Article

Feature Learning Based Approach for Weed Classification Using High Resolution Aerial Images from a Digital Camera Mounted on a UAV

Calvin Hung *, Zhe Xu and Salah Sukkarieh

The Australian Centre for Field Robotics, University of Sydney, The Rose Street Building J04, Sydney, NSW 2006, Australia; E-Mails: z.xu@acfr.usyd.edu.au (Z.X.); salah@acfr.usyd.edu.au (S.S.)

* Author to whom correspondence should be addressed; E-Mail: c.hung@acfr.usyd.edu.au; Tel.: +61-2-9036-7178; Fax: +61-2-9351-7474.

External Editors: Arko Lucieer and Prasad S. Thenkabail

Abstract: The development of low-cost unmanned aerial vehicles (UAVs) and light weight imaging sensors has resulted in significant interest in their use for remote sensing applications. While significant attention has been paid to the collection, calibration, registration and mosaicking of data collected from small UAVs, the interpretation of these data into semantically meaningful information can still be a laborious task. A standard data collection and classification work-flow requires significant manual effort for segment size tuning, feature selection and rule-based classifier design. In this paper, we propose an alternative learning-based approach using feature learning to minimise the manual effort required. We apply this system to the classification of invasive weed species. Small UAVs are suited to this application, as they can collect data at high spatial resolutions, which is essential for the classification of small or localised weed outbreaks. In this paper, we apply feature learning to generate a bank of image filters that allows for the extraction of features that discriminate between the weeds of interest and background objects. These features are pooled to summarise the image statistics and form the input to a texton-based linear classifier that classifies an image patch as weed or background. We evaluated our approach to weed classification on three weeds of significance in Australia: water hyacinth, tropical soda apple and serrated tussock. Our results showed that collecting images at 5–10 m resulted in the highest classifier accuracy, indicated by F1 scores of up to 94%.

1. Introduction

The development of low-cost unmanned air vehicles (UAVs) and light-weight imaging sensors in the past decade has resulted in significant interest in their use for remote sensing applications. These applications include vegetation mapping [1–5], archeology [6–8], meteorology [9] and high risk activities, such as volcano monitoring [10] and forest fire mapping [11]. This interest arises due a number of advantages of small UAVs over other remote sensing platforms. Firstly, small UAVs require minimal infrastructure and can thus be operated from remote or rugged locations. Secondly, small UAVs can be low cost, allowing them to be used as disposable data collection platforms in high risk environments. Further, they can be operated frequently (weather permitting) at low altitudes, allowing the collection of data at high spatial and temporal resolutions.

While significant attention has been paid to the collection, calibration, registration and mosaicking of remote sensing data, the interpretation of these data into useful information for the above-mentioned scientific, ecological management and other activities can still be a laborious task. Common approaches use the object-based image analysis (OBIA) framework [12–14]. The use of OBIA for any given application requires an extensive manual design process, comprising segment size tuning, feature selection and rule-based classifier design. Though this rule set can be re-used for different datasets, the process must be repeated for each new application.

Conventional approaches to image analysis require the selection of appropriate features to deliver high classification accuracy. These features can be the spectral bands at which data are collected and analysed or derivatives of these data. This feature selection process is a laborious process, as it is highly dependent on the application (e.g., vegetation classification and geological mapping) and the types of data being used (e.g., multi-spectral and hyper-spectral).

Features can be either designed by experts or learned from the imagery itself via feature learning algorithms. In practice, the design and selection process requires a large amount of experience from experts in the field and an understanding of the physical process that governs how light is reflected from the materials. Examples of designed features commonly used in remote sensing applications include spectral indices (for example, NDVI [15], EVI [16] and NDWI [17]), spatial features (for example, texture and shape [18], wavelets [19] and Gabor texture features [20]. While hand-designed features are a proven approach to the feature selection problem, there are two disadvantages to the designed approach. Firstly, the feature design process can be time consuming and application dependent. Secondly, a given feature may be parameterised on a large set of tuning parameters, each of which may be application specific and, thus, difficult to generalise. Classification performance has been shown to be sensitive to these parameters [21].

In this paper, we propose an alternative approach to minimise the effort in feature design by coupling a remote sensing UAV with feature learning (via the sparse autoencoder [22–24]). Feature learning has

been applied successfully in computer vision; the technique has achieved state-of-the-art performance in object detection [25], image classification [26] and object recognition [27] tasks. We demonstrate the utility and generality of such a system by using it for automated classification of three invasive weed species on the north-west slopes of New South Wales, Australia.

Invasive weeds species have major economic and environmental impacts in Australia and around the world, including competition with native plant species, choking irrigation infrastructure, reducing agricultural yields and affecting the health of livestock. The targets of this study are three weeds of national significance in Australia: (1) water hyacinth [28]; (2) tropical soda apple [29]; and (3) serrated tussock [30]. We study the use of small UAVs for the classification of these weeds, as UAVs provide capabilities that are complementary to both traditional remote sensing assets, such as satellites and manned aircraft, on one end, and field work, on the other. In contrast to satellites and manned aircraft, small UAVs can operate at lower altitudes, allowing end users to collect higher spatial resolution data, which is critical in the classification of small or localised weed presence. Further, small UAVs are typically cheaper and easier to own and operate, making their use more available to local organisations that manage these weeds. While their ability to cover large areas is limited when compared against manned aircraft or satellites, small UAVs can cover larger areas in less time than ground-based field work.

While work-flows for the collection, calibration and mosaicking of aerial imagery from small UAVs are well established, a human expert is still required to examine the data products in order to find the weeds of interest. On the other hand, in an automated weed classification pipeline, an expert provides positive examples of the weed to train a computer vision algorithm. The algorithm then assigns semantic meaning to patches in the image: either a weed of interest or not. The computer vision algorithm applies feature learning to generate a bank of image filters that allows for the extraction of features that discriminate between the weeds of interest and background objects. These features are pooled to summarise the image statistics and form the input to a texton-based linear classifier that classifies an image patch as weed or background.

There are two major contributions of this paper:

1. Theoretical: We develop a general semi-supervised approach for weed classification using remote sensing data. The approach couples feature learning, the use of filter banks and the use of textons for weed classification.
2. Experimental: We contribute a comprehensive evaluation of the proposed approach using real data. We analysed the performance of the classification results as a function of a key tuning parameter, the classifier window size, for each altitude. This is useful in defining the requirements of sensor resolution and flight altitude for future studies.

2. Methodology

Our methodology can be decomposed into three steps: (1) data collection; (2) image pre-processing; and (3) weed classification. In the data collection step, aerial imagery and geo-registration data are gathered at each site. The aerial imagery for each site is then mosaicked and geo-registered into a single ortho-photo. The mathematical parameters for projecting each aerial image into the ortho-photo are also

determined. In parallel, a weed classifier has been developed and tested in this paper. In this study, we have experimented with different altitude and patch window size settings for water hyacinth, tropical soda apple and serrated tussock. Our methodology is summarised in Figure 1.

Figure 1. Overview of our methodology. A multi-rotor UAV was used to collect aerial images. The images were mosaicked and registered into an ortho-photo. Further, a weed classifier has been developed to classify patches of water hyacinth, serrated tussock and tropical soda apple.
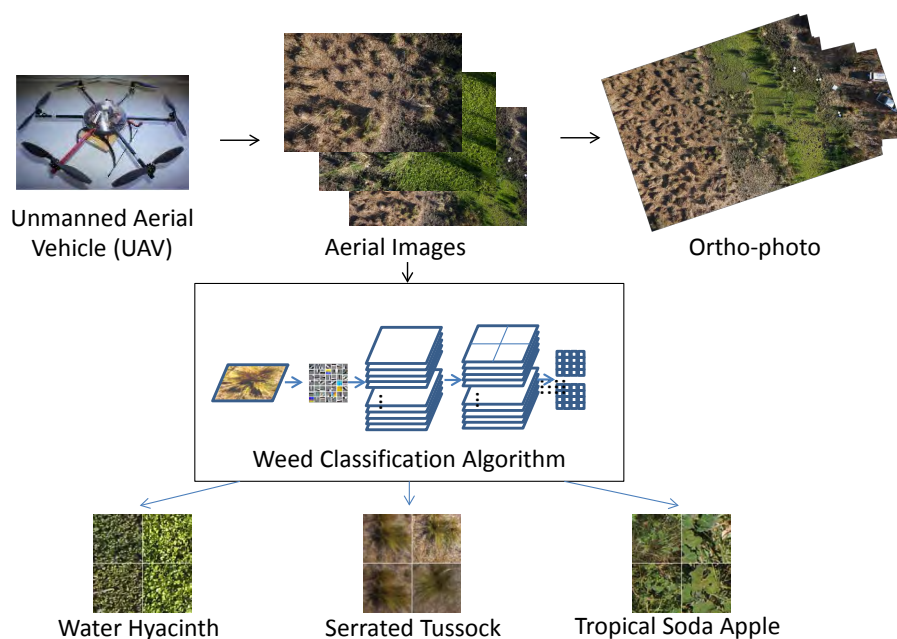
2.1. Data Collection

A multi-rotor UAV (Mikrokopter Hexacopter) was used to gather all of the imagery used in this study. The UAV is illustrated in Figure 1, and its technical specifications are presented in Table 1. We collected aerial images with a downward pointing, off-the-shelf camera (Sony NEX-7). The specifications of this camera are presented in Table 2. The UAV was flown over sites of interest at various altitudes between 5 and 30 m and captured an image approximately every 6 s.

Table 1. UAV technical characteristics.

Make, Model	Mikrokopter Hexacopter
Gross weight	1.5 kg
Dimensions	80 cm × 80 cm
Endurance	6 min
Typical speed	1.0 m/s
Typical operating altitude	20 m
Typical range	<100 m

Table 2. Camera technical characteristics.

Make, Model	Sony NEX 7
Resolution	6000×4000 pixels
Lens	16 mm $f2.8$
Angular field of view	76×55 deg.
Typical foot-print size (at 20 m altitude)	30×20 m
Typical spatial resolution (at 20 m altitude)	5 mm/pixel

2.2. Data Pre-Processing

In the data pre-processing step, we mosaic the aerial images to form a coherent picture of the survey sites. We also estimate the altitude at which each image frame was captured to support subsequent analysis on the effect of altitude on classifier performance.

2.2.1. Mosaicking

In the data pre-processing stage, we mosaic and geo-register the images into a single ortho-photo. The aerial images were geo-registered using ground control points. In this study, the markers were laid out around the weed(s) of interest, such that there is typically at least one marker visible in each image frame. The positions of these makers were surveyed using a hand-held GPS receiver.

Image mosaicking and geo-registration was performed using off-the-shelf software, Agisoft PhotoScan. The geo-registration used information from: (1) the ground control point locations; and (2) the scale provided by the (known) size of the ground control point markers.

2.2.2. Altitude Estimation

Individual image frames were also grouped by the altitude at which they were taken. While the UAV was equipped with GPS localisation, the relative inaccuracy of GPS (±20 m) compared against the altitude at which the missions were flown (5–30 m) means the GPS altitude cannot be used for the purposes of this study. Instead, we used the known size of the ground control point markers (1×1 m) and their size in the image frames to estimate the slant range at which the image was taken. We use the slant range as a proxy for the altitude in this study, as the roll and pitch angles of the UAV were small. For example, a marker that is 750 pixels wide indicated that Figure 2 was taken at 5.1-m altitude.

Not all image frames included a visible ground control point marker. To estimate the altitude of the UAV for these frames, we applied first order interpolation between the altitudes of the previous and subsequent frames. This interpolation assumes that the altitude of the UAV changes linearly between each frame.

With an altitude estimate per image, the images were then grouped into different altitude brackets for further analysis. The altitude brackets are listed in Table 3.

Figure 2. Estimating the altitude of the UAV using the dimensions of a ground control point marker. Knowing the dimensions of the marker and the optical parameters of the camera allowed us to estimate the altitude at which each image was captured.

Table 3. In this study, we grouped the images in the following altitude brackets.

Altitude Group	Range	Corresponding Pixel Size
5 m	2.5–7.5 m	1.3 mm
10 m	7.5–15 m	2.6 mm
20 m	15–25 m	5.2 mm
30 m	25–35 m	7.8 mm

2.3. Classification

In this paper, we extracted patches containing either weeds and non-weed background from the collected images to evaluate the weed classifier. Feature learning was applied to generate the filter bank, followed by pooling to summarise the image statistics before passing to the texton-based linear classifier.

2.3.1. Feature Learning and Pooling

It has been shown that in classification problems, the algorithms perform better on meaningful features instead of classifying the raw (noisy) data. For example, in RGB image segmentation problems, it is standard to perform classification using colour, texture and shape features.

The state-of-the-art approach to generating representative features for image classification is via feature learning. Feature learning is a type of deep learning [23,24,31,32] in which a set of trainable modules implementing complex non-linear functions are stacked together to capture the underlying patterns in unlabelled data. These patterns are represented as filter banks, which are convolved with the image to extract feature responses.

In this paper, we use a standard approach for learning the filter banks. In this approach, a sparse autoencoder learns the filter bank from an established training dataset, CIFAR10 [33]. The sparse autoencoder is an algorithm based on neural networks that minimises the squared reconstruction error

with an extra sparsity constraint [23]. Examples of the learned features used in this paper is shown in Figure 3.

Figure 3. Filter bank obtained using feature learning. The learning algorithm has automatically generated filters that extract edge information (the grey-scale filters), colour information (the coloured filters) and texture information.
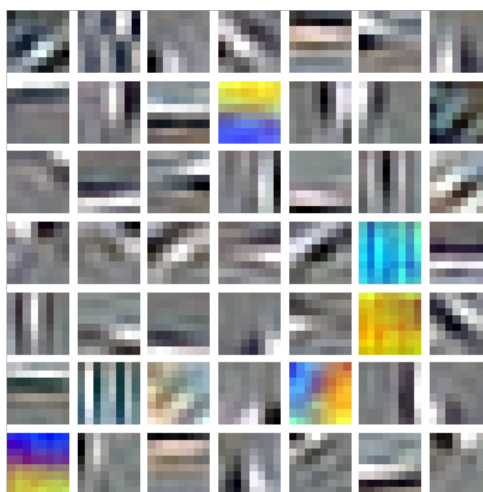
While feature responses can be used with the classifier directly, this approach can be computationally expensive. In the paper, we use the standard practice of pooling [34,35] to summarise the feature responses before applying the classifier. Pooling is an operation that divides an image into regions and collects statistics for each region. In this paper, we use max-pooling, which summarises each region by the maximum feature response. The feature learning, feature extraction and pooling pipeline are shown in Figure 4.

Figure 4. Feature extraction is performed by convolving the learned feature bank over the image. The high dimensional feature responses are then summarised using pooling before classification.

Original Image

Convolved with Filter Bank

High Dimensional Feature Responses

Pooling to summarise the statistics

Final Feature Vector
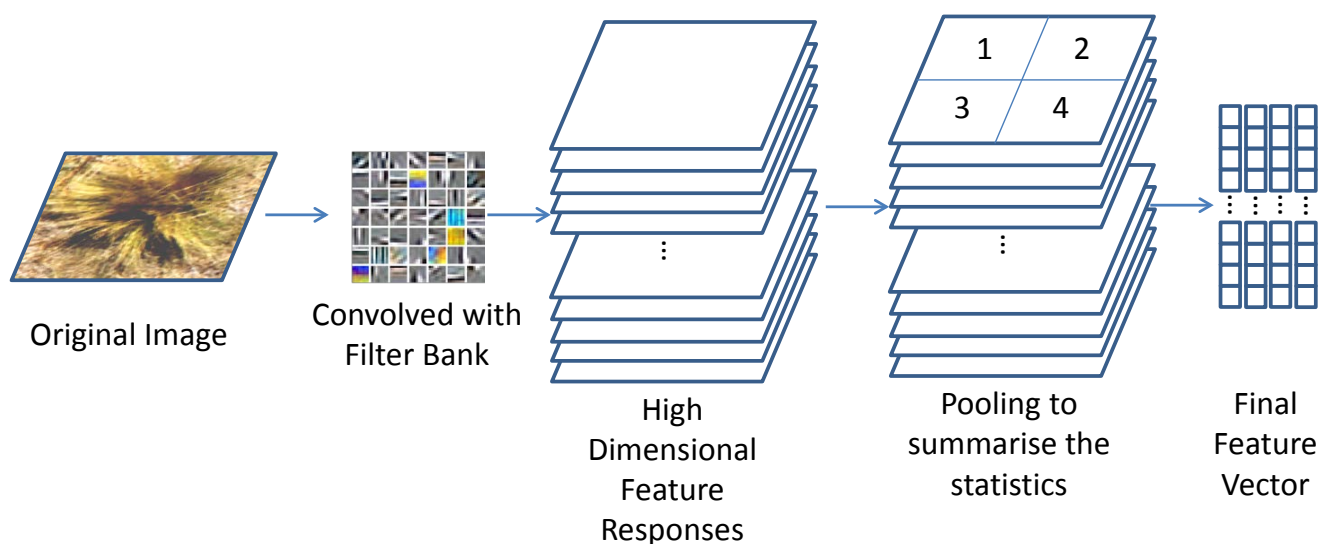
2.3.2. Classification

We applied the texton approach to classifying feature responses, as it has been shown to result in a high accuracy for image classification tasks [36,37]. During the training phase, K-means clustering is performed over the pooled feature responses from both the weed and non-weed class. The centroids of each cluster are called textons. A class is modelled by one or more textons. Examples of textons that model the water hyacinth class and non-weed class are shown in Figure 5. The histograms are visually distinct, demonstrating that the classes are easily separable. During the prediction stage, the input feature responses are compared against the texton(s) modelling each class. The most similar model becomes the predicted class.

Figure 5. The histogram of the texton representing the water hyacinth and non-weed class. The histograms are visually distinct, indicating that the classes are easily separable.
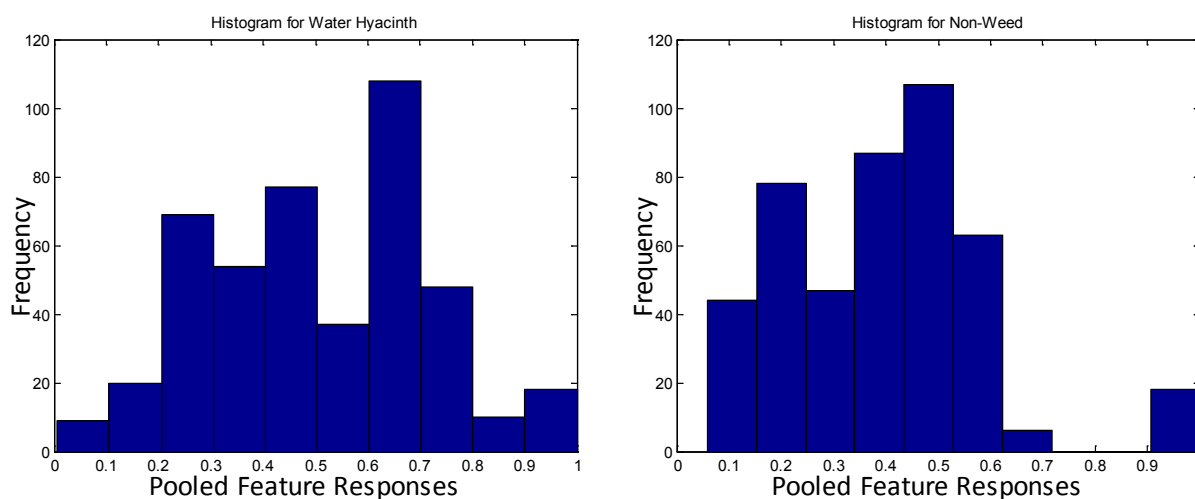
3. Experimental Setup

The experimental setup is defined by four components: (1) the survey sites; (2) the approach used to extract training and evaluation data; (3) the classification algorithm parameter selection; and (4) the definition of the accuracy metric.

3.1. Survey Sites

The survey sites were concentrated around northern-inland New South Wales, Australia. Multiple flights were performed at the sites to collect data at different times of the day. The water hyacinth sites were at Moree (29°14′33.647″S 149°17′51.552″E). The total survey area was approximately 1500 m^2, of which 300 m^2 was infested with water hyacinth. The serrated tussock site was at Armidale (30°36′44.881″S 151°42′2.845″E). The total survey area was approximately 2400 m^2, and 16 serrated tussock plants at different growth stages were discovered on this site. The tropical soda apple sites were at Carrai (30°45′50.434″S 152°20′25.250″E) and comprised a total survey area of approximately 4800 m^2 in which 18 tropical soda plants were discovered.

3.2. Training and Evaluation Data

The image collection was divided into patches, and each patch was labelled as either weed or non-weed. The labelling process was done visually for weed species that were easily distinguishable from the background, for example water hyacinth and serrated tussock.

Tropical soda apple plants were more difficult to distinguish visually. To provide reliable ground truth data, we performed two flights for each site. The first flight was used to collect the raw image data used for evaluating the classification algorithm. Weed experts then marked tropical soda apple plants with red tape. A second flight was performed to collect aerial images with these markings. An example of the raw and marked images is shown in Figure 6. The class labelling was guided by these markings. The images from the second flight were not used as part of the evaluation, as we did not want to bias the classifier by associating the unnatural red tape markings with tropical soda apple plants.

Figure 6. Two flights were performed at the tropical soda apple sites. The first flight was used to collect the raw image data for algorithm evaluation (**left**) and the second flight to identify the tropical soda apple plants, which had been marked by weed experts (**right**).

Raw Image   Labelled Image
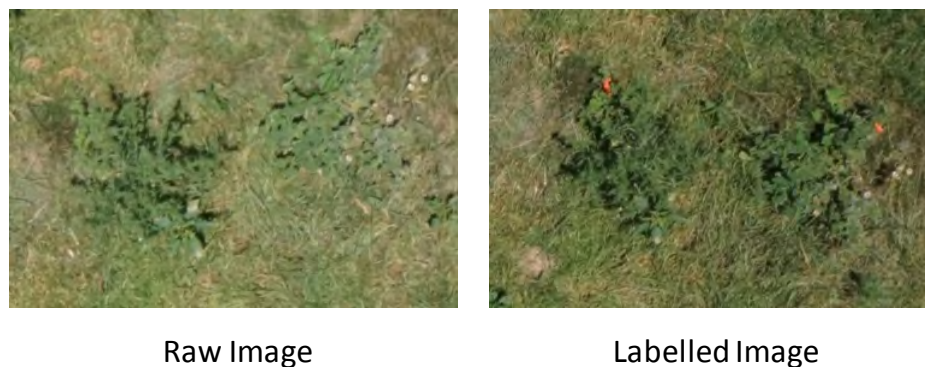
Note that there were replications in the sampled plants, because some were observed across different images; we assume the variations (different orientations, view angles, lighting condition and flight altitudes) between each observation were adequate for us to treat them as individual samples. On average, 100 image patches of positive and negative examples were collected for each species at each altitude setting for training and evaluation.

3.3. Algorithm Parameters

For feature learning, the size of the filter bank was 50. For pooling, each image was divided into 9 regions (3×3). Thus, the final feature vector has 450 dimensions. For classification, 5 textons were used to model each class (weed and non-weed).

3.4. Accuracy Metric

Two-fold cross-validation was used to evaluate the classification algorithm. In two-fold cross-validation, half of the patches were randomly selected as the training set and the remainder as the test set. The classification accuracy was measured with precision, recall and F1 score. Precision

is the ratio between the number of accurately classified weeds and the total number of image patches (including false positives). Recall is the ratio between the number of accurately classified weeds and the total number of weeds that should have been correctly classified. The F1 score is the harmonic mean of precision and recall.

Experiments were performed for 4 different altitudes and 4 different window size settings. For each setting, 20 runs were performed and the accuracy metric was averaged. The F1 scores are shown for each setting, and the detailed precision and recall breakdowns are shown for the best window size setting.

4. Results

Mosaics of the three survey sites are shown in Figure 7. The quantitative classifier performance with different pixel and window sizes for each weed species and the performance scores for the best window size are also shown in Table 4. Examples of correctly classified weed, non-weed and misclassified patches are shown in Figure 8.

4.1. Water Hyacinth

Our results showed that the classification accuracy improved with larger window sizes. We hypothesise that this is because water hyacinth grows in large patches at this site. A larger window size increases the area of observation, which leads to more consistent colour and texture features. This consistency, in turn, leads to higher classification accuracies. Further, we found that the accuracy did not change significantly with increasing altitude. The algorithm was able to distinguish water hyacinth from other classes at different altitudes with F1 scores greater than 90%. The main reason is that these data are amenable to classification; most of the green area was infested by water hyacinth.

4.2. Serrated Tussock

The classification accuracy improved with increasing the window size up to 384×384 pixels. Interestingly, this is the dimension at which most of the serrated tussock plants can be fully observed (shown in Figure 9). The classification accuracy reduced slightly with increasing altitude, as fewer details of the serrated tussock plant could be resolved.

4.3. Tropical Soda Apple

The results showed that the window size did not affect the classification performance significantly. Further, our results showed that the altitude had a significant impact on the classification accuracy; the best results were obtained at 10 m. Figure 8 shows that only images collected below 10 m had enough resolution to see the prickly leaves. We hypothesise that the ability to resolve this distinctive feature aids classification performance. It becomes increasingly difficult to resolve the distinctive leaves of the tropical soda apple from above 10 m.

Figure 7. Image mosaic of the surveyed weed sites. Images taken from different flight altitudes were mosaicked, with images captured at higher altitudes providing coverage, whereas images captured at lower altitudes had higher resolution.
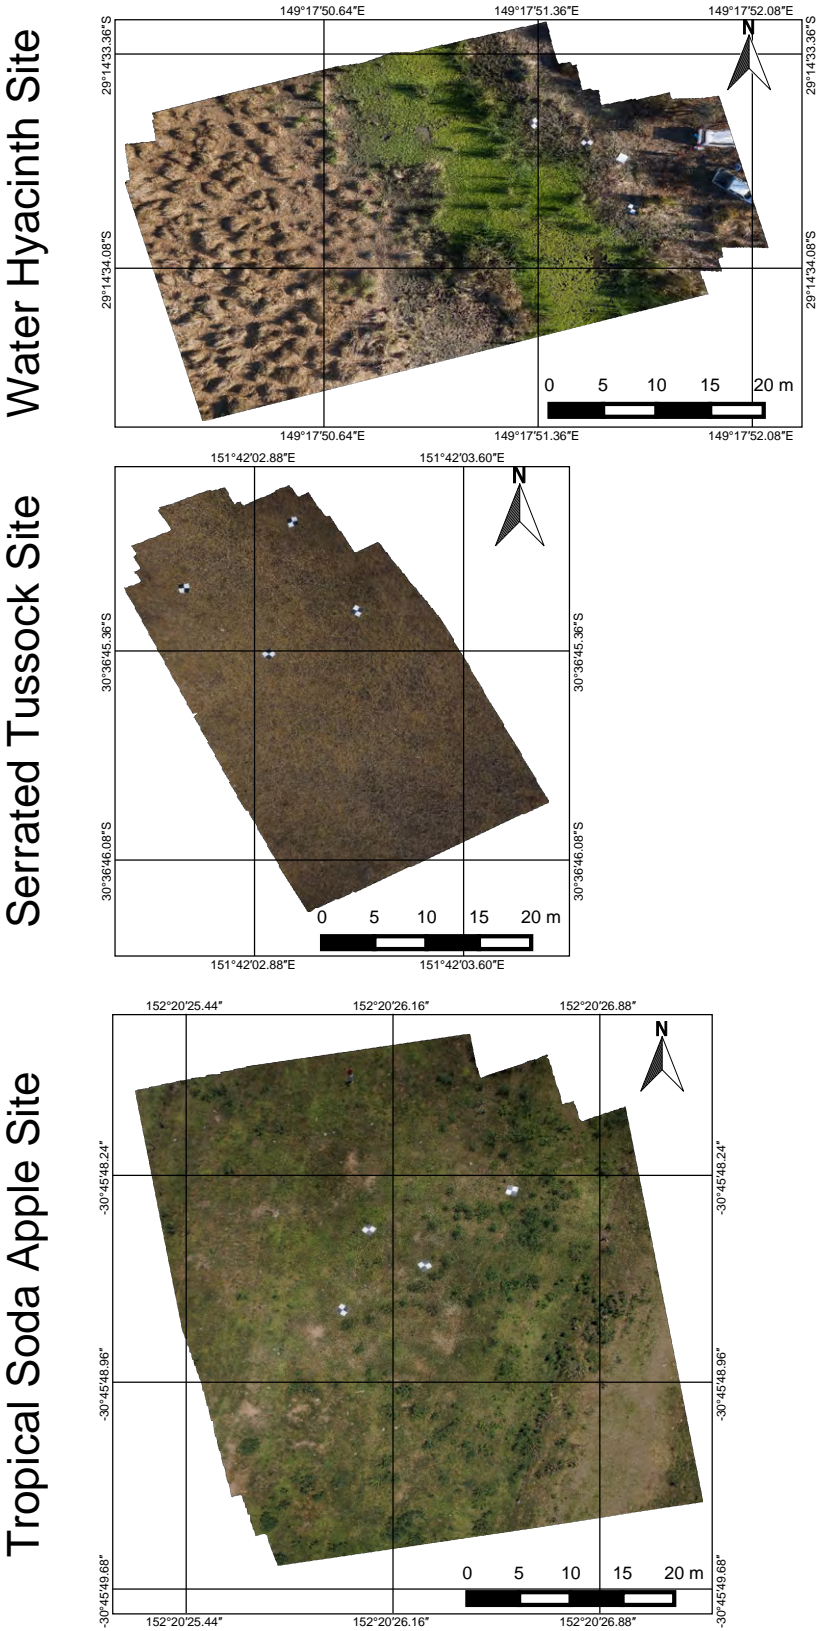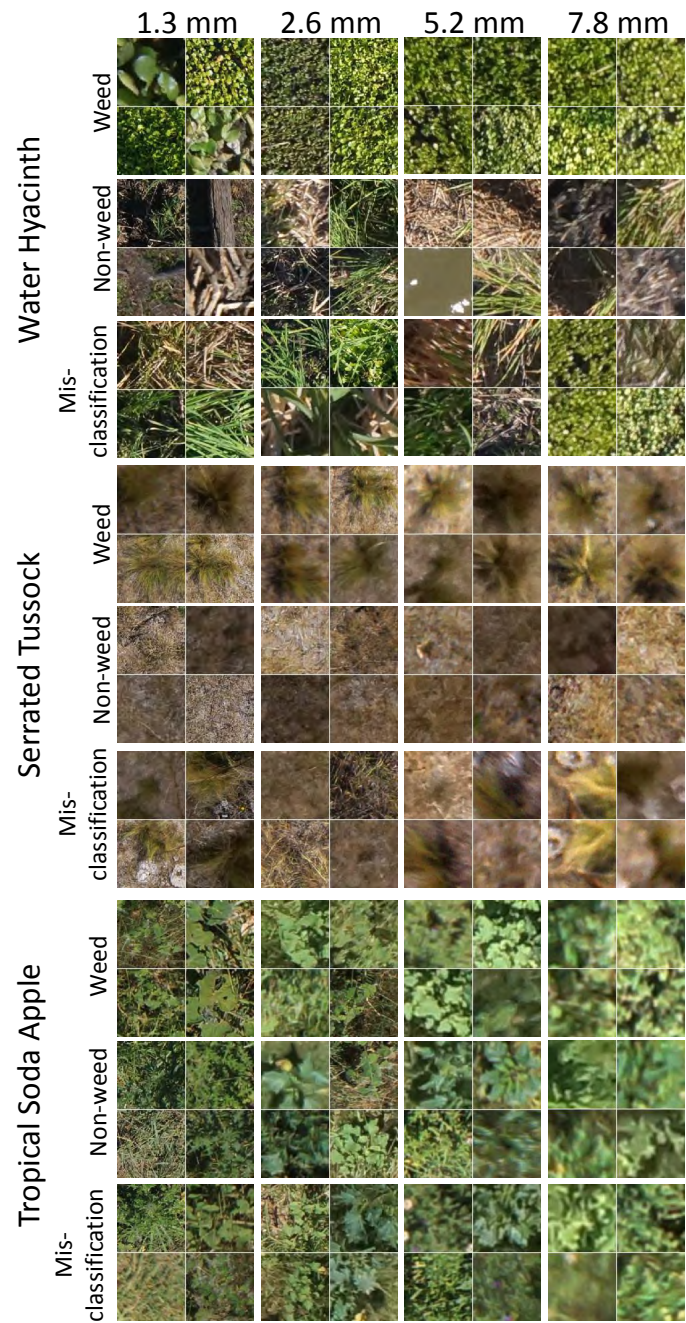
Figure 8. Qualitative results: examples of correctly classified weeds, non-weeds and mis-classification for water hyacinth, serrated tussock and tropical soda apple at different pixel size.


The classification results in Table 4 showed that the classifier did not perform as well as for the other weed species. The reason was that the current feature bank is populated with mostly colour and texture filters. While these cues were useful for water hyacinth and serrated tussock, they were not enough to describe the subtle differences between tropical soda apple and other plants of similar appearance. Further developments of a specialist classifier to incorporate the extra shape information of the leaves similar to the work on flower classification [38] can potentially improve the classifier performance.

Table 4. Quantitative results: classification performance at different window pixel sizes for each weed species. The precision and recall values for the best window size setting are also shown. These results show that window size is more important than altitude in determining classifier performance for water hyacinth, while both window size and altitude were important for serrated tussock and tropical soda apple.

Pixel Size/Window Size	128	256	384	512	Precision	Recall
Water hyacinth						
1.3 mm	80.28	82.39	92.17	**91.77**	88.59	95.19
2.6 mm	81.33	81.24	90.90	**94.31**	91.79	96.98
5.2 mm	80.00	81.63	89.90	**91.45**	86.59	96.89
7.8 mm	79.98	80.59	87.99	**90.00**	84.39	96.42
Serrated Tussock						
1.3 mm	91.73	91.86	**92.90**	87.60	93.02	92.79
2.6 mm	87.46	89.21	**92.08**	92.98	93.13	91.07
5.2 mm	84.35	87.49	**90.70**	90.01	92.02	89.42
7.8 mm	83.95	82.54	87.23	**88.54**	88.24	86.25
Tropical Soda Apple						
1.3 mm	65.95	62.74	**62.91**	60.85	63.09	62.39
2.6 mm	71.33	**72.19**	71.54	72.05	70.18	74.31
5.2 mm	69.29	**69.82**	69.54	68.20	67.26	72.59
7.8 mm	68.17	**69.29**	68.86	69.10	68.42	70.17

Figure 9. Serrated tussock patches sampled with different window sizes. The classification accuracy improved with increasing the window size up to 384 pixels. This is the dimension at which most of the serrated tussock plants can be fully observed.
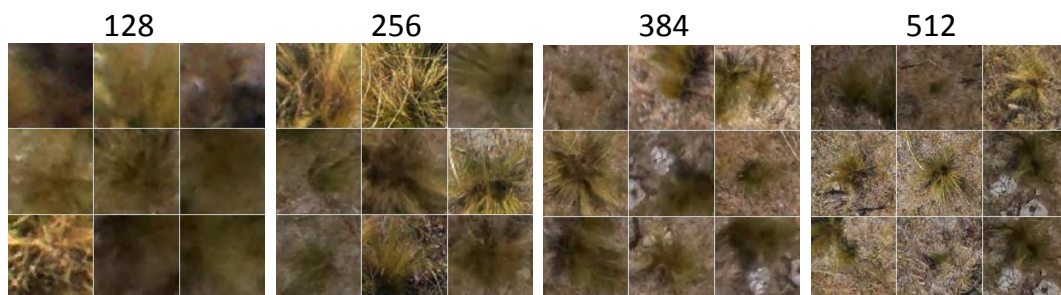
5. Discussion

The two objectives of this study are to demonstrate the effectiveness of the proposed feature learning-based weed detection algorithm and to determine the minimal flight altitude (and the

corresponding image resolution) that resulted in the highest classifier performance for three weed species.

Identical algorithm settings were used for all three weed species, without segment size tuning, feature selection and rule-based classifier design, typically required in the standard approach [12–14]. The models learned from the data and training labels, thus minimising the manual effort in algorithm design across different applications. The results showed that the proposed feature learning weed classification algorithm was effective over different weed species, with a classification accuracy of 94.3% for water hyacinth, 92.9% for serrated tussock and 72.2% for tropical soda apple.

This study has also established the resolution that resulted in the highest classifier performance for three weed species. For area infestations, such as water hyacinth, the classification algorithm was less sensitive to the altitude setting and more sensitive to the window size. The large window size allows a large area of observation, and the more consistent colour and texture features leads to higher classification accuracies.

For the classification of individual plants, such as serrated tussock and tropical soda apple, both the window size and altitude settings are important. The altitude dictates the amount of detail that can be resolved in the images, which is important, because certain species can appear similar without those details. It can also be important if additional features that describe the specified properties of the plants (such as leaf shape and spikes of tropical soda apple) are incorporated. The window size should be selected according to the physical size of the target plants; it should be large enough to include the entire plant, but not too large to introduce noise from the neighbouring plants. Interestingly, the scale dependency of the classification performance has also been observed in OBIA studies [39], where different performances were achieved with different segmentation sizes.

Our results showed that classification accuracy depended highly on how distinguishable the weed was from its surroundings. Thus, the survey strategy is as important as the classification algorithm. There are three strategies that can be applied to maximise the distinction between the weeds and its surroundings. The first (and potentially best) approach is to time the survey to coincide with seasonal changes that maximise the difference in appearance, for example during the flowering season. Secondly, we can perform the survey at lower altitudes to obtain more detail at the cost of lower coverage. The third approach is to use sensors with higher resolution to obtain the same amount of detail at higher altitudes. One disadvantage of flying at higher altitudes is that the observations become more sensitive to the vehicle's motion. Any perturbations in the platform pose at higher altitudes will cause large displacements in the sensor's orientation and may cause motion blur. Therefore, to obtain the same amount of detail at higher altitudes will require a higher resolution sensor and a more stable vehicle.

5.1. Limitations and Future Work

This paper presented a proof-of-concept and feasibility study in weed detection using the ultra-high resolution aerial images processed by the proposed image analysis pipeline. There are two main limitations needed to be addressed in future studies for the system to become operational.

The first limitation relates to the UAV platform. While the light-weight Hexacopter is a useful platform for concept studies, it has a short flight duration and is unable to cover large areas. In our

future work, we envision collecting data with a larger fixed-wing UAV that has a longer flight duration and is able to cover larger areas.

The second limitation relates to the algorithm. Two simplifications were made in this study. The first simplification was that we performed patch classification on the pre-selected example patches instead of having the algorithm search the entire image. A truly automated image processing pipeline requires a vegetation/background segmentation algorithm to automate this task. The second simplification was that the training/evaluation dataset was balanced in the sense that it contained an equal number of positive and negative examples. In reality, the dataset is more likely to be dominated by the non-weed class. This simplification was made because this study focused on validating the effectiveness of the patch classifier. In our future work, we will collect additional training data to address this imbalance. Our future work will revolve around relaxing some of the simplifications made in this study.

6. Conclusions

In this paper, we developed a learning-based algorithm for weed classification. While work-flows for the collection, calibration and mosaicking of aerial imagery from small UAVs are well established in the field of remote sensing, the amount of manual effort required to select features and design rule-sets remains a significant bottleneck in the data processing pipeline. In this paper, we addressed this bottleneck by applying a learning-based algorithm that was trained using positive and negative examples of weed images. This algorithm applied feature learning to generate a bank of image filters that allowed the extraction of relevant features that discriminate between weeds of interest and the background objects. These features were pooled to summarise the image statistics and formed the input to a texton-based linear classifier that classified each image patch as weed or background. This learning approach minimised the manual effort required in algorithm design for different weed species.

We demonstrated that this image classification algorithm was able to correctly classify weeds of interest from remote sensing data collected from a Sony Nex-7 digital camera mounted on a small UAV. This system was evaluated on three different weed species: water hyacinth, serrated tussock and tropical soda apple. As part of this evaluation, we collected data from altitudes of 5–30 m and experimented with classifier parameters to determine the best altitude and classifier tuning. Our results showed that for area infestations, such as water hyacinth, the classifier worked well at all altitudes and that larger classifier window sizes resulted in more accurate results. With the best window size, F1 scores in excess of 90% were achieved. For individual plants, such as serrated tussock and tropical soda apple, both altitude and window size settings were important. With the best window size, a F1 score of more than 90% was achieved for serrated tussock. A lower F1 score of 70% was achieved for the tropical soda apple dataset because of the similarity between the weed and surrounding plants.

This study is unique because: (1) it is the first study to use a feature learning-based approach on weed classification; and (2) we have performed a comprehensive evaluation of the algorithm using real data. Our findings are useful in defining the requirements of sensor resolution and flight altitude for future studies. For future work, we would like to extend this study to additional weed species to improve the robustness of the learning algorithm.

Acknowledgments

Author Contributions

Zhe Xu and Salah Sukkarieh organised the data collection campaign. Zhe Xu and Calvin Hung performed the UAV operation. Calvin Hung and Salah Sukkarieh developed the classification algorithm, Calvin Hung, Zhe Xu and Salah Sukkarieh wrote the manuscript.

Conflicts of Interest

The authors declare no conflict of interest.

References

1. Turner, D.; Lucieer, A.; Malenovskỳ, Z.; King, D.H.; Robinson, S.A. Spatial co-registration of ultra-high resolution visible, multispectral and thermal images acquired with a micro-UAV over Antarctic Moss Beds. *Remote Sens.* **2014**, *6*, 4003–4024.
2. Bryson, M.; Reid, A.; Ramos, F.; Sukkarieh, S. Airborne vision-based mapping and classification of large farmland environments. *J. Field Robot.* **2010**, *27*, 632–655.
3. Hung, C.; Bryson, M.; Sukkarieh, S. Multi-class predictive template for tree crown detection. *ISPRS J. Photogram. Remote Sens.* **2012**, *68*, 170–183.
4. Zhang, N.; Wang, M.; Wang, N. Precision agriculture a worldwide overview. *Comput. Electron. Agric.* **2002**, *36*, 113–132.
5. Grenzdörffer, G.; Engel, A.; Teichert, B. The photogrammetric potential of low-cost UAVs in forestry and agriculture. *Int. Arch. Photogram. Remote Sens. Spat. Inform. Sci.* **2008**, *31*, 1207–1214.
6. Sauerbier, M.; Eisenbeiss, H. UAVs for the documentation of archaeological excavations. *Int. Arch. Photogram. Remote Sens. Spat. Inform. Sci.* **2010**, *38*, 526–531.
7. Lambers, K.; Eisenbeiss, H.; Sauerbier, M.; Kupferschmidt, D.; Gaisecker, T.; Sotoodeh, S.; Hanusch, T. Combining photogrammetry and laser scanning for the recording and modelling of the Late Intermediate Period site of Pinchango Alto, Palpa, Peru. *J. Archaeol. Sci.* **2007**, *34*, 1702–1712.
8. Eisenbeiss, H.; Zhang, L. Comparison of DSMs generated from mini UAV imagery and terrestrial laser scanner in a cultural heritage application. *Int. Arch. Photogram. Remote Sens. Spat. Inform. Sci.* **2006**, *36*, 90–96.

9. Spiess, T.; Bange, J.; Buschmann, M.; Vorsmann, P. First application of the meteorological Mini-UAV "M2AV". *Meteorol. Z.* **2007**, *16*, 159–169.

10. Caltabiano, D.; Muscato, G.; Orlando, A.; Federico, C.; Giudice, G.; Guerrieri, S. Architecture of a UAV for volcanic gas sampling. *Journal name* **2005**. *1*, 739–744.

11. Casbeer, D.; Kingston, D.; Beard, R.; McLain, T. Cooperative forest fire surveillance using a team of small unmanned air vehicles. *Int. J. Syst. Sci.* **2006**, *37*, 351–360.

12. Laliberte, A.S.; Rango, A. Image processing and classification procedures for analysis of sub-decimeter imagery acquired with an unmanned aircraft over arid rangelands. *GISci. Remote Sens.* **2011**, *48*, 4–23.

13. Laliberte, A.S.; Herrick, J.E.; Rango, A.; Winters, C. Acquisition, orthorectification, and object-based classification of unmanned aerial vehicle (UAV) imagery for rangeland monitoring. *Photogram. Eng. Remote Sens.* **2010**, *76*, 661–672.

14. Peña, J.M.; Torres-Sánchez, J.; de Castro, A.I.; Kelly, M.; López-Granados, F. Weed mapping in early-season Maize fields using object-based analysis of Unmanned Aerial Vehicle (UAV) images. *PLoS One* **2013**, *8*, e77151.

15. Tucker, C. Red and photographic infrared linear combinations for monitoring vegetation. *Remote Sens. Environ.* **1979**, *8*, 127–150.

16. Huete, A.; Justice, C.; Liu, H. Development of vegetation and soil indices for MODIS-EOS. *Remote Sens. Environ.* **1994**, *49*, 224–234.

17. Gao, B. NDWI a normalized difference water index for remote sensing of vegetation liquid water from space. *Remote Sens. Environ.* **1996**, *58*, 257–266.

18. Davis, S.; Landgrebe, D.; Phillips, T.; Swain, P.; Hoffer, R.; Lindenlaub, J.; Silva, L. *Remote Sensing: The Quantitative Approach*; McGraw-Hill International Book Co.: New York, NY, USA, 1978; Volume 1, p. 405.

19. Ouma, Y.O.; Tetuko, J.; Tateishi, R. Analysis of co-occurrence and discrete wavelet transform textures for differentiation of forest and non-forest vegetation in very-high-resolution optical-sensor imagery. *Int. J. Remote Sens.* **2008**, *29*, 3417–3456.

20. Reis, S.; Taşdemir, K. Identification of hazelnut fields using spectral and Gabor textural features. *ISPRS J. Photogram. Remote Sens.* **2011**, *66*, 652–661.

21. Carleer, A.; Debeir, O.; Wolff, E. Assessment of very high spatial resolution satellite image segmentations. *Photogram. Eng. Remote Sens.* **2005**, *71*, 1285–1294.

22. Goodfellow, I.; Le, Q.; Saxe, A.; Lee, H.; Ng, A. Measuring invariances in deep networks. *Adv. Neural Inform. Process. Syst.* **2009**, *22*, 646–654.

23. Lee, H.; Ekanadham, C.; Ng, A. Sparse deep belief net model for visual area V2. *Adv. Neural Inform. Process. Syst.* **2008**, *20*, 873–880

24. Marc Aurelio Ranzato, C.; Chopra, S.; LeCun, Y. Efficient learning of sparse representations with an energy-based model. *Adv. Neural Inform. Process. Syst.* **2006**, *19*, 1137–1144.

25. Coates, A.; Ng, A. The importance of encoding versus training with sparse coding and vector quantization. In Proceedings of the 2011 International Conference on Machine Learning, Bellevue, WA, USA, 28 June 2011.

26. Yang, J.; Yu, K.; Huang, T. Efficient highly over-complete sparse coding using a mixture model. *Comput. Vis. 2010* **2010**, *5*, 113–126.

27. Bo, L.; Ren, X.; Fox, D. Unsupervised feature learning for rgb-d based object recognition. *Experimental Robotics* **2013**, 387–402.

28. J. Burton, E.V.O.; Ensbey, R. Water Hyacinth Eichhornia Crassipes. http://www.dpi.nsw.gov.au/__data/assets/pdf_file/0013/330322/Water-hyacinth-web.pdf (accessed on 29 November 2014).

29. Csurhes, S. Risk Weed Assessment: Tropical Soda Apple. Water Hyacinth Eichhornia Crassipes. https://www.daff.qld.gov.au/__data/assets/pdf_file/0014/72140/IPA-Tropical-Soda-Apple-Risk-Assessment.pdf (accessed on 29 November 2014).

30. Ayres, L.; Leech, F. Serrated Tussock—Identification and Control. http://www.dpi.nsw.gov.au/__data/assets/pdf_file/0017/362411/Serrated-tussock-identification-and-control-web.pdf (accessed on 29 November 2014).

31. Hinton, G.; Osindero, S.; Teh, Y. A fast learning algorithm for deep belief nets. *Neural Comput.* **2006**, *18*, 1527–1554.

32. Bengio, Y.; Lamblin, P.; Popovici, D.; Larochelle, H.; Montreal, U. Greedy layer-wise training of deep networks. *Adv. Neural Inform. Process. Syst.* **2007**, *19*, 153–160.

33. Krizhevsky, A.; Hinton, G. *Learning Multiple Layers of Features from Tiny Images*; Technical Report; Computer Science Department, University of Toronto: Toronto, ON, Canada, 2009.

34. Serre, T.; Wolf, L.; Poggio, T. Object recognition with features inspired by visual cortex. In Proceedings of the 2005 IEEE Computer Vision and Pattern Recognition, San Diego, CA, USA, 20–26 June 2005; Volume 2, pp. 994–1000.

35. Boureau, Y.L.; Ponce, J.; LeCun, Y. A theoretical analysis of feature pooling in visual recognition. In Proceedings of the 2010 International Conference on Machine Learning, Haifa, Israel, 21–24 June 2010; pp. 111–118.

36. Leung, T.; Malik, J. Representing and recognizing the visual appearance of materials using three-dimensional textons. *Int. J. Comput. Vis.* **2001**, *43*, 29–44.

37. Varma, M.; Zisserman, A. A statistical approach to texture classification from single images. *Int. J. Comput. Vis.* **2005**, *62*, 61–81.

38. Nilsback, M.E.; Zisserman, A. A visual vocabulary for flower classification. In Proceedings of the 2006 IEEE Conference on Computer Vision and Pattern Recognition, New York, NY, USA, 17–22 June 2006; Volume 2, pp. 1447–1454.

39. Kim, M.; Warner, T.A.; Madden, M.; Atkinson, D.S. Multi-scale GEOBIA with very high spatial resolution digital aerial imagery: Scale, texture and image objects. *Int. J. Remote Sens.* **2011**, *32*, 2825–2850.