Original papers

A rapidly deployable classification system using visual data for the application of precision weed management

David Hall[a,*], Feras Dayoub[b], Tristan Perez[a], Chris McCool[a]

[a] School of Electrical Engineering and Computer Science, Queensland University of Technology (QUT), Brisbane, Australia
[b] ARC Centre of Excellence for Robotic Vision, Queensland University of Technology (QUT), Brisbane, Australia[1]

ABSTRACT

In this work we demonstrate a rapidly deployable weed classification system that uses visual data to enable autonomous precision weeding without making prior assumptions about which weed species are present in a given field. Previous work in this area relies on having prior knowledge of the weed species present in the field. This assumption cannot always hold true for every field, and thus limits the use of weed classification systems based on this assumption. In this work, we obviate this assumption and introduce a rapidly deployable approach able to operate on any field without any weed species assumptions prior to deployment. We present a three stage pipeline for the implementation of our weed classification system consisting of initial field surveillance, offline processing and selective labelling, and automated precision weeding. The key characteristic of our approach is the combination of plant clustering and selective labelling which is what enables our system to operate without prior weed species knowledge. Testing using field data we are able to label 12.3 times fewer images than traditional full labelling whilst reducing classification accuracy by only 14%.

1. Introduction

Farmers have seen a steady increase in herbicide resistance from various species of weeds over the years (Gilbert, 2013). This has led to an increased research focus on precision weed management strategies where each weed is treated individually using a treatment which best suits that plant. To do so manually is laborious and costly when covering large areas. This, in combination with other factors, has led to a growing interest in the potential of agricultural robotics capable of performing autonomous precision weeding such as the AgBot II shown in Fig. 1.

Achieving autonomous precision weeding has been a focus of research for many years yet weed classification still remains a critical problem which is generally considered unsolved within this field (Slaughter et al., 2008). Weed classification is usually done using either species-specific classification (Gerhards and Oebel, 2006) or weed-vs-crop classification (Åstrand and Baerveldt, 2002). While successful for their specifically targeted tasks, these techniques either assume that they know exactly what species are present in the field or assume that species-specific knowledge is not necessary. Neither assumption is valid if we want species-specific weed management which is easily deployable to any given field. To the best of our knowledge, the challenge of creating rapidly deployable weed management systems that can perform intra-row weeding without needing any prior training has only been considered a few times before and these only consider weed-vs-crop classification systems (De Rainville et al., 2012; Strothmann et al., 2017; Wendel and Underwood, 2016). Neither one tackles the challenge of providing a species-specific weed management system which can be deployed rapidly in a field without knowing the species within in advance. This challenge of rapidly deployable species-specific precision weeding is what our research works to enable.

The main contribution of this work is the creation and demonstration of a rapidly deployable species-specific weed classification system for use in autonomous precision weeding. This system is able to operate without assuming that it knows what weed species are present prior to deployment and can be trained on a specific field with minimal human effort. The system is novel in many ways from what is currently found in autonomous weed classification literature. Unlike the current literature we:

- Use unsupervised clustering to summarise weed species.
- Use selective labelling based on the clustering results to fully label scouted plant data rapidly with minimal human effort.
- Use the data labelled from selective labelling to train classifiers able

* Corresponding author.
  E-mail addresses: d11.hall@hdr.qut.edu.au (D. Hall), feras.dayoub@qut.edu.au (F. Dayoub), tristan.perez@qut.edu.au (T. Perez), c.mccool@qut.edu.au (C. McCool).
[1] http://www.roboticvision.org/.

Fig. 1. AgBot II agricultural robotics platform.

to operate in the given field after one scouting operation.
- Use highly descriptive learnt features with low dimensionality (128 vs 1024 dimensions) to improve clustering results.

We evaluate several methods for clustering and selective labelling for our classification system to evaluate which approaches work best for our task. We show that on a dataset of weed images collected in a field we could label 12.3 times fewer images than full labelling and still achieve labelling accuracy of 79% and a decrease in classification accuracy of only 14% from full labelling. Through our analysis we identify the strengths and weaknesses of the different methods tested and propose directions for future work.

1.1. Literature review

There have been several classification approaches for autonomous weeding presented over the years, using a variety of plant descriptors and classifiers. These approaches try to solve either a binary weed-vs-crop, or a multi-class species-specific classification task. Both approaches tend to use either shape, reflectance, or texture features - and often combinations of thereof (Slaughter et al., 2008).

Weed-vs-crop classification methods can be seen as the simpler of the two options, being a more easily implemented weeding procedure for specific crops as it considers only a two class problem and can take advantage of the structured manner that crops are usually sown. One technique used in weed-vs-crop classification is to perform only inter-row weeding without any need for a specific plant-wise classification. Such a technique was by Emmi et al. where the main focus was on calculating weed densities between the crop rows in order to utilise resources more efficiently (Emmi et al., 2014). This simplifies the classification problem by only needing to detect plants and crop rows rather than needing to distinguish a specific species within the detected plants. In some applications, the weed-vs-crop classification task is made easier by the size difference between crop and weed (Blasco et al., 2002). Blasco et al. focused on cabbage crops which are a transplanted crop and most likely always bigger than the weeds around them (Blasco et al., 2002). This allowed for a simple size threshold to identify crop and weed.

Other works deal with a harder crop-vs-weed plant-wise classification procedure with crops of similar size to the weeds (Åstrand and Baerveldt, 2002; Lottes et al., 2017). This harder process for classification can however be aided by the use of positional information. This takes advantage of known prior information about how the crop should be planted such as knowing that the crop is sown in rows and with an expected spacing between them. Lottes et al. calculated distances between query objects or keypoints and plants previously classified as crops (Lottes et al., 2017). These distances, combined with other information are used to calculate the probability that a given plant is a crop based on these distances. This probability was then used as a

feature within their classification system.

Weed-vs-crop methods which utilise plant-wise classification are typically only designed for a single crop and would require retraining before they could be used for different scenarios, which can be a time-consuming process as crop image data needs to be collected and fully manually annotated. While these crop-vs weed systems are undoubtedly useful and easily implementable they do have one major drawback which is that they treat all weeds as the same. This does not allow for species-specific treatment which may be required by some farmers.

In order to achieve species-specific treatment, a multi-class classification approach is required. Lin demonstrated the use of an SVM classifier with shape features to achieve 75.00% and 82.85% average classification accuracy in field and greenhouse tests respectively across 6 pre-defined species (Lin, 2009). The precision weeding system presented by Bawden et al. demonstrated a classifier trained to identify five different weed species which successfully identified up to 98.8% of one of the classes correctly. While impressive, the system did struggle distinguishing between different grass species due to the large visual similarity between such plants, in one case identifying only 47.5% examples of one species correctly (Bawden et al., 2017). A unique system was developed by Haug et al. for classifying overlapping plants by classifying a grid of small patches across segmented plant regions and then interpolating the results of each patch until whole plant regions were classified, achieving an accuracy of 93.8% (Haug et al., 2014). It should be noted that in this case, of the three defined classes, only two were defined plant species and the third class was simply labelled as "other weeds".

The use of a single class to define all other weed species found in a field is not uncommon within this field of research. Sometimes this "other" class is split into "other grass" and "other broadleaf" classes such as was done by Gerhards and Oebel (2006). In that work, Gerhards and Obel classified three different weed species as well as "other broadleaf" and "grass weeds" and tested the system in the field. This system managed to provide herbicide reductions of up to 81% and increased weeding efficacy between 85% and 98% using this classification system. In a more recent work, Lottes et al. evaluated both a weed-vs-crop classification system as well as a species-wise plant classification system with three pre-defined plant species as well as an "other weeds" class (Lottes et al., 2017). Using shape, reflectance, texture and position features they achieved an overall accuracy of 86% of predicted objects for their species-wise system and 96% accuracy for crop vs weed classification. The main detractor for the precision of their system was stated as being due to the performance of the "other weeds" class. This was hypothesised as being because this class has a small number of samples and a high intra-class variance as it represented every other weed species not previously defined. This use of an "other weeds" class highlights a problem inherent to classification approaches for automated weed management. This problem is that they need prior information about which species are to be expected and cannot be adapted for different species if they are transferred to fields which do not meet with the prior assumptions being made. The algorithms need previously created, manually labelled datasets in order to be retrained which can be a laborious and slow process which won't necessarily meet with a farmer's immediate needs.

One of the few works that approaches the weed classification task without making prior assumptions about the plant species in the field is work done by De Rainville et al. (2012). They performed intrarow weed-vs-crop weed management where the crop species was not pre-defined. The only assumption made was that the crop was planted in rows. All plants not within the crop rows were considered to be weeds. Using this assumption they could infer a model for the weeds and crop and achieved a crop classification accuracy on average of 94% across two different crops. A limitation of this work is that it is only applicable once the crop has grown to a state where crop rows can be detected. This means that it cannot be applied in the fallow period before the crop