

DETECTING *RUMEX OBTUSIFOLIUS* WEED PLANTS IN GRASSLANDS FROM UAV RGB IMAGERY USING DEEP LEARNING

J. Valente^{1,2,*}, M. Doldersum¹, C. Roers³, L. Kooistra¹

¹ Laboratory of Geo-Information Science and Remote Sensing, Wageningen University & Research, 6708 PB, Netherlands -
(joao.valente, lammert.kooistra)@wur.nl, m.e.doldersum@students.uu.nl

² Information Technology Group, Wageningen University & Research, Wageningen 6700 EW, Netherlands

³ Naturschutzzentrum im Kreis Kleve e.V., 46459 Rees-Bienen, Germany – roers@nz-kleve.de

KEY WORDS: DJI Phantom, Grasslands, Machine vision, Deep learning, Rumex, Weeding, Plant detection, Aerial surveying

ABSTRACT:

Broad-leaved dock (*Rumex obtusifolius*) is a fast growing and spreading weed and is one of the most common weeds in production grasslands in the Netherlands. The heavy occurrence, fast growth and negative environmental-agricultural impact makes Rumex a species important to control. Current control is done directly in the field by mechanical or chemical actuation methods as soon as the plants are found *in situ* by the farmer. In nature conservation areas control is much more difficult because spraying is not allowed. This reduces the amount of grass and its quality. Rumex could be rapidly detected using high-resolution RGB images obtained from a UAV and optimize the plant control practices in wide nature conservation areas. In this paper, a novel approach for Rumex detection from orthomosaics obtained using a commercial available quadrotor (DJI Phantom 3 PRO) is proposed. The results obtained shown that Rumex can be detected up to 90% from a 6 mm/pixel ortho-mosaic generated from an aerial survey and using deep learning.

1. INTRODUCTION

Grasslands have a positive environmental and economic impact. In case of agricultural grasslands, they promote the production of food through dairy farming and help to improve soil health. However, there are some injurious weeds that reduce the quality of the fodder and the amount of grass. One of those weeds is broad-leaved dock or *Rumex obtusifolius* as shown in Figure 1. Rumex is a weed that competes with grass for resources such as water, nutrients and light and thereby it reduces crop yield (Hiremath et al., 2013). It is not only competing with grass but is also poisonous when eaten by livestock in large amount. It is listed as an ‘injurious weed’ under the UK Weeds Act 1959 (Weeds Act, 1959).




Figure 1. Rumex (*Rumex obtusifolius*) with an approximated area of 200 cm².

Rumex is difficult to eradicate because of its deep roots and easily dispersed seeds. The roots of Rumex can reach up to 1.5 meters into the ground and its seeds can travel through both

wind and water. The structure of the seeds also makes that they cling easily to animals or machinery (Poison Diaries, 2013).

Rumex control and management is typically done via chemical spraying. However, when spraying herbicides, not only the weeds, but also the surrounding crops or grasslands are sprayed and this has negative effects on the environment. Herbicides for example contribute to ground water pollution and they contribute to a selective bias of herbicide-resistant weeds (Kempenaar et al., 2007; Hiremath et al., 2013). There are several other methods that do not include the spraying of chemicals. These methods are for example crop rotation or thermal and biological control (van Evert et al., 2009; Bond & Grundy, 2001). Another method is the manual removal of weeds. All these methods are either time-consuming, expensive or labour intensive, so there is need of automatic detection approaches that will be less expensive, time efficient, and more reliable (van Evert et al., 2009; Bond & Grundy, 2001).

Deep learning has proven to be an effective way of object detection in images for various applications such as autonomous driving and robotics (Huval et al., 2015; Goodfellow et al., 2016). Also, in the field of detecting plant species, the use of deep convolutional neural networks has shown promising results. Dyrmann et al. (2016) used a deep convolutional network to detect 22 plant species and managed to do so with an accuracy of 86.5%. Ghazi et al. (2017) have examined the use of deep convolutional neural networks in detecting plant species captured in photographs. They aimed to find the difference in performance between transfer learning and building a model from scratch. Ghazi et al. (2017) examined different algorithms and concluded that the results of transfer learning are significantly better than the results of newly built models from scratch.

* Corresponding author

In this work a deep learning approach using Convolutional Neural Networks (CNN) and transfer learning, based on a pre-trained deep convolutional networks model is applied to high-resolution aerial imagery acquired from a Unmanned Aerial Vehicle (UAV) over wide grasslands to detect Rumex. Moreover, it attempts to introduce a more cost-effective, fast, and non-invasive method for detecting Rumex in high-resolution images obtained with off-the-shelf UAV technology.

2. RELATED WORK

The most common challenges using deep learning in agriculture are the identification of weeds, land cover classification, plant recognition, fruit counting and crop type classification (Kamilaris & Prenafeta-Boldú, 2018).

In the context of object detection or recognition (in agriculture) a deep learning model requires large volumes of training data, preferably thousands or more instances (Pan & Yang, 2010). Since this kind of input data is often not available, many researchers rely on the use of transfer learning (Lu et al., 2017; Lee et al., 2015; Christiansen et al., 2016). With transfer learning, an existing deep learning model, often trained on millions of images, is used as a basis and is then retrained with a smaller dataset in order to fit the specific problem at hand (Pan & Yang, 2010).

Although not with deep learning, there are also a number of articles regarding the detection of Rumex in grasslands. Gebhardt et al. (2006) used digital image processing to detect Rumex. They used image segmentation, local homogeneity calculation, a homogeneity threshold and morphological opening with detection rates from 71% to 95%. Misclassification rates were between 9% and 24% (Gebhardt et al., 2006). In a following research, Gebhardt and Kühbauch (2007) introduced a new algorithm for Rumex detection in mixed grasslands. Using local homogeneity and morphological operations, homogeneous regions were segmented. Additional texture and colour features were defined using stepwise discriminant analysis and maximum-likelihood classification. This resulted in an improved accuracy of up to 83% and Rumex detection rates of 93% (Gebhardt & Kühbauch, 2007). However, a drawback of this method is that the results were depending on the growth stage of the Rumex plants and relies on images taken by hand at close range.

Even higher detection rates than those of Gebhardt and Kühbauch (2007) were obtained by Hiremath et al. (2013). By segmenting texture features based on Markov random fields, a detection rate of 97,8% for Rumex detection was achieved (Hiremath et al., 2013). For this method, the input data is set as image locations containing only a part of a Rumex plant, but without containing any background grassland. This kind of input data is less suitable for long-range imagery and is part of the reason for exploring an alternative approach in this research.

Van Evert et al. (2010) created a ground-based robot to detect and control Rumex in grassland. This robot detects Rumex through a camera and real-time image processing. The image processing method is based on the differences in size, shape and colour between grassland and Rumex plants. During testing in a real world setting the robot managed to detect 93% (124 out of 134) of the present Rumex plants and reported eight false positives (van Evert et al., 2010). Although the robot can cover one hectare in three hours' detection based on UAV imagery

might contribute to speeding up this process and with that offer new possibilities.

Binch and Fox (2017) examined all currently available research on the detection of Rumex and compared them to each other by testing the methods on a new independent dataset under a large range of environmental conditions. In their examination, they found that the linear binary patterns method together with a support vector machine from Ahmed et al. (2013) scored the best with an accuracy of 83%. This is significantly lower than the accuracy when tested on the original dataset, but still a large amount of the Rumex was detected. What is noteworthy is that all articles discussed regarding the detection of Rumex rely on close range photography (Binch and Fox (2017)).

During this research, UAV imagery will be used. Unlike close range cameras, UAV imagery allows for larger areas to be evaluated at once. If the processing time does not increase substantially, the evaluation of larger areas increases the speed of the process and thus allows for a cheaper and more effective use in large grassland areas as well as additional value for real time applications. If the performance is similar to that of close range cameras this can contribute to a faster and more widely used method for weed management in grassland. Finally, it is the first time that a study using deep learning for Rumex detection in high-resolution UAV imagery is provided.

3. METHODOLOGY

The methodology employed is based on transfer learning and a pre-trained CNN (AlexNet, explained in the coming sections). The implementation was based on four main steps as illustrated in Figure 2.

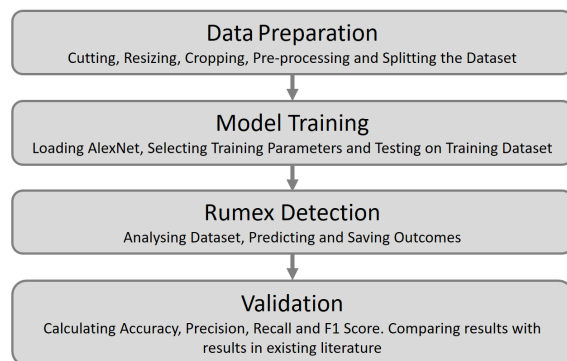


Figure 2. Flowchart of the main steps taken to execute and validate the approach proposed.

3.1 Data acquisition

The small quad-rotor UAV used in this research was the Phantom PRO 3 with a 12-megapixel camera from an elevation of 10m and duration of flight was 3 min. The flights took place on 17th of April, 2018 at 11:00 am in the Salmorth reservoir (Germany) very close to the Rhine river shore. The exact location can be seen in Figure 3. The area selected is a grassland field used for livestock. The UAV surveyed a grassland area from 0.5 ha where several Rumex patches were identified *in situ* by experts and farmers.



Figure 3. Grassland field locate in the Salmorth reserve with Rumex infestation located along the river Rhine to the West of the city of Emmerich next to the Dutch-German border.

3.2 Data post-processing and labelling

The imagery obtained from the survey (54 images) was processed using Agisoft Photoscan and an orthophoto with a spatial resolution of 6 mm/pixel was generated. The image was afterwards, divided in four quadrants to be easier to inspect. Experts of Naturschutzzentrum im Kreis Kleve e.V. located Rumex plants in this imagery and labelled them. In total 631 Rumex plants were found and labelled in the dataset. The resulting orthomosaic is shown in Figure 4.




Figure 4. Orthomosaic generated from the aerial survey over the grassland area selected.

3.3 Deep learning approach

In this section, the transfer learning based deep learning approach and its validation metrics are explained.

3.3.1 Alexnet CNN

AlexNet is a pretrained convolutional neural network model that was trained on a dataset of 1.300.000 high resolution images and can recognize a thousand different objects in images (MathWorks, 2018). The original AlexNet network consists of eight layers, of which five are convolutional layers and three are fully connected layers (Vedaldi & Lenc, 2015). An overview of the AlexNet architecture is shown in Figure 5.




Figure 5. Schematic overview of the AlexNet architecture (Shehata, 2016)

The AlexNet architecture consists of roughly 650,000 neurons and 60 million parameters (Ghazi et al., 2017). By using the AlexNet model as a starting point, transfer learning can be applied, which allows for solving other object recognition or detection tasks effectively even with more limited datasets. AlexNet is an open source model and by using the Deep Learning Toolbox Model for AlexNet Network it can be imported into MatLab directly (MathWorks, 2018).

3.3.2 Training data pre-processing

The data needs to be prepared in such a way that is it compatible with the input expectations of the model. The model, based on AlexNet, expects a set of images of size 227x227x3. Based on the coordinates of the bounding boxes, all Rumex locations will be selected, cut out into a set of smaller images and rescaled to the desired input size. These images will be divided into Rumex and grasslands bounding boxes (without any or part of Rumex plant).

These pictures are then divided into a training set and a testing set. The training set will be used to train the model. The testing set will be used to test how well the model performs. After that, the model robustness is improved by training it with partially occluded plants or plants displaced from the bounding box centre.

3.3.3 Model training

The AlexNet training options that can be selected are for example the batch size (the number of images used per iteration), the number of iterations, the learning rate, the momentum (how strongly a new iteration influences the result) or the number of epochs (how often the entire training dataset is used). There are a large number of options and combination that can be chosen from. The effectivity of which depends per case and is difficult to predict beforehand.

For this research, the default options for the SGDM (stochastic gradient descent with momentum) optimizer is used as a starting point. From there on the options will be changed systematically to see how the changes influence the outcome. Based on this approach, the most suitable set of training options will be selected. By using cross-validation the model will be trained multiple times on different subsets of the training data in order to improve performance and reduce overfitting.

The orthomosaic from Figure 4 was split in four quadrants to aid the image annotation task and processing. The training set was made up 628 Rumex plants, and 659 non Rumex plants, e.g., others species and pure grassland.

3.3.4 Validation metrics

The validation is done through the analysis of the accuracy, precision and recall. The accuracy is a measure which refers to the proportion of correct predictions. It can be defined as follows:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

Where TP = True Positives, TN = True Negatives, FP = False Positives, and FN = False Negatives.

However, accuracy alone can give a wrong indication, for example when a certain group only consists of a few entries. Additionally, to accuracy there are more measures that validate the results of a model. Two of them are precision and recall. Precision states which proportion of positive identifications was correct (Kamilaris & Prenafeta-Boldú, 2018). Precision is defined as follows:

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

Recall refers to the proportion of actual positives that was detected (Kamilaris & Prenafeta-Boldú, 2018). Recall is defined as follows:

$$Recall = \frac{TP}{TP + FN} \quad (3)$$

These two measures can sometimes contradict to each other in values but are both important to evaluate. A measure which combines precision and recall is the F1 score. The F1 score is the weighted average of precision and recall and can be calculated as following (Joshi, 2016):

$$F1\ Score = 2 \cdot \frac{Recall \cdot Precision}{Recall + Precision} \quad (4)$$

However, because a false positive and a false negative have different impacts, precision and recall will also be evaluated separately in this research.

4. EXPERIMENTS AND RESULTS

In this section, we present the results regarding the experiments with three different Rumex datasets: 1) basic (plants centred in the training input images), 2) shifted (plants that are not centred in training input images), and 3) cut-off (plants that are partially seen on the image).

These experiments aim of to give an insight of the performance of different input images that will be generated from a sliding window procedure over the orthomosaic, i.e. the orthomosaic is automatically scanned with a bounding image (227x227 pixels) from the size.

4.1 Basic model

The first model was trained with images that depict Rumex plants and grassland background. Figure 7 shows some input images examples.



Figure 7. Standalone 227x227 pixels input pictures obtained from the orthomosaic generated from the UAV flight: a) Rumex centred (top-left), b) Grassland (top-right), c) Grassland and other plant species (bottom-left), and d) Rumex centred in an earlier stage (bottom-right).

To train the model, a randomized selection of 75% of both groups has been used as training data. Using a randomized cross-validation, 30 validation rounds were performed. Based on the average of these 30 runs, a confusion matrix was created (Figure 8).

		Predicted Values	
		No Rumex	Rumex
Actual Values	n = 250 No Rumex	True Negatives 88	False Positives 5
	Rumex	False Negatives 9	True Positives 148

Figure 8. Confusion Matrix – Classification with 75% training data for basic dataset.

To account for the influence of the amount of training data, the same process was done using 80% of the data for training. Again 30 runs were performed, resulting in the confusion matrix shown in Figure 9.

		Predicted Values	
		No Rumex	Rumex
Actual Values	n = 201 No Rumex	True Negatives 69	False Positives 6
	Rumex	False Negatives 7	True Positives 119

Figure 9: Confusion Matrix – Classification with 80% training data for basic dataset.

Based on these confusion matrices the accuracy, precision, recall and F1-score of the models were calculated. These average values of these measures are shown in Table 1.

	75% training	80% training
Accuracy	94,3%	94,0%
Precision	96,9%	95,8%
Recall	94,2%	94,8%
F1-score	95,4%	95,1%

Table 1. Validation measures with different shares of training data

The results show that the differences between a training dataset of 75% and a set of 80% are not significant. The largest difference between the two sets lies in the precision. Remarkable is that a larger training dataset has a larger variance and a lower average score for precision. However, these differences are only minimal. Overall the model performs very well, with average values for all tests being above 94%.

4.2 Shifting plant location

In the initial model, the Rumex was always located in the centre of the image. To increase the robustness and its real application, another dataset is created in which the Rumex is not centred in the image and there are images with partly cut off plants. Several examples of centroid shift and partial plants on the image are shown in figure 10.



Figure 10. Decentralized and partly cut off Rumex obtained from orthomosaic generated from the UAV flight: a) plant displaced from the centre (top-left), b) plant partially cut-off (top-right), c) plant displaced from the centre (bottom-left), and plant partially cut-off (bottom-right).

The confusion matrices of using the previous model, trained on images with Rumex basic model (plant in the centre), on images with Rumex on a random location or partly cut off are shown in Figures 11 and 12. It is shown that the number of false positives stays very low, but the number of false negatives increases substantially. When the Rumex location is moved, only 63 out of 103 Rumex locations are found. When the Rumex is partly cut off only 23 out of 103 are found.

		Predicted Values	
		No Rumex	Rumex
Actual Values	n = 219		
	No Rumex	True Negatives 112	False Positives 4
Rumex	False Negatives 40	True Positives 63	

Figure 11. Confusion Matrix – Classification of Rumex location moved using the basic trained model.

		Predicted Values	
		No Rumex	Rumex
Actual Values	n = 212		
	No Rumex	True Negatives 106	False Positives 3
Rumex	False Negatives 80	True Positives 23	

Figure 12. Confusion Matrix - Classification of Rumex partly cut off using the basic trained model.

The results of the models are summarized and compared in Table 2. The results shown that in all datasets, the model has a high score on precision. Meaning that locations that are classified as Rumex are most of the time correct. As expected, based on the confusion matrices, the recall score decreases drastically. Meaning that many actual Rumex locations are not found.

	Rumex Center	Rumex Moved	Rumex Cut off
Accuracy	94,3%	79,9%	60,9%
Precision	96,9%	94,0%	88,4%
Recall	94,2%	61,1%	22,3%
F1-score	95,4%	74,1%	35,6%

Table 2. Validation measures of the basic model

To make the model more robust, it has been retrained using the images with shifted Rumex locations as input data. The same method has been used as for the initial model. The model was trained on 75% of the images and was validated by a 30-fold random cross-validation. The averaged results of the model with shifted Rumex locations are shown in the confusion matrix in Figure 13.

		Predicted Values	
		No Rumex	Rumex
Actual Values	n = 249		
	No Rumex	True Negatives 78	False Positives 14
Rumex	False Negatives 17	True Positives 140	

Figure 13. Confusion Matrix – Classification using the shifted trained model.

Figure 14 shows the confusion matrix for the model with images containing partly cut off Rumex as both training and validation data.

		Predicted Values	
		No Rumex	Rumex
Actual Values	No Rumex	True Negatives 66	False Positives 22
	Rumex	False Negatives 17	True Positives 140

n = 245

Figure 14. Confusion Matrix – Classification using the partly cut off trained model.

The validation measures for these models were calculated and are compared to the measures of the initial model in Table 3. All validation measures decrease with a few percent per model. However, it is noteworthy that even in the model where Rumex is only present in the image partly, the model still performs well.

	<i>Rumex Central</i>	<i>Rumex Moved</i>	<i>Rumex Cut off</i>
<i>Accuracy</i>	94,3%	87,5%	83,8%
<i>Precision</i>	96,9%	91,8%	87,9%
<i>Recall</i>	94,2%	89,0%	88,9%
<i>F1-score</i>	95,4%	89,9%	87,5%

Table 3: Validation measures of the re-trained model

The confusion matrix and validation measures for the final model, trained and tested on both moved and partly cut off Rumex plants, are shown in figure 15 and table 4.

		Predicted Values	
		No Rumex	Rumex
Actual Values	No Rumex	True Negatives 208	False Positives 17
	Rumex	False Negatives 18	True Positives 188

n = 431

Figure 15. Confusion Matrix– Final model: Rumex moved and partly cut off

<i>Final Model - Rumex Moved and Cut off</i>	
<i>Accuracy</i>	91,9%
<i>Precision</i>	91,7%
<i>Recall</i>	91,3%
<i>F1-score</i>	91,5%

Table 4. Validation measures of the final model

In the confusion matrix in figure 15 it can be seen that the number of false positives and false negatives is very similar. This results in similar values for precision and recall, which can be seen in table 4. Table 4 also shows that the overall validation measures for the final model, containing both moved and cut off Rumex plants, are higher than the measures of the previous models.

4.3 Discussion

The final model shows that the retrained CNN is capable of detecting over 90% of all Rumex weeds in the test images, with a relatively low number of false positives. This accuracy is higher compared to some of the previous works, such as, Binch and Fox (2017) and Gebhardt & Kühbauch (2006), where the overall detection rate is less than 85%.

The work of Gebhardt & Kühbauch (2007), and Hiremath et al. (2013), performed better but the dataset used was obtained with close-range cameras and invasive data collection practices (human observers and ground robots). The proposed approach proposed is faster and provides a non-invasive way to collect data because plants are identified from an aerial high-resolution image obtained with small UAV.

The input images were selected in such a way that the results can give an indication of the results of using a sliding window approach over a whole field. This procedure will enable to automate all the Rumex detection and positioning workflow. The proposed approach is also robust to images with partly cut off Rumex plants in the classification window.

While interpreting the results, it should be enhanced that the images used have all been acquired under the same weather conditions and acquired from 10m elevation (surveying time 3 min.). It will be interesting to study if the model is robust enough to work under different weather or light conditions, as well as, different flying heights.

Next to that, the raw input images were used in this research. The only edits that were made were the cropping of Rumex locations from the orthomosaic and shifting the location. It might be interesting to see if increasing the size of the dataset using other methods of data augmentation can increase the quality of the model predictions.

5. CONCLUSIONS

This novel Rumex detection strategy based on transfer learning proves once again the adaptability of pre-trained deep convolutional networks for plant detection. However, in this study high-resolution images were acquired from UAVs instead of close-range images as has been shown in previous studies.

It was shown that the classification results obtained are a good starting point for Rumex detection within high-resolution aerial images, i.e., orthomosaics, using transfer learning, and a deep learning network specifically AlexNet.

In a next phase the training and testing dataset will be updated with images with increased heterogeneous characteristics, such as, spatial resolutions, light conditions, and other morphologic transformations.

We are confident that this novel approach for Rumex detection using UAV technology will benefit grasslands management and livestock practices.

ACKNOWLEDGEMENTS

This work was supported by the SPECTORS project (143081) which is funded by the European cooperation program INTERREG Deutschland-Nederland.

REFERENCES

- Ahmed, F., Kabir, M.H., Bhuyan, S., Bari, H., & Hossain, E., 2013. Automated weed classification with local pattern-based texture descriptors. *Int. Arab J. Inf. Technol.*, 11(1), 87-94.
- Binch, A., & Fox, C.W., 2017. Controlled comparison of machine vision algorithms for Rumex and Urtica detection in grassland. *Computers and Electronics in Agriculture*, 140, 123-138.
- Bond, W., & Grundy, A.C., 2001. Non-chemical weed management in organic farming systems. *Weed research*, 41(5), 383-405.
- Dyrmann, M., Karstoft, H., & Midtiby, H.S., 2016. Plant species classification using deep convolutional neural network. *Biosystems Engineering*, 151, 72-80.
- Gebhardt, S., & Kühbauch, W., 2007. A new algorithm for automatic Rumex obtusifolius detection in digital images using colour and texture features and the influence of image resolution. *Precision Agriculture*, 8(1-2), 1-13.
- Gebhardt, S., Schellberg, J., Lock, R., & Kühbauch, W., 2006. Identification of broad-leaved dock (Rumex obtusifolius L.) on grassland by means of digital image processing. *Precision Agriculture*, 7(3), 165-178.
- Ghazi, M.M., Yanikoglu, B., & Aptoula, E., 2017. Plant identification using deep neural networks via optimization of transfer learning parameters. *Neurocomputing*, 235, 228-235.
- Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y., 2016. *Deep learning* (Vol. 1). Cambridge: MIT press.
- Hiremath, S., Tolpekin, V.A., van der Heijden, G., & Stein, A., 2013. Segmentation of Rumex obtusifolius using Gaussian Markov random fields. *Machine vision and applications*, 24(4), 845-854.
- Huval, B., Wang, T., Tandon, S., Kiske, J., Song, W., Pazhayampallil, J., & Mujica, F., 2015. An empirical evaluation of deep learning on highway driving.
- Kamilaris, A., & Prenafeta-Boldú, F.X., 2018. Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147, 70-90.
- Kempenaar, C., Lotz, L.A.P., Van der Horst, C.L.M., Beltman, W.H.J., Leemans, K.J.M., & Bannink, A.D., 2007. Trade off between costs and environmental effects of weed control on pavements. *Crop Protection*, 26(3), 430-435.
- Lu, H., Fu, X., Liu, C., Li, L.G., He, Y.X., & Li, N.W., 2017. Cultivated land information extraction in UAV imagery based on deep convolutional neural network and transfer learning. *Journal of Mountain Science*, 14(4), 731-741.
- MathWorks, 2018. *Introducing Deep Learning with MatLAB*.
- Pan, S.J., & Yang, Q., 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10), 1345-1359.
- Poison Diaries, 2013. *Poisonous Weeds – Rumex Obtusifolius (Broadleaf Dock)*. Visited on 04-12-2018 through <http://thepoisondiaries.tumblr.com/post/40849549707/poisonous-weeds-rumex-obtusifolius-broadleaf>
- Evert, van, F.K., Polder, G., Van Der Heijden, G.W.A.M., Kempenaar, C., & Lotz, L.A.P., 2009. Real-time vision-based detection of Rumex obtusifolius in grassland. *Weed Research*, 49(2), 164-174.
- Evert, van, F.K., Samsom, J., Polder, G., Vijn, M., Dooren, H. J.V., Lamaker, A., & Lotz, L.A., 2010. A robot to detect and control broad-leaved dock (Rumex obtusifolius L.) in grassland. *Journal of Field Robotics*, 28(2), 264-277.
- Vedaldi, A., & Lenc, K., 2015. Matconvnet: Convolutional neural networks for matlab. In *Proceedings of the 23rd ACM international conference on Multimedia* (pp. 689-692). ACM.
- Weeds Act, 1959. Visited on 04-12-2018 through https://www.legislation.gov.uk/ukpga/1959/54/pdfs/ukpga_19590054_en.pdf

Revised January 2019