TropWATER
Centre for Tropical Water and Aquatic Ecosystem Research

JAMES COOK UNIVERSITY
AUSTRALIA

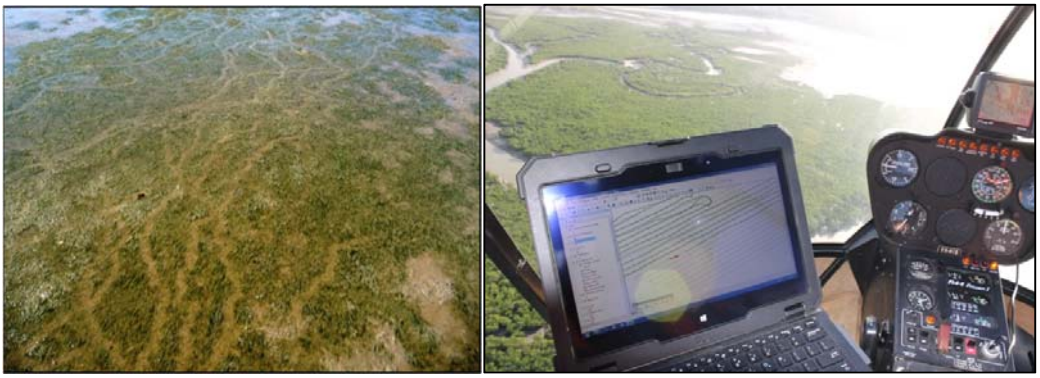

Dugong Feeding Ecology and Habitat Use on Intertidal Banks of Port Curtis and Rodds Bay – Final Report

Michael Rasheed, Damien O'Grady, Emma Scott, Paul York, and Alex Carter

Report No. 17/13

March 2017

Dugong Feeding Ecology and Habitat Use on Intertidal Banks of Port Curtis and Rodds Bay – Final Report

A Report to the Ecosystem Research and Monitoring Program Advisory Panel

Report No. 17/13

March 2017

Prepared by

Centre for Tropical Water and Aquatic Ecosystem Research
(TropWATER)
James Cook University
Townsville
Phone : (07) 4781 4262
Email: TropWATER@jcu.edu.au
Web: www.jcu.edu.au/tropwater/

This report should be cited as:
Rasheed, M.A., O'Grady, D., Scott, E., York, P.H., and Carter, A.B. (2017). Dugong Feeding Ecology and Habitat Use on Intertidal Banks of Port Curtis and Rodds Bay – Final Report. Report produced for the Ecosystem Research and Monitoring Program Advisory Panel as part of Gladstone Ports Corporation's Ecosystem Research and Monitoring Program. Centre for Tropical Water and Aquatic Ecosystem Research (TropWATER) Publication 16/14, James Cook University, Cairns, 68 pp.

This report has been produced for the Ecosystem Research and Monitoring Program Advisory Panel as part of Gladstone Ports Corporation's Ecosystem Research and Monitoring Program. The study was undertaken through a Consultancy Agreement (CA14000187) between Gladstone Ports Corporation and James Cook University

This Publication has been compiled by:
Centre for Tropical Water and Aquatic Ecosystem Research (TropWATER)
James Cook University
seagrass@jcu.edu.au
PO Box 6811
Cairns QLD 4870

Acknowledgements:
This project is funded by the Gladstone Ports Corporation's Ecosystem Research and Monitoring Program. We wish to thank the staff at TropWATER that have assisted with data collection and analysis including Skye McKenna, Mark Leith, Lloyd Shepherd and Alysha Sozou.

Table of Contents

KEY OUTCOMES ii

1 INTRODUCTION 1

1.1 Background 1
1.2 Limitations 2
1.3 This report 2

2 METHODS 3

2.1 Sites and surveys 3
2.2 Data acquisition and processing methods 5
2.2.1 Equipment preparation 5
2.2.2 Flight planning 7
2.2.3 Data capture 8
2.2.4 On-site verification 8
2.2.5 Preprocessing 8
2.3 Identification of Dugong Feeding Trails 10
2.4 Quantification of Dugong Feeding Trails 11
2.4.1 Visual Census of Dugong Feeding Trails 11
2.4.2 Algorithm Development for Automated Extraction of Dugong Feeding Trails 14
2.5 Seagrass and Environmental Data 17
2.6 Statistical Analyses 18

3 RESULTS 19

3.1 Orthomosaic Generation 19
3.2 Longevity of feeding trails between surveys 20
3.3 Visual Census of Dugong Feeding Trails 25
3.4 Algorithm Extractions of Dugong Feeding Trails 33

4 DISCUSSION 36

5 REFERENCES 43

6 ABBREVIATIONS 46

7 APPENDICES 47

Appendix 1: Matrix of correlations for seagrass analysis 47
Appendix 2: Methods to automate mapping of dugong feeding trails 48
Introduction 48
Problem defined 48
A typical working sequence 54
Potential approaches to the feature extraction problem 57
Results and discussion 65
Conclusion and further work 66

KEY OUTCOMES

This report details the findings of the Ecosystem Research and Monitoring Program (Project CA14000187): *"Dugong Feeding Ecology and Habitat Use on Intertidal Banks of Port Curtis and Rodds Bay"* for Gladstone Ports Corporation. In this project, key seagrass areas within Port Curtis and Rodds Bay were identified and surveyed quarterly to assess temporal and spatial feeding activity by dugongs from May 2015 to November 2016. Key outcomes include:

- Results highlighted the importance of intertidal seagrass meadows as foraging habitat for the dugong population of Port Curtis and Rodds Bay both throughout the year and the survey region. The program provides a method and baseline information on the use of seagrass meadows by dugongs against which the impact of future developments in the region can be assessed.

- A new method was successfully developed to quantify dugong feeding trails using low-level aerial photography and next generation photogrammetry (structure from motion) techniques and software. This enabled the production of orthomosaics of the target areas with less than 5 cm pixel resolution suitable for identifying dugong feeding scars or trails. For the first time this allowed the effective examination of dugong feeding over large "meadow" scales and a tool to address knowledge gaps around the temporal and spatial use of intertidal seagrass habitat and its relationship to weather, season and changes to seagrasses.

- All meadows sampled showed dugong feeding activity throughout the year with no consistent temporal patterns among sites. Sites closer to the port (Wiggins Island, South Trees and Pelican Banks) had higher levels of feeding than the two Rodds Bay meadows.

- In general, Dugong Feeding Trails (DFTs) were positively correlated with increased seagrass presence and above ground biomass. The main driver of seagrass abundance was linked to the amount of rainfall (and linked river flow) received in the months prior to the survey although this varied among sites.

- An analysis of the longevity of DFTs indicated that trails were unlikely to persist between quarterly sampling events with the majority of trails indistinguishable between 2 and 10 weeks after being first recorded, therefore quarterly surveys are likely to identify only new DFTs.

- Several approaches were developed and trialled to produce algorithms that would allow reliable automated or semi-automated extraction of DFTs from the imagery. While the approaches produced good results at small scales the complexity of the DFT features and the background landscape mosaic on which they occur in Port Curtis and Rodds Bay, meant that at larger meadow scales, they were unsuitable for automated extraction. At larger spatial scales a manual visual classification approach was more effective.

- The algorithm development work however resulted in substantial technical advances and highlighted potential avenues to further develop the techniques for future application.

This project provides the basis for an ongoing assessment of intertidal dugong feeding activity in Port Curtis.

It was not possible to ascertain in a quantitative way potential impacts of the Western Basin Dredging and Disposal Project on dugong feeding as this study has only collected DFT data post-dredging. However, qualitative assessments from observational data collected during seagrass surveys that spanned pre, during, and post dredging confirm that dugong feeding activity assessed in 2015-16 occurred in the same meadows where feeding was observed prior to (2009 – 2011) and during capital dredging (2011-13). For future activities, the results of this project provide a baseline and method for effective quantitative comparisons of intertidal dugong feeding to assess potential impacts, if any.

1 INTRODUCTION

1.1 Background

Dugongs (*Dugong dugon*) are large herbivorous marine mammals restricted in range to tropical and sub-tropical locations of the Indo-west Pacific region. Dugong habitats generally correspond to shallow water (< 10 m) seagrass meadows as their diet consists predominately of seagrass, although they can also eat macro-invertebrates and algae (Marsh et al. 2011). Globally, dugongs are listed as 'vulnerable' by the International Union for the Conservation of Nature (IUCN). In Australia, they are protected under the *Environment Protection and Biodiversity Conservation Act 1999* (EPBC Act) as a migratory species and are considered a Matter of National Environmental Significance (MNES). Between the 1960s and the late 1990's, the dugong populations on the east coast of Queensland (between Brisbane and Cooktown) have undergone serious declines (Marsh et al. 2005). In Queensland, dugong are listed as a 'vulnerable' species in the *Nature Conservation (Wildlife) Regulation 2006* under the *Nature Conservation Act 1992*. Within the central Queensland region, a relatively small dugong population (in the low hundreds at most) utilises a large area of Port Curtis and Rodds Bay where extensive seagrass meadows (Figure 2) support dugong feeding (Sobtzick et al. 2013; Bryant et al. 2014a). As there are no other known major areas of seagrass between Shoalwater Bay and Hervey Bay, the Port Curtis and Rodds Bay area is a potentially important connecting habitat for dugong populations in southern Queensland (Sobtzick et al. 2013, 2017, Cleguer 2015b). This region has been recognised as an area important for dugong conservation through the declaration of the Rodds Bay Dugong Protection Area under Queensland Fisheries legislation, however, the seasonality and extent of dugong feeding is not well understood for this area.

The feeding ecology of dugongs has been assessed by various methods including direct observation, analysis of stomach contents and faecal matter and measurements of food abundance combined with feeding trail analysis (Marsh et al. 2011). Within dugong populations there is no clear preference for one seagrass species as a primary food source (Marsh et al. 2011). Dugongs use two primary modes of feeding depending on seagrass morphology and the nature of the sediment. When feeding on large dense seagrass genera such as *Amphibolis* species in Shark Bay, Western Australia, they have been observed to crop only the above ground leaves and stems (Anderson 1982). However, when feeding on structurally smaller species such as those in the *Halophila*, *Halodule*, and *Zostera* genera, particularly in soft sediments (conditions that occur in our study region), dugongs excavate the plant community to consume both above and below ground parts of the plant (Anderson and Birtles 1978). When feeding in an excavation mode, each feeding effort can leave long serpentine furrows (10-25 cm wide, usually several metres long and up to 6 cm deep) within the seagrass meadows known as dugong feeding trails (DFTs – Figure 1). These feeding trails offer some of the best physical evidence of intertidal habitat use and are common in dugong foraging habitats in the intertidal regions of the Great Barrier Reef (GBR) (Preen 1992).

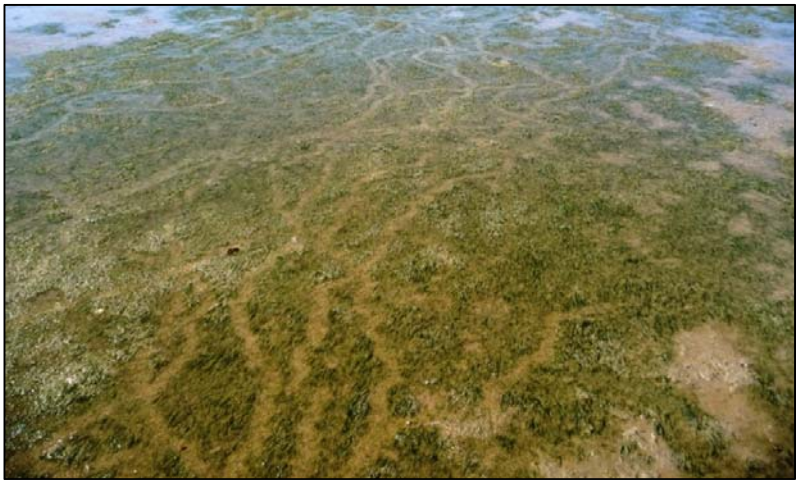

Figure 1: Dugong feeding trails on intertidal seagrass meadows.

Little quantitative information exists on how dugongs utilise seagrass habitats over time within Port Curtis and Rodds Bay. One of the major challenges to studying feeding behaviour is that it is incredibly difficult to observe the animals in the field. The large areas of intertidal seagrass meadows in the region are dominated by small structural seagrass species that dugongs are prone to excavate when feeding and provides an ideal opportunity to examine the potential of using DFTs as a proxy for intertidal habitat use by dugongs.

The use of remote sensing was identified and trialled as a sampling tool to provide a census of DFTs across a large study area over multiple years. Given the spatial and temporal scale of the project and the known variability of the trails, the development of an automated extraction of DFTs was considered desirable and formed part of the goals of the project.

1.2 Limitations

The study has some inherent limitations that need to be recognised when assessing the potential to provide information on habitat use by dugongs in the Port Curtis and Rodds Bay region. Firstly, remote sensing from the air in turbid estuaries like Port Curtis is only useful for intertidal areas as it is not possible to detect and map DFTs in subtidal areas through the water column in most conditions. Fortunately, in the case of Port Curtis and Rodds Bay, large areas of the available seagrass meadows for dugong feeding are in intertidal areas.

Secondly, as we are using only linear features in seagrass meadows that we can positively identify as dugong feeding trails, we are likely to be missing activities such as cropping of seagrass, total removal of small irregular patches of seagrass through intensive excavation, and foraging for burrowing invertebrates that can leave circular scars and craters. Bearing this in mind, our analysis of feeding activity is very likely to be an underestimate of the use of intertidal seagrass by dugongs and provides no data on subtidal habitat use. Finally, without additional data such as aerial surveys of dugong numbers it is not possible to use this report to make meaningful inferences regarding the size of local dugong populations. This is because similar densities of DFTs can be created through either intense feeding by a small number of dugongs or by intermittent feeding by a large group.

1.3 This report

This is the final report detailing the findings of aerial assessments of DFTs in Port Curtis and Rodds Bay in 2015-16 (project CA14000187). This report presents results from visual analysis of DFTs at five sites sampled quarterly from May 2015 to November 2016. The analysis looks at spatial and temporal patterns in dugong feeding in intertidal seagrass meadows in Port Curtis and Rodds Bay.

The project investigates potential correlation between dugong feeding patterns and seagrass meadow characteristics and environmental and climate variables. A finer time-scale analysis was also conducted from August to October 2015 to assess the length of time that trails can be detected from aerial photographs in seagrass meadows after creation. Additionally, this report also details the research and progress toward development of a means to automate the extraction of DFTs from the imagery into a vector feature set against which individual trail characteristics could be evaluated and monitored over time.

The report details a new method to describe intertidal dugong feeding activity using excavation mode at relatively large spatial scales with the aim of providing valuable information on the use of intertidal seagrass beds for feeding by dugongs in the Port Curtis and Rodds Bay region. This method provides a useful tool for investigating excavation foraging and habitat use by dugongs in intertidal habitats throughout their geographical range. The goal of the project is to provide information to increase understanding of the role of dugong in ecosystem processes and to add fundamental information to assist in the effective conservation of these important animals and their habitats. The project was also designed to meet approval conditions as part of the Western Basin Dredging and Disposal Project (WBDDP- see Discussion).

2 METHODS

2.1 Sites and surveys

Methods were developed to assess evidence of dugong feeding in intertidal seagrasses in a quantitative way so that comparisons through time and across locations could be made in an effort to understand how dugong feeding activity occurs and varies in the Port Curtis and Rodds Bay region and what the major influencing factors may be.

TropWATER has historical data from over a decade of annual and sub-annual monitoring of seagrass sites in Port Curtis and Rodds Bay where DFTs have been qualitatively recorded during surveys (Bryant et al. 2014 a and b). This information was used to narrow the focus to five target seagrass meadows where DFTs had regularly been observed. During the initial sampling trip in May 2015, a reconnaissance was done to check that both seagrass and dugong feeding activity were present at the sites. The three initial sites in Port Curtis all contained seagrass and evidence of dugong feeding (South Trees, Pelican Banks and Wiggins Island – Figure 2). Only one suitable area was found in Rodds Bay during the May survey (Rodds Bay North) but a second area was identified and added to the program in the subsequent survey in August (Rodds Bay South – Figure 2).

Surveys were carried out at the five target meadows in 2015 as part of the quarterly assessments on 16th and 17th May, 29th and 30th August, and 27th and 28th November and in 2016 on 9th and 10th February, 7th and 8th May, 18th and 19th August and 14th and 15th November (Table 1).

In order to examine the longevity of trails between the regular quarterly surveys, a sub-set of smaller areas where DFTs had been mapped in August were reassessed on three occasions corresponding to 2, 4 and 8 weeks after the August 2015 survey (13th September, 28th September, 28th October – Table 1). These areas were located at Rodds Bay North, Rodds Bay South, South Trees and Pelican Banks (Figure 2).

Most surveys were carried out from flights at a height above ground of around 175 metres (see 2.2.2 below). In order to test the relative performance of identifying DFTs and feature extraction with different Ground Sampling Distances (GSDs, which equate to pixel sizes in resulting mosaics), an area of the South Trees meadow was also carried out at a lower altitude of 90 metres during five of the 2015-16 surveys.

Table 1: Dugong Feeding Trail sampling meadows and dates surveyed for the quarterly surveys in 2015/2016 and interim surveys assessing DFT longevity in 2015.

Monitoring sites (see Figure 2)	2015						2016			
			Interim surveys to examine DFT longevity between quarterly sampling							
	May	Aug	Sept 13	Sept 28	Oct 28	Nov	Feb	May	Aug	Nov
RODDS BAY										
Rodds Bay North	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Rodds Bay South	**N/A**	✓	✓	✓	✓	✓	✓	✓	✓	✓
PORT CURTIS										
South Trees	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Pelican Banks	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Wiggins Island	✓	✓				✓	✓	✓	✓	✓

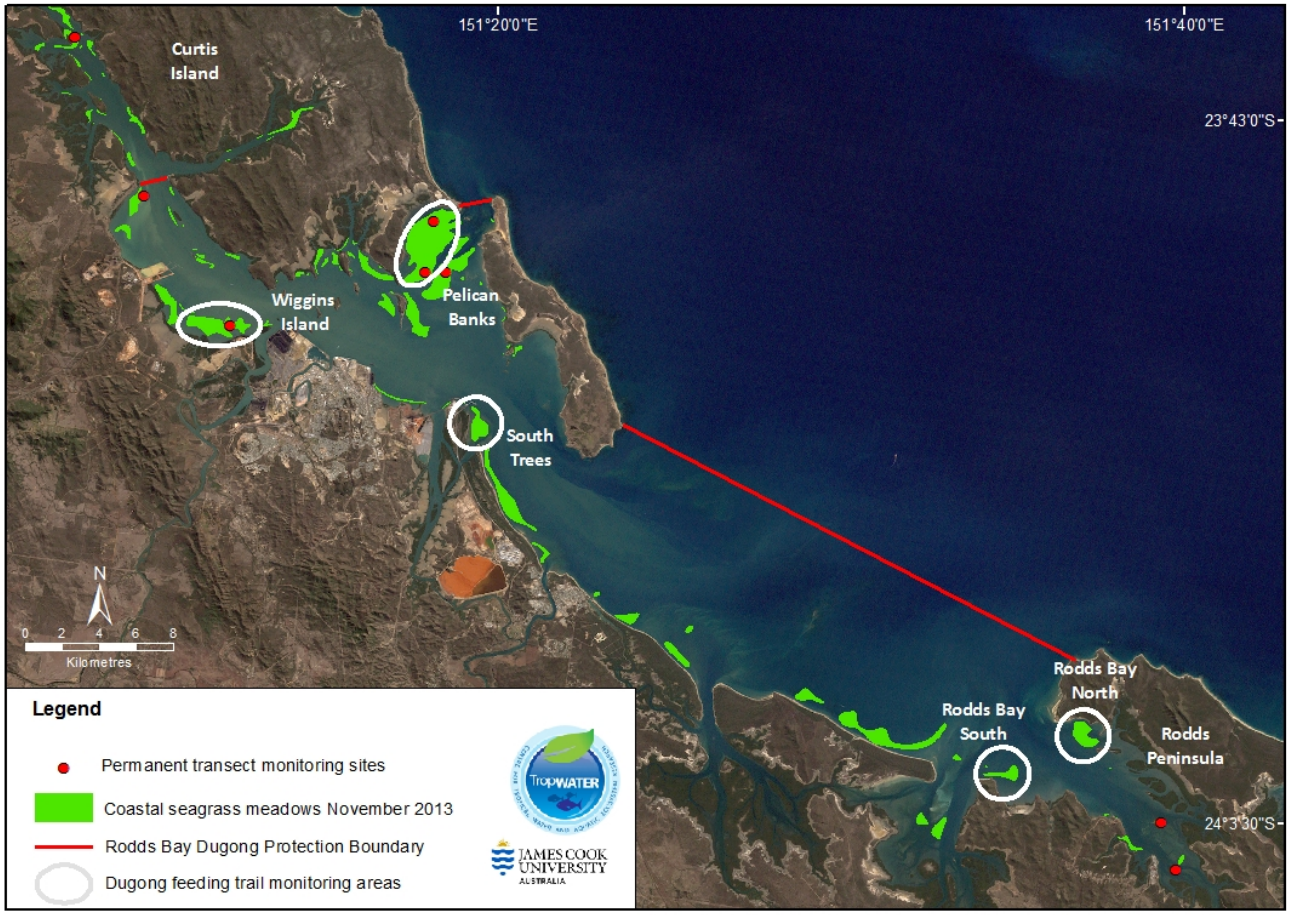

Figure 2: Port Curtis and Rodds Bay region showing the location of the five meadows assessed quarterly for dugong feeding surveys and the shallow subtidal and intertidal seagrass meadows present in November 2013.

2.2 Data acquisition and processing methods

In order to produce suitable imagery over the relatively large spatial scale required for this project (largest meadow over 680 hectares), a cost effective technique with sufficient resolution to identify DFTs (pixel size 5 cm) was required to be developed. Recent advances in photogrammetric structure from motion software allow for the development of high resolution mosaics from relatively simple methods that essentially pixel match a series of individual overlapping photos. The following section details how this was developed and applied in Port Curtis and Rodds Bay to produce the imagery for DFT analysis.

2.2.1 Equipment preparation

Field equipment for the surveys included the vehicle (helicopter), the camera (with spares), a camera staff, a computer and GPS units (multiple).

Helicopter

A four-seater light helicopter was used (Figure 3). Although other platforms such as Unmanned Aerial Vehicles (UAV's) could be utilised, for our purposes, the faster speed and longer range of helicopters made them more suitable to cover the large meadow scale areas we were required to assess within the small tidal window available for the surveys. In addition, using helicopters allowed us to carry a heavier payload than typically available in most UAV's and we were therefore able to use larger cameras (see below) with better optics than available for most small research UAV's. For the helicopter survey one researcher sat in the front and assisted the pilot to navigate the prescribed course, maintaining the height and speed within the required parameters (see below) and liaising with the camera operator. The camera operator occupied the rear seat. The port-side door was removed, and the camera operator steadied the staff in its position with the camera triggering at a fixed rate from below the skids of the aircraft.

Figure 3**:** Helicopter used in DFT surveys.

Camera

The camera used was a Canon EOS 1100D (Figure 4a), a 12 megapixel digital SLR, with an 18-55mm Canon Zoom lens. The Canon cameras are easily customised using freely available *Canon Hackers Development Kit* (CHDK) if necessary, and the wired remote control trigger requires no particular protocol, and is easily wired to the camera. We used a Hähnel Giga T Pro intervalometer (fires camera at set intervals) and wireless actuator (Figure 4b). Images were stored on SD cards, with a transfer rate of 40 MB/s.

(a)

(b)

Figure 4: The (a) Canon EOS 1100D camera similar to those used in these surveys and the (b) Hähnel Giga T Pro intervalometer and wireless actuator.

Camera staff

During the survey, the camera was positioned below the level of the helicopter skids. This was achieved using a purpose-made aluminium staff (Figure 5), built to firmly hold the camera aray, and to safely contain the wiring. The staff included a padded bracket which transferred the weight of the system to the floor of the helicopter without contact with the outer shell of the vehicle to avoid unnecessary vibration. Two lanyards ensured the safe containment of the staff, one at the top fixed to the back of the operator's chair, with the other close to the camera, under tension, countering the force of the air resistance on the camera to keep it steady. These tethers were arranged so that at no time was it possible for the staff to fall into a position that compromised the safety of the aircraft or the passengers.

Figure 5: Camera staff as used in helicopter (installed below helicopter skids after take off).

Computer

A Dell Toughbook with ArcGIS software was used to plot the flight pattern, with the in-built GPS allowing a real-time position indicator on the screen to show the progress of the aircraft along the flight pattern, and to maintain a trail of the actual path flown. This allowed the researcher and pilot to monitor the deviation of the aircraft from the planned track, so that remedial action could be taken, and once the prescribed flight pattern was completed go back and fill in any gaps that resulted from sudden wind shifts or pilot error. A constant dialogue was maintained between the navigating researcher and the pilot, to ensure that horizontal and vertical position, as well as speed, was maintained according to the planned route and within the prescribed specifications for future image processing.

GPS

The cameras used for the surveys did not have in-built GPS capabilities. As such, the position of the aircraft had to be known for each camera actuation. The bundle adjustment process used to develop the orthomosaic from the photographs is capable of producing a resultant mosaic with geopositional accuracy of a few centimetres. For such a level of accuracy, it is necessary to have a GPS, preferably Real-Time Kinematic (RTK), at a very close position to the camera, with an RTK ground station set up and with sufficient visible ground control points being surveyed from which the imagery can be georeferenced. For our purposes, it was sufficient that the survey placed the imagery within a few metres of the correct location, so that we could subsequently co-register repeat surveys to each other based on visible features. The bundle adjustment process itself manages the relative position of pixels within the mosaic to a high degree of accuracy. This method eliminates the need for expensive ground surveys to be carried out.

On each flight, there were typically three GPS units tracking the current location of the aircraft, any one of which was sufficient to be able to match photograph frames with a time. These were the pilot's navigation GPS, the GPS embedded with the Toughbook and an additional hand-held GPS.

2.2.2 Flight planning

Flight planning involved many considerations, including timing, cost and imagery requirements. Overall timing was governed by the lowest tides that fell into the approximate quarterly survey schedule. It was important to produce flight patterns which adequately covered the targeted meadows within the tight window of opportunity that tidal exposure provided— typically one or two hours either side of a spring low tide.

Choice of parameters

There are many factors which play a part in the results of an aerial survey, determining the final spatial resolution, geolocation accuracy and extent of coverage. These include:

- Height
- Speed
- Image overlap along an individual leg (the flight path flown Figure 6))
- Image overlap between parallel legs (side overlap)
- Camera trigger frequency

For photogrammetric techniques employing Structure-from-Motion (SfM) bundle adjustment software to produce structure and geo- and ortho-rectification from multiple GPS-tagged images, it is desirable to overlap consecutive images by 80% along the azimuth, and then by 60% along each adjacent path. It was decided that the camera would be triggered at a fixed frequency, as this provided the simplest operation of the camera, without requiring location feedback from GPS to instigate the data capture. This removed a complication that may increase the chance of error on site. As a result of this, the speed of the helicopter must be constant. It was necessary to choose a speed which could safely be maintained in all reasonably foreseeable conditions as well as a speed that allowed the completion of the survey area within the available tidal window, which was around 60 knots.

A focal length of 18 mm was chosen for the camera, which balanced field of view with minimal distortion. A targeted pixel dimension of 5 cm was chosen, as the DFTs were observed to have a minimum width of about 10 cm. The flight height was then determined using

$$H(m)=\frac{W_{im}\times P_x\times F}{W_s\times 100}$$

where H(m) is the height in metres, W_{im} is the image width (4272 pixels), P_x is the pixel size, F is the focal length (mm) and W_s is the sensor width (22 mm). This equates to a flight height of 174 m, or 573 feet. With 5 cm pixels, the image height on the ground will be 2848×5/100=142 m. For an 80% overlap along the azimuth, we therefore require the camera shutter to trigger every 28 m flown. We can comfortably trigger the camera at 1 second intervals, which equates to 28 m at around 55 knots. The image width on the ground is 4272×5/100= 214 m. Thus for a 60% overlap between legs (adjacent parallel flight runs Figure 6), we require a distance of 128 m between each leg.

Local test flight

An initial test flight, to assess the suitability of the specifications above, was flown at a seagrass location close to Cairns Airport. The results were successful, except that it was found in some instances that the distance between legs was too great to ensure sufficient overlap. It was therefore decided that the distance between parallel legs for the actual DFT surveys would be reduced from 128 m to 100 m.

Gladstone flight plan

Flight patterns were designed to meet the configuration described above, for each of the five seagrass meadows to be visited. A typical example, Wiggins Island, is shown in Figure 6. The actuation of the camera was checked frequently, and the adequacy of the images captured was confirmed after the conclusion of the flight on each day, to ensure that the data capture was successful.

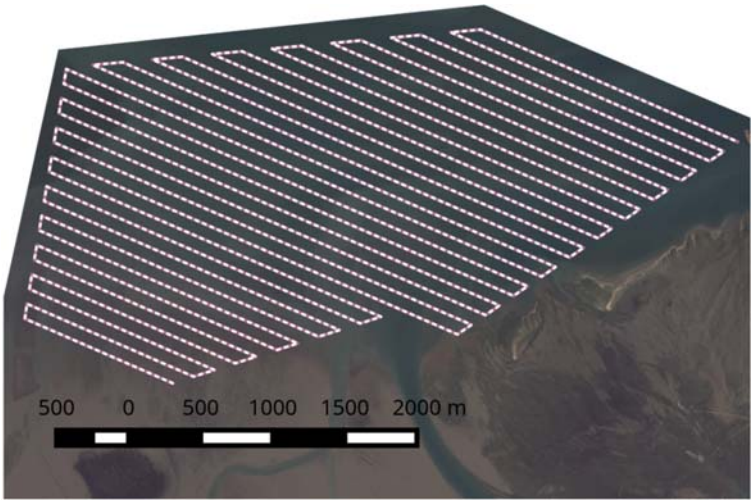

Figure 6 : Wiggins Island flight pattern showing planned layout of parallel legs spaced 100m apart to generate sufficient image overlap for the generation of suitable orthomosaics.

2.2.3 Data capture

Helicopter surveys were carried out on the dates shown in Table 1. Site alternatives stated as *high* and *low* indicate that *high* was flown at the standard height of 175 m and *low* denotes a flight done at half of the usual height and half of the usual speed to examine the benefits of having twice the spatial resolution in the resultant imagery. Flights commenced two hours prior to low tide. At a time close to each survey, the exact time on the camera was recorded in relation to the exact time determined by an atomic clock using the website *time.is*[1]. This time was identical to the time on the GPS units.

2.2.4 On-site verification

At the completion of a day's survey, the photographs were run through an initial stage of the photogrammetry software (Pix4Dmapper Pro) to produce a quality report to ensure that sufficient coverage had been achieved for the bundle adjustment process to work. Verification of imagery allowed sufficient time to take remedial action the following day if discrepancies arose.

2.2.5 Preprocessing

Before the extraction and analysis of seagrass and DFTs could begin, the data needed to be converted from sets of individual photographs to georeferenced orthographic images, to ensure spatial integrity. This involved multiple steps, including the elimination of irrelevant photographs, the insertion of geopositional information into the photograph EXIF tags, the SfM bundle adjustment process and finally the coregistration of each of the images corresponding to the same locations in a time series.

Filtering of images

One survey of a single seagrass meadow resulted in the capture of between 1,000 and 5,500 photographs. The bundle adjustment process which is used to turn these into useful imagery is very

[1] http://time.is/

computationally intensive, so a manual elimination of obviously irrelevant photographs was performed (e.g. shot over water and mangroves) prior to processing.

Correlation of GPS data with individual photographs

As the camera used for the surveys was not GPS-enabled, each photograph was tagged with its corresponding horizontal and vertical position using the open source software *GPS Correlate*[2]. The software reads the time tag on each image, and looks for the closest time on the GPS track recorded from the independent GPS unit, interpolating or extrapolating to adjust the GPS track position to match the time exactly, before feeding the geo-positional data back into the EXIF tag of each photograph.

Bundle adjustment

The bundle adjustment process aggregates the GPS coordinate grid, and knowing the camera parameters, is able to more precisely calculate the likely position and poise of the camera when each shot is taken. This is done by matching clusters of patterns identified on multiple images, and performing many triangulation calculations. This adjustment can be seen in the "ray cloud" image shown in Figure 7a. The blue dots represent the location logged by the GPS, and the green dots represent the position as corrected by the SfM software. The image shows the ends of a series of runs, in which the helicopter is being turned from each run to the next. The poise of the camera as the helicopter banks can clearly be seen through the turns.

From these calculations, a 3-D point cloud is constructed, which is used to ortho-rectify the image mosaic. A snapshot of the point cloud at South Trees is shown in Figure 7b. The bundle-adjustment process, taking the data from geo-tagged photographs through to an ortho-mosaic and digital surface model (DSM), was carried out using the proprietary software Pix4Dmapper Pro[3].

(a) 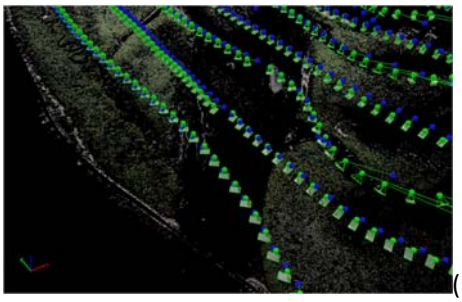 (b) 

Figure 7: a) Ray cloud from the SfM software, showing camera location and poise calculations at the end of a series of survey legs and b) South Trees point cloud.

Coregistration

Once the orthomosaics were produced, the time series of images for each location were coregistered in relation to each other to ensure that the geo-positioning of each mosaic corresponds as closely as possible. Prior to co-registration the process results in the position of imagery from different times to be within a few metres. Coregistration involves the identification of invariant features located on each image, which allows GIS software to determine a polynomial function to apply to the coordinates of each pixel in-order to bring them into alignment with the reference image. Having done this, the software then uses a bilinear or convolution method to resample the pixel values in order to fit them into the transformed image. We used the GIS package QGIS for coregistration (Figure 8). For each location, the initial sampling image was used as a base for coregistration, and this was built upon with consecutive orthomosaics being georeferenced to the most recent image to minimise the propagation of spatial errors.

[2] https://github.com/freefoote/gpscorrelate

[3] https://www.pix4d.com/product/pix4dmapper-pro/

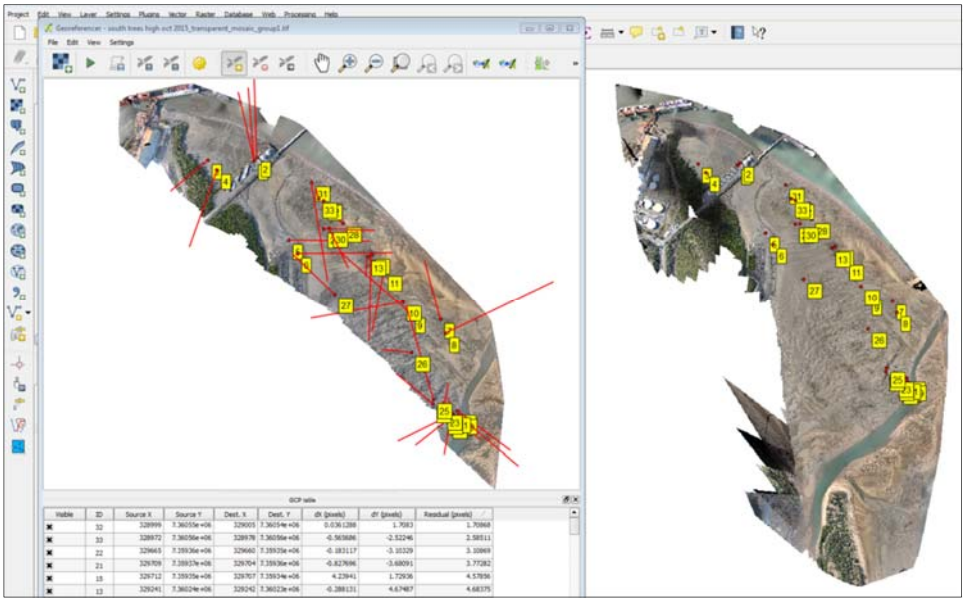

Figure 8: Screen shot of the QGIS coregistration interface.

2.3 Identification of Dugong Feeding Trails

DFTs can be identified generally as sinuous paths of clear substrata through patches of seagrass. Their width is recognisable as being between 100 and 250 mm. However dugong feeding can create larger cleared areas which may only be identified by their connectivity with the recognisable thin trails (Figure 9). For the purposes of this study only linear features that clearly occurred within seagrass and could be identified as feeding trails were used in analysis. Larger cleared areas that may have multiple sources were excluded. DFT linear features were quite distinct from other linear features in seagrass such as propeller scars and feeding or excavation by other animals such as rays. Examples of these are presented in the results.

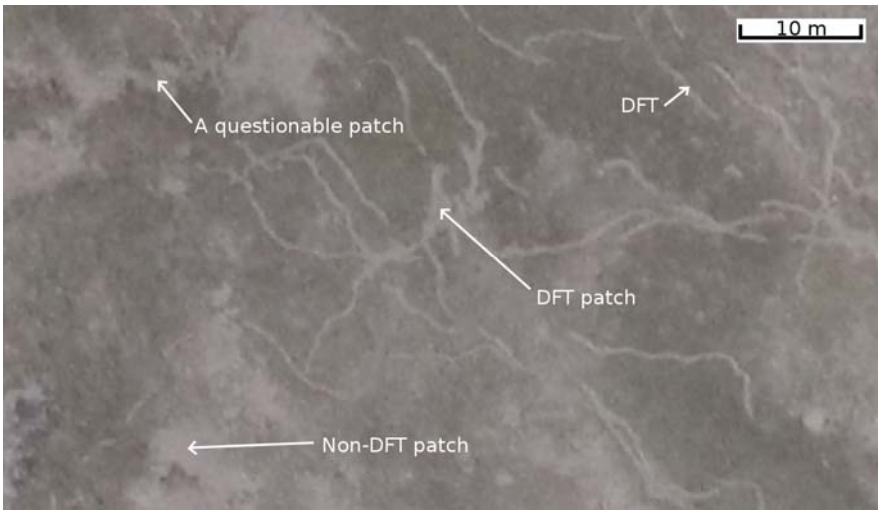

Figure 9: An example of distinct DFTs as well as patches that were definitely not DFTs or where there was uncertainty (questionable patch). Only distinct DFTs were included in the analysis.

2.4 Quantification of Dugong Feeding Trails

Two approaches were used to classify DFTs from the orthomosaic imagery. The first was a manual visual census and classification of the imagery by experienced observers. As this was a time consuming method the second approach was to investigate the development of supervised and semi-supervised algorithms to automate the process as much as possible. For both methods, ground-truthing of potential trails seen during the high level photographic runs was performed in the field by periodically bringing the helicopter down to ground level where observers could examine features in more detail.

2.4.1 Visual Census of Dugong Feeding Trails

A visual census of DFTs was conducted on coregistered orthomosaic images at each location and at the seven sampling periods during the study (N.B. Image quality at Pelican Banks in May 2015 was not considered adequate and was left out of the assessment, and sampling at Rodds Bay South commenced in August 2015 so both of these sites had six sampling periods). To achieve this, a fixed grid template with 50 m x 50 m cells was created to overlay on the targeted sampling area for each location which excluded obvious vegetated land areas and deep water (Figure 10). The number of grids at each location varied due to differences in the size of seagrass meadows, with totals as follows: Rodds Bay North – 929; South Trees - 989; Rodds Bay South - 1688; Wiggins Island - 2447 and Pelican Banks - 3637. A 50 m x 50m grid cell size was selected as this provided both a level of resolution that was appropriate to observe feeding trail features before they became too pixelated for visual discrimination as well as a fine enough scale to determine spatial changes in DFTs across the meadow

Each individual grid cell was assessed by TropWATER seagrass ecology group staff with experience of DFTs in the field. Assessment noted seagrass presence or absence and DFT presence or absence on an LCD computer monitor at a scale of 1:300 or less. DFTs were considered present only if apparent in seagrass and did not include the assessment of old trails in areas with no above ground seagrass biomass. Prior to assessment, a series of images containing dugong feeding trails, ray trails, propeller scars and footprints (see Figure 11 for examples) were sent to dugong experts (Helene Marsh and Chris Cleguer) for verification. DFTs identified had to be obvious linear features and did not include areas where grazing appeared to have occurred in concentration and lacked distinct lines. Trails were carefully and conservatively assessed and marked with a confidence score of 1 or 2 with 1 being very confident and 2 being fairly confident. Any grid cells that were scored 2 were reassessed at a later time by a second observer for confirmation and validation as part of a quality assurance procedure. Cells that showed image distortion or water coverage which prevented assessment of potential trails were given an ‘unable to assess’ classification and were discarded from overall values.

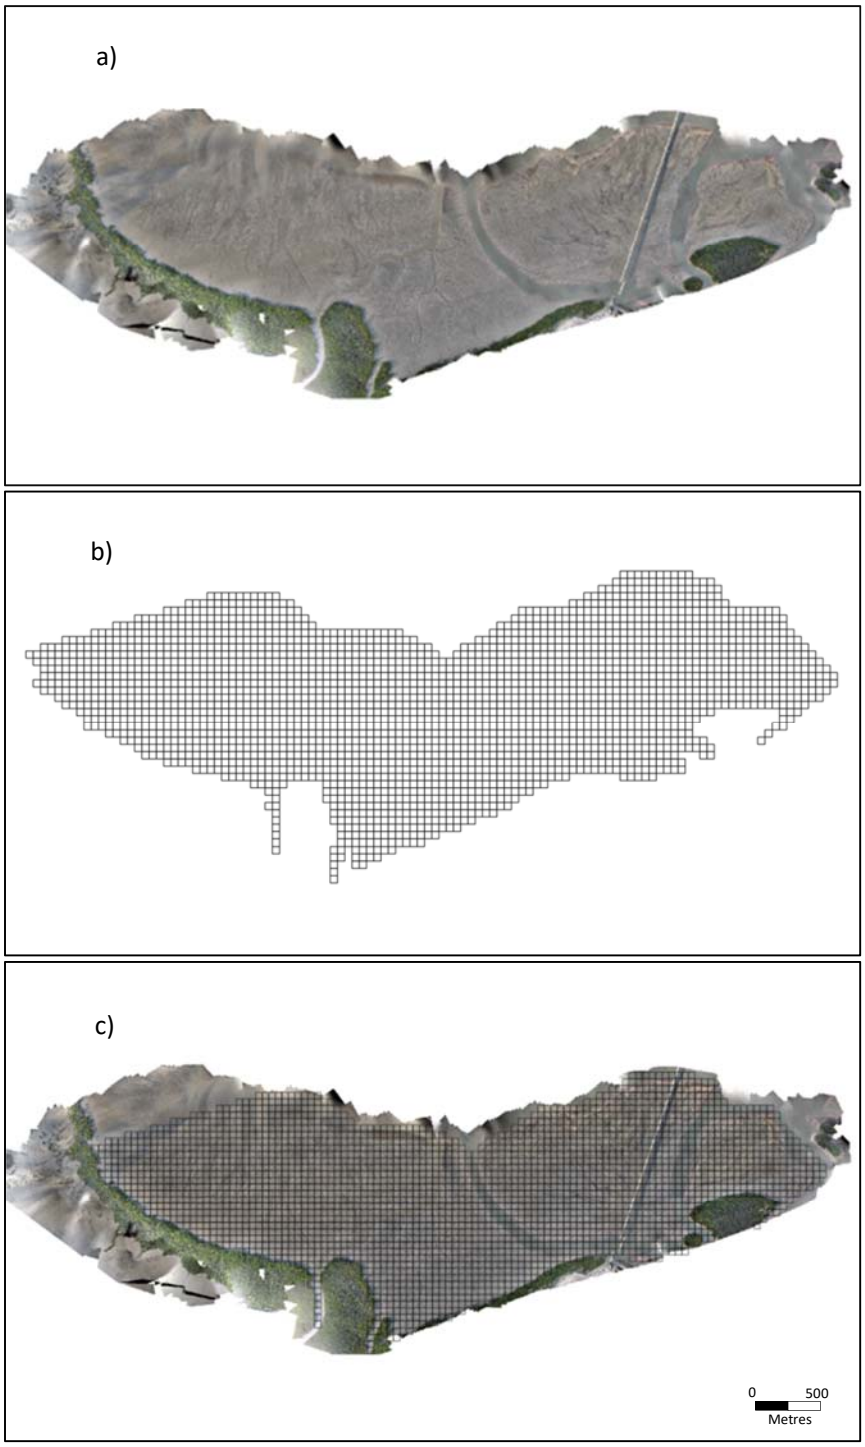

Figure 10: Example of (a) Wiggins Island mosaic, (b) grid template for Wiggins Island and (c) template applied to mosaic ready for assessment.

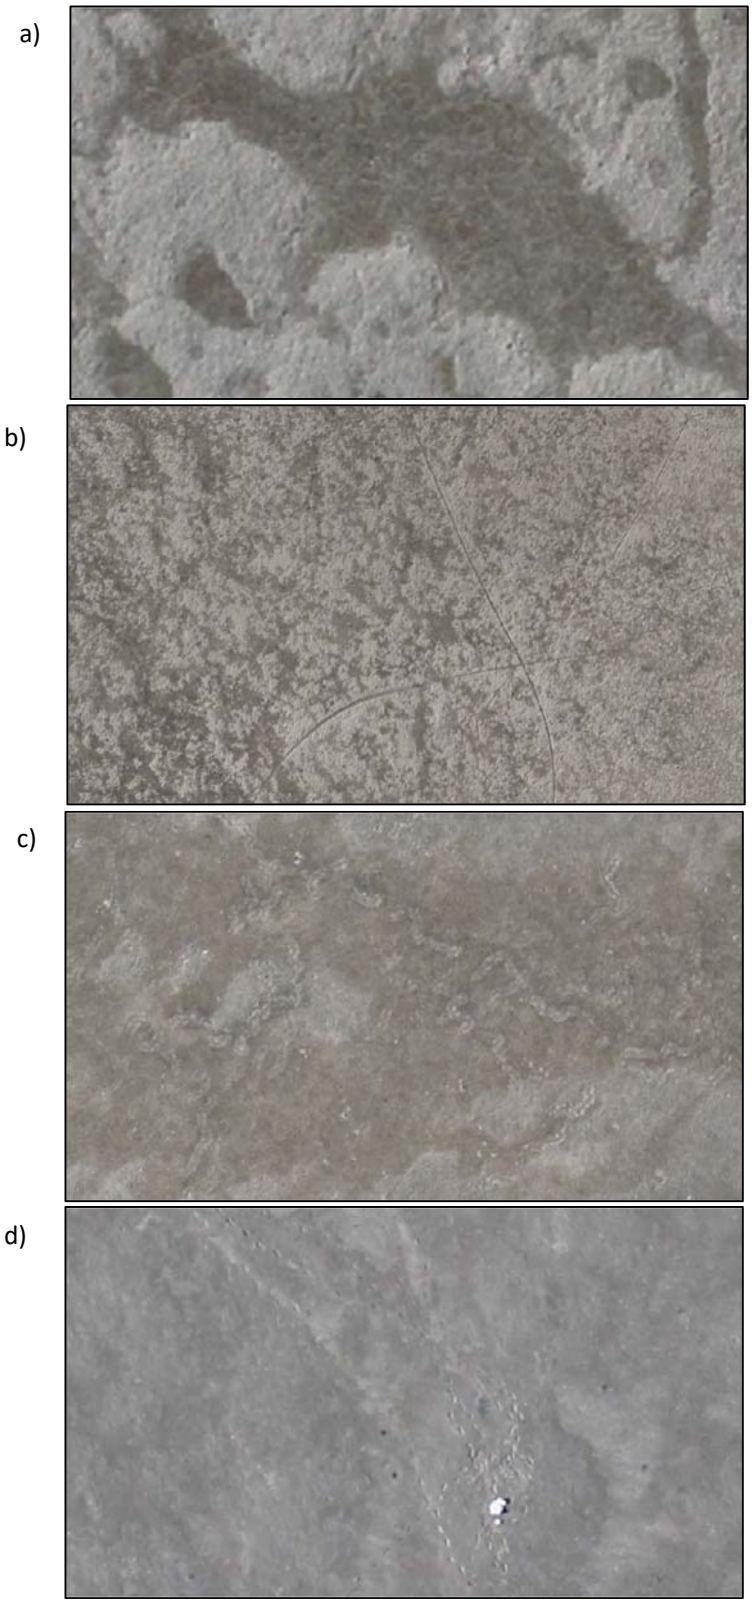

Figure 11: Examples of: (a) DFTs in seagrass (pale trails in darker area), (b) boat propeller scar, (c) sting ray trails and (d) footprints and buckets from on ground sampling team.

2.4.2 Algorithm Development for Automated Extraction of Dugong Feeding Trails

There was an iterative approach to the development and testing of automated and trained algorithms for extracting DFTs. A more complete description of the process and pathways investigated can be found in Appendix 2.

The nature of DFTs as well as the background imagery presents a range of challenges to automated feature extraction. The lengths of individual trails vary from 30 cm to more than 6 m, however, these can sometimes be interconnected forming single continuous features up to 20 m in length. Their length is often discontinuous, increasingly so over time with seagrass regrowth into the DFT, and the relationship between disconnected components of the same trail may only be recognised by our brains' ability to "join the dots", something which humans are very good at, and which computers are only beginning to succeed at, and only after considerable periods of algorithm training.

As there is a considerable amount of variability in the length and continuity DFTs, to be able to extract the features there needs to be a clear distinction of what comprises a DFT. Once separation of individual features based on pixel values (segmentation) has taken place, then classification of these linear features based on their characteristics, such as shape, size, texture and spectral signature needs to be addressed.

The spectral signature of the trails was determined by using a principal components transformation, and selecting the first component (PCA1). In order to normalise the spectral variance further an additional layer was created to represent a standardised score from the PCA1 values bringing the background contrast of feeding trails into similar ranges to clarify the DFT boundaries (Figure 12).

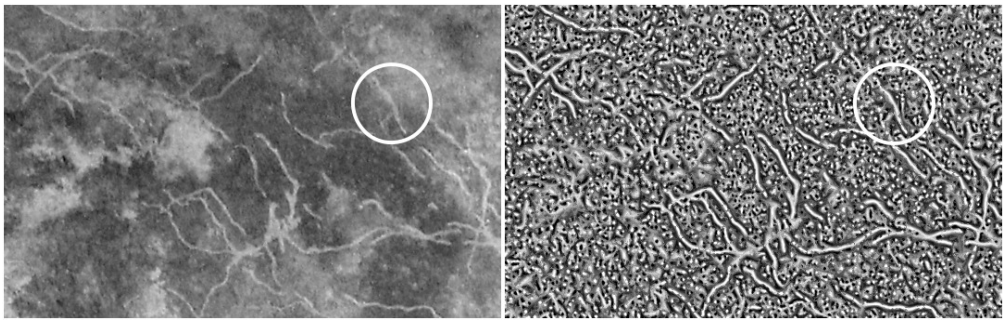

Figure 12: First principal component of the red, green and blue reflectance values of a section of a seagrass meadow (a), and the standardised score of those values (b). The highlighted region encircles a trail whose contrast with the background is much less than those other trails towards the centre of the image. The standardised score brings the contrast into a similar range of values.

Following separation of individual features based on pixel values, shape characteristics can be determined. Metrics including compactness, fractal dimension and sinuosity were used to ascertain the optimal algorithm for this extraction. However, as these metrics produce fixed thresholds and there is high variability in length and width of trails (i.e. trails turn back on themselves or cross over) this approach runs the risk of increasing errors of omission and inclusion. To normalise trails and decrease this variability a combination of skeletonisation of the features (thinning until 1 pixel width) and distance transforms were used. This, however, introduced the possibility of inhibiting original connectivity of trails and therefore required some tolerance to be built into the measure (Figure 13). For the delineation of DFT boundaries and the association of fragmented parts of a trail, a number of methods to define their morphology and shape were investigated including edge detection measures, textural predictors, binary and greyscale morphology operations, and active contours. The most effective operators for DFT extraction were applied in the algorithm (see Appendix 2).

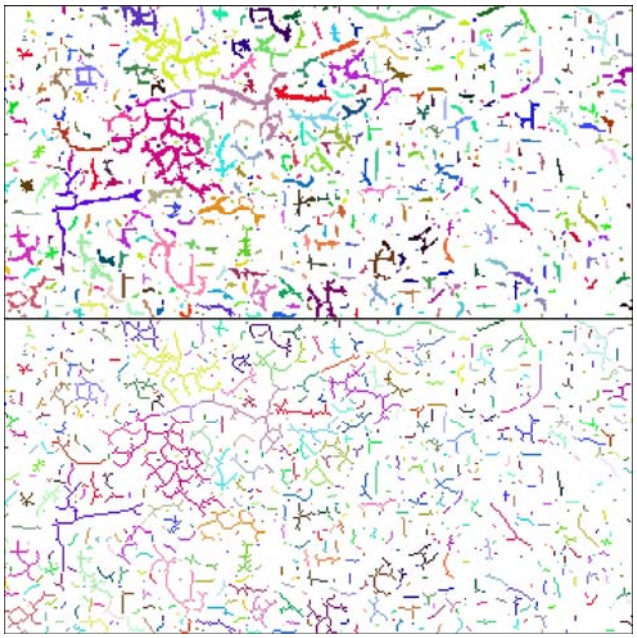

Figure 13: Original features (above), thinned by morphological "skeletonisation" until 1 pixel thick (below).

To optimise the parameters required in the algorithm for trail extraction, two initial approaches were trialled. The first approach was the training of a machine learning algorithm, which was used as a trial on a small area of the South Trees meadow. In order to reach an acceptable level of accuracy, a great deal of training was necessary. Training involved the selection of hundreds of features identified as DFTs, and the identification of linear features such as wave fronts and gullies that were not DFTs. The training required a substantial investment of time that presented challenges which may have been overcome by increased workforce but ran the risk of increased inconsistency due to the variability of all DFT types.

Given the variability in DFTs and the volume of data, a second approach to optimise the parameters required for the algorithm to extract DFTs was used in which ranges were set for the major parameters of determining a trail and were then individually varied with the others fixed to determine a 'best fit'. To reduce computation times, parameters were applied to small sample areas of the sites considered representative of the whole study area and individual parameters were altered in reasonably large increments. This approach produced multiple sets of potential parameter-set images, termed 'swatches', from which to determine the best results (Figure 14)

Due to the variability in image characteristics between sites and also sampling times (exposure, light, sediment characteristics etc.) swatches were required to be produced for each meadow and each sampling time. Once the best swatch results were determined for the test area for a particular mosaic, those parameters were then applied to the full meadow to extract DFT features from the image.

For a detailed narrative of DFT extraction and algorithm development see Appendix 2.

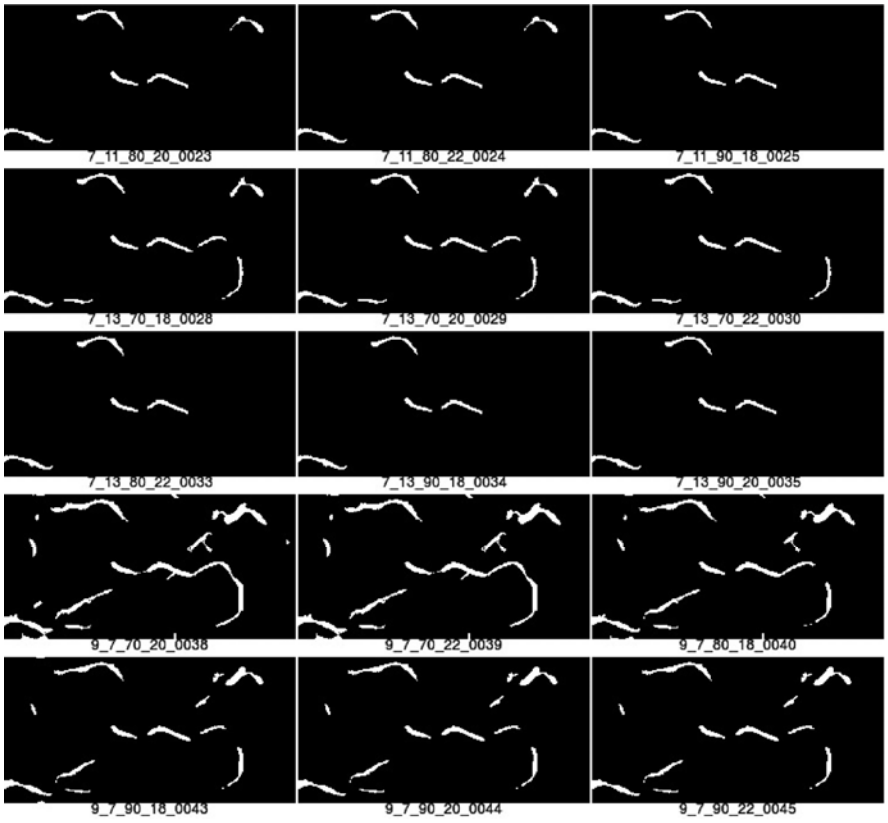

Figure 14: Extract from a swatch. The whole swatch might contain hundreds of sample images with potential DFT features to help in parameter selection.

2.5 Seagrass and Environmental Data

Seagrass change and available environmental data were examined to establish what links they may have in explaining dugong feeding activity changes in Port Curtis and Rodds Bay.

Seagrass monitoring in Port Curtis is conducted by the research team as part of two separate research programs. Full details and results of those programs can be found in Davies et al. (2016) and Bryant et al. (2016). As part of monitoring linked to the Western Basin Dredging and Development Project (WBDDP), a series of sensitve receptor seagrass assesment sites were established for quarterly monitoring of seagrass change. Five of these sites coincided with areas assessed in this DFT project and were used as part of correlative analysis with the DFT results.

Permanent transect sampling design

Seagrass was sampled quarterly between May 2015 and November 2016 at five intertidal permanent transect sites within the monitored DFT meadows: four at the Pelican Banks meadow (GH1 – GH4) and one at Wiggins Island (WW2). At each site three parallel transects laid 50 m apart were established. Above ground biomass was estimated from 11 quadrats (0.25 m^2) per transect (n = 33 quadrats per site) (see Bryant et al. 2016). Seagrass sampling was conducted within a day or two of the DFT surveys.

Seagrass biomass

Seagrass above-ground biomass and species composition were estimated for each quadrat using the "visual estimates of biomass" technique (Mellors 1991; Rasheed and Unsworth 2011). This technique involves an observer ranking seagrass biomass in the field in each quadrat while referring to a series of photographs of quadrats with similar seagrass habitats for which the above-ground biomass has been measured previously. The relative proportion of the above-ground biomass of each seagrass species within each quadrat was recorded. Field biomass ranks were then converted to above-ground biomass estimates in grams dry-weight per square metre (g DW m^{-2}). At the completion of sampling, each observer ranked a series of photographs of calibration quadrats across the range of seagrass biomass observed in the survey. Seagrass in these calibration quadrats had been harvested and the actual seagrass above ground biomass (in g DW m^{-2}) determined in the laboratory. A separate linear regression of ranks and measured biomass from these calibration quadrats was then generated for each observer. The regression was then applied to field biomass ranks to account for differences in observers between sampling periods to determine above-ground biomass estimates.

Environmental conditions

Environmental conditions in the seagrass meadow were recorded using two separate *in situ* loggers:

1. Autonomous iBTag™ submersible temperature loggers recorded water temperature (°C) within the seagrass canopy every 15 minutes; and
2. Submersible Odyssey™ photosynthetic irradiance (light) autonomous loggers recorded irradiance (measured as photosynthetically active radiation, PAR) every 15 minutes.

Tidal exposure to air for the meadow was calculated by summing the total daylight hours that tidal height was $\leq$ 1.0 m (the tidal height at which seagrass meadows expose at this location). Tidal data was provided by Maritime Safety Queensland (© The State of Queensland (Department of Transport and Main Roads) 2016, Tidal Data) for Gladstone at Auckland Point (MSQ station # 052027A; www.msq.qld.gov.au). Total rainfall (mm) and air temperature were obtained for the nearest weather station (Gladstone Airport, station # 039123) from the Australian Bureau of Meteorology website (http://www.bom.gov.au/climate/data/). Calliope River water flow data (total monthly megalitres (ML)) was obtained from the Department of Natural Resources and Mines (station # 132001A; www.watermonitoring.derm.qld.gov.au).

For each of the sampling periods total rainfall, total river flow from the Calliope River (ML, megalitres), total daytime tidal exposure, mean total daily PAR, mean maximum daily water temperature (°C), and mean maximum daily air temperature (°C) were determined for one month and three months prior to the day that seagrass biomass and DFTs were measured.

2.6 Statistical Analyses

We used binomial generalized linear models with a logit link function (Zuur et al. 2013) to model the effects of (1) sampling period (*p*) and location (*l*) on the probability of DFT presence (*DFT* ~ *p***l*); and (2) the proportion of grid cells with seagrass present (*s*) for a given *p* and location (*l*) on the probability of DFT presence (*DFT* ~ s_p**l*). The same method was also used to model the effects of (1) seagrass biomass from nearby seagrass monitoring sites (Bryant et al. 2016) (*b*) and location (*l*) on the probability of DFT presence (*DFT* ~ *b***l*); and (2) the effects of seagrass percent cover from nearby seagrass monitoring sites (*c*) and location (*l*) on the probability of DFT presence (*DFT* ~ *c***l*). Two meadow locations were included in these analyses: Pelican Banks and Wiggins Island. Analysis was conducted using the lme4 package (Bates et al. 2015) and Figures 21-23 were created using the ggplot2 package (Wickham 2009) in R (R Core Team 2016).

Generalized additive mixed models (GAMM) were used to examine the effects of environmental variables on seagrass biomass using the "mgcv" package for R (Wood 2014). GAMMs fit a non-parametric model to the data where the functional form is not specified *a priori*, but instead additive non-parametric functions are estimated using smoothing splines (*s*) to model covariates (Zuur et al. 2014). GAMMs were used because data exploration indicated a non-linear response of biomass to PAR.

Prior to fitting models, environmental covariates were tested for collinearity using variance inflation factors (VIFs). Sepreate VIFs were calculated for environmental data collected one and three months prior to measuring seagrass biomass. Collinearity was high between environmental variables for both data sets. For the 1 and 3 month data sets, air temperature and water temperature were excluded from the analysis (negative relationship with exposure, see Appendix 1), as was river flow (positive relationship with rainfall, Appendix 1). The VIFs of PAR, rainfall, and tidal exposure were <1.1 and <2.2 for the one and three month data sets, respectively, indicating that collinearity was within reasonable limits and would not substantially inflate the standard errors of the model's parameter estimates (Zuur et al. 2009). The response variable biomass was averaged across 11 quadrats per transect before analysis. Transect means were analysed to reduce zero-inflation inherent in the quadrat data set which led to unstable models, with zero counts reduced from 55% of quadrats in the total data set to 8% of transects.

Separate global models for the one and three month data sets were run to determine optimal models for final analysis. The global models were:

$$B = s(P \times S) + R \times S + E \times S + \beta transect + \varepsilon$$

Where *B* is biomass, *P* is PAR, *R* is rainfall, *E* is tidal exposure, *S* is Site, β_{transect} is the random effect of transect, and ε is the random error term. Sub-model sets of the global model were generated using the dredge function in the MuMIn package (Bartoń 2013). The best-fit model was considered to be the simplest model with the lowest Akaike's Information Criterion corrected for small sample sizes (AICc) that fell within two of the lowest AICc (Burnham and Anderson 2002). A "weights" function was used to correct residual variance heterogeneity among sites (varIdent (form = ~1|Site)) (Zuur et al. 2009). Normalised residuals were inspected for the best-fit final model using residual plots and qq-plots for violations of the assumptions of homogeneity of variance and normality (Zuur et al. 2014).

3 RESULTS

3.1 Orthomosaic Generation

Individual photo images were taken from 175 m altitude and were of sufficient resolution and quality to visually identify DFTs in the seagrass meadows as well as remnant feeding scars in areas where seagrass had been lost. The images were used to successfully generate orthomosaics for the entire area of the investigation meadows (see Figure 15 for an example).

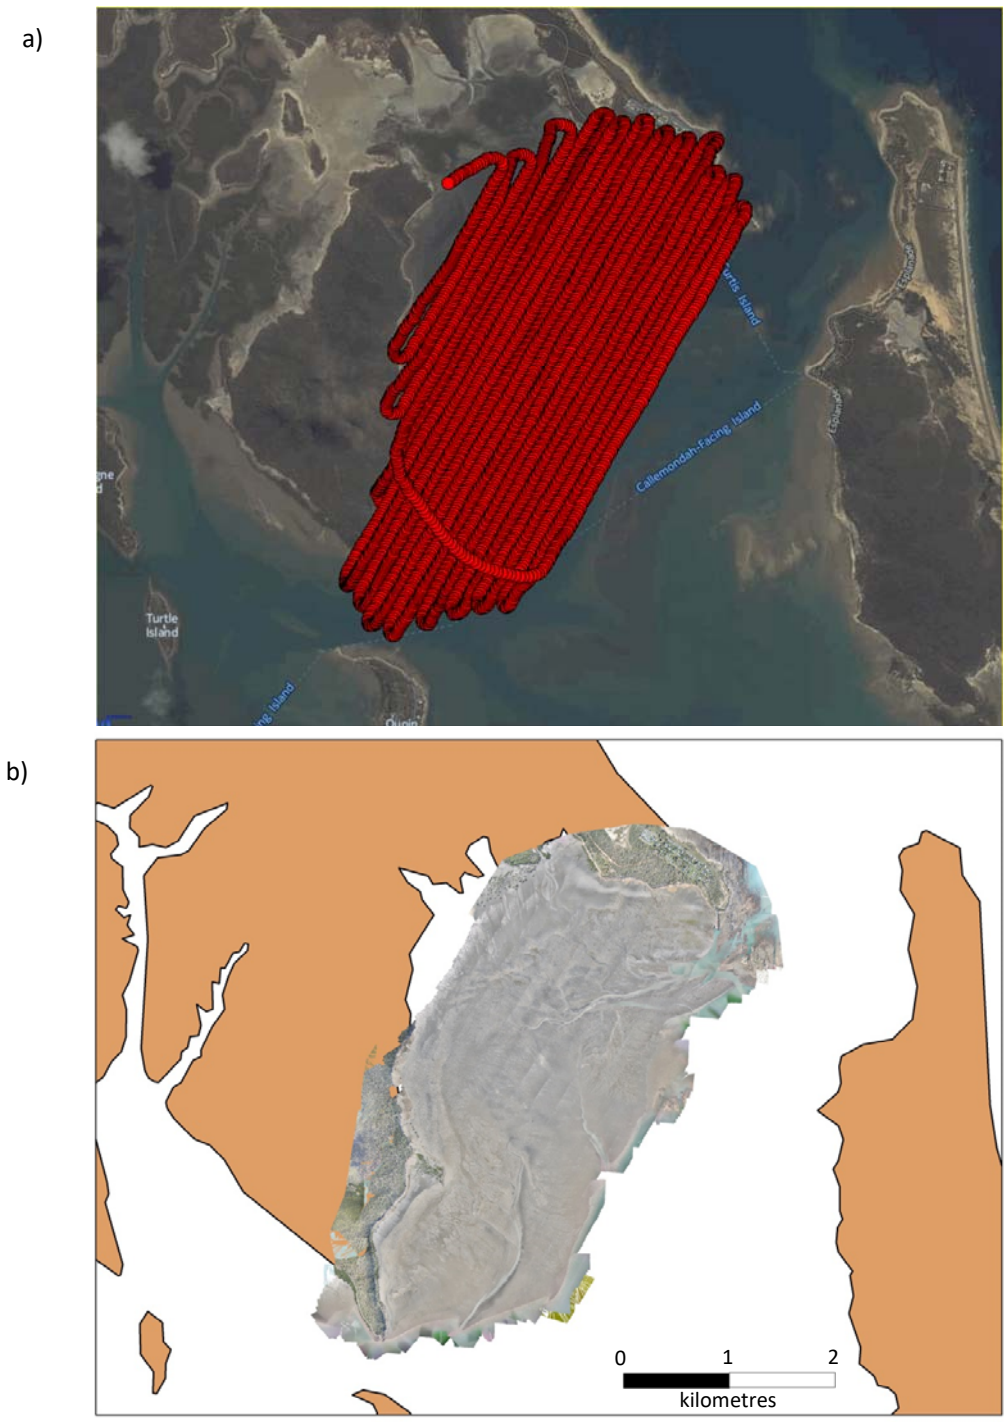

Figure 15: Example of (a) individual image locations before processing the orthomosaic for May 2016 survey at Pelican Banks, and (b) the final orthomosaic photo image after processing on Pix4D software.

Comparison with lower altitude flight

Imagery from the lower level flights (90 m) of the test area at South Trees (13th September 2015 to 7th February 2016) was also processed with orthomosaics generated (Table 1). These flights led to images with a reduced pixel size from ~5 cm to <2.5 cm compared with the standard 175 m altitude imagery. Figure 16 shows a comparison of the high and low level orthomosaics for the same area from the September 13th survey. While the smaller pixel image is visually sharper, the imagery didn't reveal any great improvement in DFT detection. The extra flight and processing time required for these images would have significantly added to the time and cost of the image collection making it infeasible for a project of this scale.

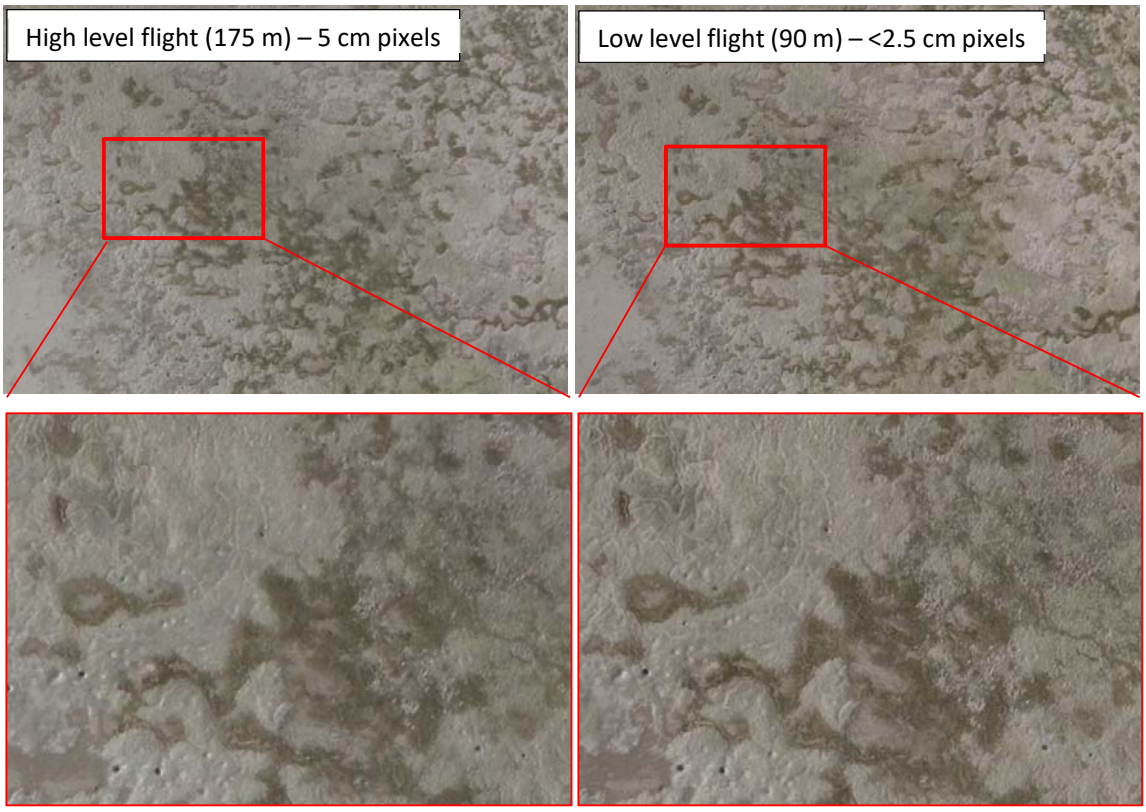

Figure 16: Comparison of high level (175 m) and low level (90 m) flights and resultant orthomosaic generation from the September 13th survey.

3.2 Longevity of feeding trails between surveys

Qualitative observations of a random selection of new DFTs identified in the interim surveys during spring 2015 at several locations showed that DFTs within seagrass were unable to be seen in the imagery from between two to six weeks later and were not present after ten weeks (Pelican Banks: Figure 17, South Trees: Figure 18, Rodds Bay North: Figure 19). Surveys on 13th September 2015 identified new DFTs that had not been present at the previous sampling period two weeks prior (29th August) at Pelican Banks and South Trees. These feeding trails were tracked over time and many had disappeared by the time of the following survey on 28th September 2015 and they had all disappeared apart from a small section of one trail at Pelican Banks after 6 weeks (28th October sampling). The best time series for DFT tracking available in Rodds Bay North showed new trails appearing on the 28th September 2015 sampling that were not present two weeks prior (13th September). Only a small proportion of these trails were still discernible four weeks later. Tidal inundation prevented further assessment of these trails at the next sampling period on 28th October 2015. While the more frequent longevity sampling was only conducted during the 2015 growing season, examination of images between quarterly sampling in February 2016 and May 2016 indicates that even during the senescent season for seagrass growth, trails did not persist during the three month time period between surveys, and feeding trails observed in each quarterly survey were likely to be novel (Figure 20).

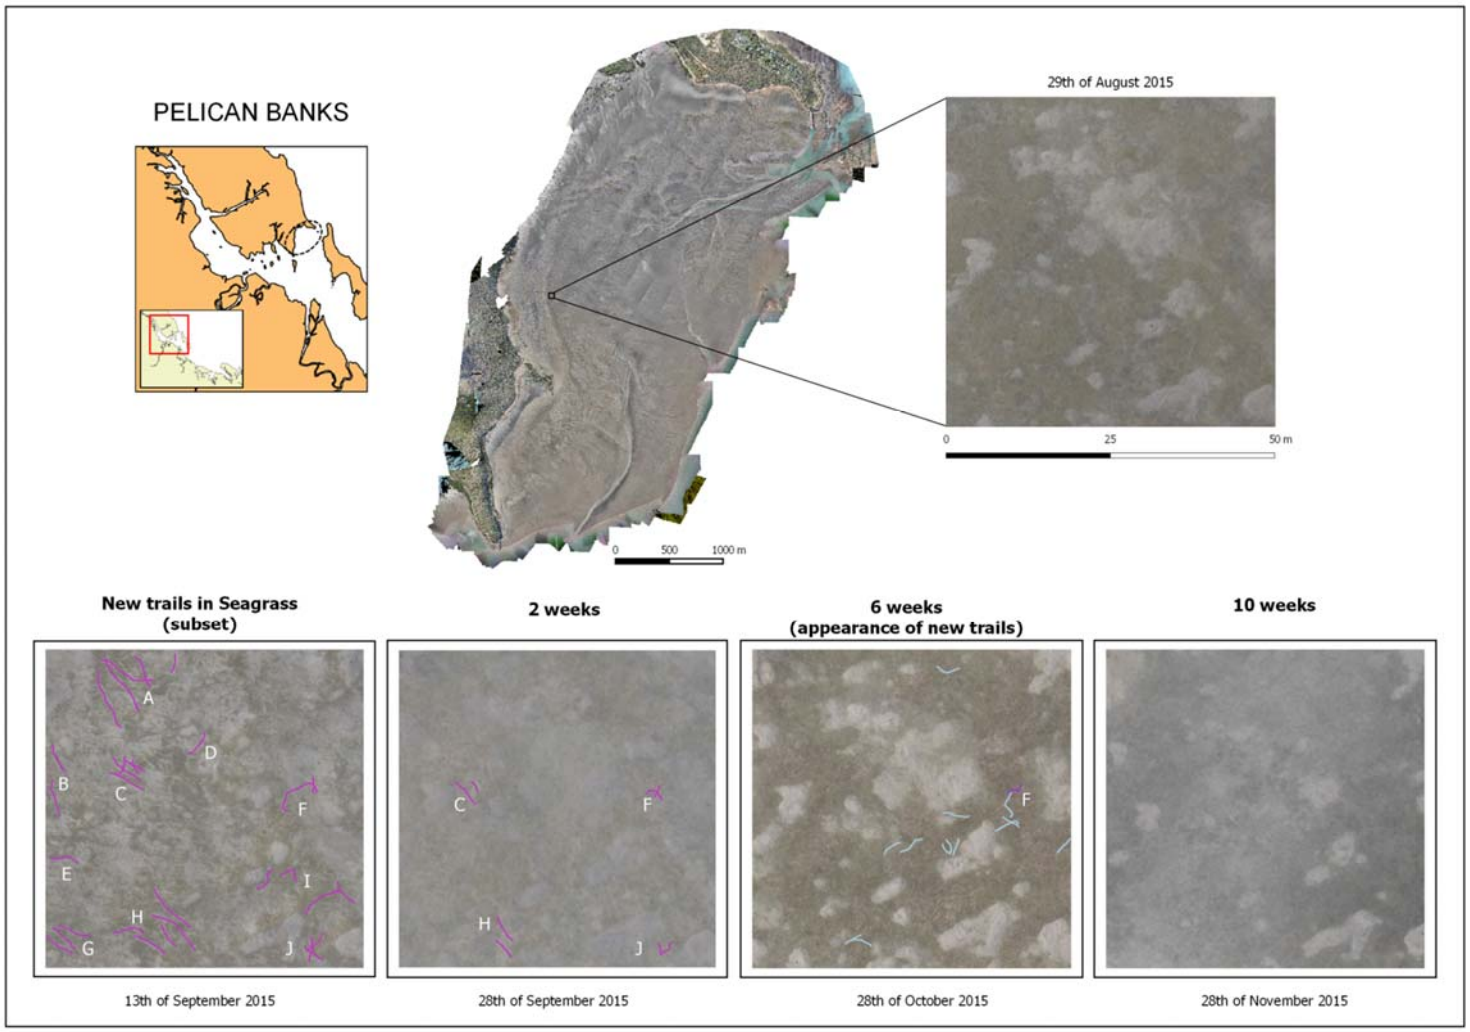

Figure 17: A time series tracking the appearance and disappearance of feeding trails was conducted over a 10 week period in spring 2015. At Pelican Banks, new trails (less than 2 weeks old) were identified on 13th September and only a few remained visible 2 weeks later. Part of one of the original trails was visible when resampled six weeks later, however new trails had appeared. None could be detected after 10 weeks (28th November).

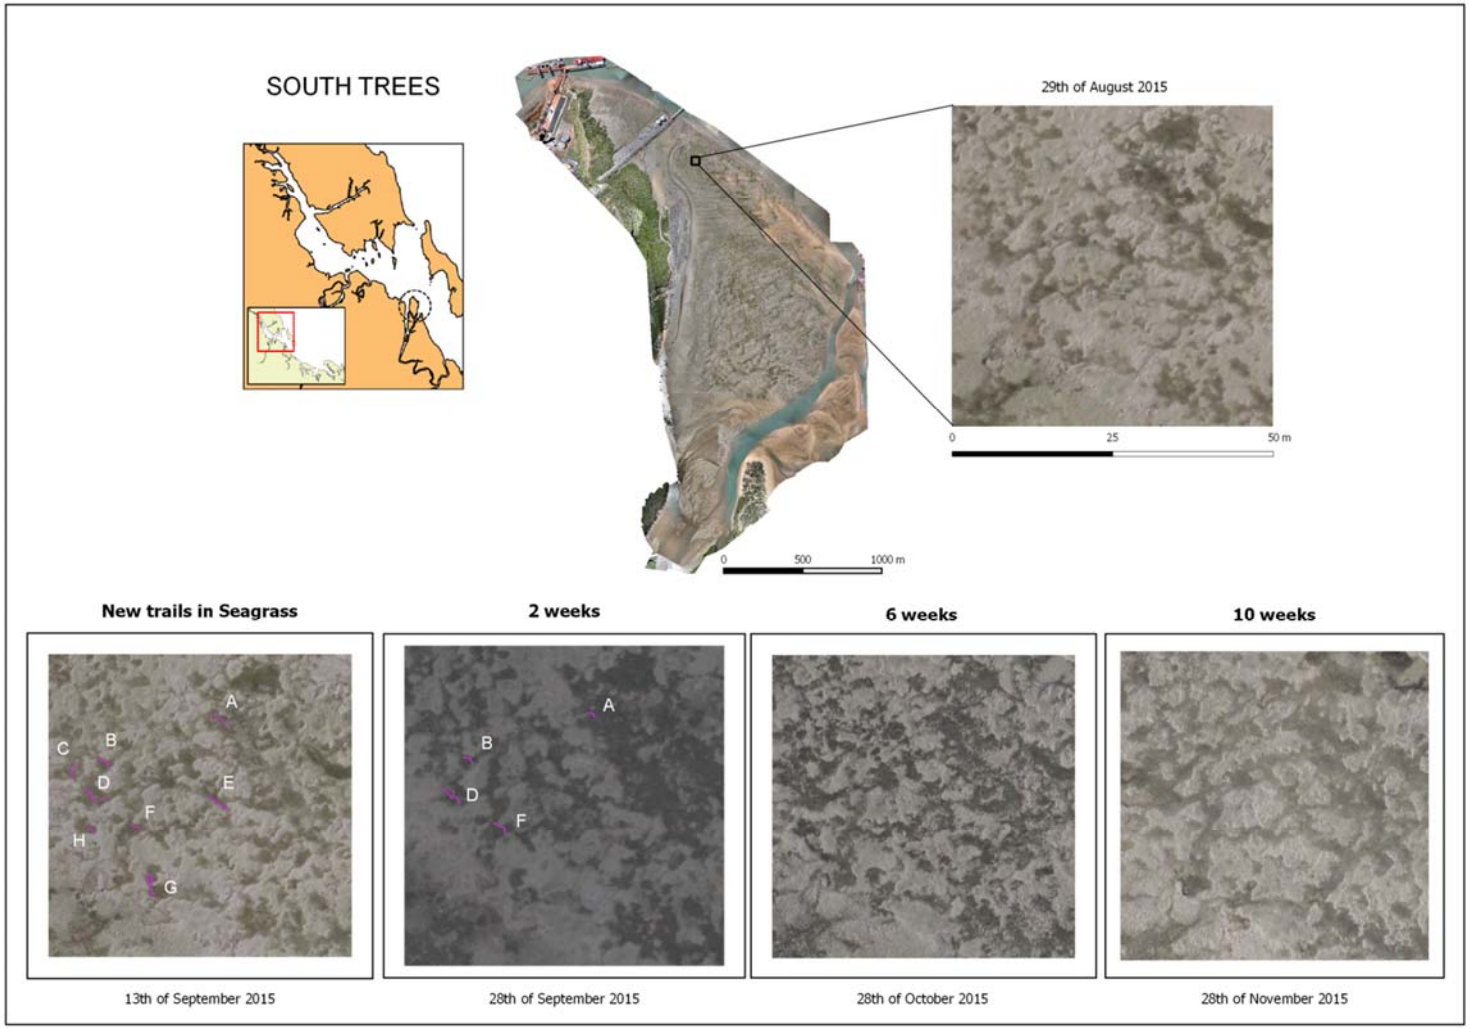

Figure 18: A time series tracking the appearance and disappearance of feeding trails was conducted over a 10 week period in spring 2015. At South Trees, new trails (less than two weeks old) were identified on 13th September and only a few remained visible 2 weeks later. None of these trails were visible when resampled six weeks later on the 28th October.

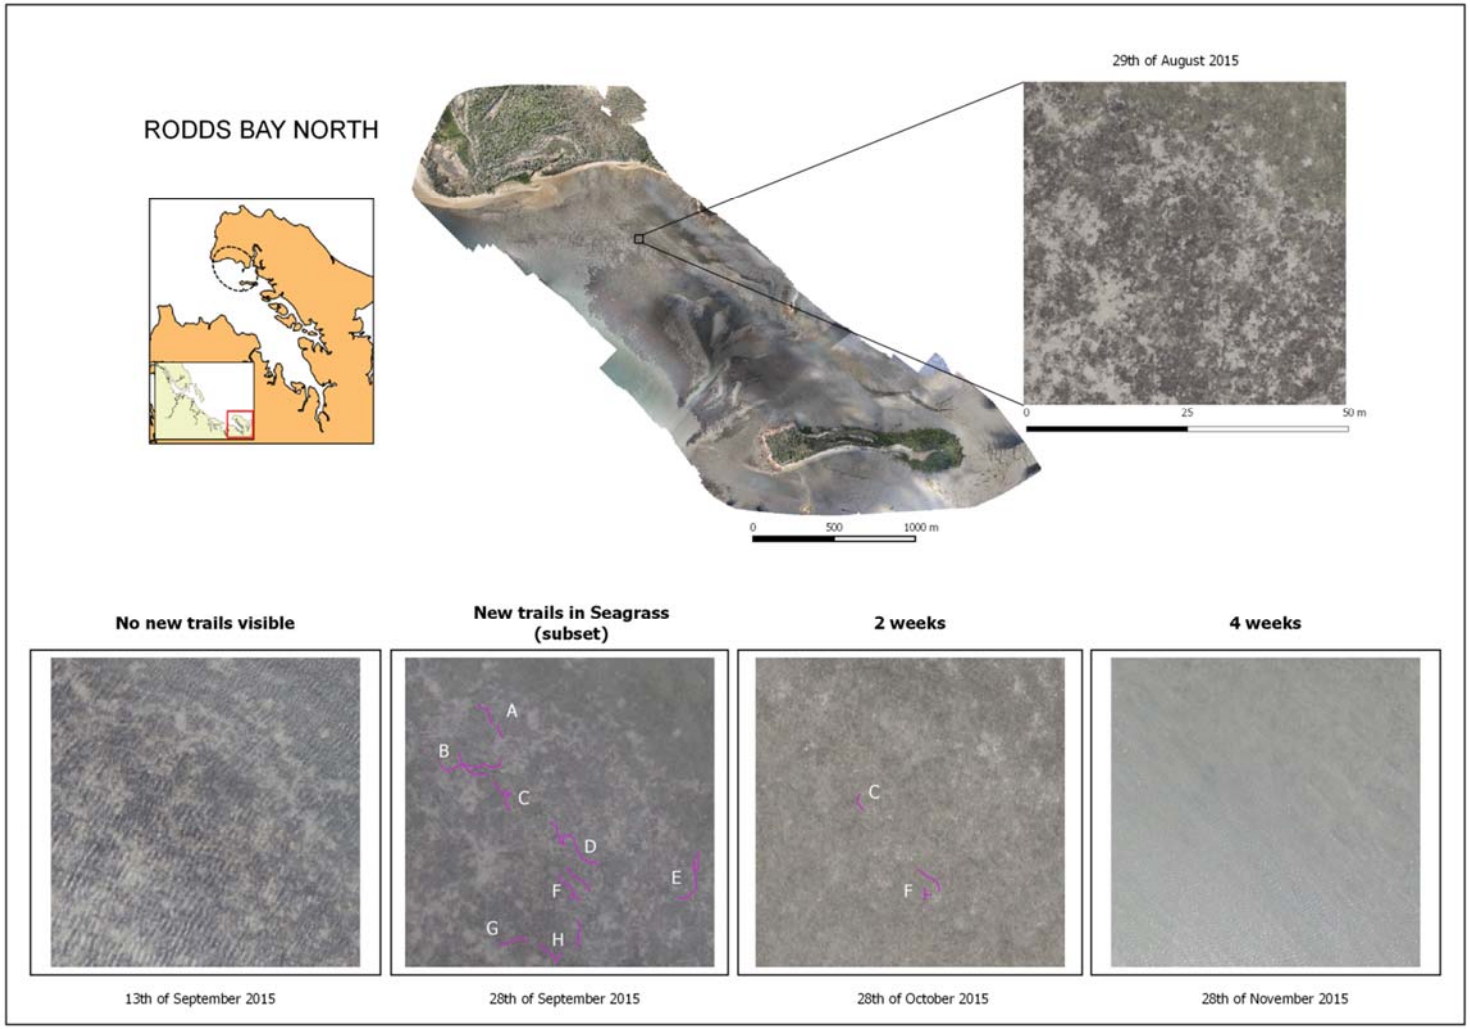

Figure 19: A time series tracking the appearance and disappearance of feeding trails was conducted over a 10 week period in spring 2015. At Rodds Bay North, new trails (less than two weeks old) were identified on 28th September and only a few remained visible four weeks later on the 28th October. Tidal inundation prevented assessment of these trails two months after initial detection.

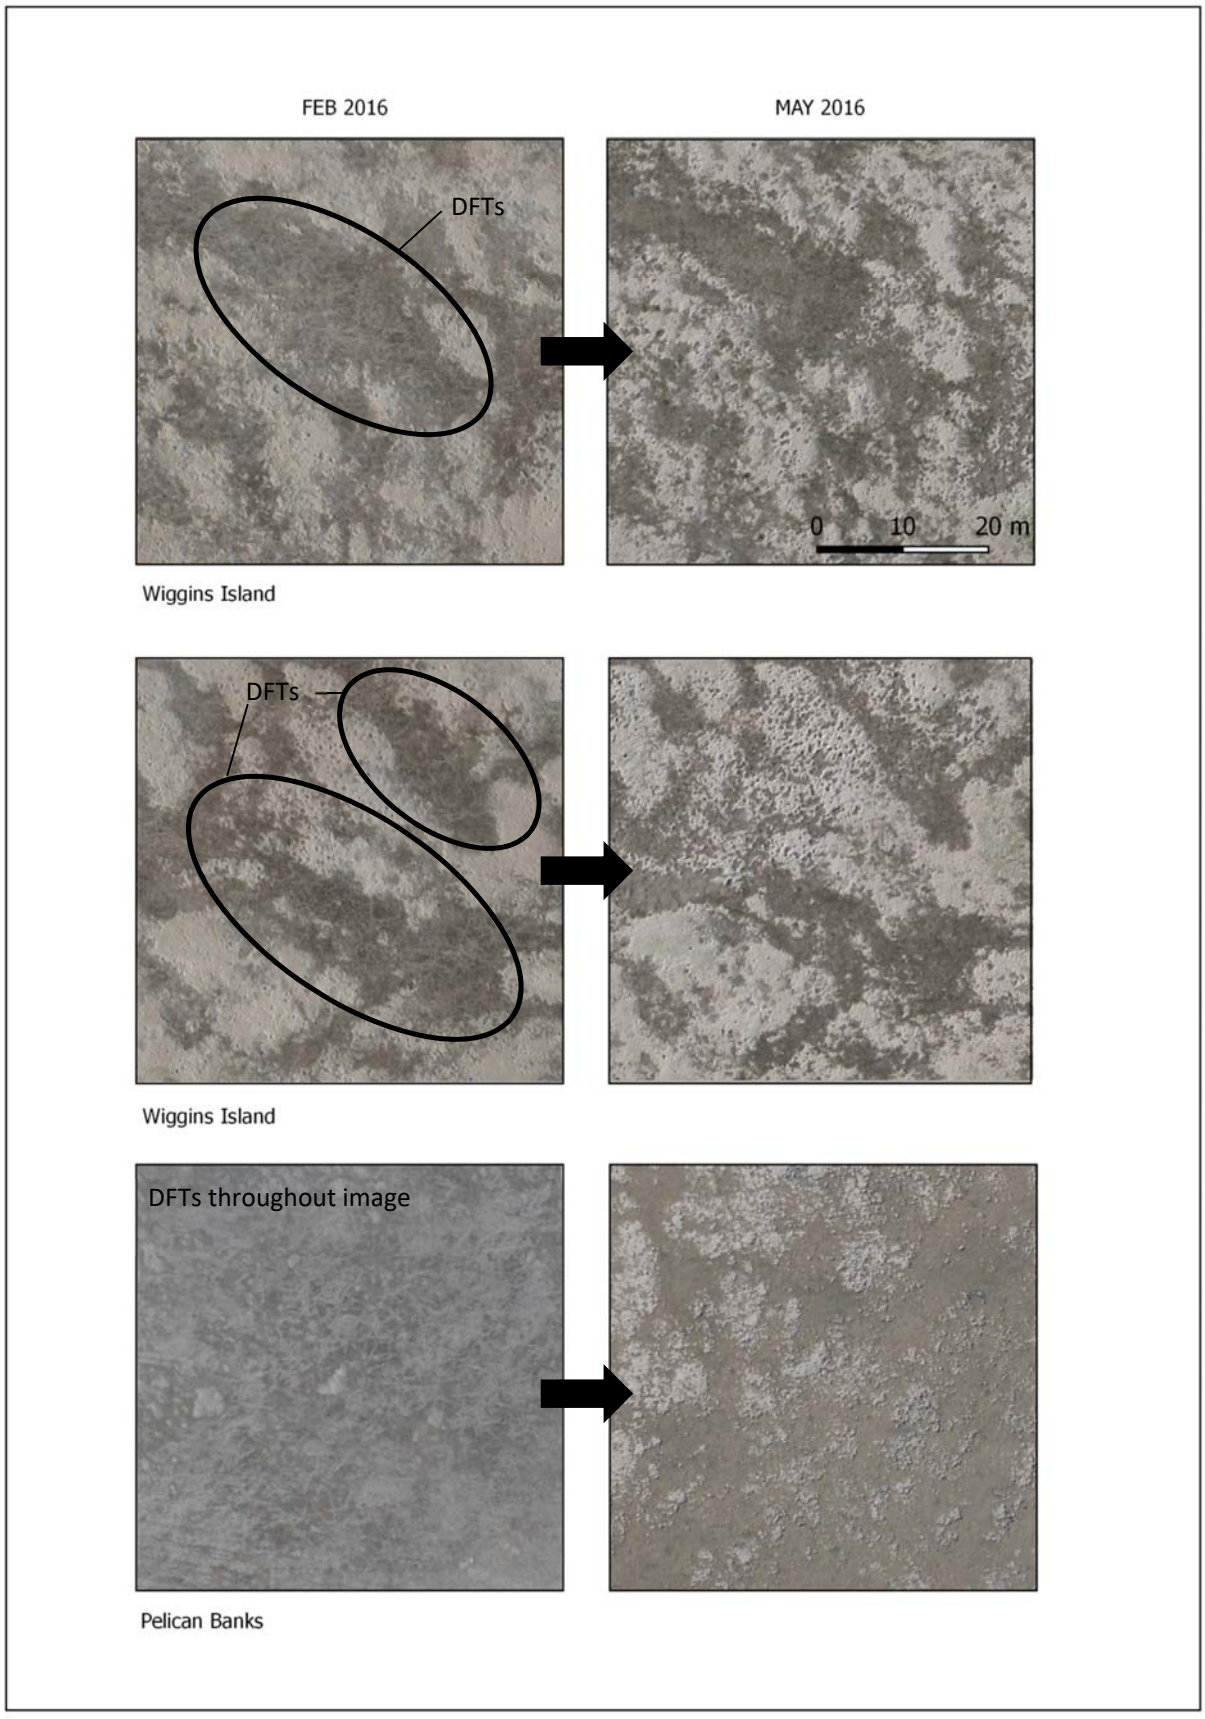

Figure 20: Examples of longevity of feeding trails between senescent season sampling events February 2016 to May 2016. No trails observed in February were still detectable in May 2016.

3.3 Visual Census of Dugong Feeding Trails

Examination of visual census data showed a significant interaction among the five locations and the six sampling periods analysed (Figure 21; Table 2, Model 1). This result highlights that while dugong feeding occurred across all locations and at most times there were no clear temporal or spatial patterns. Pelican Banks, South Trees and Wiggins Island in general had a greater probability of dugong feeding trails occurring in available seagrass compared to the two locations in Rodds Bay (Figure 21). These three locations all showed peaks in feeding in November 2015, however in Pelican Banks and Wiggins Island this was followed by a reduction to very little feeding in May 2016 and then an increase through August and December, whereas South Trees declined in February 2016 and then increased through May and August 2016 before declining sharply again in November 2016 (Figure 21, Wiggins Island: Figure 26, South Trees: Figure 27, Pelican Banks: Figure 28). Rodds Bay North and South meadows both had moderate probabilities of DFTs in August and November 2015, which declined to very low probabilities in February 2016 and then increased through May and August 2016 before declining again in November 2016 (Figure 21, Rodds Bay South: Figure 24, Rodds Bay North: Figure 25) .

Table 2: Summary of binomial generalised linear models (GLM) examining the effects of (1) sampling period and location on the probability of DFT presence; and (2) proportion of grids with seagrass present and location on the probability of DFT presence. Significant values are in bold.

Model	Source	*Df*	Deviance	Resid. *df*	Resid. Dev.	*P*(>Chi)
(1)	Null			39095	26048	
	Sampling period	4	625.27	39091	25422	**<0.001**
	Location	5	1600.96	39086	23821	**<0.001**
	Sampling period * Location	20	1152.35	39066	22669	**<0.001**
(2)	Null			42210	27422	
	Location	4	509.27	42206	26913	**<0.001**
	Proportion seagrass	1	1246.77	42205	25666	**<0.001**
	Location * Proportion seagrass	4	260.66	42201	25405	**<0.001**

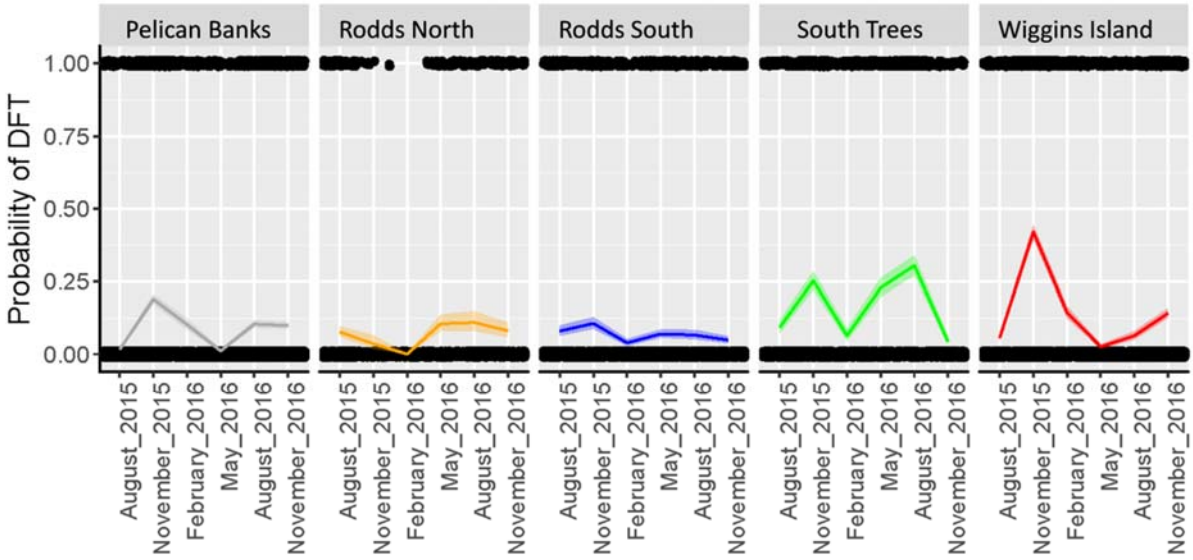

Figure 21: Fitted values from binomial generalised linear model of probability of DFT presence among sampling periods (August 2015 – November 2016) and seagrass meadows (Pelican Banks, Rodds Bay North, Rodds Bay South, South Trees and Wiggins Island). Shaded areas are ±95% confidence intervals.

There was a significant increase in the probability of DFT presence as the proportion of grids with seagrass present increased (Table 2; Figure 22). This relationship varied among locations, with both Wiggins Island and South Trees showing steep exponential increases in DFTs with increases in the proportion of seagrass present. Pelican Banks showed a similar pattern, however the relationship was not as strong as the other two locations and there were much weaker relationships in the Rodd's North and Rodd's South meadows with only a small linear increase in DFTs with increased seagrass presence at the site (Figure 21).

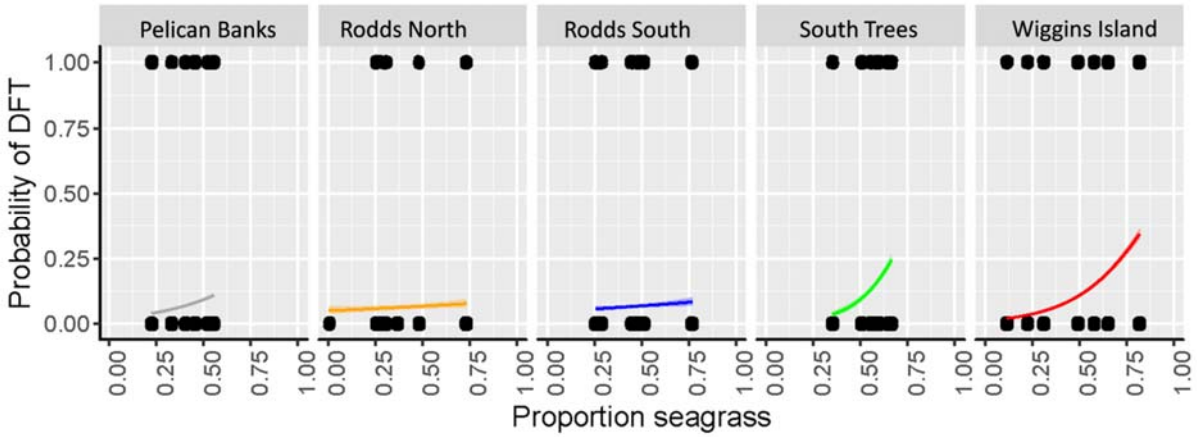

Figure 22: Fitted values from binomial model of the relationship between probability of DFT presence and proportion of grids with seagrass present for a given meadow (Pelican Banks, Rodds Bay North, Rodds Bay South, South Trees and Wiggins Island). Shaded areas are ±95% confidence intervals.

The probability of DFT presence significantly increased both with biomass and percent cover of seagrass recorded in seagrass permanent monitoring sites within the meadows (Figure 22; Table 3). These increases differed between locations and showed very similar patterns to the proportion of seagrass at the two sites from the visual census of orthomosaics discussed above. For both of the locations where seagrass transect data coincided with the DFT assessments, there was an increase in the probability of DFTs with an increase in seagrass biomass and percent cover at the site, however these increases were much steeper and showed more of an exponential than a linear pattern at Wiggins Island compared with Pelican Banks. As biomass and percent cover of seagrass are highly correlated, similar relationships existed between the two, as expected, and we have shown the graphical relationship for biomass only in Figure 23.

Table 3: Summary of binomial generalised linear models (GLM) examining the effects of (1) biomass and meadow on the probability of DFT presence; and (2) percent cover and meadow on the probability of DFT presence. Significant values are in **bold**.

Model	Source	*Df*	Deviance	Resid. *df*	Resid. Dev.	*P*(>Chi)
(1)	Null			27161	17895	
	Biomass	1	105.84	27160	17789	**<0.001**
	Location	1	519.59	27159	17270	**<0.001**
	Biomass * Location	1	733.76	27158	16536	**<0.001**
(2)	Null			27161	17895	
	Percent cover	1	0.21	27160	17895	<0.65
	Location	1	415.10	27159	17480	**<0.001**
	Percent cover * Location	1	53.89	27158	17426	**<0.001**

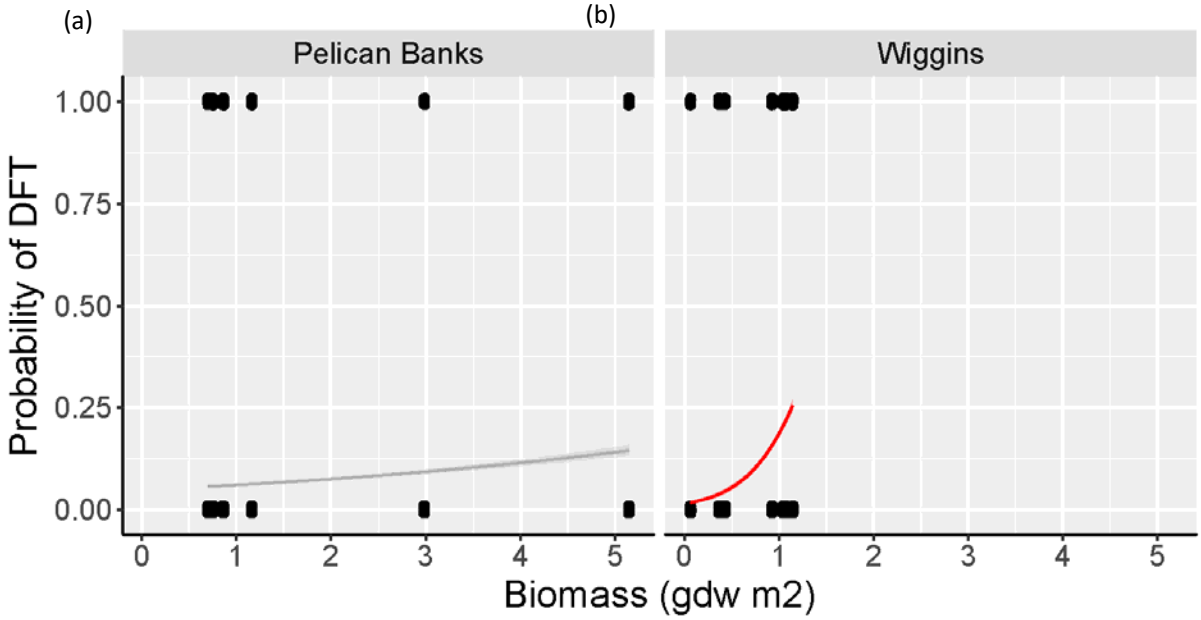

Figure 23: Fitted values from binomial generalised linear model of probability of DFT presence with above ground biomass measured at nearby seagrass monitoring sites at Pelican Banks and Wiggins Island meadows. Shaded areas are ±95% confidence intervals.

Site and rainfall for one month and three months prior to each sampling event had a significant effect on seagrass above ground biomass (Table 4), but these models explained only 15% and 23% of the variance, respectively (based on the adjusted R square statistic). Seagrass above ground biomass was greatest at the Pelican Banks North meadow relative to other sites, and at all sites above ground biomass declined with total rainfall. As rainfall was correlated with river flow (Appendix 1), these variables are likely to be interchangeable within each model.

Table 4: Overall fit of selected best models of above ground biomass (gdw m^{-2}) for one and three month periods, including degrees of freedom (*df*), *F*-statistic and *p*-values. Covariates are site (*S*) and total rainfall (*R*).

Model	Model terms	Source	*df*	*F*	*p*-value
1 month	Parametric terms	*rainfall*	1	7.739	< 0. 01
		site	4	5.432	< 0. 01
3 months	Parametric terms	*rainfall*	1	23.963	< 0. 001
		site	4	5.553	< 0. 001

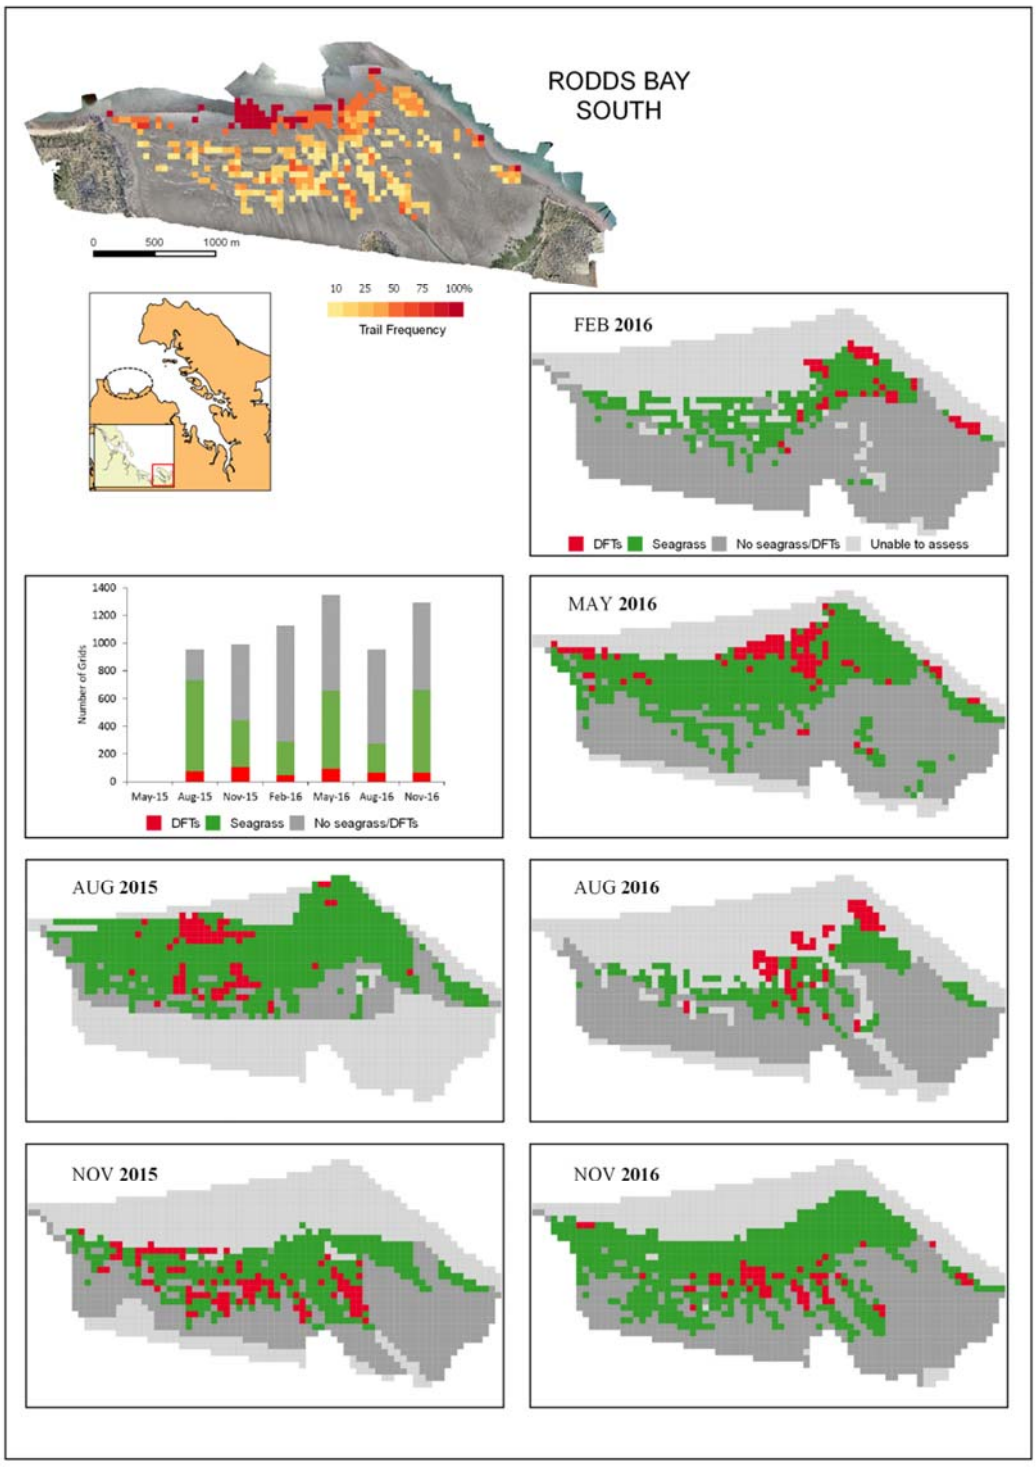

Figure 24**:** The heatmap (top of page) shows the frequency that DFTs were present in each grid over all sampling periods that assessment was possible for that grid at Rodds Bay South. The column graph in the top left panel shows the total number of grids that were able to be assessed at each sampling period. These are divided into those that contained no seagrass (grey), those that contained seagrass (green) and those that contained seagrass with DFTs (red). Spatial configuration of these assessments at each sampling event is shown on the maps in each of the remaining panels with colours representing the grids that were able to be assessed but contained no seagrass (dark grey), grids assessed that contained seagrass (green) and grids assessed that contained seagrass with DFTs (red). The remaining grids could not be assessed due to poor image quality from the presence of water or distortion from photogrammetry processing (light grey).

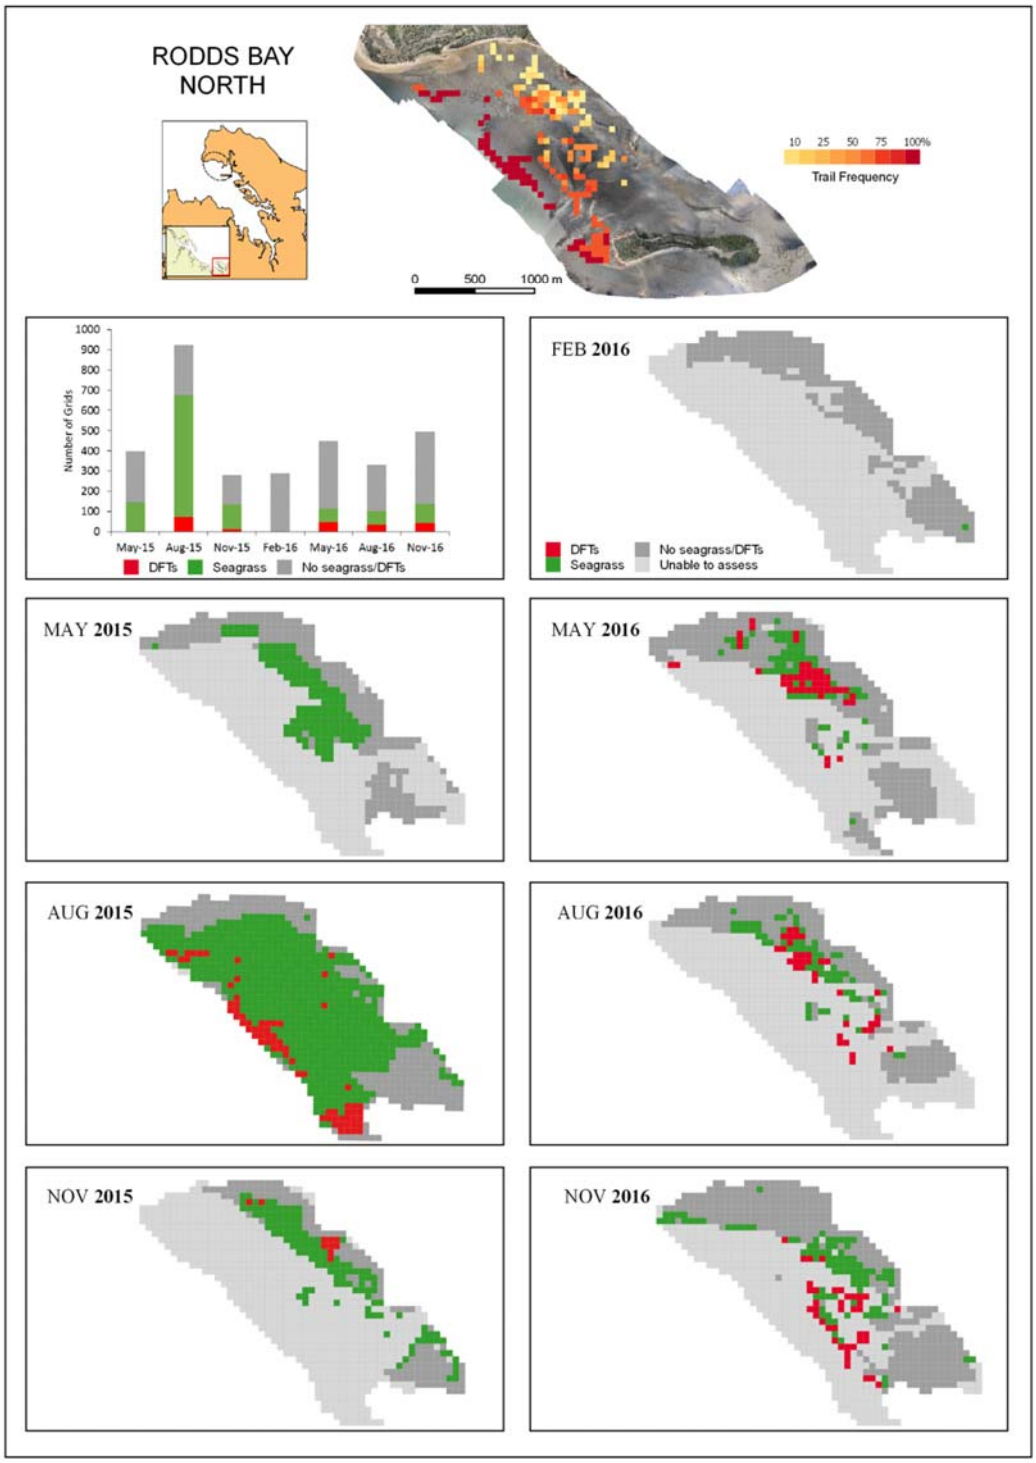

Figure 25**:** The heatmap (top of page) shows the frequency that DFTs were present in each grid over all sampling periods that assessment was possible for that grid at Rodds Bay North. The column graph in the top left panel shows the total number of grids that were able to be assessed at each sampling period. These are divided into those that contained no seagrass (grey), those that contained seagrass (green) and those that contained seagrass with DFTs (red). Spatial configuration of these assessments at each sampling event is shown on the maps in each of the remaining panels with colours representing the grids that were able to be assessed but contained no seagrass (dark grey), grids assessed that contained seagrass (green) and grids assessed that contained seagrass with DFTs (red). The remaining grids could not be assessed due to poor image quality from the presence of water or distortion from photogrammetry processing (light grey).

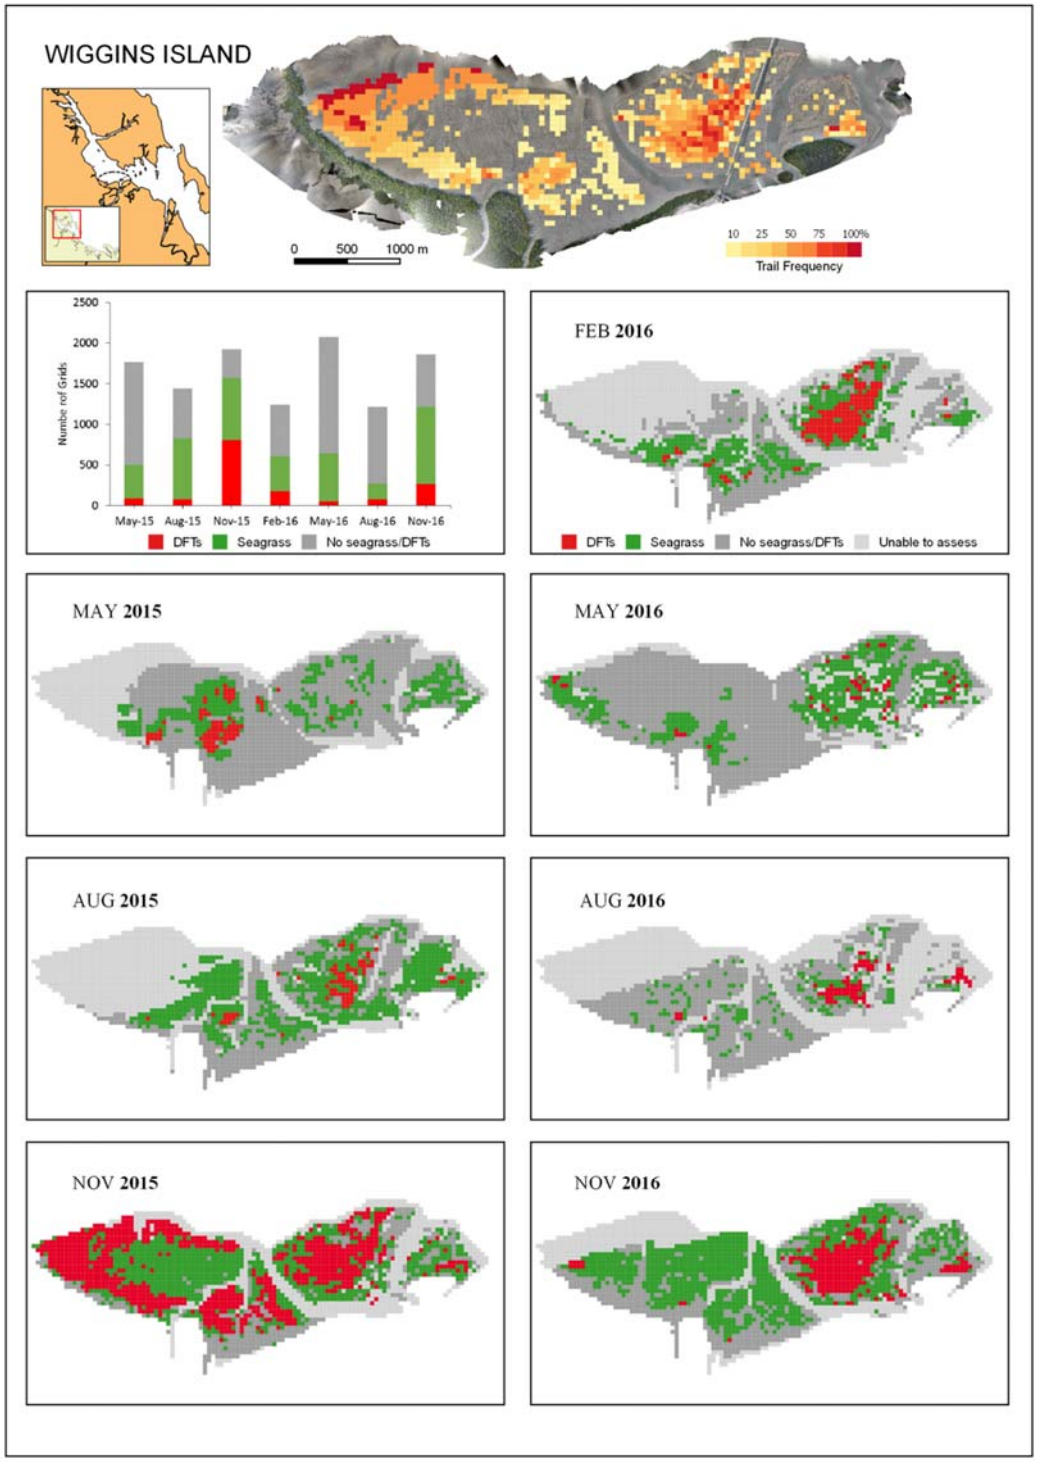

Figure 26**:** The heatmap (top of page) shows the frequency that DFTs were present in each grid over all sampling periods that assessment was possible for that grid at Wiggins Island. The column graph in the top left panel shows the total number of grids that were able to be assessed at each sampling period. These are divided into those that contained no seagrass (grey), those that contained seagrass (green) and those that contained seagrass with DFTs (red). Spatial configuration of these assessments at each sampling event is shown on the maps in each of the remaining panels with colours representing the grids that were able to be assessed but contained no seagrass (dark grey), grids assessed that contained seagrass (green) and grids assessed that contained seagrass with DFTs (red). The remaining grids could not be assessed due to poor image quality from the presence of water or distortion from photogrammetry processing (light grey).

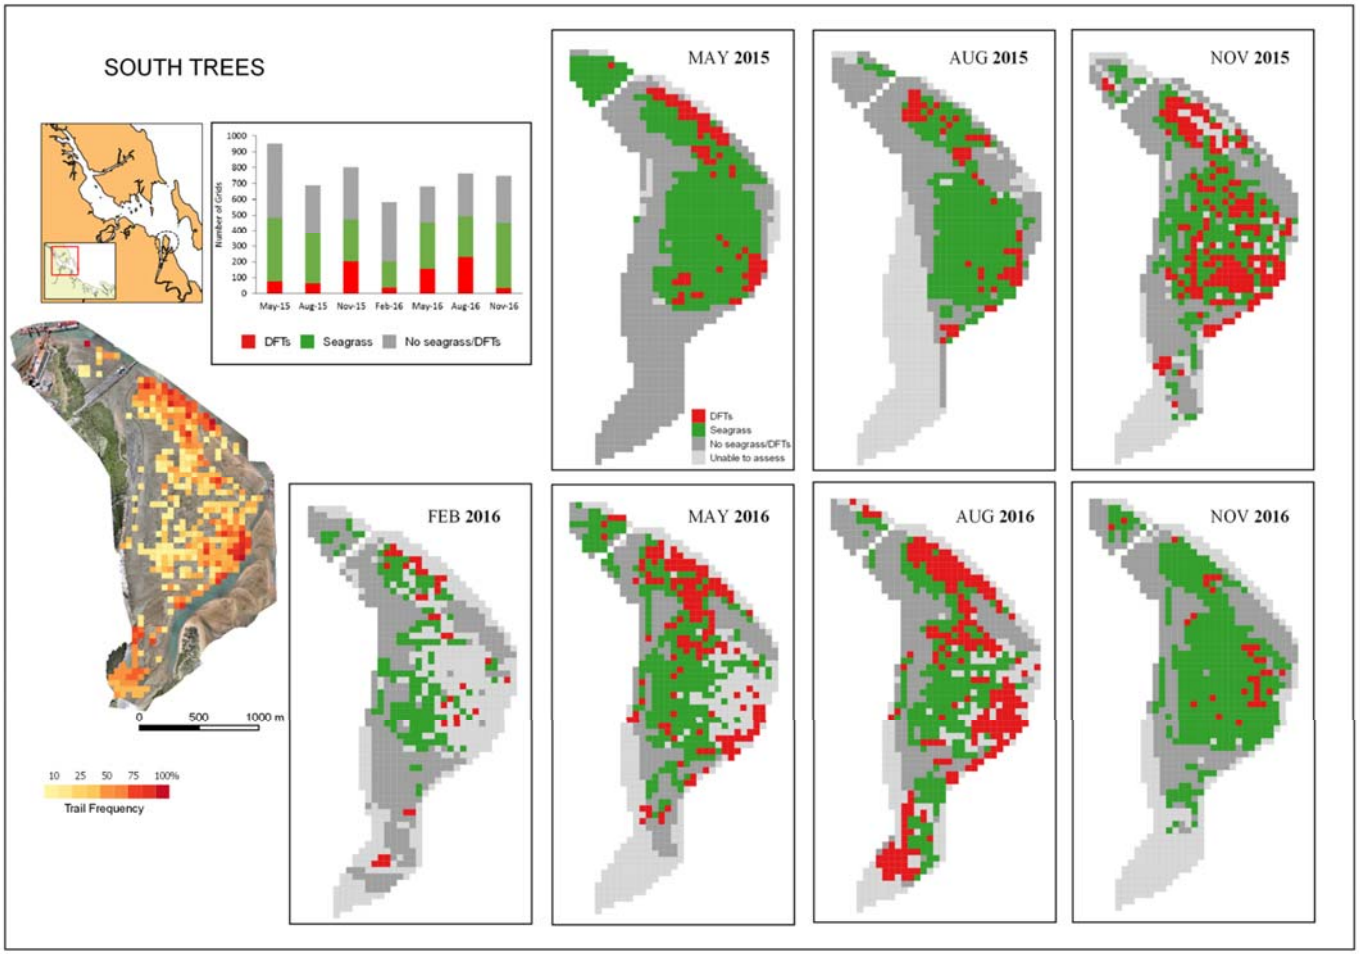

Figure 27**:** The heatmap (top of page) shows the frequency that DFTs were present in each grid over all sampling periods that assessment was possible for that grid at South Trees. The column graph in the top left panel shows the total number of grids that were able to be assessed at each sampling period. These are divided into those that contained no seagrass (grey), those that contained seagrass (green) and those that contained seagrass with DFTs (red). Spatial configuration of these assessments at each sampling event is shown on the maps in each of the remaining panels with colours representing the grids that were able to be assessed but contained no seagrass (dark grey), grids assessed that contained seagrass (green) and grids assessed that contained seagrass with DFTs (red). The remaining grids could not be assessed due to poor image quality from the presence of water or distortion from photogrammetry processing (light grey).

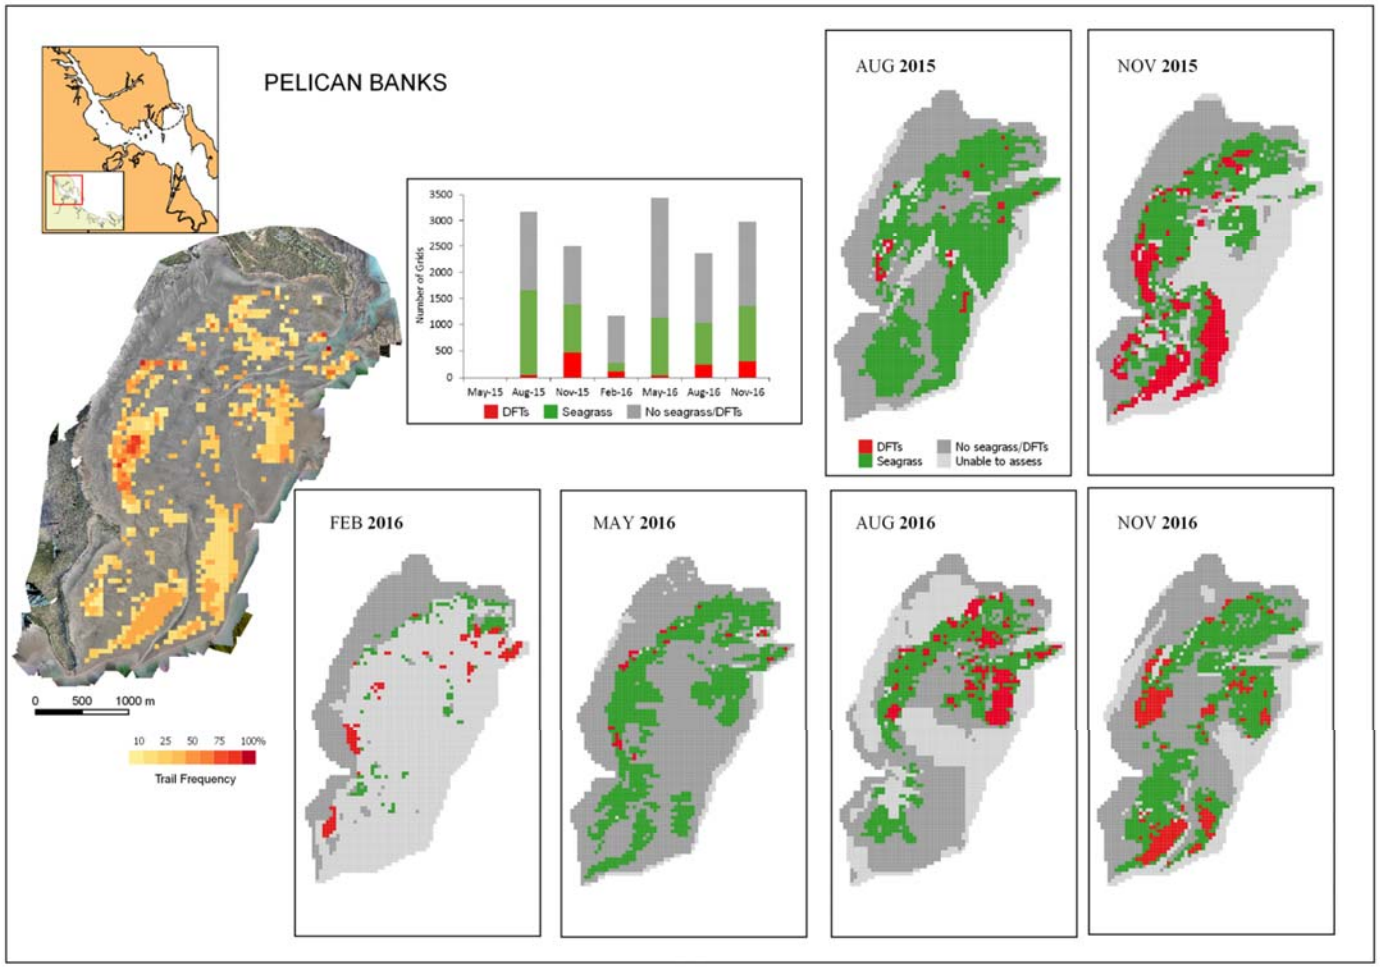

Figure 28**:** The heatmap (top of page) shows the frequency that DFTs were present in each grid over all sampling periods that assessment was possible for that grid at Pelican Banks. The column graph in the top left panel shows the total number of grids that were able to be assessed at each sampling period. These are divided into those that contained no seagrass (grey), those that contained seagrass (green) and those that contained seagrass with DFTs (red). Spatial configuration of these assessments at each sampling event is shown on the maps in each of the remaining panels with colours representing the grids that were able to be assessed but contained no seagrass (dark grey), grids assessed that contained seagrass (green) and grids assessed that contained seagrass with DFTs (red). The remaining grids could not be assessed due to poor image quality from the presence of water or distortion from photogrammetry processing (light grey).

3.4 Algorithm Extractions of Dugong Feeding Trails

An extensive range of approaches, parameters and techniques were investigated and trialled to automate the extraction of DFT features from the imagery using algorithms. A full description of these and the issues and recommended approaches can be found in Appendix 2. In this section we summarise the results of the two main themes that were investigated, training of a machine learning algorithm, and parameter range optimisation (swatch approach).

Using the first approach of a machine learning algorithm, we were able to train and classify DFTs in a small test area of the South Trees meadow to a high degree of accuracy. To accommodate potential differences in interpretation for uncertain features, and to allow adjustment at a later stage, we used the ability of machine learning predictors to output results as probabilities, rather than binary classes. An extract of the results is shown in Figure 29. As described in the methods and Appendix 2, this approach needed a much greater amount of training to provide a level of accuracy when applied to the larger meadow scale area to be acceptable. This approach was therefore deemed unsuitable due to the time constraints of the project and the likely unrealistic number of hours that would be required to effectively train the process.

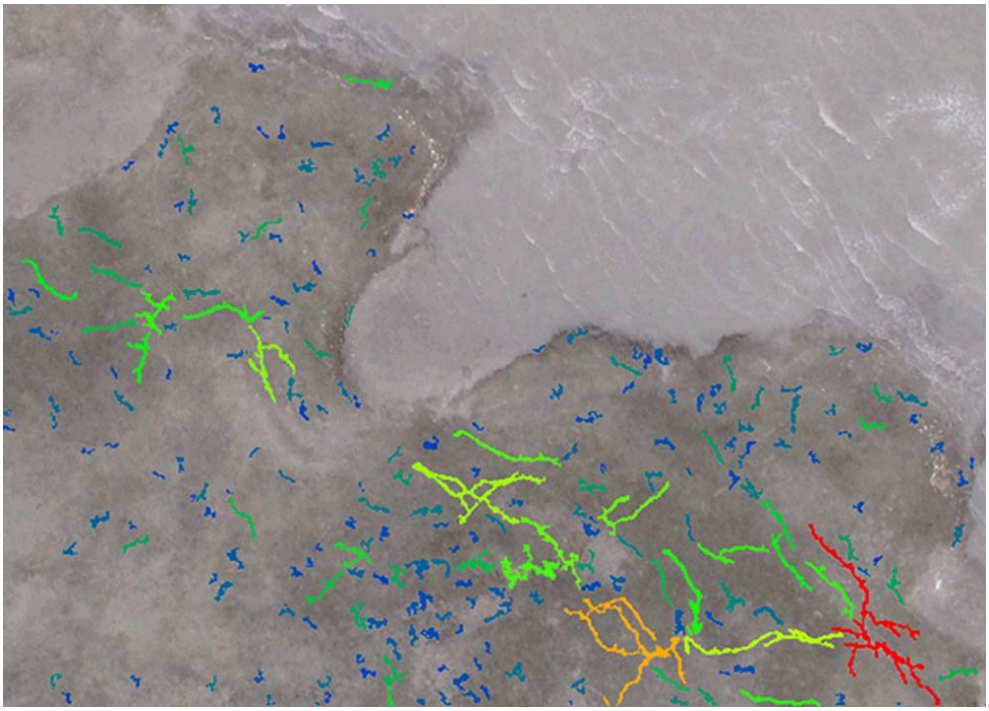

Figure 29: DFTs extracted using a training / prediction process with the Random Forest algorithm. Colours represent the probability of the feature being a DFT from blue (least likely) to red (most likely).

Following the second approach of using parameters to find a 'best fit', swatches were produced for 26, 100 m x 100 m areas across the five meadows with 120 different combinations of parameters for each one (see Appendix 2). The best match for extraction of DFT accuracy was then chosen by visual comparison with the raw image (e.g. Figure 30). In some cases there was a high degree of accuracy from the swatch method in extracting DFT features from the test area. In other cases the range of parameter values chosen for the swatches did not cover the optimal values particular to the environmental conditions of the imagery and so produced no satisfactory samples. Significant computational time was required for every re-run of swatches with a parameter range shift (in the order of 48 hours per run for a given parameter set and area). Therefore once we had established enough suitable cases where swatches produced parameter ranges that sufficiently matched the known DFTs, these were tested for their suitability across the entire meadow. Although these parameters provided a highly acceptable level of accuracy in the sample region, these results did not extrapolate well to larger areas. Results produced false negatives and false positives throughout and were not representative of the entire landscape (Figure 31). In addition, as the degree of success for the

parameters could not be seen until all areas of the location had been processed it put limitations on how many were able to be assessed.

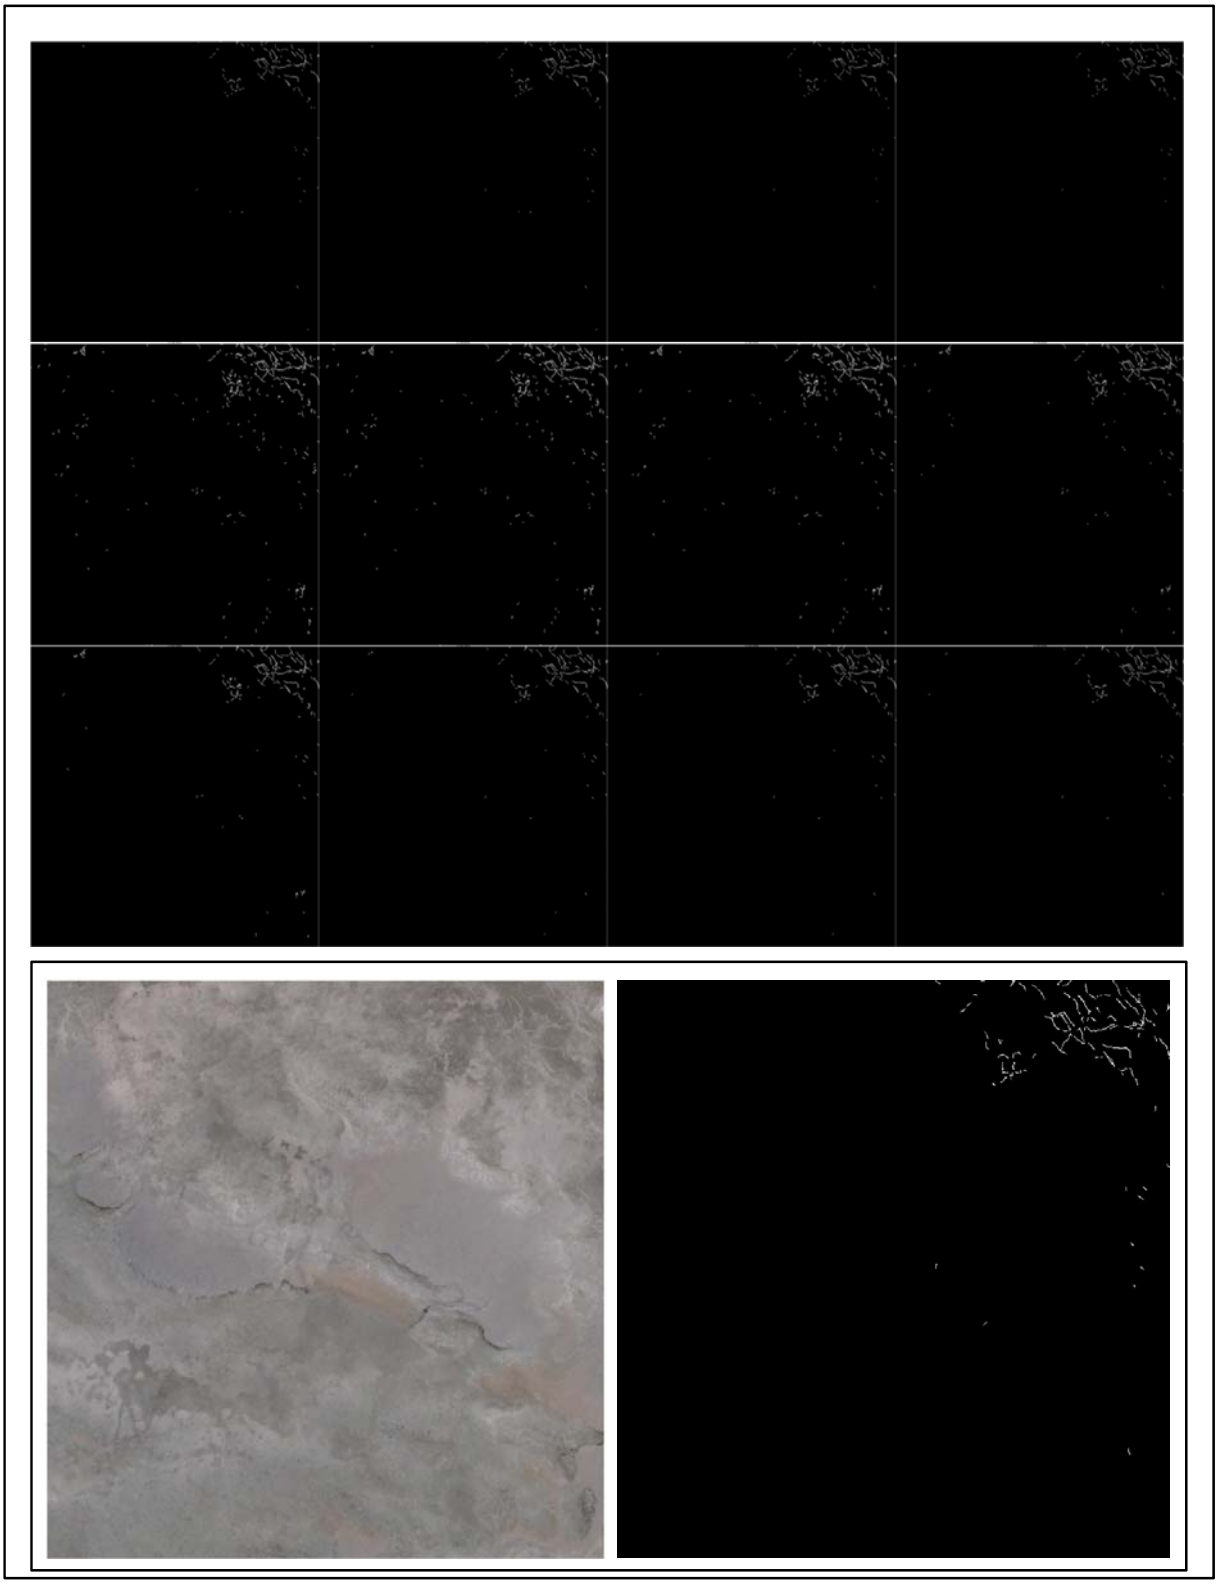

Figure 30: An example of a subset of 100 m x 100 m swatches from South Trees on May 2015 showing the extraction of DFTs using different parameter ranges. The “best fit” model for this area can be seen at the bottom of the figure.

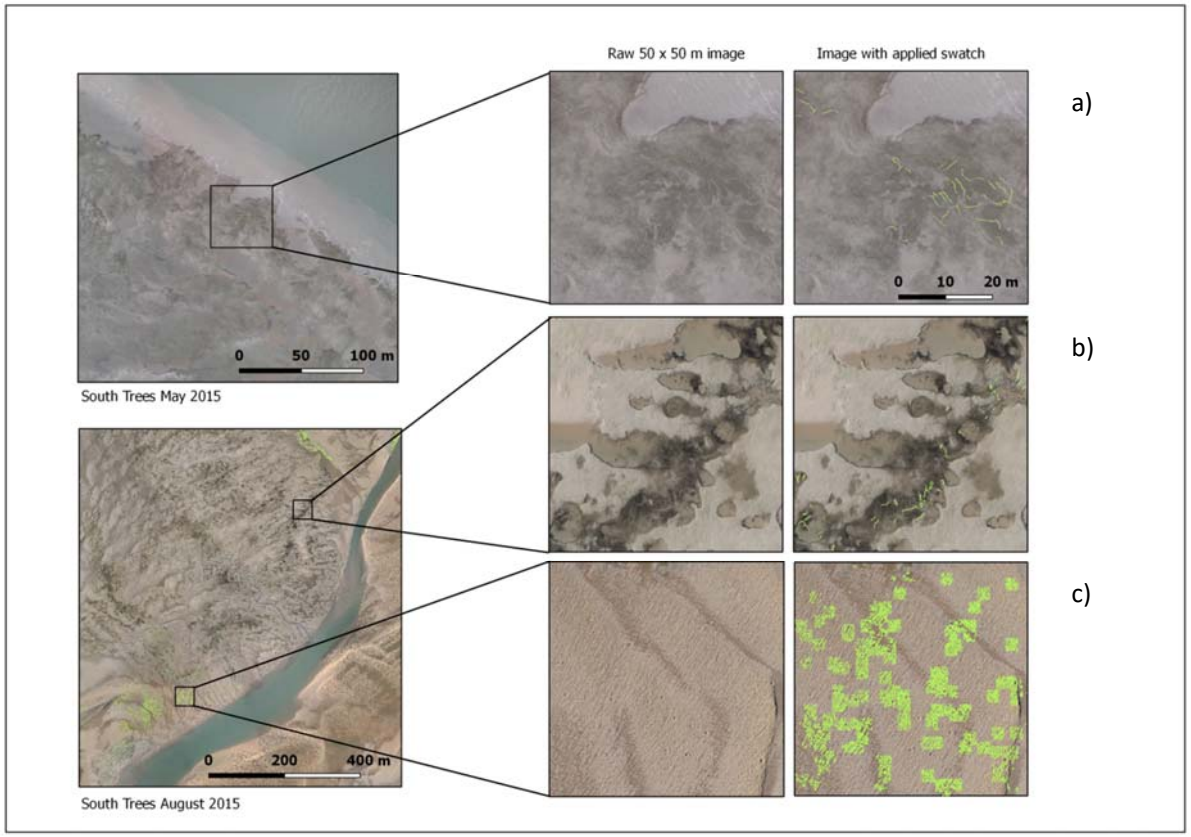

Figure 31: Results showing the extrapolation of the "best fit" parameters that were successful in 100 m x 100 m grid swatches to the entire meadow landscape. Examples of both (a) false negatives (b) good matches and (c) false positives for DFTs can be seen.

4 DISCUSSION

This study presents a positive demonstration of a new technique for sampling dugong feeding activity and seagrass change at a large spatial scale. Using aerial photography and new photogrammetry techniques we were able to successfully capture and compile orthomosaic images of entire intertidal seagrass meadows that provided the level of detail required to distinguish linear DFTs from other features such as propeller scars from boats, and similar but larger trails made by other biota such as stingrays. This is the first time that dugong feeding activity has been able to be assessed at this scale and frequency and the approach has great promise for future application, both in Port Curtis, and other intertidal areas where dugongs feed largely by excavation of seagrasses.

The techniques can be performed at relatively low cost as very little in the way of specialised equipment is required in order to construct suitable orthomosaics for DFT assessments. The most expensive outlay is the structure from motion software required, in our case Pix4Dmapper Pro. The camera can be any camera capable of being remotely fired with a high enough resolution. While in this project we chose to use helicopters to overfly the areas to be photographed, other platforms such as fixed wing and multi-rotor drones could easily be used. For us, using helicopters enabled the large areas of investigation to be covered quickly within the narrow tidal window we had available. In other applications, particularly projects investigating smaller areas, relatively in-expensive off the shelf drones could be utilised.

Using these techniques quarterly surveys of five representative intertidal seagrass meadows in Port Curtis and Rodds Bay were successfully carried out to determine the feeding patterns and use of these habitats by the dugong population within this region. A visual census was conducted of each orthomosaic image overlaid with a sampling grid to detect the presence and distribution of both seagrass and DFTs.

Our results show that while there were no strong seasonal patterns across all sites, all of the meadows were used by dugongs to varying extents throughout the year. This indicated that all of the intertidal meadows were of value for dugong feeding. Results also highlighted that there were large variations in dugong feeding activity between times and meadows and although there were links to feeding activity and changes in seagrass extent and coverage, a large part of the variation remains unexplained, suggesting there is a lot we still don't understand about what drives dugong foraging choices.

Different locations had peaks and declines in the amount of DFTs at different times of the year. Higher levels of DFTs were generally observed in Port Curtis meadows (Pelican Banks, Wiggins Island, and South Trees) throughout the study and particularly at the culmination of the seagrass growing season (November), compared to the two Rodds Bay meadows. This may be due to larger areas of subtidal seagrass meadows available in Rodds Bay compared to Port Curtis making dugongs less reliant on the intertidal meadows we assessed. On many of the sampling surveys in Rodds Bay dugongs were observed at the deeper areas of the meadows outside of the survey limits which were constrained to exposed areas of the meadows. Evidence of this was borne out in the analysis from August 2015 in Rodds Bay North and August 2015 and May 2016 at Rodds Bay South when lower tides and greater exposure allowed for surveys to extend over larger areas of the meadow. During these periods, a high presence of DFTs occurred on the deeper areas of the meadow where it was not possible to assess trails at other sampling times when tides were less favourable.

The tidal cycle can play an important role in structuring dugong feeding. A study of dugong movements across tidal cycles in Hervey Bay (Sheppard et al. 2010) suggested that dugongs use high tide to feed on intertidal seagrass species with high energy content (e.g., starch - in this case *H. uninervis*). At low tide, when high energy seagrasses are not accessible, dugongs may feed on areas of high biomass for foraging efficiency (Sheppard et al. 2010). As all the meadows monitored in Port Curtis and Rodds Bay were dominated by *Z. muelleri* subsp. *capricorni* and there was sufficient biomass of subtidal seagrass at Rodds Bay it is likely there was little need for dugong to venture into shallower waters at high tide to forage there. Such preferences for feeding in deeper seagrass areas when available may be linked to reduced risk of predation, with studies suggesting that feeding in deeper areas provides better escape routes from sharks (Wirsing et al. 2007, Wirsing et al. 2008).

In general, there was a significant trend of more DFTs as the presence and amount of seagrass increased. This trend was consistent whether analysing the relationship between DFTs with the probability of seagrass presence in sampling grid cells from the aerial photo mosaics, or the relationship with seagrass above ground biomass and percentage cover data collected within the meadows during on ground seagrass monitoring surveys at Pelican Banks and Wiggins Island. This result concurs with a recent study of dugong feeding preference in intertidal seagrasses to the north, around the Cairns and Townsville region. In that study, Tol et al. (2016) found that seagrass biomass was the primary factor that determined the presence of dugong feeding and that species composition and nutritional value (nitrogen content of the plant) influenced feeding to a lesser degree. These findings contradict some previous results from subtidal meadows in south-east Queensland. Preen (1995) found feeding preferences in Moreton Bay were thought to be based on the nutritional value of the plant material and pioneer species like *Halophila ovalis* were targeted because of this. Another study in Hervey Bay found *Halodule uninervis* to be targeted by dugongs, again based on nutritional value; however this was also the species with the highest biomass in the study area (Sheppard et al. 2007). In our study, all of the intertidal meadows were dominated by *Zostera muelleri* subsp. *capricorni* with very small percentages of *H. ovalis* and *H. uninervis.* Given the relative uniformity in species assemblages and the inability to distinguish between seagrass species from the aerial photographs, no assessment could be made of dugong feeding preference in regards to seagrass species or nutritional value in our study.

The relationship between the increase in seagrass presence and above ground biomass and the probability of dugong feeding varied among locations. In meadows with sparser seagrass abundance and more variable seagrass distributions such as Wiggins Island and South Trees, DFTs increased markedly and in an exponential relationship with seagrass. In contrast, more stable meadows like Pelican Banks, DFTs increased in a more gradual, linear fashion. This suggests that for the sparser more ephemeral meadow types, when seagrasses increase substantially, they reach a critical point that encourages dugong to spend significantly more time and effort grazing there. Such a response makes sense as at these times there is likely a substantially increased reward for grazing effort and previous studies suggest that dugong grazing tends to target seagrass areas with continuous cover (D'Souza et al. 2013). In meadows where seagrass coverage is more stable and consistent there is not such a noticeable change in reward for effort.

Seasonal change in seagrass biomass and distribution is characterised by an increase during the growing season (July-December) and decrease during the senescent season (January-June) in Port Curtis (Chartrand et al. 2016). These changes are primarily driven by available light, which is higher in the dry (growing) season due to less run-off from precipitation and river flow which increases turbidity in the water column. This was highlighted in the analyses of relationships between seagrass above ground biomass and environmental and climatic variables at each location. Rainfall (and strongly linked river flow) in the months prior to each survey was the best predictor of seagrass abundance with a significantly negative relationship between precipitation and seagrass above ground biomass. However, the relationship between seagrass and environment is complex and rainfall only explained a small portion (15-25 %) of this variability on its own.

Satellite tagging of dugongs in the Port Curtis area has been conducted on an opportunistic basis as part of the Gladstone Ecosystem Research and Monitoring Plan (ERMP) (Cleguer et al. 2015a, 2015b). As numbers of tagged animals were low (two in 2014 and one in 2015 - coinciding with timing of our study), it is difficult to draw much in the way of correlation between those studies and ours. The two dugong that were tagged in 2014 did spend their time in the inner harbour (Cleguer 2015a), and collectively in the meadows where we recorded most of the DFT activity (Cleguer et al. 2015a). One of the tagged animals spent its time moving between the main meadows we assessed dugong feeding at in the inner harbour, such as Pelican Banks, South Trees and Wiggins Island, supporting our findings that all of these intertidal meadows are important for dugong feeding. The dugong tagged in 2015 moved out of the Gladstone area, north to Shoalwater Bay (Cleguer et al. 2015b).

Our assessment of DFT longevity found that DFTs recorded in each quarterly survey are likely to be unique to those surveys, with trail longevity measured in our study much shorter than 3 months. These findings are largely in line with previous research examining recovery of seagrasses into feeding trails. Recovery of

simulated DFTs took three months in meadows of *Zostera/ Cymodocea* and *Halophila ovalis* in experiments conducted to the north near Cairns (Aragones and Marsh 2000). Recovery of *H. ovalis* in intertidal dugong feeding trails were much faster (less than one month) in Thailand (Nakaoka and Aioi 1999) and Moreton Bay (McMahon 2005). While these studies report full recovery of DFTs to levels of surrounding un-grazed seagrass, the recovery of trails to the point where they would no longer be recognisable as distinct DFTs from our more remote aerial surveys is likely to be much less than this.

There is likely to be some temporal/seasonal variation to DFT longevity, as in a large part, visible trail longevity is determined by the rate of recolonisation of seagrasses into the excavated region. Seagrass recolonisation and growth, whether through asexual colonisation from the surrounding meadow or from seeds, is known to be highly seasonal (Rasheed 1999, 2004, Rasheed et al. 2014). Recovery rates were slower in the slow growth season than in spring and summer high growth season for simulated DFT experiments in Moreton Bay (Preen 1995). Our detailed assessments of longevity were performed during the high season for seagrass growth in this region (August-November) when seagrass recovery would be expected to be the most rapid. Residence times of DFTs during the senescent season, when seagrass asexual colonisation is slower, are likely to be longer. However, while we didn't examine DFTs at intervals more frequently than quarterly in the senescent season, it was clear from examination of areas of DFTs between February and May 2016 (low growth season) that individual trails did not endure long enough to be identified in subsequent surveys.

In the circumstance where seagrass around a DFT disappears after the trail was created (due to seasonal declines in overall seagrass distribution for example), the residence time for the DFTs remaining visible in the imagery may be substantially longer. This is as expected as trail longevity is determined by seagrass re-growth. In these cases trails may remain in place between the quarterly sampling events as the excavated furrows remain distinct for longer without seagrass re-growing into them. For this reason we excluded such trails from our analysis of new dugong feeding activity.

Although DFT trail lengths were not systematically measured as part of our assessment, individual trails were detected from 30 cm to over 6m in length which is in the range found in literature from other locations (e.g., from 1 m to 5 m in Heinsohn et al. 1977, and varying from a mean of ~3 m to 8 m in Shoalwatwer Bay, Anderson and Birtles 1978).

Automated feature extraction

To generate a more automated approach to DFT quantification from imagery, we focused on a sequence of operations which were likely to be achievable, and yield good results within the time and resource constraints of the project (see Appendix 2). The approaches we used over any given confined region extracted DFTs well, once parameters were set accordingly. On a small scale (10's of metres) it seems the process of DFT extraction is relatively straight forward. Similar positive results have recently been achieved for small scale (20 m x 20 m) sub-tidal DFT assessments in clear water in the Philippines using similar techniques to our study (Mizuno et al. 2017). Successful extrapolation of small scale approaches to larger meadow scale extraction (10's of kilometres) however, relies on the parameters used being relatively consistent across the broader region. This was not the case for the Port Curtis and Rodds Bay meadows, and any one smaller area where parameters are set to adequately extract features, represented an over-fit that is not transferable to the broader data coverage. Such an extrapolation may be possible for other locations where DFTs and the background substrata on which they occur are more uniform in time and space than was the case for the Port Curtis and Rodds Bay meadows. However, based on our assessments, the development of useful algorithms for automated or semi-automated DFT feature extractions, over meaningful spatial scales is likely to require a more complex approach.

The work we have done as part of this project has gathered and tested the tools likely required to produce such a successful DFT extraction method. We now understand the variability of the trails and of the substratum on which they appear. Having once decided to avoid the machine-learning option for fear of insufficient resources to train the algorithm, we limited ourselves to a handful of parameters by which to

adjust the model. We now know that this did not afford the complexity required to account for the enormous variability in background and DFT features within a single mosaic and between sites and times. With hindsight, persevering with the machine learning technique may have provided the best way to manage the large number of variables on which the model needs to be based.

From the work we have done to date, we have the means to segment the DFTs to reach the training stage of a machine learning approach. Using broad ranges of parameter values, we can apply the current model to extract all possible features inclusively into vector format. These are then ready for operators to sub-sample for the identification of DFTs and non-DFTs. We also have the expertise within the research team to work with machine learning algorithms. In addition, computational resources, which were a significant consideration throughout the project, have undergone significant advances since the project started. More cloud-based distributed networks have become available on which to process massive amounts of data. Amazon's AWS[4], for example, allows the choice of multiple graphics processing units (GPU), which are much more efficient than central processing units (CPUs) in the handling of imagery processes such as photogrammetry (Christophe et al. 2011). Google's *TensorFlow*[5] offers an open source library through a single API that allows us to deploy operations to many CPUs and GPUs as required. The network is designed to allow deep learning research, which is essentially a multi-layered machine learning approach.

For the analysis presented in this report, we have had to rely on a manual classification approach to analyse the location and density of DFTs, due to the problems and limitations encountered in this research into automated extraction. While this has allowed a viable method to achieve the goals of the project, manual classification is highly labour intensive at large spatial scales, and has limits in what can be achieved in terms of finer scale analysis at the individual trail level. We believe that the significant work we have done to date in developing an algorithm extraction approach and the pathway outlined above is worth pursuing further. Once developed it would offer the potential for rapid assessments of imagery over very large spatial scales, which would provide an excellent tool for understanding and managing dugongs and their feeding behaviour over large regional scales, as well as detailed assessments at small scales.

Management implications

The mapping of DFTs in Port Curtis and Rodds Bay shows consistent and often high levels of usage of intertidal seagrass meadows in the region. The definition of DFTs used in this study is inherently conservative and only incorporates distinct linear features that can be positively identified as DFTs while missing other forms of use such as cropping and intensive grazing in non-linear patches. This invariably leads to an underestimate in the level of use in these areas. The results therefore clearly show that dugong are relying heavily on these intertidal meadows for food throughout the year. This is not surprising given a large proportion of the seagrass meadows in Port Curtis are intertidal (Davies et al. 2016). Meadows are generally dominated by *Z. muelleri* subsp. *capricorni* which has high light requirements and the high levels of turbidity in the estuary and low light levels limit seagrass growth to shallow areas. The continued health and widespread distribution of these seagrass meadows may therefore be crucial in supporting dugong populations in central Queensland. Extensive seagrass declines due to extreme rainfall and cyclone events in 2010-2011 were responsible for high mortalities (Meager and Limpus 2014) and lower fecundity (Fuentes et al. 2016) of dugong along the coast of Queensland adjacent to the southern Great Barrier Reef. This highlights the need to protect and maintain healthy seagrass meadows as one of the key components in dugong conservation programs.

This project was triggered as part of conditions associated with a major capital works program in Port Curtis that occurred from 2011-13 to expand port facilities to support development of emerging industries in the region such as Liquid Natural Gas. This project, the Western Basin Dredging and Disposal Project (WBDDP), involved dredging activities to deepen and widen existing channels and to create new channels and swing basins as well as a new reclamation area. The activities were deemed to have the potential to impact seagrass habitat and the marine megafauna that used these habitats for foraging through poor water quality,

[4] Amazon Web Services, `https://aws.amazon.com/`

[5] TensorFlow™, `https://www.tensorflow.org/`

increased vessel movement around the development and loss of foraging habitat. As such the WBDDP came under the guidelines of the EPBC Act due to the potential to effect Matters of National Environmental Significance (MNES) as the project was likely to impact on listed migratory species including dugongs. Conditions of the approval for the WBDDP required the development and implementation of a Port Curtis and Port Alma Ecosystem Research and Monitoring Program (ERMP) with an advisory panel that reported to the Commonwealth Department of the Environment. An aim of the ERMP is to acquire a detailed ecological understanding of the marine environment of Port Curtis that can be used to monitor, manage and/or improve the regional marine environment and to offset potential impacts from the WBDDP on listed threatened and migratory species. This report and project, in part, addresses Conditions 27, 31 and 33 of the WBDDP approval, particularly 33(d) by "identifying potentially suitable habitats for key megafauna in the region at an ecological scale appropriate for megafauna species".

This DFT study has only collected detailed data on the distribution of DFTs in the post dredging environment (2015-2016) and therefore we cannot use it to make quantitative comparisons of the potential impact of the development on the use of intertidal seagrass meadows by dugongs. While no comparable DFT data was available pre- and during the dredging campaign, qualitative assessments from observational data collected during seagrass surveys spanning pre, during and post dredging are available. Seagrass monitoring surveys in Port Curtis conducted by TropWATER (see Davies et al. 2016 for most recent report) noted the presence of DFTs in intertidal seagrass survey sites providing qualitative evidence of dugong usage in these areas over the last six years, including the period before the WBDDP. Those observations show that the three meadows surveyed for DFTs during the current study were also consistently used for dugong feeding prior to the WBDDP (Figure 32). This indicates that following the WBDDP, dugong were at least accessing similar meadows to those they were using prior to WBDDP. It is difficult to identify any potential impacts to dugong feeding activity that may have occurred during the dredging program as the program overlapped with a period of seagrass loss in central and northern Queensland due to extreme weather events, flooding and cyclones in 2010-2011 and many of the inner harbour sites experienced dramatic declines in seagrass cover and biomass during that period (Davies et al. 2016). When we look at observations of DFTs between 2012 and 2013 the Wiggins Island area stands out as having no feeding trail observations (Figure 32). However, while the Wiggins Island area was the closest assessed meadow to the dredging operations, the loss of almost the entire seagrass meadow linked to the extreme weather events of 2010-11 make it difficult to attribute what additional impacts the development may have had on dugong, if any. This meadow is also the closest to the Calliope River mouth and as such most immediately impacted by flooding events that occurred at the time.

For future activities, the results of this project provide a baseline and method for effective quantitative comparisons of intertidal dugong feeding to assess potential impacts with much more certainty. The power of the program to achieve this would be significantly enhanced if some ongoing assessments of a sub-set of the areas investigated in this study, were to be incorporated into a longer term monitoring strategy.

Future of the Program and Potential Design

This project has now captured seven quarterly surveys of dugong feeding on representative intertidal seagrass meadows through Port Curtis and Rodds Bay. The continuation of the project would provide a greater understanding of seasonal changes in dugong usage within the region and longer term trends that are not possible with a < 2 year data set. Extending the study would also provide a better basis for detecting any environmental or climatic drivers that may be affecting the use of these intertidal meadows by dugong by incorporating a greater range of data for each parameter to better inform the models in the analysis. We suggest that this could be achieved by continuing surveys on a quarterly basis for a further two years on smaller areas that capture the known hotspots of feeding activity from the initial two years of the study (Figure 33). This reduced survey footprint would result in substantial savings in field and analysis costs. DFTs would be analysed using the visual census of the sampling grids for continuity of method to allow comparison of results. This could be enhanced with further investigation of automated extraction techniques.

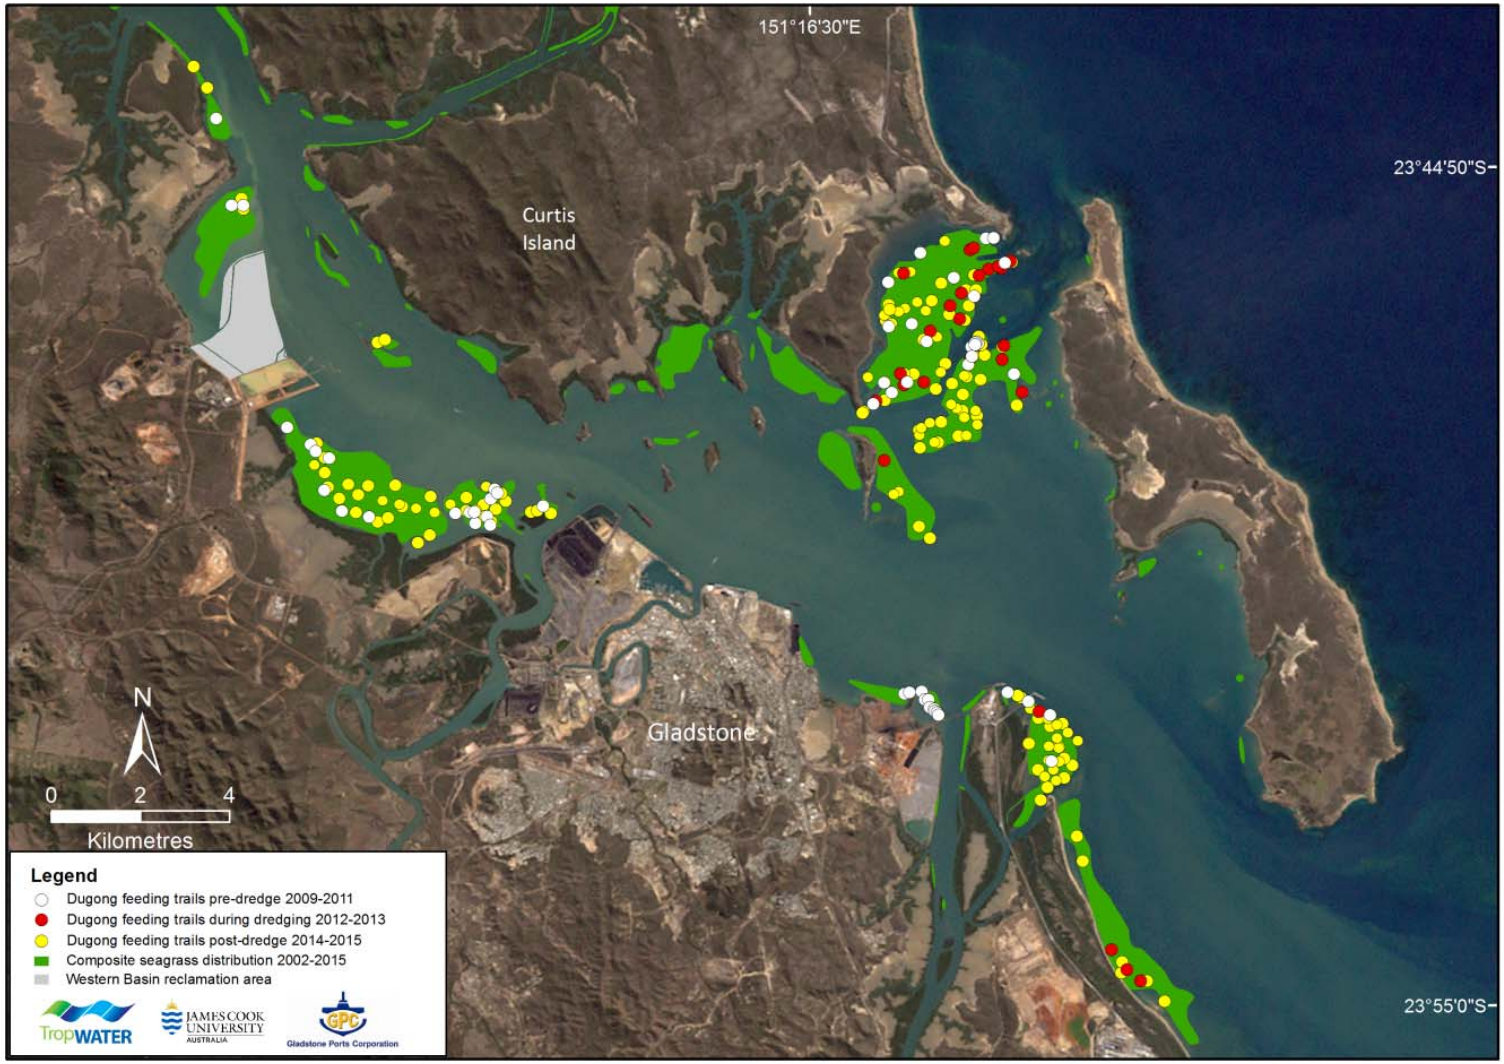

Figure 32: Presence of DFTs recorded during seagrass monitoring surveys in Port Curtis from 2009-2015 incorporating periods before (white dots), during (red dots) and after dredging (yellow dots) for the WBDDP.

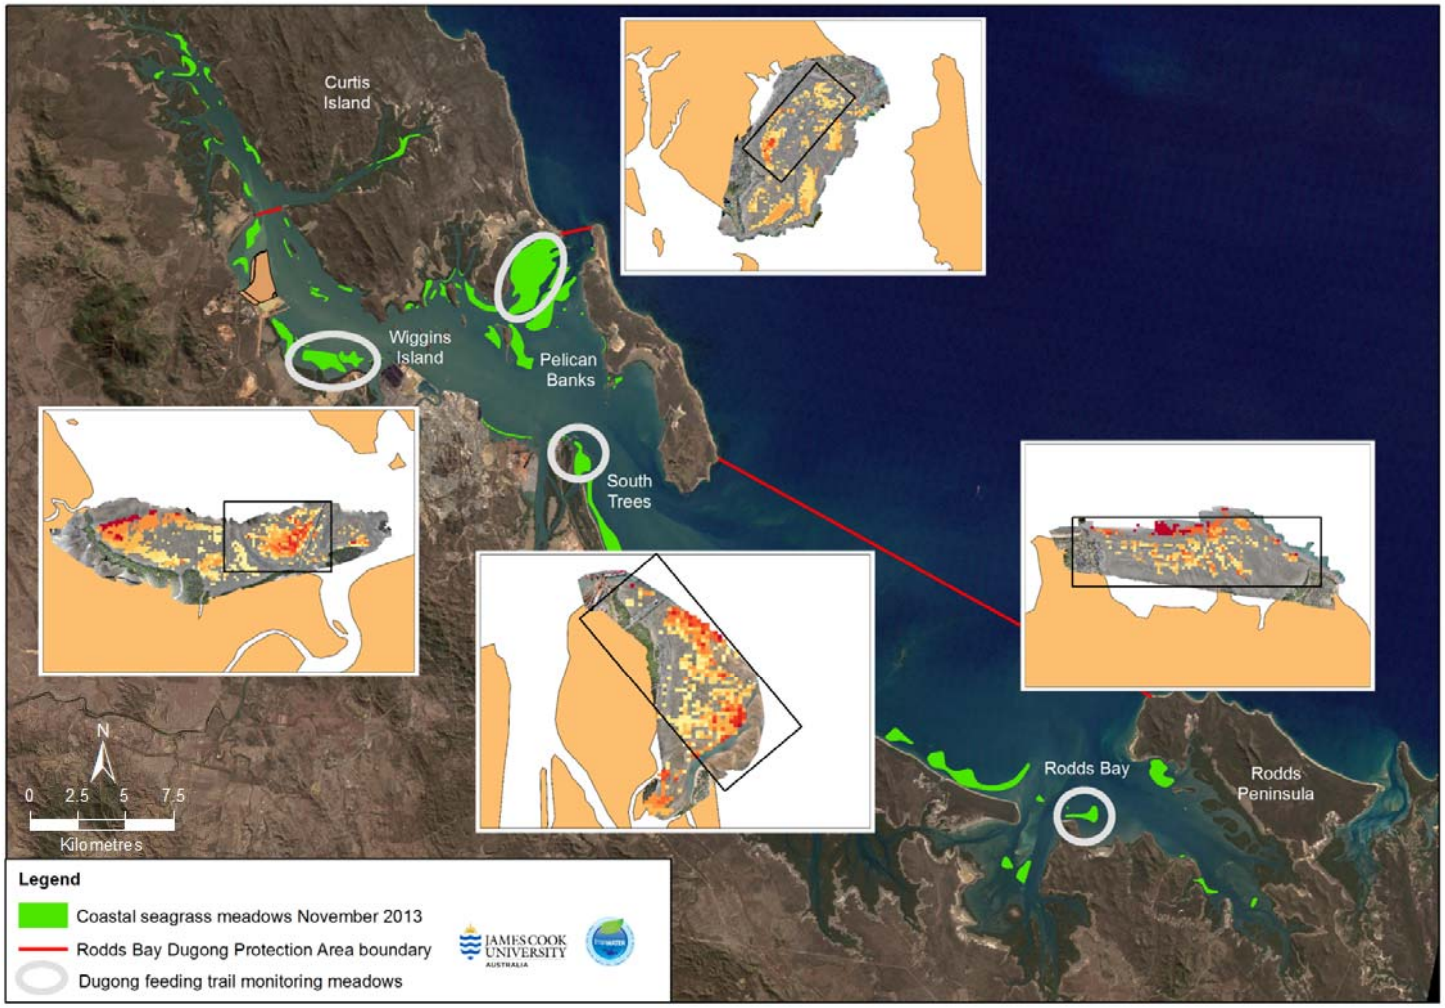

Figure 33: Suggested locations (black rectangles) for ongoing DFT monitoring in Port Curtis and Rodds Bay.

5 REFERENCES

Anderson PK (1982). Studies of dugongs at Shark Bay, Western Australia. II. Surface and sub-surface observations. Australian Wildlife Research, 9: 85-100.

Anderson PK and Birtles A (1978). Behaviour and ecology of the dugong, *Dugong dugon* (Sirenia): Observations in Shoalwater and Cleveland Bays, Queensland. Australian Wildlife Research, 5: 1–23.

Aragones LV, and Marsh H (2000). Impact of Dugong grazing and turtle cropping on tropical seagrass communities. Pacific Conservation Biology 5:277-288.

Bartoń K (2013). MuMIn: Multi-model inference. R package version 1.9. 13.

Bates D, Maechler M, Bolker B and Walker S (2015). Fitting linear mixed-effects models using lme4. Journal of Statistical Software, 67: 1-48.

Bryant CV, Davies JD, Jarvis JC, Tol S and Rasheed MA (2014a). Seagrasses in Port Curtis and Rodds Bay 2013: Annual Long Term Monitoring, Biannual Western Basin Surveys and Updated Baseline Survey, Centre for Tropical Water & Aquatic Ecosystem Research (TropWATER) Publication 14/23, James Cook University, Cairns, 71pp.

Bryant CV, Davies JN, Sankey T Jarvis JC and Rasheed MA (2014b). Long Term Seagrass Monitoring in Port Curtis: Quarterly Seagrass Assessments and Permanent Transect Monitoring Progress Report 2009 to 2013', Centre for Tropical Water & Aquatic Ecosystem Research (TropWATER) Publication 14/18, James Cook University, Cairns, 84pp.

Bryant CV, Davies JN, Sankey T and Rasheed MA (2016). Long-Term Seagrass Monitoring in Port Curtis: Quarterly Permanent Transect Monitoring Progress Report 2009 to 2015, Centre for Tropical Water & Aquatic Ecosystem Research (TropWATER) Publication 16/34, James Cook University, Cairns, 61pp.

Burnham, KP, and Anderson DR (2002). Model selection and multi-model inference: a practical information-theoretic approach. Springer, New York.

Cleguer C, Limpus CG, Gredzens C, Hamann M and Marsh H (2015a). Annual report on dugong tracking and habitat use in Gladstone in 2014. Report produced for the Ecosystem Research and Monitoring Program Advisory Panel as part of Gladstone Ports Corporation's Ecosystem Research and Monitoring Program, 13pp.

Cleguer C, Limpus CG, Hamann M and Marsh H (2015b). Annual report on dugong tracking and habitat use in Gladstone in 2015. Report produced for the Ecosystem Research and Monitoring Program Advisory Panel as part of Gladstone Ports Corporation's Ecosystem Research and Monitoring Program, 17pp.

Chartrand KM, Bryant CV, Carter AB, Ralph PJ and Rasheed MA (2016). Light thresholds to prevent dredging impacts on the Great Barrier Reef seagrass, *Zostera muelleri* ssp. *capricorni*. Frontiers in Marine Science 3: 106.

Christophe E, Michel J and Inglada J (2011). Remote sensing processing: From multicore to GPU. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 4: 643–652.

Davies JD, Bryant CV, Carter AB and Rasheed MA (2016). 'Seagrasses in Port Curtis and Rodds Bay 2015: Annual long-term monitoring', Centre for Tropical Water and Aquatic Ecosystem Research (TropWATER) Publication 16/02, James Cook University, Cairns, 66pp.

D'Souza E, Patankar V, Arthur R, Alcoverro T and Kelkar N (2013). Long-term occupancy trends in a data-poor dugong population in the Andaman and Nicobar Archipelago. Plos One 8:e76181.

Fuentes MMPB, Delean S, Grayson J, Lavender S, Logan M and Marsh H (2016). Spatial and Temporal Variation in the Effects of Climatic Variables on Dugong Calf Production. Plos One 11:e0155675.

Heinsohn GE, Wake J, Marsh H, and Spain AV (1977). The dugong (*Dugong dugon* (Müller)) in the seagrass system. Aquaculture 12: 235-248

Marsh H, De'ath G, Gribble N and Lane B (2005). Historical marine population estimates: Triggers or targets for conservation? The Dugong Case Study. Ecological Applications 15: 481-492.

Marsh H, O'Shea T and Reynolds III J (2011). Ecology and conservation of the sirenia: dugongs and manatees. Cambridge, Cambridge University Press.

McMahon K (2005). Recovery of subtropical seagrasses from natural disturbance. PhD Thesis. University of Queensland, Brisbane.

Meager JJ and Limpus C (2014). Mortality of inshore marine mammals in eastern Australia is predicted by freshwater discharge and air temperature. Plos One 9:e94849.

Mellors JE (1991). An evaluation of a rapid visual technique for estimating seagrass biomass. Aquatic Botany, 42: 67-73.

Mizuno K, Asada A, Matsumoto Y, Sugimoto K, Fujii T, Yamamuro M, Fortes MD, Sarceda M and Jimenez LA (2017). A simple and efficient method for making a high-resolution seagrass map and quantification of dugong feeding trail distribution: A field test at Mayo Bay, Philippines. Ecological Informatics, 38: 89-94.

Nakaoka M and Aioi K (1999). Growth of seagrass *Halophila ovalis* at dugong trails compared to existing withi-patch variation in a Thailand intertidal flat. Mar Ecol Prog Ser 184:97-103.

Preen A (1992). Interactions between dugongs and seagrasses in a subtropical environment. PhD Thesis, James Cook University, Townsville, Australia.

Preen A (1995). Impacts of dugong foraging on seagrass habitat: observational and experimental evidence for cultivation grazing. Marine Ecology Progress Series 124:201-213.

R Core Team (2016). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

Rasheed MA (1999). Recovery of experimentally created gaps within a tropical *Zostera capricorni* (Aschers.) seagrass meadow, Queensland Australia. Journal of Experimental Marine Biology and Ecology, 235: 183-200.

Rasheed MA (2004). Recovery and succession in a multi-species tropical seagrass meadow following experimental disturbance: the role of sexual and asexual reproduction. Journal of Experimental Marine Biology and Ecology, 310: 13-45.

Rasheed MA, McKenna SA, Carter AB and Coles RG (2014). Contrasting recovery of shallow and deep water seagrass communities following climate associated losses in tropical north Queensland, Australia. Marine Pollution Bulletin, 83: 491-499.

Rasheed MA and Unsworth RKF (2011). Long-term climate-associated dynamics of a tropical seagrass meadow: implications for the future. Marine Ecology Progress Series, 422: 93-103

Sheppard JK, Marsh H, Jones RE and Lawler IR (2010) Dugong habitat use in relation to seagrass nutrients, tides, and diel cycles. Marine Mammal Science, 26: 855-879.

Sheppard JK, Lawler IR and Marsh H (2007). Seagrass as pasture for seacows: Landscape-level dugong habitat evaluation. Estuarine, Coastal and Shelf Science 71:117-132.

Sobtzick S, Cleguer C, Hagihara R and Marsh H (2017). Distribution and abundance of dugong and large marine turtles in Moreton Bay, Hervey Bay and the southern Great Barrier Reef. A report to the Great Barrier Reef Marine Park Authority. Centre for Tropical Water & Aquatic Ecosystem Research (TropWATER) Publication 17/21, James Cook University, Townsville.

Sobtzick S, Grech A, Coles R, Cagnazzi D and Marsh H (2013). Status of the dugong population in the Gladstone area. Report produced for the Ecosystem Research and Monitoring Program Advisory Panel as part of Gladstone Ports Corporation's Ecosystem Research and Monitoring Program. 37 pp.

Tol SJ, Coles RG and Congdon BC (2016). Dugong dugon feeding in tropical Australian seagrass meadows: implications for conservation planning. PeerJ 4:e2194.

Wickham H (2009). ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York. URL: http://ggplot2.org

Wirsing AJ, Heithaus MR, and Dill LM (2007). Living on the edge: dugongs prefer to forage in microhabitats that allow escape from rather than avoidance of predators. Animal Behaviour 74:93-101.

Wirsing AJ, Heithaus MR, Frid A, and Dill LM (2008). Seascapes of fear: evaluating sublethal predator effects experienced and generated by marine mammals. Marine Mammal Science 24:1-15.

Wood S (2014). Package 'mgcv', version 1.8-3.

Zuur AF, Hilbe JM and Ieno EN (2013). A beginner's guide to GLM and GLMM with R. Highland Statistics: Newburgh, UK.

Zuur AF, Ieno EN, Walker NJ, Saveliev AA and Smith GM (2009). Mixed effects models and extensions in ecology with R. Springer, New York.

Zuur AF, Saveliev AA and Ieno EN (2014). A beginner's guide to generalized additive mixed models with R. Highland Statistics, Newburgh, UK.

6 ABBREVIATIONS

DFT	Dugong feeding trail
GBR	Great Barrier Reef
API	Application Program interface
EPBC	Environment Protection and Biodiversity Conservation Act 1999
MNES	Matter of National Environmental Significance
GPS	Global Positioning System
RTK	Real-Time Kinematic
SfM	Structure-from-Motion
EXIF	Exchangeable image file format
DSM	Digital surface model
PCA	Principal Components Analysis
WBDDP	Western Basin Dredging and Development Project
GAMM	Generalized additive mixed models
VIF	variance inflation factor
AICc	Akaike's Information Criterion corrected
ERMP	Environmental
GPU	Graphics processing units
CPU	Central Processing Unit
GLCM	Grey-Level Co-occurrence Matrix

7 APPENDICES

Appendix 1: Matrix of correlations for seagrass analysis

Scatterplot matrix of correlations between environmental variables one month and three months prior to seagrass sampling at Port Curtis permanent transect sites. Values in bottom left portion of figure are Pearson correlation coefficients.

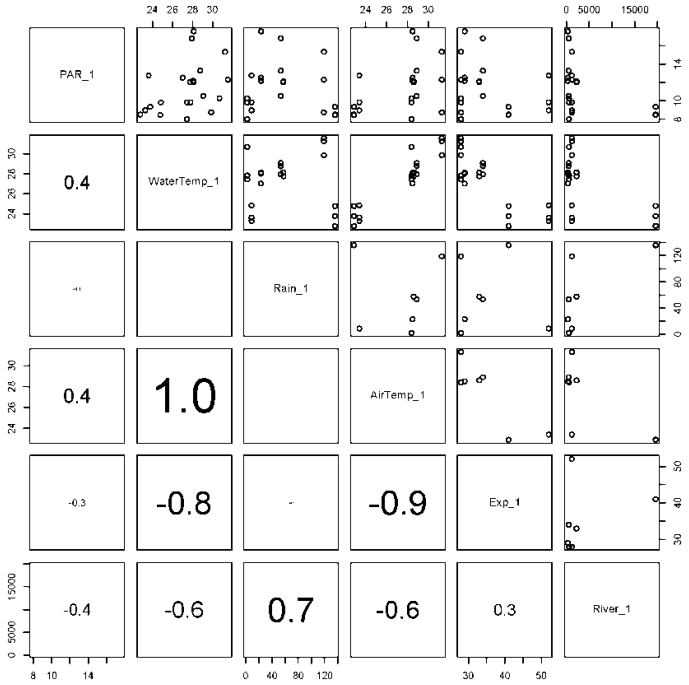

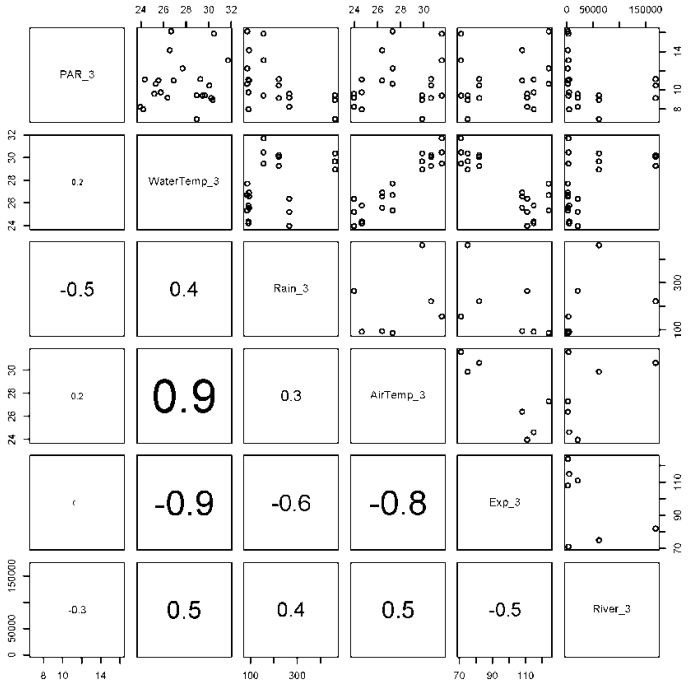

Appendix 2: Methods to automate mapping of dugong feeding trails in intertidal seagrass meadows

Introduction

The potential to use remote sensing to provide a census, as opposed to a sampled survey, of DFTs across the study area over several multiple time steps was identified. For this purpose, imagery was captured from a helicopter over approximately 20 km^2 in 10 repeat operations over a 2 year period.

In order to understand seasonal changes in DFTs and how they disturb seagrass and effect regrowth, fine scale spatial analysis of the trails is required. Given the spatial and temporal scale of the project and the known variability of the trails, the automated extraction of DFTs as opposed to manual classification would be a distinct advantage. This appendix details the research and development we applied to automate the extraction of DFTs from the imagery into a vector feature set against which individual trail characteristics could be evaluated and monitored over time.

The nature of the task and research into contemporary methods into related remote sensing problems is outlined to demonstrate available ways in which to approach the task. Ultimately, a work flow had to be settled upon, the success of which could not be known until the first pass through the process. The development of the approach was necessarily iterative with small scale trials of techniques necessary to determine likely success. With the survey itself, and the intense preprocessing of image data taking place in parallel, and using the same computational resources, we were faced with the fact that the methodology, and the preprocessing required for the first pass through the data at full scale, finished at around the same time, very close to the planned completion date of the project as a whole. The result is that while success was achieved at smaller test scales the aim of being able to fully automate the identification and mapping of DFTs at larger scales was not achieved. However, the segmentation required to train a machine learning algorithm to do the task, and likely methods by which to achieve success, have been learned through this work. The chosen procedure and outcomes are described below as well as the issues that created the complexities in automating DFT extraction from the Port Curtis and Rodds Bay data.

Problem defined

The objective was to extract DFTs from imagery collected during multiple aerial surveys over around 20 km^2 . The intended results would provide a 4-dimensional matrix of DFTs in space and time, allowing the team to know when each trail first appeared, how its size and continuity changed over time until it finally disappeared. The results would also allow for the analysis of the distribution of DFTs over time in relation to seagrass metrics being measured on the ground and aerially within the same survey periods.

There were several ways to approach the task. As with many complex feature extraction methods, the optimal approach was always likely to involve the fusion of different techniques in stages, taking advantages of the strengths of each. The overall project window was a major factor governing the approach. Although the duration of two years seems significant, the reality was that many data acquisition, processing, and experimentation tasks had to be conducted in parallel. It was necessary, for example, for the surveys to commence without having determined the best configuration under which to conduct the flights, in terms of pixel size, relating to flight height and velocity, camera shutter speed, image overlap and light conditions. Ultimately such parameters therefore relied on some research into contemporary work, although no such work related directly to DFTs. Allowance had to be made for the large demand on time and computational resources to pre-process imagery, on training and testing the data and on running the final algorithm.

Nature of dugong feeding trails

DFTs can be identified generally as sinuous paths of clear sand through patches of seagrass within the intertidal range. Their width is recognisable as being between 150 and 350 mm, but there exist larger areas evidently cleared by dugongs, which may only be identified by their connectivity with the recognisable thin trails, and these patches must be observed in the context of seagrass meadows which are, by their nature, discontinuous and "patchy" (Supplementary Figure 1). The length of the trails varies from nothing to several tens of metres. Their length is often discontinuous, increasingly so over time with seagrass regrowth, and the relationship between disconnected components of the same trail may only be recognised by our brains' ability to "join the dots", something which humans are very good at, and which computers are only beginning to succeed at, and only after considerable periods of algorithm training.

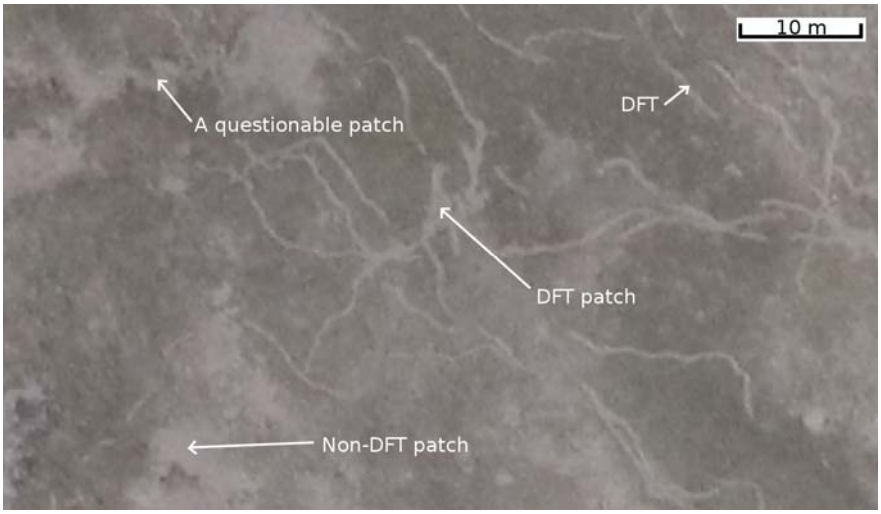

Supplementary Figure 1: An example of distinct DFTs as well as patches that were definitely not DFTs or where there was uncertainty (questionable patch). Only distinct DFTs were included in the analysis.

The ability to distinguish DFTs depends on the ability to extract features by their contrast with the immediate surroundings. Once some sort of segmentation has taken place, then classification can involve characteristics, such as shape, size, texture and spectral signature. The problem is made complex by the varied spectral response of the seagrass itself, the varied contrast of the trails against the seagrass and by the fact that some trails manifest in complete reverse - i.e. they appear as dark lines through a brighter background, as with some examples in Supplementary Figure 2. The segmentation and classification process must accommodate these differences.

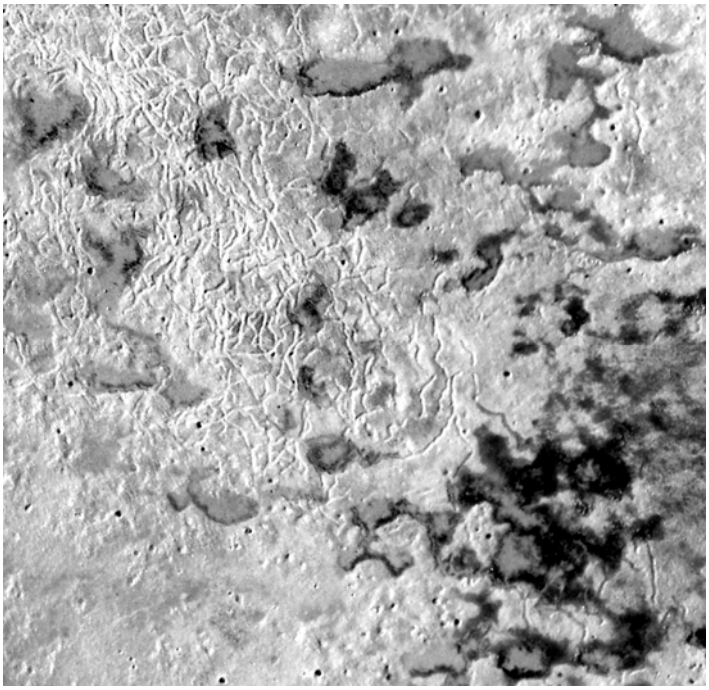

Supplementary Figure 2: DFTs represented by channels in bare sand, the seagrass having been stripped bare.

A range and combination of methods was tested. This section describes how each of the principal challenges was tackled.

How to account for varied spectral response

Classification of remote sensing imagery often involves the determination of a spectral signature of the target features, made up of values from different parts of the visual spectrum, such as red, green and blue. Very early on it became clear that with the images of the seagrass meadows, there was very little distinction between the signals in the green and blue channels in particular (Supplementary Figure 3). It was therefore decided to extract as much information from the three channels as possible into one channel, which would substantially reduce computation time. This was done using a principal components transformation, and selecting the first component (generally referred to as PCA1).

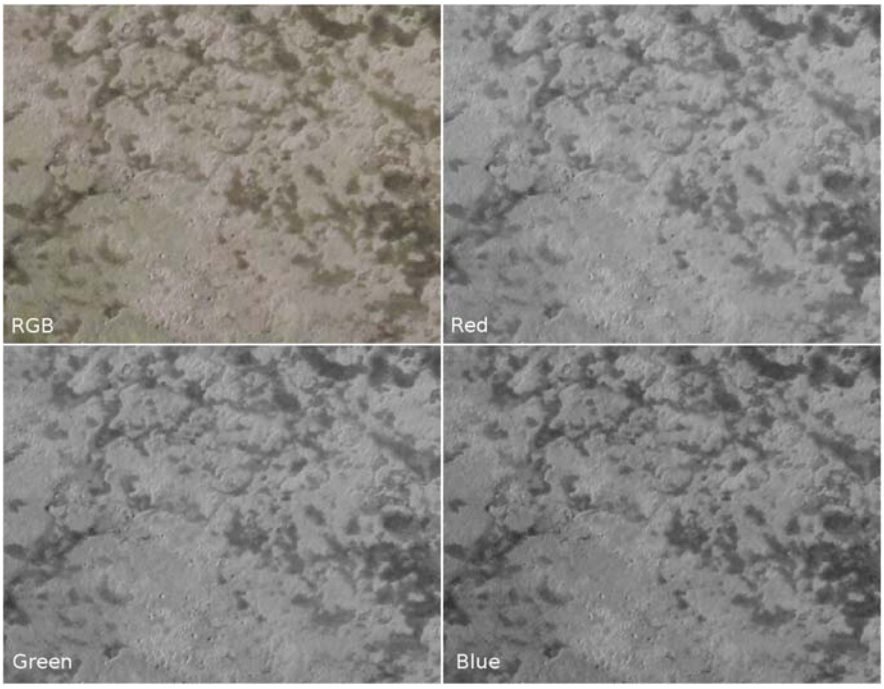

Supplementary Figure 3: RGB image of a section of the intertidal zone (top left), with the red, green and blue channel values shown separately.

Further to this, in order to normalise the variance of feeding trails such that broad-scale variation in spectral reflectance was largely discarded, an additional layer was created representing the *standard score* or *Z-score* of the original PCA1 values. Supplementary Figure 4 shows the result of such a calculation. For each pixel in the PCA1 image on the left, the mean μ and standard deviation σ of the values of all pixels in a 7×7 pixel window centred on the target pixel were used to create the Z-score, where $Z=(\text{PCA1}-\mu)/\sigma$. The highlighted region encircles a trail whose contrast with the background is much less than those other trails towards the centre of the image. The Z-score brings the contrast into a similar range.

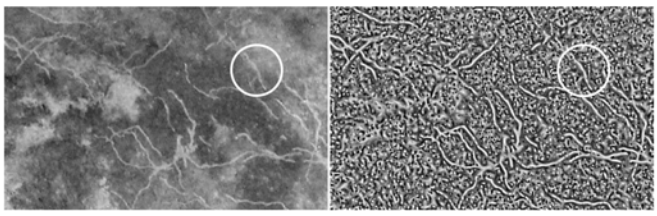

Supplementary Figure 4: First principal component of the red, green and blue reflectance values of a section of intertidal zone (left), with the Z-score of those values on the right.

Another approach which reduces the disruptive influence of variation in absolute values across an image is to focus on the rate of change of pixel values, using first and second order derivatives. Various well known tools exist to extract edges from imagery in this way, such as the Sobel and Canny operators (Danielsson and Seger 1990; Canny 1986), the latter of which is shown in Supplementary Figure 5. From this image it is clear that the mitigation of the influence of varied pixel values in this way is accompanied by another problem — the difficulty in differentiating between edges of interest and noise. Many such edge detection techniques were tested for their potential in the extraction of DFTs. Although it is quite possible to achieve impressive levels of success in mapping DFT edges in this way, the problem comes with the need to *close* a DFT border to identify one DFT. Even the use of more sophisticated techniques, such as with curvelets as described by Gebäck and Koumoutsakos (2009), improve the connectivity of detected edges, but fail to "complete the loop" as needed here.

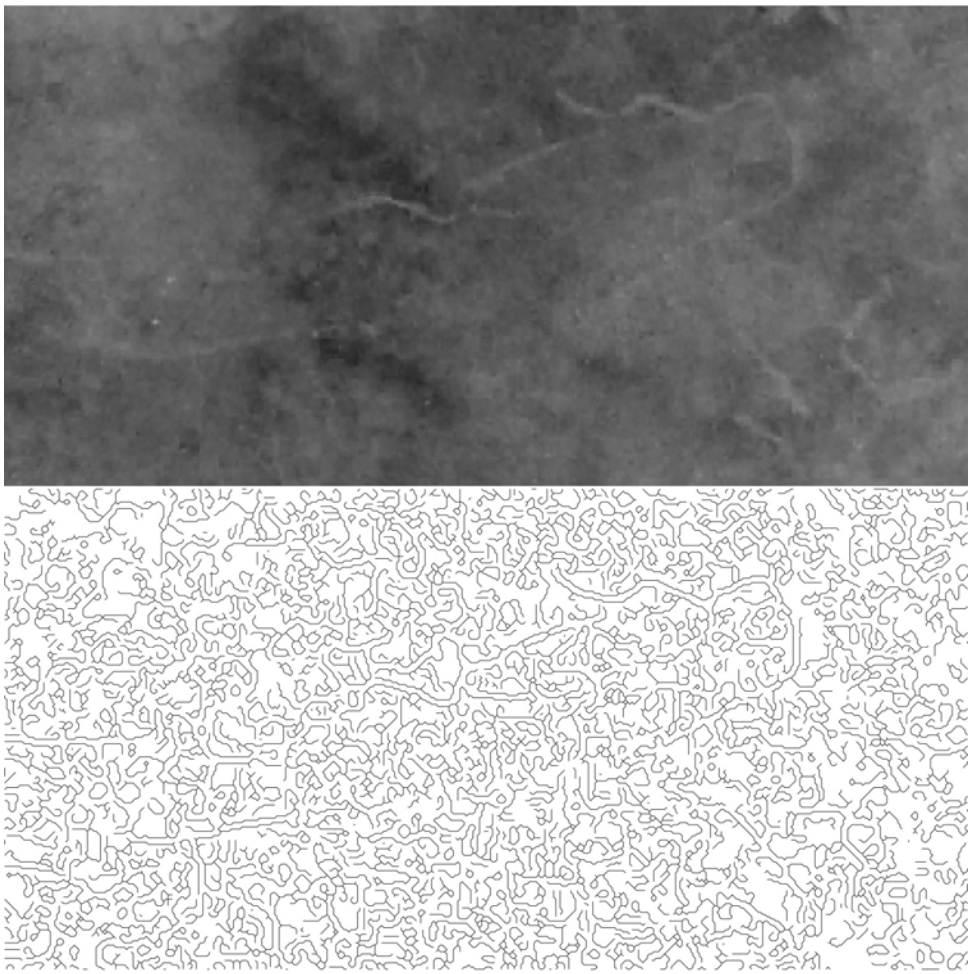

Supplementary Figure 5: PCA1 image including some DFTs (above), and the edges detected using the Canny operator.

Where local variation in values with respect to DFTs and their surrounds are consistent, but contain bias due to broad scale variation in values, one approach is to decompose the image using Fourier transforms or wavelets, which result in several layers of scale or frequency components, the sum of which produces the final image. The method is then to remove the layer or layers representing the scale at which the unwanted variations exist, to produce a normalised surface. This is effectively a spatial frequency low pass filter.

In Supplementary Figure 6, the original image (top left) has been decomposed into ten wavelet layers (right). The broad-scale background variations are contained in the larger scales. The lower-left image shows the recombination of scales 1–6 only, which tends to normalise the background values.

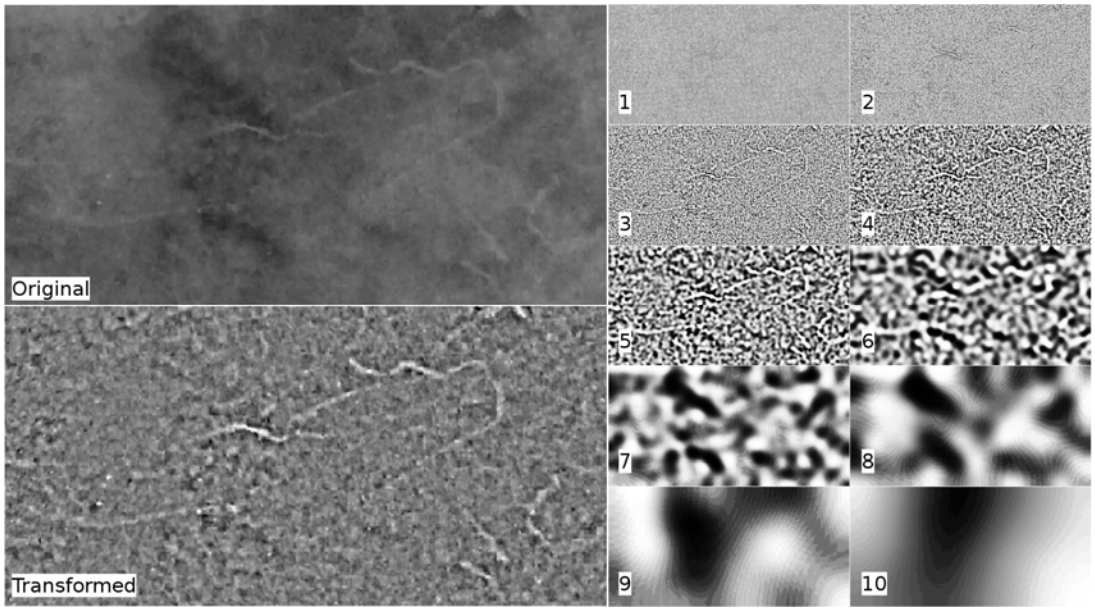

Supplementary Figure 6: The right section shows the individual wavelet scales decomposed from the original image. The "transformed" image shows the recombination of the wavelet scales 1–6 only.

How to extract DFTs based on shape and scale

The identification of DFTs through their shape assumes that segmentation, i.e. separation of individual features based on pixel values, has already taken place, given some of the methods described already. Following segmentation, the features can be considered as individual objects, for which shape characteristics may be determined. As DFTs are (on the whole) "long and thin", it seems a simple task to use some measure of this to identify them. There are many appropriate metrics to gauge this property, and all were used in the course of determining the optimal algorithm. Some of these are:

Compactness: $perimeter/(2\times\sqrt{\pi\times area})$

Fractal Dimension: $2\times\log(perimeter)/\log(area)$

Sinuosity: Centreline length / distance between end points.

The use of such measures introduces a further range of threshold choices, and the use of fixed thresholds risks errors of inclusion or omission. One of the problems with these metrics is that they do not work so well when trails have multiple branches. A large patch with a complex boundary can have a similar compactness and fractal dimension as a large DFT with a simple boundary. Sinuosity fails when a trail turns back on itself, placing the start and end points close together. We therefore adopted a slightly different approach in early classification attempts, seeking all shapes that had one dimension of the approximate width of the typical DFT (5 to 10 pixels), with the orthogonal dimension being larger by a certain factor.

We achieved this using a combination of skeletonisation of the features, and by distance transforms. Skeletonisation is a morphological process (described later) which thins features one pixel at a time until they are one pixel in width (Supplementary Figure 7).

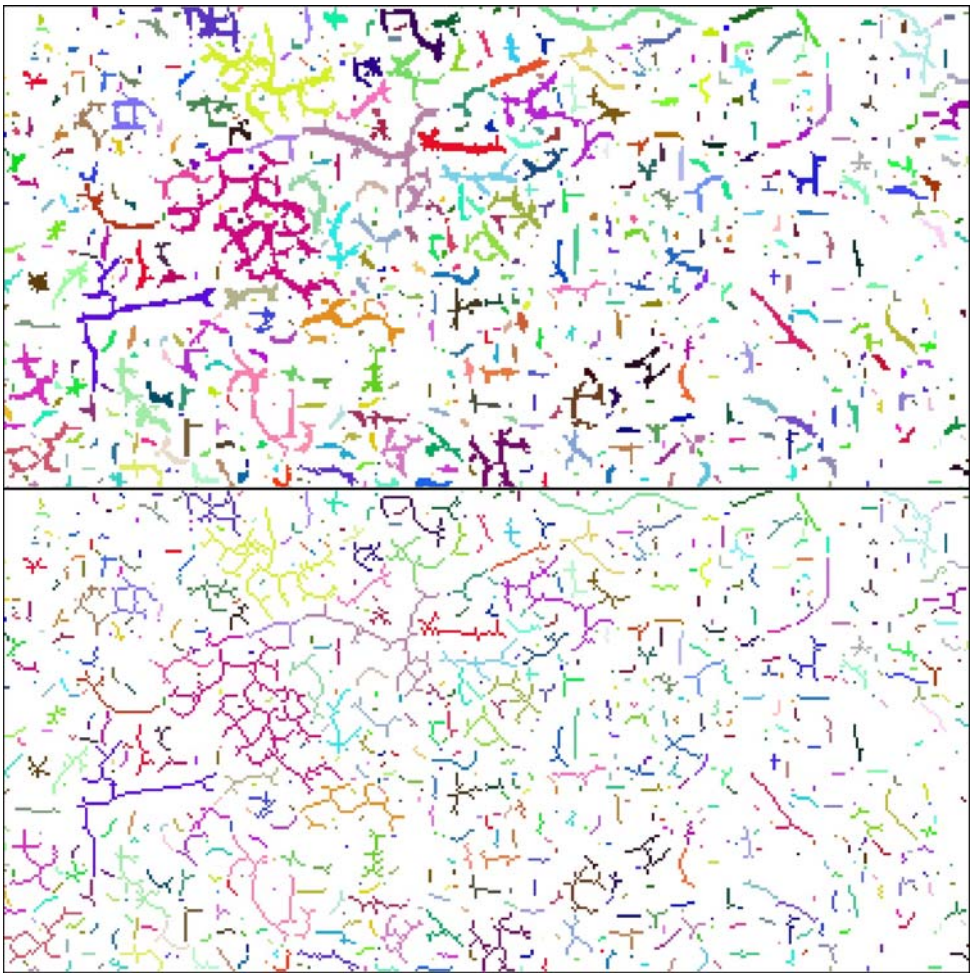

Supplementary Figure 7: Original raster features (above), thinned by morphological "skeletonisation" operator until 1 pixel thick (below).

The skeletonisation operation tends to “break” features with narrow necks, and so a further operation is necessary to retain the original connectivity. This process thins both large patches and thin trails. The way to determine that the original feature had the desired linear shape is to divide the number of pixels in the original feature by the number of pixels in the thinned feature. For a DFT, we would expect the ratio to be between 3 and 7, assuming a pixel dimension of about 5 cm. Some tolerance must be built in to this measure, and hence another threshold was introduced.

A typical working sequence

Various combinations of methods were tried. A typical work flow is given here, to demonstrate the type of parameters which need evaluating to extract DFT features in this way:

1. Determine sample region
2. Principal components transformation (PCA)
3. Calculate Z-score: **Aggregate window size = W**
4. Segmentation stage 1: **Threshold on Z-score = Z**
5. Segmentation stage 2: Associate features by connectivity
6. Vectorise features so that zonal statistics may be calculated

7. Zonal stats: area **A**, sinuosity **S**, mean (**μ**) and standard deviation (**σ**) PCA values, fractal dimension **FD**.
8. Skeletonise and determine **Pixel Width Ratio = R**

This particular work flow is one example of many that were tested, using combinations of methods incorporating a variety of approaches including wavelet decomposition, morphological operators and other computer vision techniques, some of which are described later. This example produces eight parameters which need to be optimised.

How to decide on the optimal parameter values

The possible approaches to optimising the parameters required the algorithm to fall into two distinct types. Firstly, we can invest time in the training of a machine learning algorithm. This approach has the following requisites:

1. The trainer must be certain in their mind of what is and what is not a DFT. If the task requires more that one user, such determination must be consistent.
2. The segmentation has already taken place. This is required in order to obtain the zonal statistics described above.
3. The training must include DFTs of all types found across all of the imagery, being sure to include different environments, sizes and substratum.
4. In addition to the DFTs identified in the training, there must be a comparable number of similar features positively identified as non-DFTs. This is essential.

These requisites all present a challenge. Challenges related to sheer volume of work can be mitigated with an increased workforce. However, this increases inconsistency. Experience through this project taught us that when imagery was not ideal or potential DFTs were not absolutely clear there are a lot of “maybes” when classifying DFT’s so effectively training the algorithm could be problematic.

The second broad approach addresses all of these challenges to a certain extent, at the expense of increased computational overhead. This approach is to set ranges for the major parameters and vary each incrementally with the others fixed. This could be done for selected sample areas deemed representative of the whole. Computation times could be massive, with the number of runs being the multiple of the number of increments allowed for each parameter. The trick was to keep the increments reasonably large, allowing us to hone in on the best values with smaller increments over smaller ranges.

Given the volume of data, and the uncertainty of how many different DFT types would need to be trained, the second approach was favoured. Small sample areas run across the range of parameters resulted in sets of potential parameter-set images which we called “swatches” (Supplementary Figure 8). We felt that this approach would also facilitate agreement on what is and what is not a DFT, with the optimal parameters being chosen from the swatch set offering the best compromise.

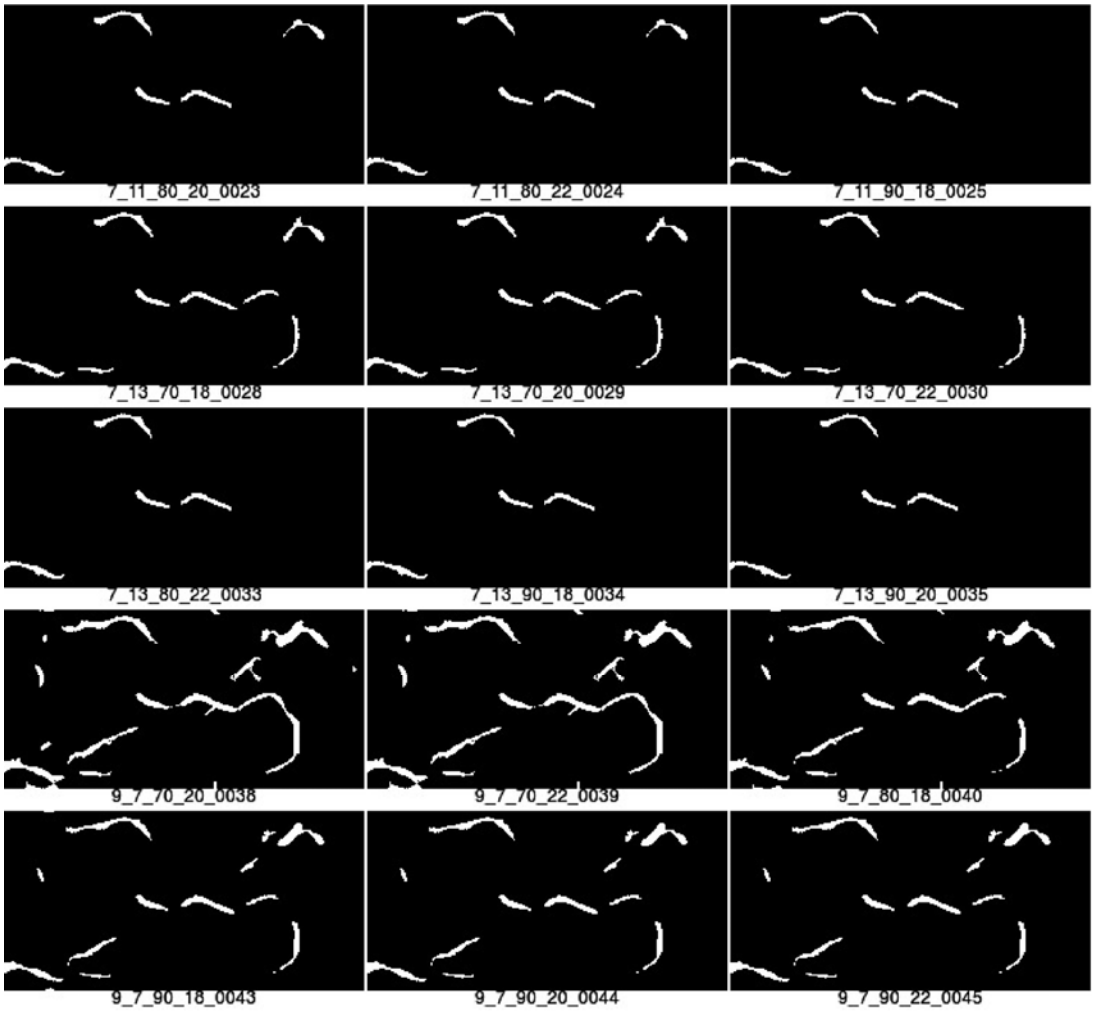

Supplementary Figure 8: Extract from a swatch. The whole swatch might contain hundreds of sample images with potential DFT features to help in parameter selection.

How to overcome computational constraints

There were two computationally intense stages of the project. One was the processing of the thousands of individual image frames into mosaics, and the other was the running of the DFT algorithm. The first was done using the proprietary software Pix4D (Pix4D, 2014). This software is capable of parallel processing using multiple cores, and relies heavily on available random access memory (RAM). As the software is expensive, we chose to spend the money on increasing the capacity of the "workhorse" computer, rather than on additional software licences, so that the improvements would serve both stages of the computation. The machine used to run the photogrammetry therefore had 128 GB of RAM and 16 coprocessors.

Data size

The substantial volume of data to be processed required an increase in internal hard-drive storage capacity to 6 TB on the working machine. Network allowances and the clearing of data from intermediate steps ensured that we were able to accommodate the space requirements of the project, whilst maintaining backup copies on external drives and across the network.

Memory capacity

The second stage computation, running the algorithm to produce swatches and maps for entire seagrass meadows, was done using code written in BASH, perl, python, and R, which made system calls to various software packages. Some of these were capable of advanced memory management and/or multi-core

processing, others were not. For this reason, it was necessary to break some of the work into batches. Apart from being a necessity, this provided the advantage that the batches could be submitted to a queue on the university's High Performance Computer (HPC) network and run in parallel.

Approach

Following photogrammetry, individual mosaics were clipped to their core region of interest, and then broken down to individual tiles of 1200 x 1200 pixels, comprising a 1000 x 1000 pixel core with a 100 pixel overlap on each side. Once processing was complete, the tiles were clipped and merged. The data were sent to the HPC, and batch scripts were run with the chosen parameters in parallel with little strain on the resources of the local workhorse.

Potential approaches to the feature extraction problem

Introduction

We cover just a few contenders here as a subset of available methods, as these broad measures are felt to encapsulate the core issue, which is the delineation of DFT boundaries and the association of fragmented parts of a DFT.

Edge detection

Ultimately, in order to delineate the boundaries between DFTs and their surroundings, the edges of the DFTs must be mapped. It therefore makes sense to make use of edge detection techniques from the outset to establish these boundaries. A successful procedure from that point will need to close boundaries and categorise pixels as being *inside* or *outside*, in order for shapes to be established, whose metrics and radiometric statistics can then be used to determine whether each shape is or is not a DFT.

There are well known simple tools to find edges in an image, and these have been built upon and refined over time. First order differential operators such as the Canny edge detector search for a value of the first-order derivative (ie. the gradient) above a set threshold, and then continue along the edge until other conditions are met. Other edge detectors search for the zero-crossing of the second-order derivative, indicating those points at which the gradient is at a maximum.

The incorporation of edge detection techniques into our methodology faces the following challenges:

1. Edges must be closed. Enforcing this in an algorithm is computationally expensive, and detracts from the main advantage of using simple edge detection in the first place - that it is computationally simple.
2. The various thresholds need to be adjusted to the local environment in order to produce a consistent output.
3. Every feature has edges, and much computation would need to be expended before the types of tests that filter for shape could be applied.

Gebäck and Koumoutsakos (2009) describe a multi-step edge detection method using the Canny edge detector (Canny 1986) to mark edges, including a method to improve connectivity between edges using the discrete curvelet transform. We incorporated this approach into trials on small areas, but found the same constraints as mentioned above.

Textural measures

Where there is little spectral variance in an image, or, as is the case with the seagrass meadows studied here, there is little distinction between radiometric response from different wavelengths, we often have a greater chance at classifying an image based on the pattern of variability of the values of a single band, which is described as *texture*. Probably the most common measures used in extracting textural predictors from an image are those that were introduced by Haralick et al. in the early 1970's, who identified a suite of measures known collectively as the Grey-Level Co-occurrence Matrix (GLCM) (Haralick et al. 1973). As these measures

work by identifying the relationship of values of a minimum of two pixels over varied distances, their applicability for the identification of DFTs with small dimensions may be limited, but GLCM measures may play a part in the pre-classification of DFT backgrounds. Calculated measures include the following:

Entropy: A measure of order, or *randomness* of values.

Angular Second Moment: A measure of local homogeneity.

Contrast: The linear dependency of pixel values on their neighbours.

There are many such measures which won't be listed here. It is crucial that the distance between pairs of analysed pixels and the size of the moving window are set appropriately for the resolution of the image as compared with the size of the features. We found that a window size of 11, with a distance of five accentuated the linearity of features at the scale of DFTs.

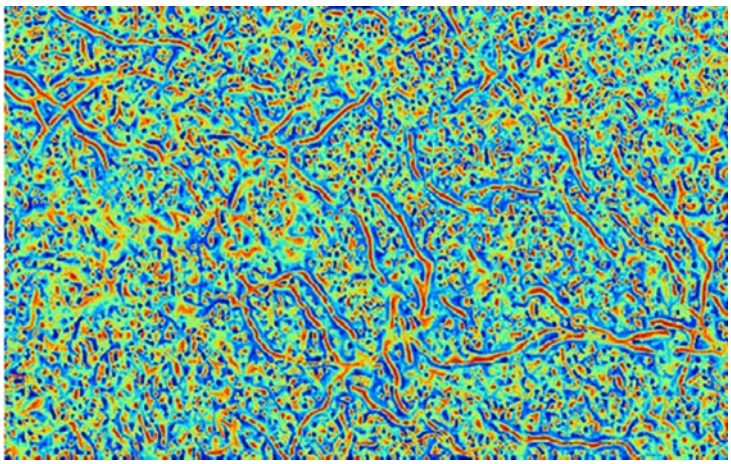

Supplementary Figure 9: The sum of Angluar Second Moment (ASM) texture filter applied to all cardinal directions over a section of beach.

Texture cannot define the boundaries of potential features, but once features have been extracted, textural measures can be aggregated as zonal statistics and associated with the features to determine their likely membership to the target DFT set.

Morphology

Morphology deals with the analysis of shapes and functions using a reference shape, called a *structuring element* (SE), which in operations is applied to every part of a matrix (such as an image), whose values are compared to the SE to direct some modification or outcome (Droogenbroeck and Talbot 1996). There are both binary and greyscale morphology operations. They serve to thin, thicken, bridge or breach features, depending on the particular operation, and the relationship between the raster features and the size and shape of the SE.

If we consider τ as a threshold function on pixels of value *v* using a threshold *t*, such that

$$\Theta(v,t) = \begin{cases} 1, & v > t \\ 0, & \text{otherwise} \end{cases}$$

and that *c* is the count of binary TRUE values inside the structuring element *s* when it is centred over the pixel under analysis, such that:

$$c = v \odot s$$

(Szeliski 2011)

Then the basic binary operations are as follows (the following descriptions and images are taken from the OpenCV documentation Bradski (2000) and Szeliski (2011):

Erosion

$$\text{erode}(v, s) = \Theta(c, S)$$

where *S* is the size of the SE.

Erode will only assign TRUE to the target pixel if all pixels covered by the SE are TRUE (Supplementary Figure 10).

Supplementary Figure 10: Left: Original binary image. TRUE or 1 is represented by white, and FALSE or 0 by black. Right: Image after an erosion filter with an SE as a 5×5 pixel box is applied.

Dilation

$$\text{dilate}(v, s) = \Theta(c, 1)$$

Dilate will assign TRUE when at least one pixel covered by the SE is TRUE (Supplementary Figure 11).

Supplementary Figure 11: Image after an dilation filter with an SE as a 5×5 pixel box is applied.

Common secondary operations combine two of these basic functions, as follows:

Opening

$$\text{open}(v, s) = \text{dilate}(\text{erode}(v, s), s)$$

An opening is an erosion followed by a dilation. It has the effect of producing the original image, having filtered out elements that are smaller than the SE (Supplementary Figure 12).

Supplementary Figure 12: Noise in the image on the left that is smaller than the SE is filtered out after an openning operation.

Connected components

The connected components operation is an important step in the conversion from pixel-based identities to those of features. Pixels with values of 1 or TRUE in a binary image which are adjacent are associated together as part of a feature, and reassigned a common unique identifier as a value. Adjacency can be defined as including pixels which share sides only (the $\mathcal{N}_4$ neighbourhood) or that include corners ($\mathcal{N}_8$).

Such binary operations have greyscale equivalents. Vincent (1993) described how greyscale reconstruction (entity selection by marking) may be applied using an efficient algorithm to filter microaneurisms from an image (Supplementary Figure 13).

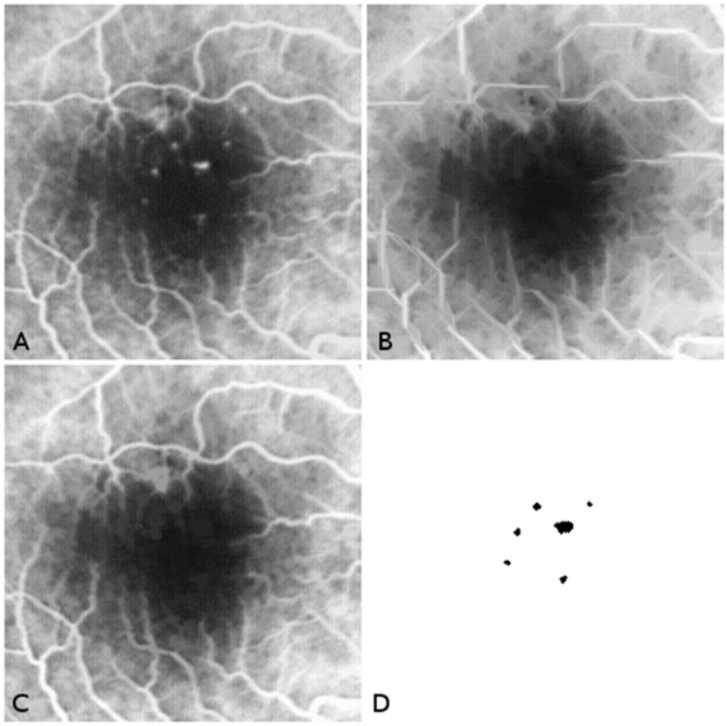

Supplementary Figure 13: Use of grayscale reconstruction for image segmentation: (a) original image of blood vessels, (b) supremum of openings by segments, (c) reconstructed image, (d) microaneurisms obtained after substraction of (c) from (a) and thresholding step (Vincent 1993)

What is interesting for us is the use of linear SEs. In Supplementary Figure 13 the aneurisms at the centre of the image have very similar brightness values to the blood vessels around the outside. They are extracted by virtue of the fact that their shape is not linear. A series of opening operations is performed on the image using an elongated SE in all orientations. The aneurisms are eliminated by this operation, but for the blood vessels, there is at least one direction in which the vessels are larger than the SE, which results in a (thinned) version in image B. This is used as a marker to select connected components of the original image (C), which can then be subtracted from the original to yield the aneurisms. The similarity of the manifestation of DFTs in our images to the blood vessels in Supplementary Figure 13 implies the possible suitability of an adaptation of this method to extract DFTs. Using similar tools, Zana and Klein (2001) also worked on extracting features from retinal images, this time concentrating on the blood vessels themselves, the work therefore being directly analogous to our own problem. They describe the use of a linear SE (Supplementary Figure 14) in a sum of top hat transforms to filter out the vessels as the last part in a chain of processes, involving noise reduction, the application of a Laplacian of Gaussian operator, evaluation of cross-curvature and linear filtering.

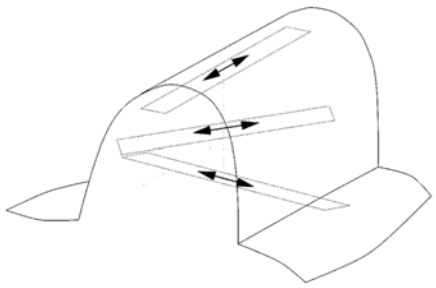

Supplementary Figure 14: Openings using linear structuring elements (Zana and Klein 2001).

The results of the algorithm proposed by Zana and Klein (2001) can be seen in Supplementary Figure 15. The top hat transform referred to above is commonly known as the white top hat transform, which is the difference between an image before and after an opening operator has been applied. It has the effect of removing features smaller than the SE. The sum of a series of multi-directional top hat operations using a linear SE will therefore tend to allow through structures that have a linear component in one direction or other.

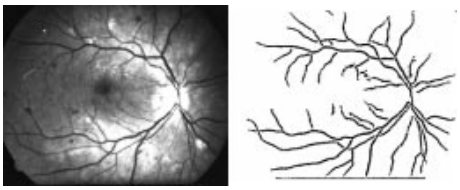

Supplementary Figure 15: Original image (left) and the results of the Zana and Klein (2001) proposed algorithm (right)

Popularity of morphological methods has led to the development of algorithms to increase their computational efficiency, which means that this method is attractive to us when considering the very large, high resolution data with which we must deal (Droogenbroeck and Talbot 1996; Gratin et al. 1993).

Active contours

Active contours are boundary features that iteratively converge around a target feature based on optimisation conditions and/or manual guidance (Szeliski 2011; Blake and Isard 1998). A couple of examples of active contour methods are *snakes* and *level sets*.

Snakes are two-dimensional splines that seek to minimize internal and external energy (Szeliski 2011). Internal energy relates to the spline energy, and external energy is a function of the presence of ridges, gradients and line terminations (Supplementary Figure 16).

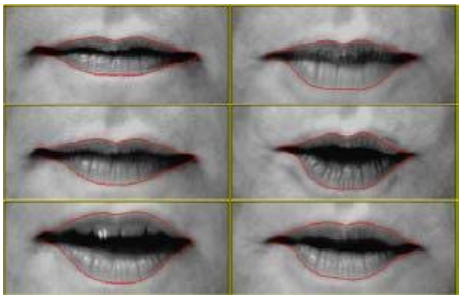

Supplementary Figure 16: Active contour for lip tracking (Szeliski 2011; Kass et al. 1988).

Snakes are initialised by drawing them outside the object of interest, as they tend to function by shrinking around the feature. This is problematic to us, as DFTs are intertwined and sinuous. In this regard, it would almost be necessary to know where the DFTs were in the first place in order to start the mapping process using snakes.

Level sets are defined by second order zero-crossings and seek to modify their underlying function in order to track objects. As this implies, they are suited to initialisation with prior knowledge of the initial state of the feature, and therefore probably less suited to our type of problem (Osher and Paragios 2003).

Chosen method

The process of determining the method of feature extraction may be considered in two parts—before and after having completely processed seagrass meadow imagery on which to test the outcomes. The first part involved research into work to date on the extraction of features under similar conditions, but only with the final data were we able to understand the variation in conditions which needed to be accommodated.

Trials

Once a single dataset including DFTs had been processed, we were able to test many of the methods discussed. On the assumption that a training/prediction approach using machine learning was to be used, we were able to train and classify DFTs in a small area of the South Trees meadow to a high degree of accuracy. To accommodate potential differences in interpretation on features that were potentially in dispute or there was a lack of certainty that they represented feeding trails, and to allow tweaking at a later stage, we used the ability of machine learning predictors to output results as probabilities, rather than binary classes. An extract of the results is shown in Supplementary Figure 17.

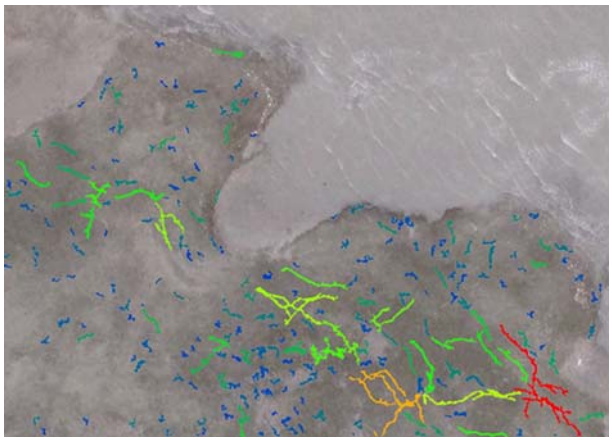

Supplementary Figure 17: DFTs extracted using a training / prediction process with the Random Forest algorithm. Colours represent the probability of the feature being a DFT from blue (least likely) to red (most likely).

Problems and constraints

To reach an acceptable level of accuracy, a great deal of training was necessary. This involved the selection of hundreds of features identified as DFTs, and the identification of other linear features such as wave fronts and gullies that were not DFTs, but which may be easily confused as such. This process was very labour intensive.

In order for the training to be done, it was necessary for segmentation to have been completed. This meant that the parameters used in the morphology and shape metric process had to be decided upon first. Sometimes it was not evident that the chosen values were not ideal until the classification had been completed, and a rerun of the segmentation meant a rerun of at least part of the training process, as there was rarely a one-to-one relationship between the features before and after the rerun. This made the process almost unworkable once several iterations were required.

Despite having tried to reduce potential false positives in the training stage, the process still required a manual "clean-up", with wave-front features having to be deleted along with DFT-like features which were present in areas known to have been too far away from the shoreline to have been feasible.

Perhaps the biggest problem with the training / prediction approach was the potential for “over-fitting”. This is where parameters are tweaked to provide a highly acceptable level of accuracy in a sample region that cannot be representative of the entire landscape, rendering the parameters ineffective elsewhere. The degree that this is a problem could not be known until all areas had been run. With this approach, we were faced with the possibility that it may be faster just to pick every DFT manually across all data sets through the time series.

Chosen course

We chose the swatch-based approach for the overall configuration. For the algorithm itself, we found that energy-based computer vision techniques such as snakes and active contours were not only too computationally intensive for our volume of data, but also suffered from the same dangers of over-fitting as the machine-learning approach. Furthermore, using swatches to optimise energy-based algorithms did not seem feasible. We therefore opted to concentrate on the flexible and widely-used morphology suite of operators, which seemed to offer good results with reasonable resource usage based on the smaller scale trails.

The algorithm used

The following sequence of morphological operators were applied to each tile in order to extract features. The results depend on the four input parameters TOPHAT, NEEDLE, SIZE_THRESHOLD and GREYSCALE_THRESHOLD.

Sum of Top Hats: The *top hat* operator will brighten a region if it has high values falling within and in the direction of the linear structuring element (SE). The parameter TOPHAT defines the length of the SE. The *sum of top hats* adds the response to the kernels of all orientations. Eliminates fat patches.

Greyscale threshold: The *sum of top hats* operation above will yield low values if a linear feature below TOPHAT length does not exist. The GREYSCALE_THRESHOLD parameter operates a high pass filter to allow only conforming shapes through.

Mark: Connects components to features and selects only those features that have their longest access greater in length that the parameter NEEDLE.

Sieve: SIZE_THRESHOLD

Reconstruct: The results of the previous operations are combined.

Closing: Bridges small gaps between features to join them.

Erosion: Uses the SE of TOPHAT dimensions to erode features that may have been increased in size by reconstruction and closing. Fragments that were connected in the closing operation will survive, those that were not connected will not survive.

Connected Components: A unique value is shared by components that remain physically connected to identify them as part of one feature.

Procedure

The general work flow to run the algorithm is shown in Supplementary Figure 18. Most stages were managed by BASH and perl scripting, operating further scripts written in R and running on the HPC. The swatches were generated locally. Once the parameters had been determined, they were written to a configuration file on the HPC. The masked meadow mosaic was tiled and grouped in sets. Each set was run as one batch on the HPC, and scripts were run which gave instructions to queue the batch tasks to be run in parallel as managed by the HPC server.

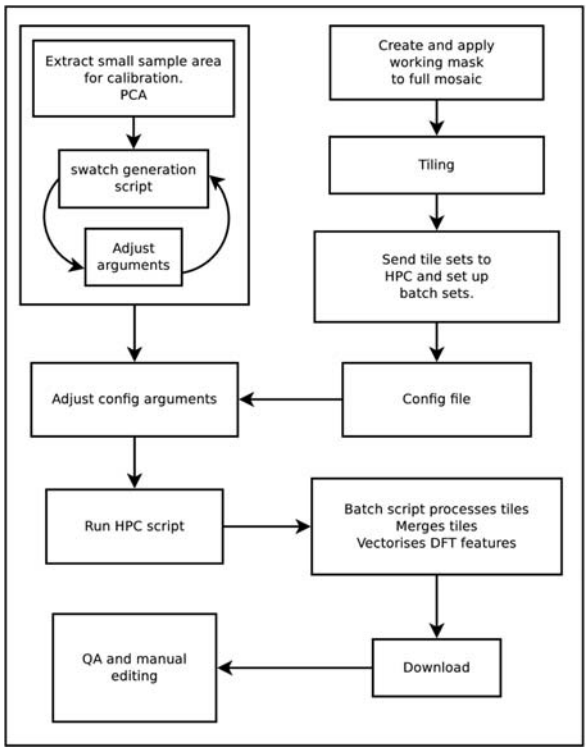

Supplementary Figure 18: DFT extraction work flow.

Once the batch tasks were complete, the resulting vector features were downloaded for analysis.

Results and discussion

The clear lesson that was drawn from the results was that the parameter set used in the final algorithm was not sufficiently complex to account for the range of characteristics found across the data. This was true for the values across a single meadow as much as it was across the time series. In some cases the swatches delivered a set of parameter values that worked well for the swatch area, but failed to distinguish DFTs elsewhere in the meadow when the algorithm was run over the complete mosaic. In other cases, the range of parameter values chosen for the swatches did not cover the optimal values peculiar to the environmental conditions of the imagery, and so produced no satisfactory samples. To date, time constraints have prevented us from being able to mitigate this problem with more swatch outputs, and it is clear that better course of action would be to change the method according to the lessons learned. We believe that accurate results are possible with a combination of the methods described in Section 2. The results achieved so far were constrained by the following problems:

1. Variation in absolute and relative neighbourhood pixel values across space and time were not adequately accounted for.
2. The range of parameter values chosen for the swatches did not contain the optimal value.
3. Ambiguity in the interpretation of a shape as a DFT.

Item 1 requires either the addition of further normalisation (such as the extraction of high wavelet scales) and/or a change in the principal method of segmentation. Another approach to this problem would be a preliminary broad-scale classification of substratum, prior to the feature extraction stage. In this way we may have parameter sets which are geared towards substratum classes. This approach was taken for the interim report. Supplementary Figure 19 shows an area of intertidal zone in which the seagrass has been classified into three spectrally distinct types, shown as different shades of green. In such a case, separate sets of parameters could be used to extract DFTs in each type.

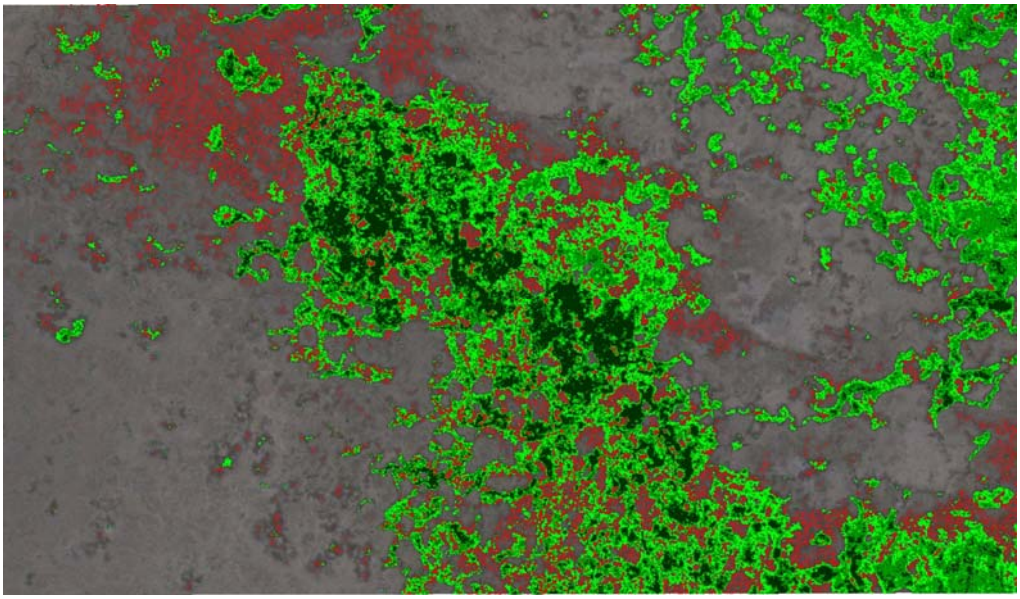

Supplementary Figure 19: DFTs (red) over a region in which vegetation has been classified into 3 types or densities.

Item 2; Ensuring an appropriate parameter range would require some *a priori* knowledge of the range of values to expect before setting the swatch parameters, which suggests a "swatch for the swatch", or at least a two-stage swatch procedure. This in itself is not unthinkable, but as more time and effort is devoted to the guessing of correct parameters, the advantages to be gained from devoting more time to training in a machine-learning approach become more attractive, especially in light of the requirement for greater complexity, and therefore more parameters. Machine learning lends itself to a situation where we have many (even hundreds) of independent variables in our model, as long as we are confident that, among them, they are capable of predicting the features we wish to extract. Background seagrass characteristics, for example, could make up some of these parameters. The beauty of the machine learning approach is that the actual effects of the substrata on the sensitivity of the other parameters do not need to be known explicitly.

Item 3, the subjectivity in the definition of a DFT, will remain with any extraction method. With a machine-learning approach, segmentation of potential features is required prior to training. In the segmentation stage, geometric and topographic decisions can be made such as length, width range and proximity to the shoreline, which can be discussed and set down at the outset. A slightly more difficult criterion might be the length of a group of short but obviously related patches that once made up a continuous trail. These patches can be connected using morphological operators and so could be filtered at a secondary stage in the segmentation process. However, as we have discovered in this project, this already involves the setting of parameters which can be problematic. The solution can only be to introduce the training early. This requires a segmentation process that tends to be inclusive, resulting in many thousands of features, only a small proportion of which would be DFTs, to be used for training.

Conclusion and further work

We have described feature extraction methods in contemporary use, many permutations of which may form the best solution for the extraction of DFTs. Given time and resource constraints, we have focused on one sequence of operations which we felt was achievable and which would be likely to yield good results. Over any given confined region, we have a method that can extract DFTs well, once parameters have been set accordingly. However, our model over-fits to an individual circumstance and is not transferable to the broader data coverage. In this respect the method fails.

We have, however, gathered and tested the tools likely required to produce a successful DFT extraction method. We understand the variability of the trails and of the substratum on which they appear. Having once decided to avoid the machine-learning option for fear of insufficient resources to train the algorithm, we limited ourselves to a handful of parameters by which to adjust the model. We now know that this did not afford us the complexity required to account for the enormous variability in background and DFT features within a single mosaic and between sites and times. Avoiding machine learning cut off, what with hindsight appears the only way to manage the large amount of variables on which we need to base the model.

From the work we have done to date we have the means to segment the DFTs to reach the training stage of a machine learning approach. Using broad ranges of parameter values, we can apply the current model to extract all possible features inclusively into vector format. These are then ready for operators to sub-sample for the identification of DFTs and non-DFTs. We also have the expertise within the research team to work with machine learning algorithms. In addition computational resources, which have been a significant consideration throughout the project have undergone significant advances since the project started. More cloud-based distributed networks have become available on which to process massive amounts of data. Amazon's AWS[6], for example, allows the choice of multiple graphics processing units (GPU), which are much more efficient than CPUs in the handling of imagery processes such as photogrammetry (Christophe et al. 2011). Google's *TensorFlow*[7] offers an open source library through a single API that allows us to deploy operations to many CPUs and GPUs as required. The network is designed to allow *deep learning* research, which is essentially a multi-layered machine learning approach.

For the analysis presented in the main body of this report, we have necessarily had to rely on a manual classification approach to analyse the location and density of DFTs, due to the problems and limitations encountered in this research into automated extraction. While this has allowed a viable method to achieve the goals of the project, manual classification is highly labour intensive at large spatial scales, and has limits in what can be achieved in terms of finer scale analysis at the individual trail level. We believe that the significant work we have done to date in developing an algorithm extraction approach and the pathway outlined above is well worth pursuing further. Once developed it would offer the potential for rapid assessments of imagery over very large spatial scales which would provide an excellent tool for understanding and managing dugong and their feeding behaviour over large regional scales, as well as detailed assessments at small scales.

[6] Amazon Web Services, `https://aws.amazon.com/`

[7] TensorFlow™, `https://www.tensorflow.org/`

References

Blake A and Isard M (1998). *Active Contours: The Application of Techniques from Graphics, Vision, Control Theory and Statistics to Visual Tracking of Shapes in Motion.* Springer Verlag, London

Bradski G (2000). The OpenCV Library. Dr. Dobb's Journal of Software Tools, citeulike:2236121.

Canny J (1986). A computational approach to edge detection. IEEE Trans. Pattern Analysis and Machine Intelligence, 8: 679–68.

Christophe E, Michel J and Inglada J (2011). Remote sensing processing: From multicore to GPU. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 4: 643–652.

Danielsson P and Seger O (1990). *Generalized and Separable Sobel Operators*. Academic Press, San Diego.

Droogenbroeck MV and Talbot H (1996). Fast computation of morphological operations with arbitrary structuring elements. Pattern Recognition Letters, 17.

Gebäck T and Koumoutsakos P (2009). Edge detection in microscopy images using curvelets. BMC Bioinformatics, 10: 1–14.

Gratin C, Meyer F and Talbot H (1993). *Visual Communications and Image Processing*, chapter Fast grey-level morphological transforms with any structuring elements. SPIE.

Haralick RM, Shanmugam K and Dinstein I (1973). Textural features for image classification. *IEEE Transactions on Syste*, SMC-3: 610–621.

Kass M, Witkin A and Terzopoulos D (1988). Snakes: Active contour models. *International Journal of Computer Vision*, 1: 321–331.

Osher S and Paragios N, editors (2003). *Geometric Level Set Methods in Imaging, Vision, and Graphics*. Springer, London.

Pix4D (2014). *Pix4Dmapper Pro*. Pix4D SA, EPFL Innovation Park, Building F 1015 Lausanne, Switzerland.

Szeliski R (2011). *Computer Vision Algorithms and Applications*. Springer, London.

Vincent L (1993). Morphological grayscale reconstruction in image analysis: Applications and efficient algorithms. *IEEE Transactions on Image Processing*, 2(2):176–201.

Zana F and Klein J-C (2001). Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE Transactions on Image Processing*, 10(7):1010–1019.