papanikolopoulos, N.: Weed detection and classification in high altitude aerial images for robot-based precision agriculture. In: 2019 27th MED. pp. 280–285. IEEE (jul 2019)
7. Buttenschøn, R., Waldispühl, S., Bohren, C.: Guidelines for management of common ragweed, Ambrosia artemisiifolia. University of Copenhagen (2010)
8. Chen, L., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. CoRR **abs/1802.02611** (2018)
9. Grigorescu, S.M., Trasnea, B., Cocias, T.T., Macesanu, G.: A survey of deep learning techniques for autonomous driving. CoRR **abs/1910.07738** (2019)
10. Haas, B., Wendt, A., Jantsch, A., Wess, M.: Neural network compression through shunt connections and knowledge distillation for semantic segmentation problems. In: AIAI. pp. 349–361. Springer International Publishing, Cham (2021)
11. Hinton, G., Dean, J., Vinyals, O.: Distilling the knowledge in a neural network. pp. 1–9 (03 2014)
12. Lu, X., Kang, X., Nishide, S., Ren, F.: Object detection based on ssd-resnet. In: 2019 IEEE 6th CCIS. pp. 89–92 (2019)
13. Mehdipour Ghazi, M., Yanikoglu, B., Aptoula, E.: Plant identification using deep neural networks via optimization of transfer learning parameters. Neurocomputing **235**, 228–235 (2017)
14. Plaščak, I., Barač, Z.: Remote detection of ragweed (ambrosia artemisiifolia l.). Tehnički glasnik (2018)
15. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. CoRR **abs/1506.01497** (2015)
16. Richter, R., Berger, U.E., Dullinger, S., Essl, F., Leitner, M., Smith, M., Vogl, G.: Spread of invasive ragweed: Climate change, management and how to reduce allergy costs. Journal of Applied Ecology **50**(6), 1422–1430 (aug 2013)