LWA project code USY13

Australian Government
Land & Water Australia

The University of Sydney

ACFR
Australian Centre for Field Robotics

Aquatic weed surveillance using robotic aircraft

Salah Sukkarieh

Introduction

Several aquatic weeds are aggressive invaders of waterways in Australia. Species such as alligator weed, cabomba and salvinia, which have been declared Weeds of National Significance, can cover entire water surfaces. Flows are prevented, channels blocked and flood patterns altered. Weed mats reduce available oxygen in waterways, resulting in increased fish kills and loss of native plant species, and adversely affecting water quality. Unchecked, aquatic weed invasions cause millions of dollars of damage to agriculture, fisheries and the environment.

One major limitation in controlling aquatic weeds is the difficulty of conducting detailed surveillance over vast areas such as irrigation schemes, or over inaccessible aquatic habitats. Satellite remote sensing has been used in the past to overcome this limitation, but is not cost effective and cannot detect small infestations, especially where overhanging foliage or environmental sensor clutter and backscattering can affect surveillance performance.

About the project

For the 2007/08 Defeating the Weed Menace program (DWM), we proposed to build and test a prototype robotic aircraft and surveillance system to detect aquatic weeds in inaccessible habitats. The know-how and technology is not new: we have been working on such systems for the aerospace industry for over a decade. However, a new application of the technology and the different test environment meant we had to devise new technologies and approaches.

The prototype aerial robot houses sensors and spray systems. The sensors take imagery of the environment the robot flies over, classify the imagery so as to detect where the weeds are (if any), and geo-reference the location of those weeds. The robot can then be tasked to go back to those weed locations and spray them with an appropriate herbicide, or be tasked to spray them at the same time that they are detected.

The project was divided up into two key areas:

1. Development of the robotic aerial platform
2. Development of novel machine learning algorithms to discriminate between different plant species.

An aquatic weed infestation in an area that would prove difficult to map and control by conventional means. The objective of this project was to develop a system which could easily conduct surveillance and control missions over such environments. Photo courtesy of Andrew Petroeschevsky, NSW DPI. Opposite: Salvina. Photo Arthur Mostead.

What are robotic aircraft?

Robotic aircraft, otherwise known as unmanned air vehicles (UAVs), have been around since the development of manned aircraft. They are essentially the same platform minus the pilot. The pilot is replaced with sensors, computers, actuators and algorithms. The sensors detect properties such as the velocity of the platform and wind speed; the actuators provide a way to drive the various control surfaces on the platform; and the computers house the algorithms which do all the "thinking" the pilot would do, such as waypoint traversal, hovering, and sensor pointing. Modern day commercial aircraft, such as 747s and the A380, are essentially UAVs because most of their flight is autonomous.

Robotic aircraft come in various shapes and sizes; from the Global Hawk which has a wing span of over 35 metres, to micro-UAVs which can easily sit on the palm of a hand. The decision as to size and shape, and whether it should be a fixed wing UAV or a rotary UAV, comes down to the payload weight one wants to carry, the flight duration and the distance to be travelled.

What is machine learning?

Machine learning is the science and implementation of computer algorithms that give the computer the ability to "learn": to improve its performance in a particular task based on the data that it receives. Machine learning has been used in a wide variety of applications, from the control of helicopters, to medical diagnosis, speech recognition, and the detection of individual faces.

There are many machine learning techniques, but the most important for our work is "supervised learning". In supervised learning you provide the algorithm with training data (input and output data). The algorithm then aims to learn a model which describes the input-output relationship. This learnt model can be used to predict the output of any new input data.

The robotic platform

We used a modified model helicopter as the platform. Using a helicopter gave full manoeuvrability, including ability to hover, making it possible to traverse large distances, move in tight situations, and hold position to take imagery or to spray herbicide. This involved development and tuning of flight control and navigation algorithms and spray mechanisms. The final system can fly for approximately two hours and carry approximately 500 ml of herbicide (water was used in the project for demonstration). Image on the front cover shows the robotic aerial vehicle over sprayed salvinia.

Surveillance system

Early in the project, many experiments were conducted testing various ideas on the type of sensors and detection algorithms we could use. We needed this information to establish the size of the surveillance system, which in turn would decide how we would mount the system on the helicopter. Colour and multi-spectral vision systems including IR and NIR were analysed. We determined that a light-weight, high resolution sensing device which provided separate RGB colour information was needed. See figure 1 below.

Figure 1. A closer look at the robotic platform. The metal box at the front holds most of the "intelligence" of the platform. The 3CCD Video camera passes the imagery information into the computer processors located within this box. Here classification of the imagery takes place. The GPS position system and other sensors provide position and velocity information. The laser range finder provides height-above-ground information. Also located at the bottom of the platform is the spray tank which housed water (herbicide in real applications), and a spray pump. When activated the spray pump transfers the liquid from the tank through the spray boom arms (shown here).

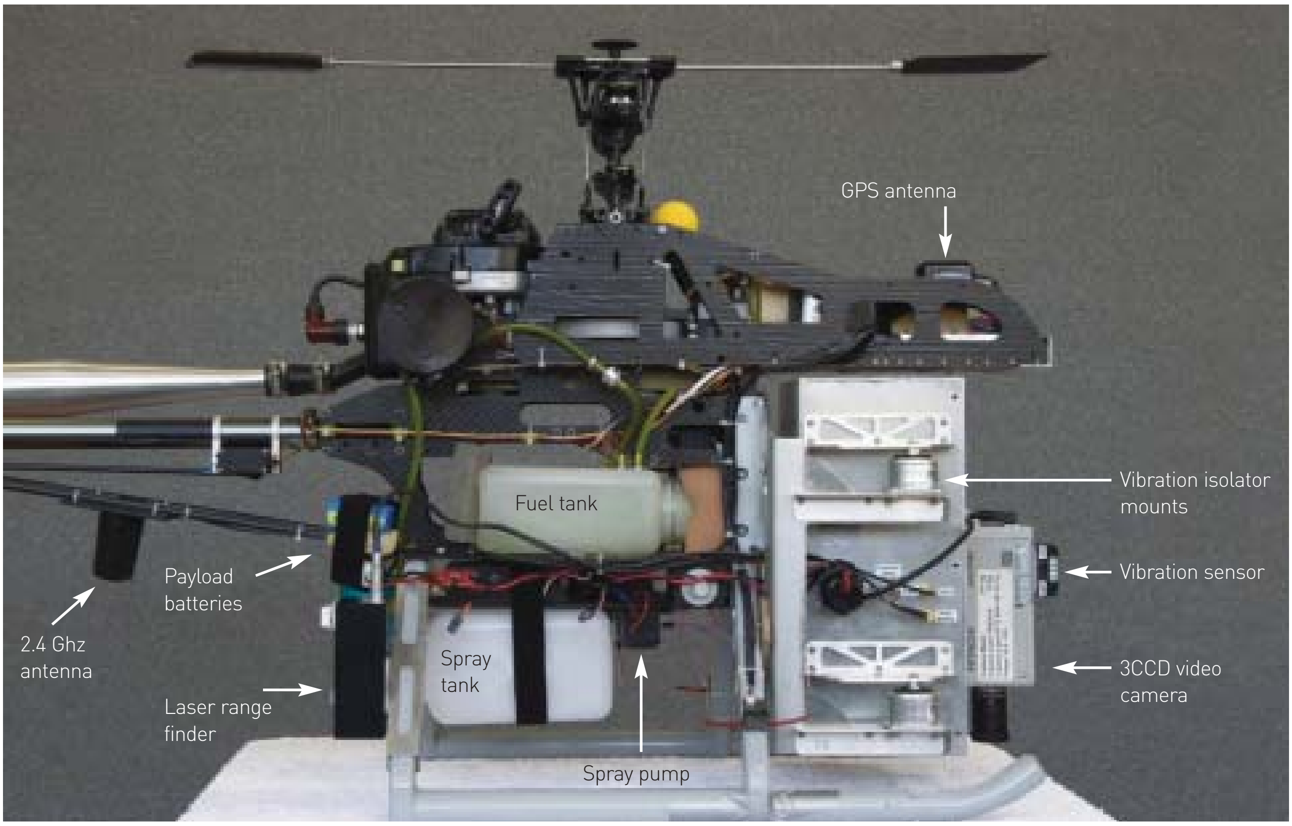

In order to classify weeds in near real-time we needed a means of learning the particular attributes of a weed so that an algorithmic model describing the weed could be developed and used on the platform. This would provide us with a computationally efficient algorithm that could reasonably quickly detect the likelihood and provide a probability measure of a particular weed being present. Using supervised learning techniques we were able to learn classification models of alligator weed and salvinia. The classification algorithm is based on maximum margin classification. A large data set of *n x n* pixels of imagery taken from the camera is collected, and the set separated manually into what is and what is not a weed. Each *n x n* pixel is in itself marked as an image of *p* dimensions. These dimensions take into consideration colour, shape and texture. The classification algorithm then tries to determine a hyperplane which separates the images into two sets. The objective of the algorithm is to determine this hyperplane, which is of maximum distance (Euclidean in the imagery space) between the two datasets. When found, this hyperplane becomes the detection model. As new data is collected it is passed to this model and will fall on either side of the hyperplane depending on whether it is or is not a weed. Depending on how far the image fits away from the hyperplane will determine the probability that the image is within that set. The algorithm proved to be very robust and reliable. See figure 2 below for a representation of this process.

Flight tests

Flight tests were conducted at the Killarney Chain of Ponds in Pitt Town near Sydney. This area is known for its spread of alligator weed and salvinia. The tests were conducted in April (alligator weed) and in August (sprayed salvinia) 2008, on one of the farms we had access to. The system worked as planned, being able to fly over the aquatic site, collect imagery, communicate to the ground station, and spray at designated locations. See figure 3 on the following page.

Where do we go from here?

The project developed a prototype system and demonstrated its effectiveness. The various approaches to weed management using the system were discussed with stakeholders of the project. This led us to discuss the potential of the system with other interested bodies, specifically council officers.

There is significant scope for this system to be actively used in weed surveillance. It provides a means of traversing large distances, accessing difficult ground operation areas, and for improving the efficiency of the weed management cycle.

There were many new challenges and these form part of the ongoing work in the area.

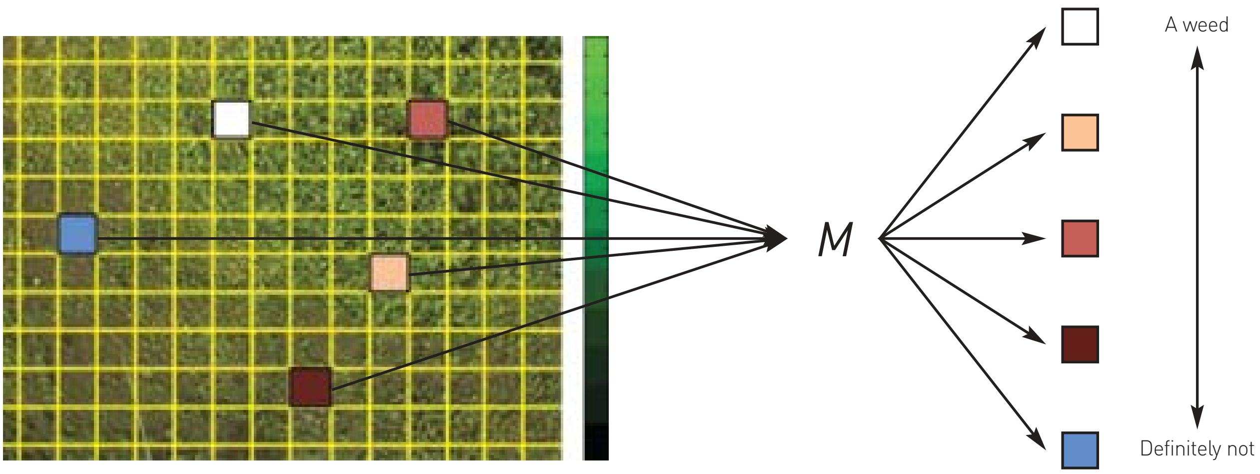

Figure 2. A representation of how the on-line classification process works. A model "M", which is an algorithmic representation of relationship between colour, texture and shape, is "learnt" off-line, and then implemented on-line. Each image taken by the camera in operation is divided up into n x n pixels, known as a "sub-image". Each sub-image is now considered to be an image in its own right, and is passed to the model M. The output of the model M is a probability measure which represents whether the sub-image contains a weed or not.

Figure 3. The top images represent the classification results of alligator weed, and the bottom that of sprayed salvinia. The images in the left hand column are those directly from the imaging system, and that in the right hand column are the output of the learnt classification model.

CASA regulations

Flying a UAV in a populated area requires the users to overcome significant regulation hurdles. Even the process of flying along a river with houses nearby will be an issue. Doing so with herbicide on board adds an extra complication. Discussions with CASA regulators suggest that gaining permission is possible, but needs to be carefully approached. Regulations for flying in remote areas, or along irrigation channels, are not as strict, and this is where we see the introduction of such a system.

Team of people

Operating the system as a UAV requires three people: two that deal with flight operations and safety, and one to deal with computing and communications. Instead of flying the system as a UAV, it could be flown remotely, but the vehicle has to remain within line-of-sight. As the technology progresses, the number of people will drop. Five years ago, six people would have been needed to operate such a system. One of the key areas for research and development is in appropriate human-machine interfacing.

Detection algorithms

The detection algorithms need to be as accurate as possible to minimise false detections. The current algorithms proved to be very effective and will become more robust as more data is collected. However, tuning is required, and so a way for the operator to easily tune and validate the results is needed. Again, research and development in the area of human/machine interfaces is required.

Decision about spraying

The "operator-in-the-loop" is a classic example of how autonomous systems interact with human ground operators in making collective decisions. Before a weed can be sprayed, there needs to be confirmation of the decision from an operator. This will probably continue to happen regardless of how accurate the detection algorithms become, because of the potential to cause widespread damage, and for OH&S reasons. Safety mechanisms need to be in place both from a hardware and software perspective and tight logic control needs to be placed around the system's decision functions for it to be used safely and effectively.

Flight control

Flight control proves to be the most delicate aspect of the system. The flight control unit needs to safely move the platform and conduct the operation at hand. Off-the-shelf flight control systems will stabilise the platform and allow you to command it to go to designated waypoints, but flight control management is a mission-specific task. Therefore the system will inevitably require a computer module that talks to the flight control computer to command it to do actions at specific instances both under normal operation and under safety critical operations. Such operations are specific to the mission at hand — you cannot buy an off-the-shelf version of such management systems. This computer module could either be autonomous and placed on the platform or could be at the ground station and managed by the operator. If the former, then extensive testing and an understanding of the limitations is required; if the latter then there needs to be a certain level of operator knowledge of the flight control system and its operation. The way forward is through human operator management and training, to better appreciate the system's limitations.

Collision avoidance

The system needs to be able to conduct collision avoidance. This is a long-term research issue. In many cases the flight path can be constructed to be safe, but the system still requires the ability to detect whether an obstacle has come within a safe bounding box. This involves sensing capabilities as well as fusion algorithms. What this implies is that in the short term, an operator needs to judge whether detection is required and how to overcome this through better path management. For many missions, such as irrigation channels and wide spaces, it is less of an issue, but caution is required for aquatic weed management along rivers, or where there are trees within the area.

The potential

The future looks very exciting for this intelligent little machine. We plan to continue the project with a specific focus on aquatic weeds, and to broaden its capabilities into other ecology management arenas, such as woody weeds and biomass measurements.

More information

More information on the project can be found at:

- www.acfr.usyd.edu.au/research/aerospace.shtml
- lwa.gov.au/weeds

Or, contact:

Salah Sukkarieh, Associate Professor, Australian Centre for Field Robotics, The University of Sydney. E-mail: salah@acfr.usyd.edu.au

Acknowledgements

Special thanks to the aerospace team: Jeremy Randle, Steve Keep, Ali Goktogan, Muhammad Esa Attia, Mitch Bryson, Todd Lupton and Calvin Hung.

We would like to thank Land & Water Australia for funding this project and recognising its potential. We would also like to thank Andrew Petroeschevsky from NSW Department of Primary Industries, Luke Joseph from Farm and Dam Control, and Sunwater for their in-kind contribution, knowledge, and constant enthusiasm. Finally, we thank the Australian Research Council for their support.

PRODUCED BY LWA'S WEEDS R&D PROGRAM. FUNDED BY THE AUSTRALIAN GOVERNMENT'S DEFEATING THE WEED MENACE PROGRAM.

Photo on this and previous page: The McKenzie River clogged with water plants, near St Lawrence, Queensland. Photo Arthur Mostead.

Printed on Revive Silk which is made in Australia from elemental chlorine free pulp with 65% derived from sustainable plantation forests and 35% recovered fibre. It is PEFC Certified and Australian Paper is an ISO 14001 certified mill which utilises renewable energy sources.

Published by: Land & Water Australia
Web: www.lwa.gov.au
ISBN print: 978-1-921544-69-9
ISBN electronic: 978-1-921544-71-2

This publication can be ordered free of charge from CanPrint Communications, freecall: 1800 776 616, or e-mail: lwa@canprint.com.au and quote product code **PN22514**

Photos: Front page, Muhammad Esa Attia
Design: Angel Ink
Print: Paragon Printers, June 2009