TropWATER
Centre for Tropical Water and Aquatic Ecosystem Research

JAMES COOK UNIVERSITY
AUSTRALIA

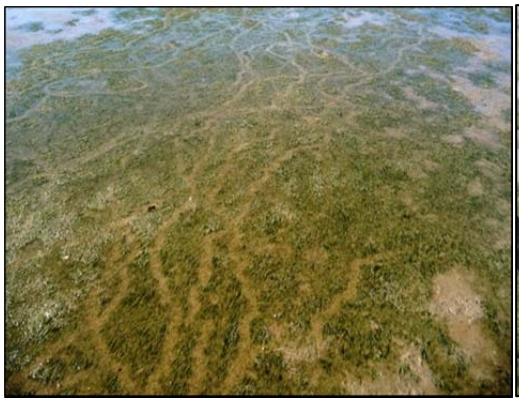


Dugong Feeding Ecology and Habitat Use on Intertidal Banks of Port Curtis and Rodds Bay – Interim Progress Report 2015

Michael Rasheed, Damien O'Grady and Emma Scott

Report No. 16/14

March 2016

Dugong Feeding Ecology and Habitat Use on Intertidal Banks of Port Curtis and Rodds Bay – Interim Progress Report 2015

A Report to the Gladstone Ecosystem Research and Monitoring Program

Report No. 16/14

March 2016

Prepared by

Centre for Tropical Water & 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. (2016). Dugong Feeding Ecology and Habitat Use on Intertidal Banks of Port Curtis and Rodds Bay – Interim Progress Report 2015. 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 & Aquatic Ecosystem Research (TropWATER) Publication 16/14, James Cook University, Cairns, 30 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 & 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.

Table of Contents

PROGRESS & KEY FINDINGS

This interim report details the progress of the first 8 months of the project: *"Dugong Feeding Ecology and Habitat Use on Intertidal Banks of Port Curtis and Rodds Bay"*. Results presented are preliminary as many aspects of the project are still in progress and methods and analysis will continue to be developed during 2016. Key highlights and progress include:

- Five seagrass areas within Port Curtis and Rodds Bay were identified for seasonal assessment of dugong feeding activity (3 in Port Curtis; 2 in Rodds Bay).

- A new method was 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 5cm pixel resolution suitable for identifying dugong feeding scars or trails.

- Successful quarterly mapping campaigns were conducted in May, August and November in 2015

- Three additional assessments were made at a smaller sub-set of areas between the August and November surveys to assess longevity of feeding trails between quarterly surveys (13th & 28th September and 28th October).

- Good progress has been made developing an algorithm to allow reliable automated extraction of dugong feeding trails from the imagery. The algorithm will undergo further refinements over the next 12 months to improve automated extraction and accuracy of extractions.

- Strong seasonal differences in seagrass and levels of dugong feeding activity were apparent from initial analysis of the imagery.

- A more detailed analysis of a section of the South Trees seagrass meadow was conducted for this interim report using a highly "trained" algorithm and has shown:
    - Seagrass coverage varied by 65% between quarterly surveys peaking in coverage between September and November 2015.
    - Dugong feeding trails were observed in the area of investigation in every survey but while initially following the increasing seagrass trend, declined in density in October & November 2015. This coincided with the appearance of new seagrass areas at nearby Wiggins Island with high levels of dugong feeding activity recorded.
    - An analysis of the longevity of feeding trails indicated that trails were unlikely to persist between quarterly sampling events with the majority of trails indistinguishable between 4 and 8 weeks after being first recorded.
    - Feeding trails that occurred in areas where all seagrasses were subsequently lost had a longer residence time.

- The small scale analysis provides a template for an expanded spatial analysis and comparison of seasonality and spatial change in the utilisation of intertidal seagrass meadows by dugong in Port Curtis and Rodds Bay during the second year of the project. During 2016 we will also
    - Continue developing the algorithms for seagrass and dugong feeding trail extraction
    - Investigate a new drone based platform for creating orthomosaics
    - Examine lower level flights with resultant smaller pixel sizes (2.5cm)

- Detailed analysis and full reporting will be presented on all of the investigation areas and the additional year of quarterly assessments in the final report at the completion of 2016 sampling (report due in February 2017).

1 INTRODUCTION

1.1 Background

Dugong (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 seagrass meadows as their diet consists predominately of seagrass, although they also eat macro-invertebrates and algae (Marsh et al. 2011). Since the 1960s, dugong populations have been declining rapidly and are currently listed as 'vulnerable to extinction' by the International Union for the Conservation of Nature (IUCN). Within the Gladstone region, a relatively small dugong population utilises a large area of Port Curtis and Rodds Bay where seagrass meadows are extensive enough to 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 Gladstone area is a potentially important connecting habitat for dugong populations in southern Queensland (Sobtzick et al. 2013). While there are some reports of preferential selection of seagrass by dugongs in Queensland (Preen 1992) the seasonality of dugong feeding within these meadows 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). Feeding preference may be related to seagrass abundance (relative biomass) and nutritional quality of their food (Searle et al. 2005, Sheppard et al. 2010). The feeding trails left behind when dugong excavate for seagrass offer some of the best physical evidence of feeding and habitat use and are common in preferred foraging habitats in the intertidal regions of the GBR (Preen 1992).

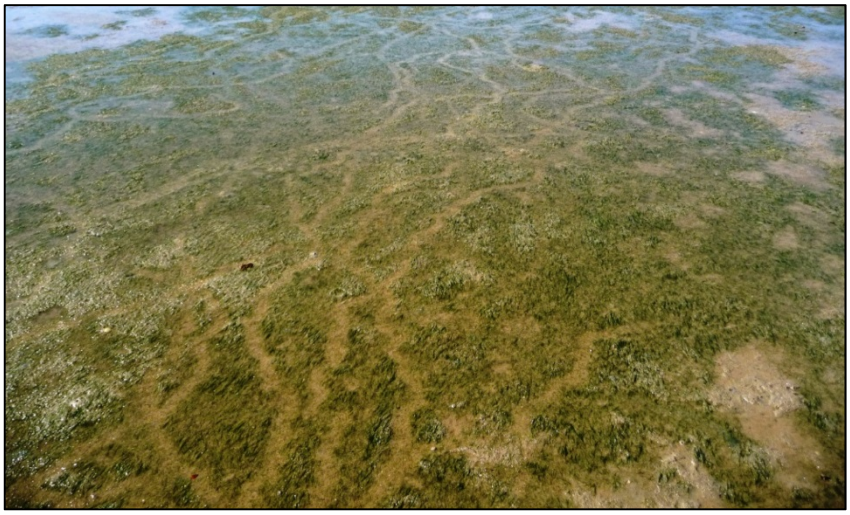

Figure 1: Dugong feeding trails on intertidal seagrass meadows.

1.2 Knowledge gaps

Little quantitative information exists on how dugong utilise seagrass habitats over time within Port Curtis and Rodds Bay. Seasonal analysis of dugong feeding trails will provide valuable information and inform the timescale that seagrasses take to recover from these disturbances. This information will increase understanding of the role of dugong in ecosystem processes and provide fundamental information for the effective conservation of these important animals and their habitats.

1.3 This report

This interim report details the progress of the first 8 months of the project. Results presented are preliminary as many aspects of the project are still in progress and methods and analysis will continue to be developed during 2016.

This report explains the process we have followed to map Dugong Feeding Trails (DFTs), from planning, data acquisition, processing through to the extraction of DFTs from the imagery. We present example results from the broader scale assessments and then focus on one DFT cluster area at one of the sites, showing the mapping of seagrass and DFTs and how they vary over time. This demonstrates the type of analysis that will be performed over the entire site and study area by the end of the project.

The extraction of DFTs from aerial survey data involves the following stages:

1. Acquisition of aerial photography by helicopter survey;
2. Processing of the raw photographs into contiguous georeferenced images (orthomosaics), using bundle-adjustment and photogrammetry methods;
3. The use of feature extraction techniques to develop a semi-automated algorithm to scan the orthomosaics and identify likely DFTs;
4. Using the algorithm to extract DFTs from each of the orthomosaics;
5. Analysis of the results.

Most of the work to date has been concerned with Stages 1 to 3.

2 METHODS

2.1 Sites and surveys

TropWATER has historical data from over a decade of annual and sub-annual monitoring of seagrass sites in Port Curtis and Rodds Bay where dugong feeding trails have been qualitatively recorded during surveys (Bryant et al. 2014 a & b). This information was used to narrow the focus to five target seagrass meadows where dugong feeding trails (DFT's) 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 1)). 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 1)).

Surveys were carried out at the 5 target meadows in 2015 as part of quarterly assessments on 16th May, 29th August, and 27th November (Table 1).

In order to examine the longevity of trails between the regular quarterly surveys a sub-set of smaller areas where DFT's had been mapped in August were reassessed on three occasions corresponding to 2, 4 and 8 weeks after the August 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 the feature extraction algorithm with different Ground Sampling Distances (GSDs, which equate to pixel sizes), an area of the South Trees meadow was also carried out at a lower altitude of 90 metres during four of the 2015 surveys (Table 1).

Table 1: Dugong Feeding Trail sampling meadows and dates surveyed for the quarterly and interim surveys in 2015

			Interim Surveys Sub Sampled			
	May	**Aug**	**Sept 13**	**Sept 28**	**Oct 28**	**Nov**
Rodds Bay North	✓	✓	✓	✓	✓	✓
Rodds Bay South	**N/A**	✓	✓	✓	✓	✓
South Trees (high- 175m)	✓	✓	✓	✓	✓	✓
South Trees (low – 90m)	**N/A**	**N/A**	✓	✓	✓	✓
Pelican Banks	✓	✓	✓	✓	✓	✓
Wiggins Island	✓	✓				✓

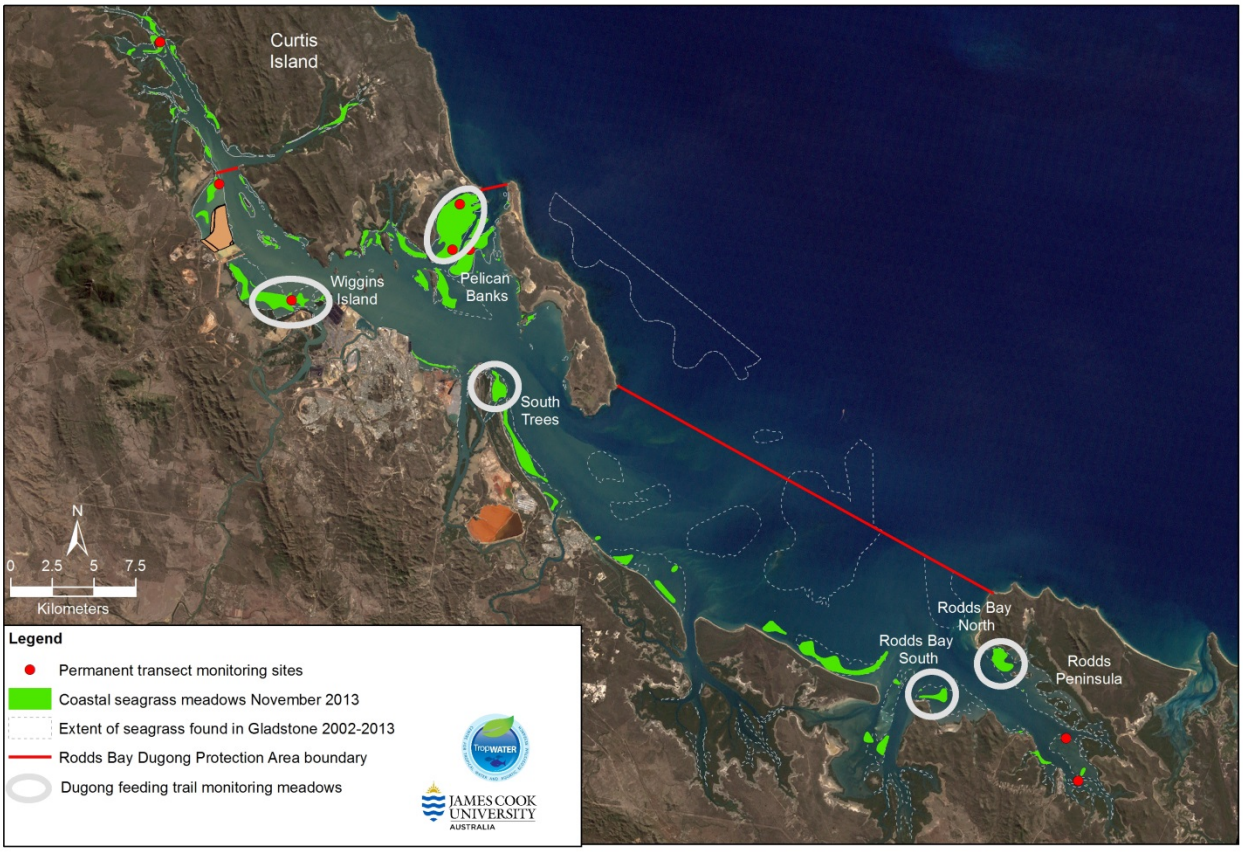

Figure 1: Port Curtis and Rodds Bay area showing location of the 5 meadows assessed quarterly for dugong feeding assessments.

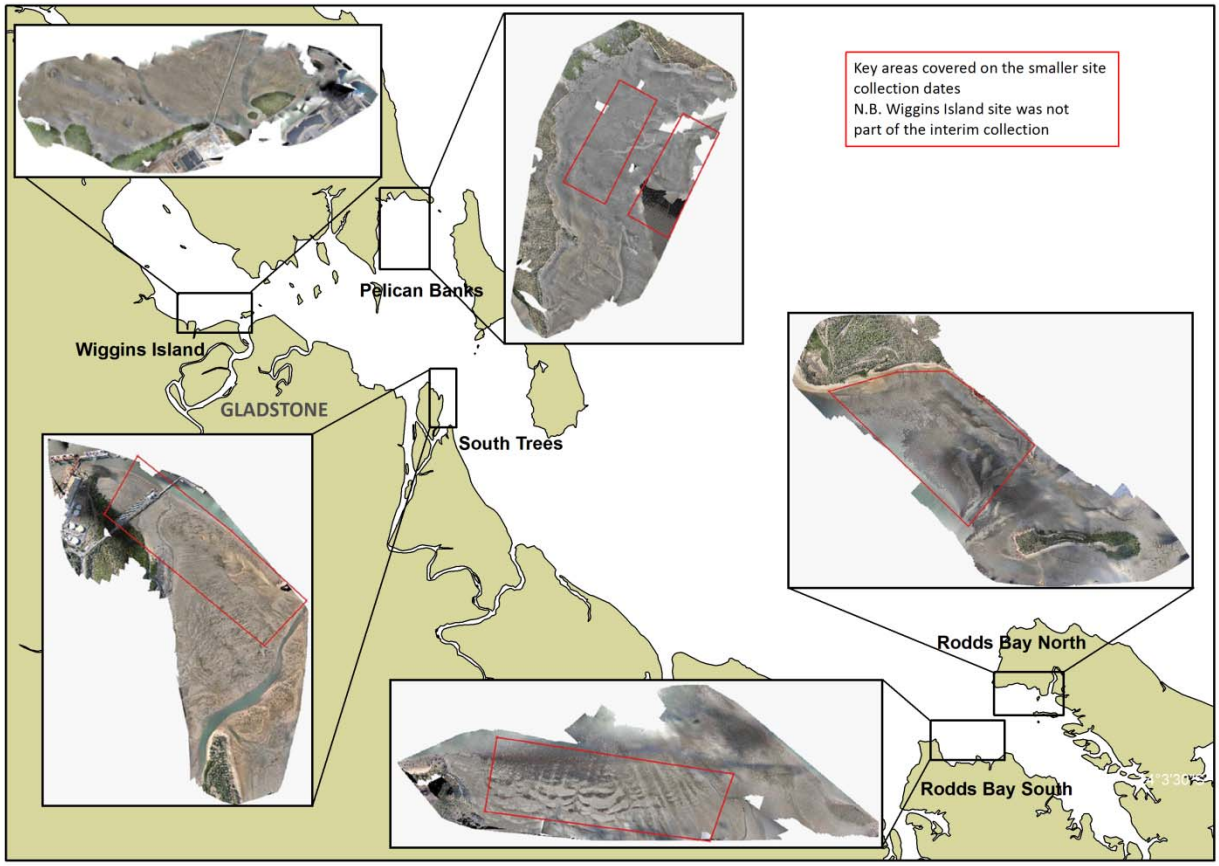

Figure 2: Location of the five seagrass and dugong feeding quarterly monitoring meadows as well as areas sub-sampled for the three interim surveys assessing feeding trail longevity between quarterly sampling (13th & 28th September, & 28th October). Target areas denoted by red polygons.

2.2 Data acquisition and processing methods

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 (Robinson 44 or Bell Jet Ranger) carried the researchers and the camera (Figure 3). On board were the pilot and two researchers. One researcher sat in the front and assisted the pilot to navigate the prescribed course, keeping an eye that the height and speed were maintained within the required parameters (see below) and liaising with the camera operator. The camera operator occupied the rear seat behind the navigator. The port-side door was removed for the survey flights, 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 4 (a)), a 12 megapixel digital SLR, with an 18-55mm Canon Zoom lens. The team have previously utilised these cameras in conjunction with intervalometers for remote firing and were therefore familiar with their capabilities and limitations, and were able to prepare and take three identical cameras with identical settings, for redundancy in the field. The Canon cameras are also 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 and wireless actuator, shown in Figure 4 (b). Images were stored on Scandisk Ultra C10 32GB SD cards, with a transfer rate of 40 MB/s.

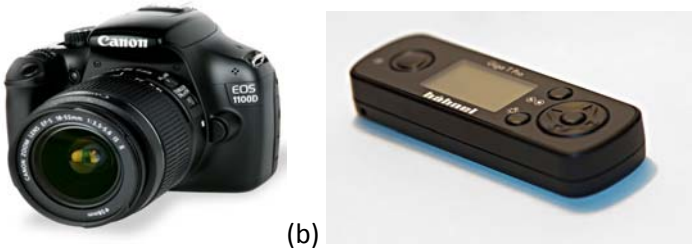

(a) (b)

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

Camera staff

The camera was positioned during the survey below the level of the helicopter skids. This was achieved using a purpose-made aluminium staff (Figure 5), built to firmly hold the camera and the actuator, and to safely contain the wiring between the two. The staff included a padded bracket which transferred the weight of the system to the floor of the helicopter next to the operator, 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

Computer

A Dell Toughbook was used by the researcher sitting in the front next to the pilot. 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, once the prescribed flight pattern was completed, to 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 as described above.

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 repeat survey schedule. It was then 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 the azimuth
- Image overlap between 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. For this reason, 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 dugong trails 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=142m. For a 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, we require a leg distance of 128 m.

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 leg distance for the actual DFT surveys would be reduced from 128 m to 100 m.

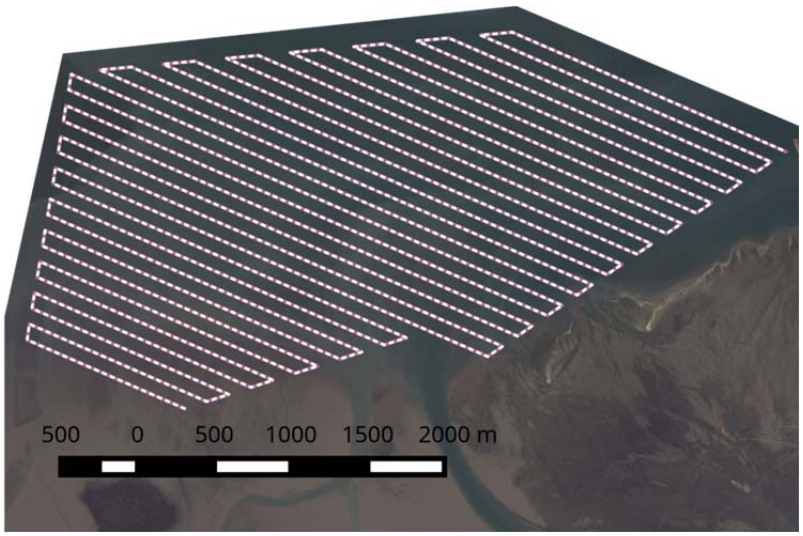

Figure 6: Wiggins Island flight pattern

Gladstone flight plan

Flight patterns were designed to meet the configuration described above, for each of the 5 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. Initially, the alignment of the legs to be flown was chosen in accordance with the direction of the known prevailing winds. After the first survey, it soon became apparent that the biggest advantage was to be gained in optimising the alignment to suit the longest legs, with the fewest turns. Despite cross winds for some sites the pilots were able to maintain the desired path successfully.

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 as described previously, but *low* denotes a flight done at half of the usual height and half of the usual speed. This was done to examine the benefits of having twice the spatial resolution in the resultant imagery, to determine to what extent DFTs may be more easily extracted at such levels of precision, in order to inform the planning of future surveys.

Flights generally took off from Gladstone Airport two hours prior to low tide. The strategy was to fly longitudinal stretches, following the direction of the tide: i.e. starting along the land-side working towards the sea on a falling tide, and vice versa on a rising tide, thus optimising the area of dry banks surveyed. The camera actuator, set at a frequency of one shot per second, was started immediately prior to the commencement of the first leg of each meadow, and stopped on completion of the last leg, with periodic checking throughout the survey to ensure that the camera was being triggered properly.

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 is identical to the time on the GPS units, which is established from the GNSS network to a very high degree of precision.

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, to produce a quality report to ensure that sufficient coverage had been achieved for the bundle adjustment process to work. This allowed sufficient time to take remedial action the following day if discrepancies arose.

[1] http://time.is/

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 computationally intensive, and so much advantage is gained by the manual elimination of obviously irrelevant photographs, mainly shot over water and mangroves at each end of the survey, plus those at the extremities of each leg. This manual exercise was carried out only in selected cases where the photograph set contained a particularly high proportion of such images

Correlation of GPS data

As the camera used for the surveys was not GPS-enabled, it was necessary to tag each photograph with its corresponding horizontal and vertical position as a separate task. This was done using open source (orphaned) software called *GPS Correlate*[2]. The software reads the time tag on each image, and looks for the closest time on the GPS track, 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. The software allows the user to feed in any known lag between the camera time and that of the GPS, which in our case had been recorded for each survey.

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 7. 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 mosaicked image. A snapshot of the point cloud at South Trees is shown in Figure 7. The bundle-adjustment process, taking the data from geo-tagged photographs through to an ortho-mosaic and DSM, was carried out using the proprietary software Pix4Dmapper Pro[3].

(a) 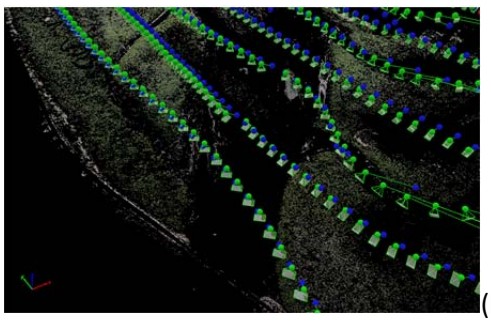 (b) 

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

[2] https://github.com/freefoote/gpscorrelate
[3] https://www.pix4d.com/product/pix4dmapper-pro/

Coregistration

Once the orthomosaics have been produced, then in order for the time series of data for each location to be analysed spatially, it is necessary to ensure that the geo-positioning of each mosaic corresponds as closely as possible. The process so far has ensured that the position of imagery from different times corresponds spatially to within a few metres, which aids the coregistration process. Coregistration involves the identification of invariant features that may be located on each image, which allows GIS software to determine a polynomial function that may be applied to the coordinates of each pixel in the target image in order to bring 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.
The software facilitates the process by displaying the target and reference images side by side, keeping their location in sync according to their approximate location, while the operator selects matching pairs of points on each image, as seen in Figure 8. The more points identified across the wider extent of the images, the more accurate the coregistration process.

For each location, a "base" image was chosen, against which each of the other images was registered, so as to minimise the propagation of spatial errors.

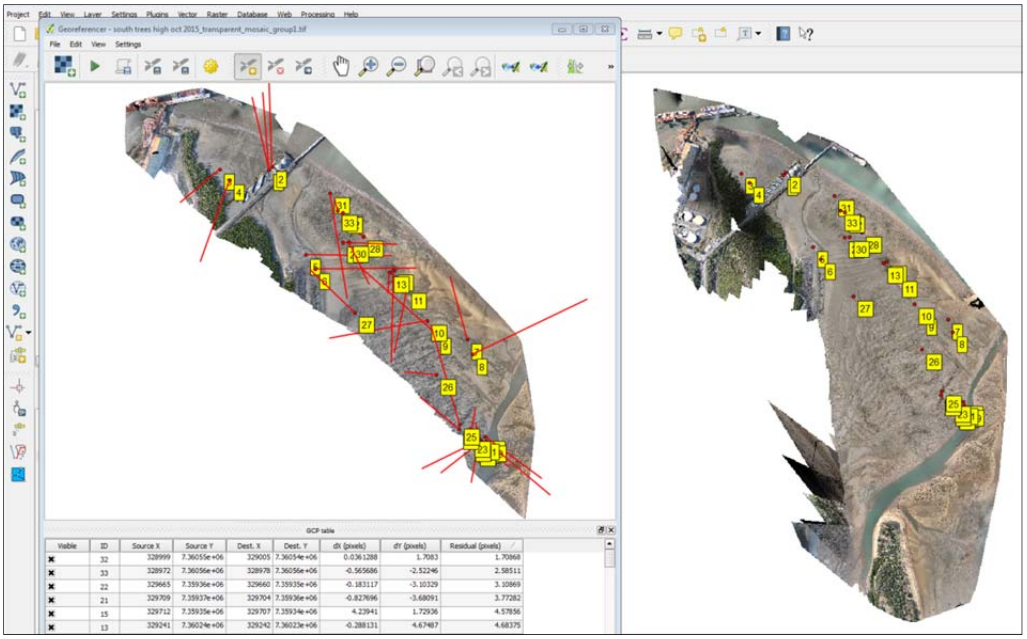

Figure 8: Screen shot of the QGIS coregistration interface

2.3 Feature Extraction

2.3.1 The algorithm

The overall process for the extraction of DFT's from the imagery involves two stages. Firstly, areas where DFT's are present must be located. Secondly, the trails themselves must be delineated and extracted from within those areas. This second stage is a classification process which is automated using a machine-learning algorithm. This itself also has two stages: training and prediction. Training machine-learning classifiers to identify clusters based on multi-band imagery often simply involves an operator digitising polygons over a sample image. This is problematic in our case, as the DFT's, by their very nature, are long and thin, and make for an extremely difficult and error-prone exercise. This coupled with the need for as much training data as possible, leads to the necessity to help automate even the training process, to some extent. This was done by carrying out preprocessing of the imagery to first extract all "long thin" features of the right scale which may possibly be DFT's. These features were then converted to vector objects. This

simplifies the training process considerably, whereby the operator needs just to click on an object and indicate whether the object is, or is not, a DFT.

The second stage involves the derivation of spatial layers representing all of the parameters that the classification algorithm can use to identify DFT's. These can be values for each of the red, green and blue image bands, as well as textural measures characterising the variation of those values within a shape. These characteristics vary with scale, and the relevant parameters are therefore calculated at many scales. Once this is done, all of these parameters are associated with each of the objects, DFT's and non-DFT's. These form the basis on which the algorithm tries to predict whether or not a shape is a DFT.

The algorithm development process is outlined in Appendix 1

2.3.2 Process used for test area in this interim report

For this interim report we selected an area of the South Trees meadow to examine in more detail as a demonstration of the type of outputs that can be generated for the final product at the end of the project (Figure 9). The area was also selected as it contained a mixture of DFT's in current seagrass patches as well as older DFT's in areas where seagrass had been subsequently lost. It was also within an area where the 3 interim surveys were conducted between quarterly sampling so that an indication of feeding trail longevity could be performed.

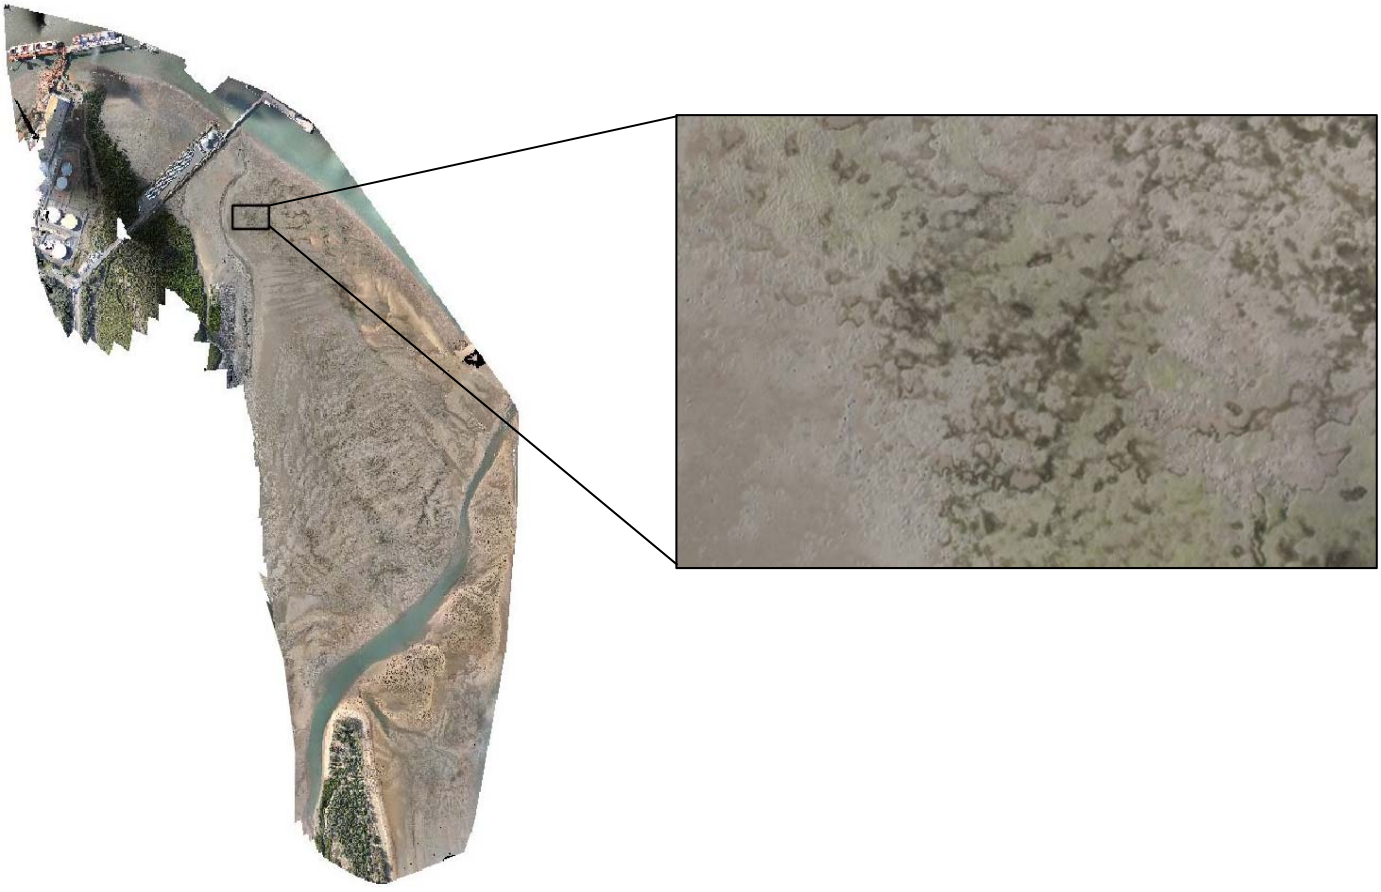

Figure 9: Area used for detailed examination of seagrass and dugong feeding activity for interim report

For this area we extracted the seagrass features at each time point in three density classes. This was done using a two-stage unsupervised classification process (eg Celik 2009). Firstly, a modification of the k-means algorithm was used to build pixel clusters based on similarities between their values for each of the red, green and blue spectral bands. It was found that the first three spectral signatures corresponded very

accurately to three distinct densities of seagrass that were visibly discernible in the image and seagrass presence confirmed from field observations. This was found to be consistent through the time series, despite radiometric differences in absolute values due to differing light conditions for each survey. This is likely due to the fact that the clustering algorithm identifies pixels based on the relationship between the mean and covariance of the values, rather than their absolute level, which tends to mitigate bias. The second stage of the classification process is to predict membership of each pixel to a seagrass category using a maximum likelihood algorithm, based on their proximity to the signatures developed in the first stage.

Dugong feeding trail features were extracted using a modified form of the full algorithm described in Appendix 1. To facilitate a more reliable extraction while the full algorithm is still undergoing development a more heavily trained extraction process was utilised than planned for the final product. This is substantially more labour intensive but achievable at the small scale of the interim assessment area. Feeding trails are generally visible as light patches of bare sediment through relatively darker patches of seagrass. A principal components transformation was carried out on the three image bands, and the first component was used for further analysis. To identify trails which stand out from their surrounds due to their higher reflectance values, a 7x7 pixel moving window convolution filter was passed over this raster (PCA_1), and the value of the central pixel of the resultant raster was calculated as the coefficient of variation of the first raster, where $C_v = \frac{PCA_1 - \mu}{\sigma}$. Following this, textural moving window filters were applied in order to highlight linear features. These were designed to match the scale of the majority of observed trails. Their effect was to "reward" pixel clusters with high values running through the centre and lower values to either side of a strip of high values, and to "penalise" clusters with higher values around the edges, rather than through their centre. From this, thresholding was used to finally segregate the linear features from the raster image.

3 RESULTS

3.1 Orthomosaic Generation

Individual photo images from 175m altitude were of sufficient resolution and quality to visually identify dugong feeding trails in the seagrass meadows as well as remnant feeding scars in areas where seagrass had been lost. Figure 10 shows an example of a typical single image from the Wiggins Island meadow from the November 2015 survey. Dugong feeding trails and a boat propeller scar are clearly visible within the image.

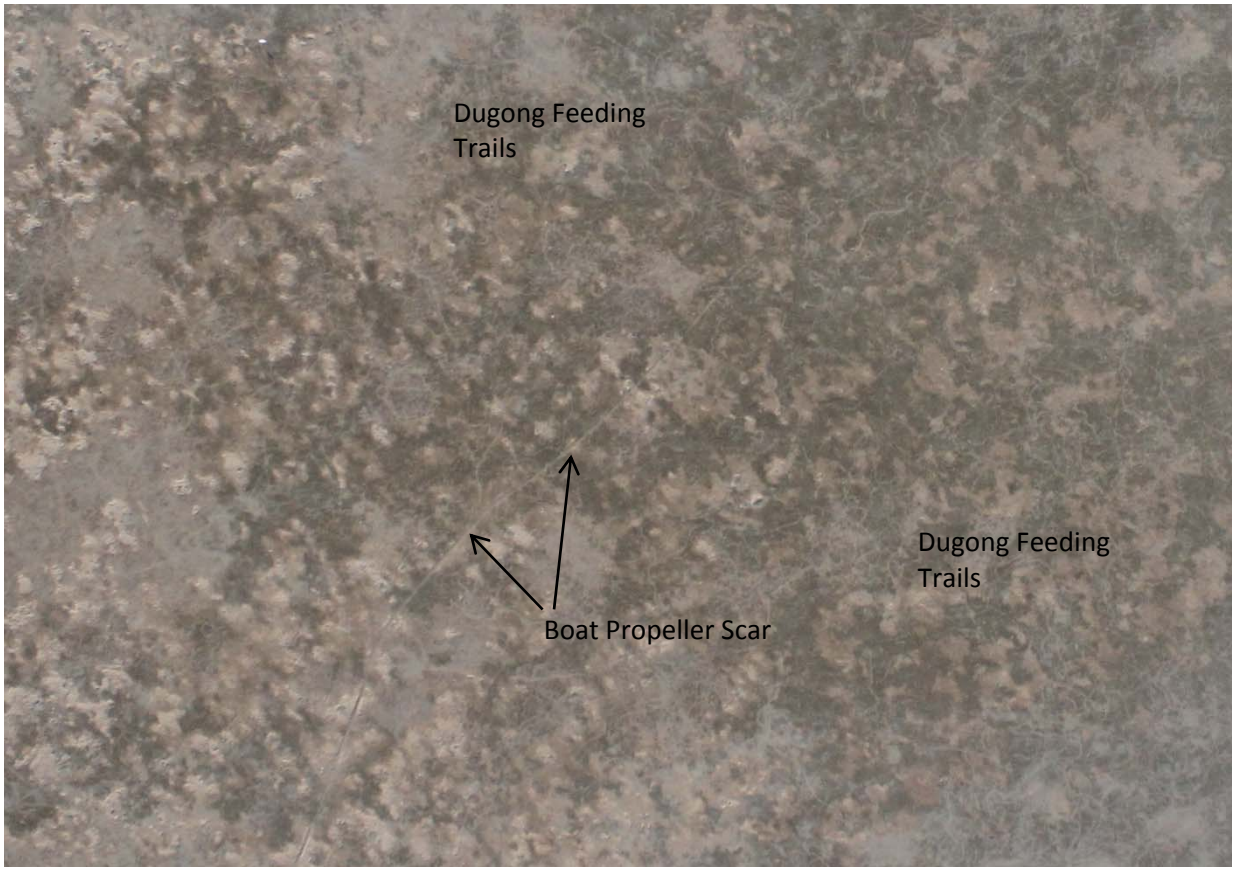

Figure 10: Section of a single photo image from November 2015 survey of Wiggins Island meadow showing dugong feeding trail and propeller scar detail

The images were used to successfully generate orthomosaics for the entire area of the investigation meadows. Orthomosaics have been generated for all sites for the first sampling (May 2015) and a complete time series of orthomosaics for three of the five sites - South Trees, Rodds Bay North and Rodds Bay South have also been completed (Figures 11-15). Processing of mosaics for the remainder of the Pelican Banks and Wiggins Island sampling events was ongoing at the time of this interim report.

Rodds Bay North

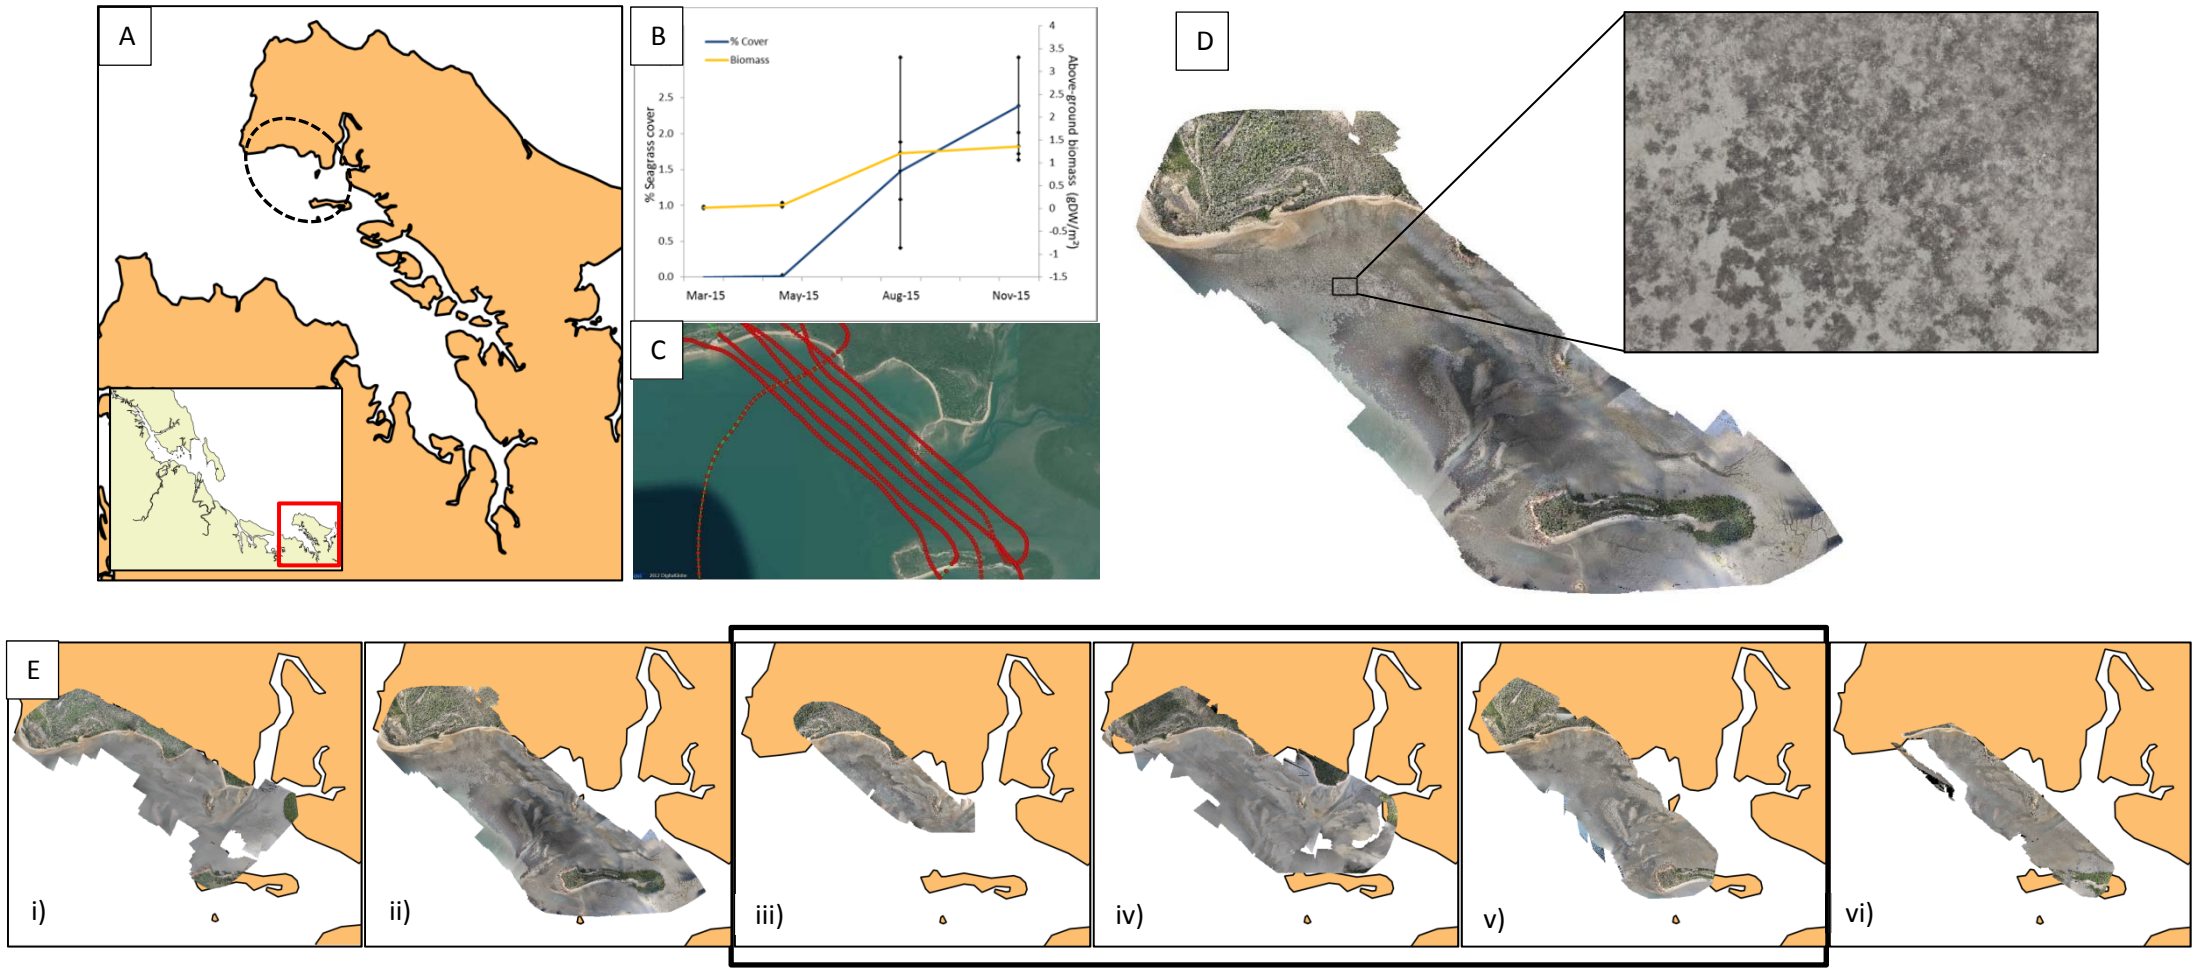

Figure 11: **A** - Rodds Bay North site location; **B**- changes in seagrass biomass and percentage cover March to November 2015 at nearby Rodds Bay permanent transect sites; **C**- individual image locations before processing the orthomosaic on Pix4D software; **D** - Example of the produced orthomosaic for the site with blow up showing dugong feeding trail detail
E - Orthomosaics produced for the Rodds Bay site across study dates in 2015; quarterly i) May, ii) Aug and vi) Nov, and the smaller interim sites iii) Sept 13, iv) Sep 28 and v) Oct 28

South Trees

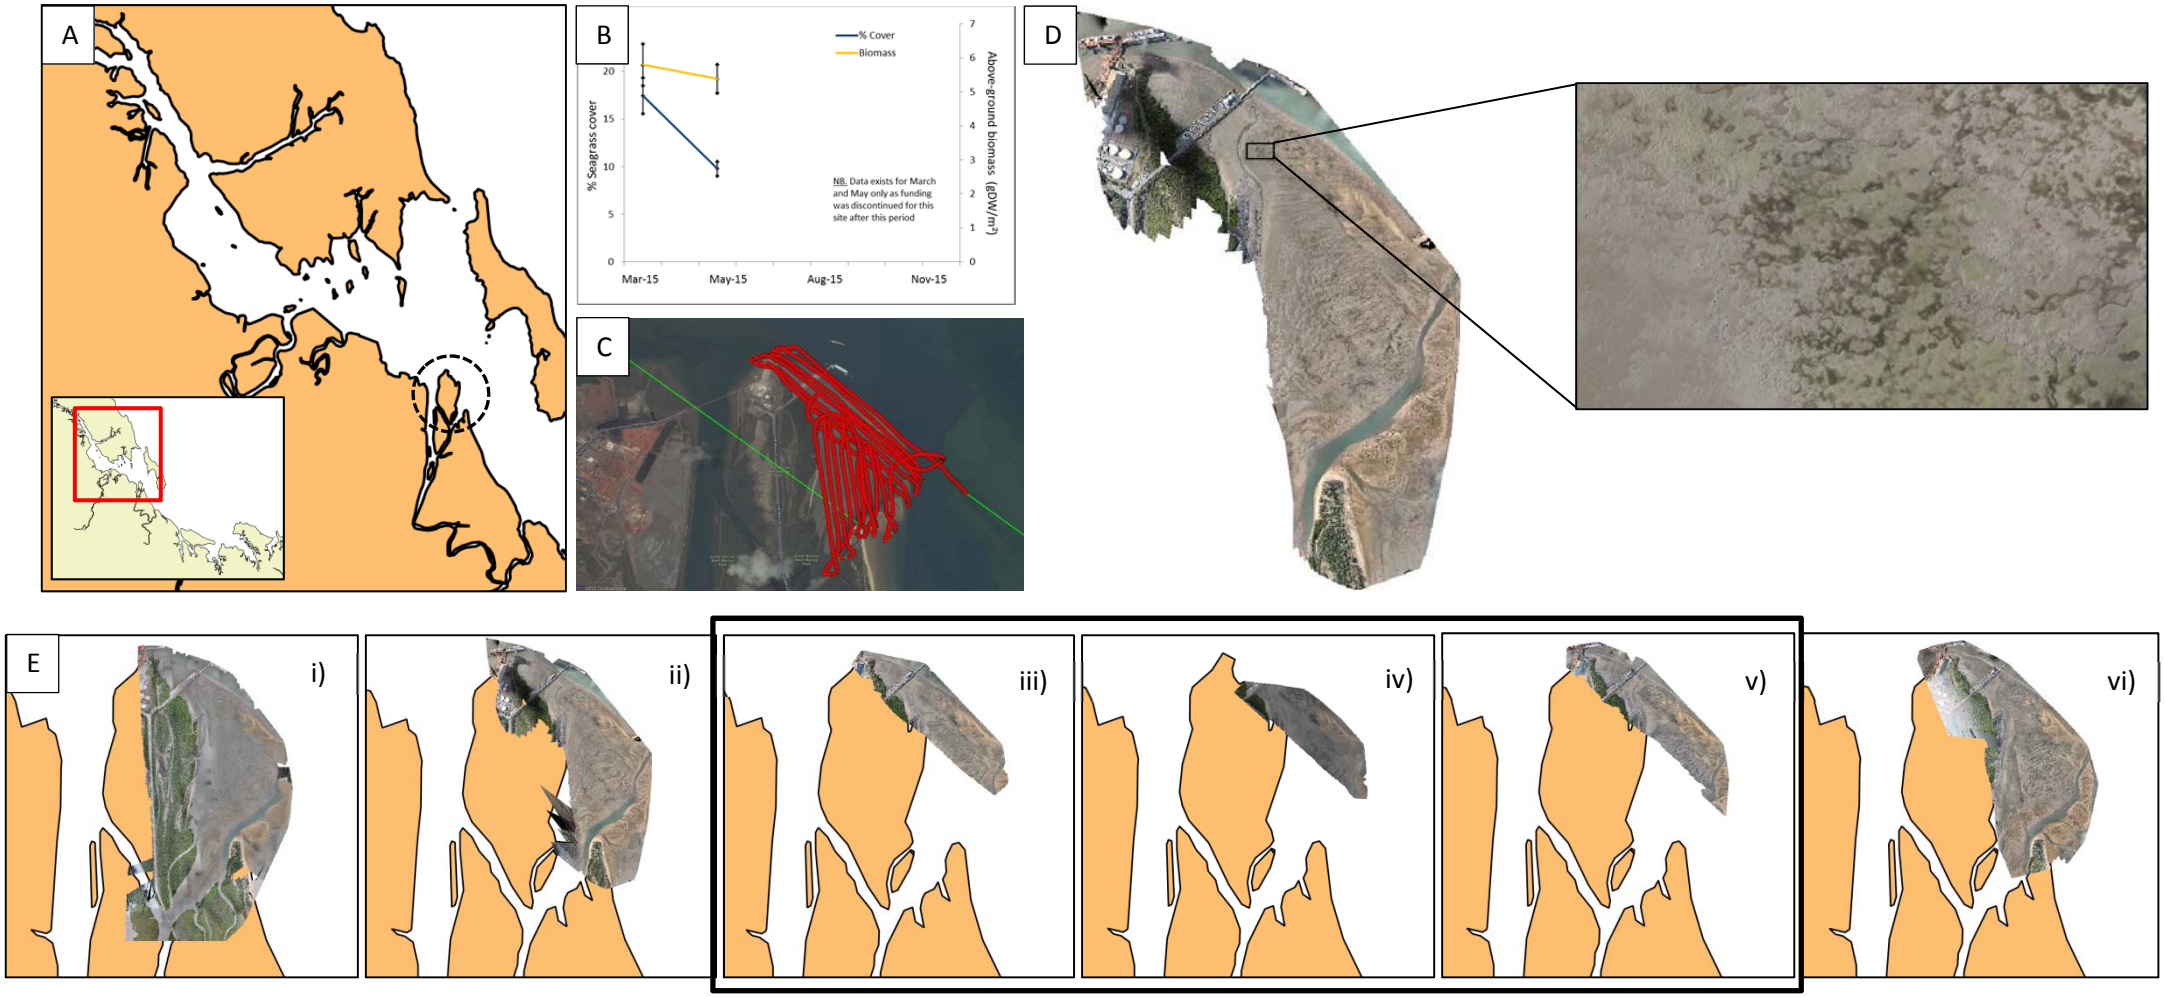

Figure 12: **A** - South Trees site location; **B**- changes in seagrass biomass and percentage cover March to November 2015 from permanent transect site (funding and data collection at site discontinued after May 2015); **C**- individual image locations before processing the orthomosaic on Pix4D software; **D** -example of the produced orthomosaic for the site with blow up showing dugong feeding trail detail.
E - Orthomosaics produced for the South Trees site across study dates in 2015; quarterly i) May, ii) Aug and vi) Nov, and the smaller interim sites iii) Sept 13, iv) Sep 28 and v) Oct 28

Rodds Bay South

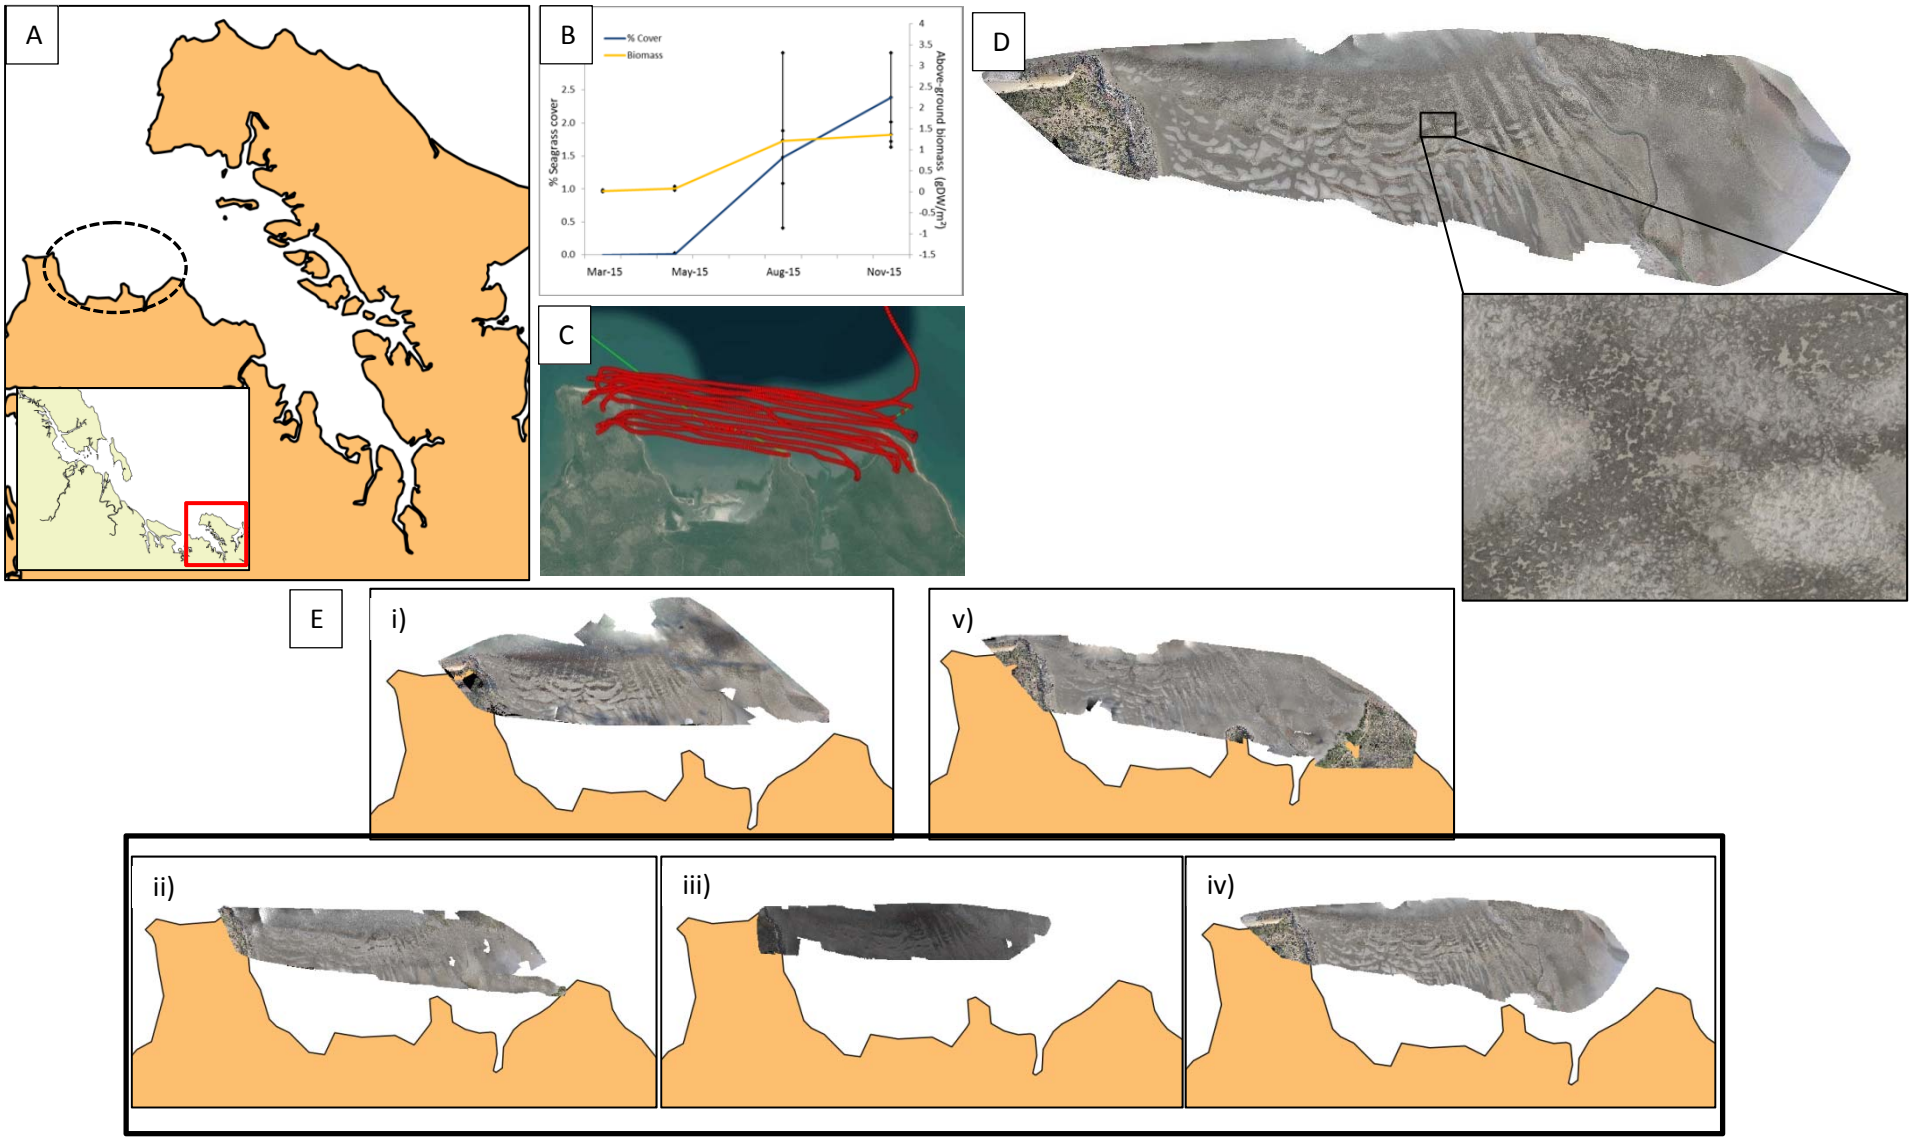

Figure 13: **A** – Rodds Bay South site location; **B**- Changes in seagrass biomass and percentage cover recorded over the period March 2015 to November 2015 from nearby Rodds Bay permanent transect site; **C**- individual image locations before processing the orthomosaic on Pix4D software; **D** -example of the produced orthomosaic for the site with blow up showing dugong feeding trail detail; **E** - Orthomosaics produced for the Rodds Bay site across study dates in 2015; quarterly i) Aug and v) Nov, and the smaller interim sites ii) Sept 13, iii) Sep 28 and iv) Oct 28

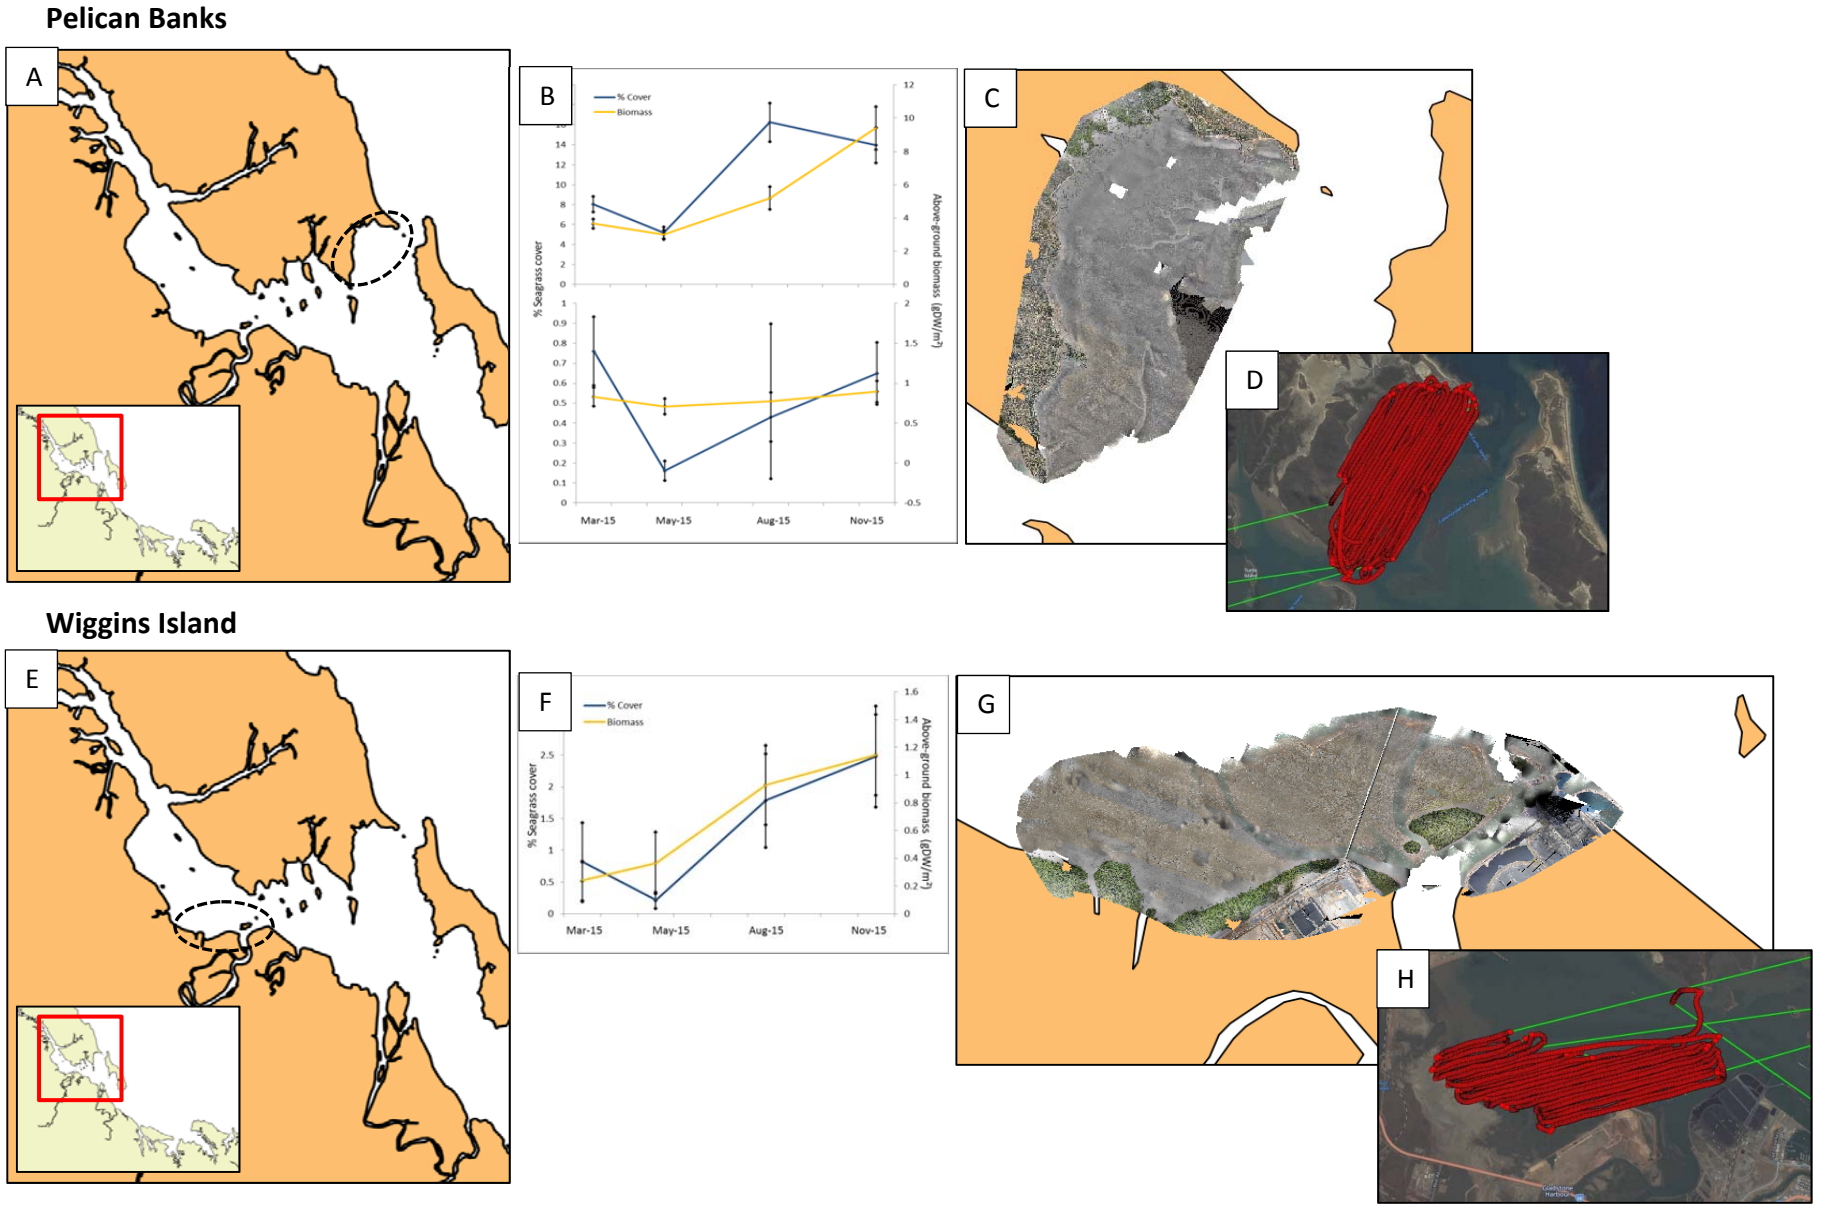

Figure 14: **A** - Pelican Banks site location; **B** - Changes in seagrass biomass and percentage cover March 2015 to November 2015; **C** – May 2015 Pelican Banks orthomosaic; **D -** individual image locations for Pelican Banks before orthomosaic processing. **E** – Wiggins Island site location; **F** - Changes in seagrass biomass and percentage cover March 2015 to November 2015; **G** – May 2015 Wiggins Island orthomosaic; **H -** individual image locations for Wiggins Island before orthomosaic processing

3.2 Dugong feeding trail extraction

Seasonal change in seagrass and dugong feeding

An examination of the South Trees investigation area using the interim image extraction algorithms showed marked patterns of change in both seagrass area and total area of Dugong Feeding Trails (DFT's) (Figures 15 & 16). Total area of seagrass increased by 65% between May and September with this area subsequently sustained through to the November sampling. Initially area of DFT's mirrored this trend, increasing between May and September, but there was a sharp decline in area of DFT's in October and again through to November. This 95% decline in area of DFT's between September and November occurred despite seagrass area remaining largely unchanged , indicating that for this small area at least, dugong feeding activity is not entirely coupled with seagrass coverage.

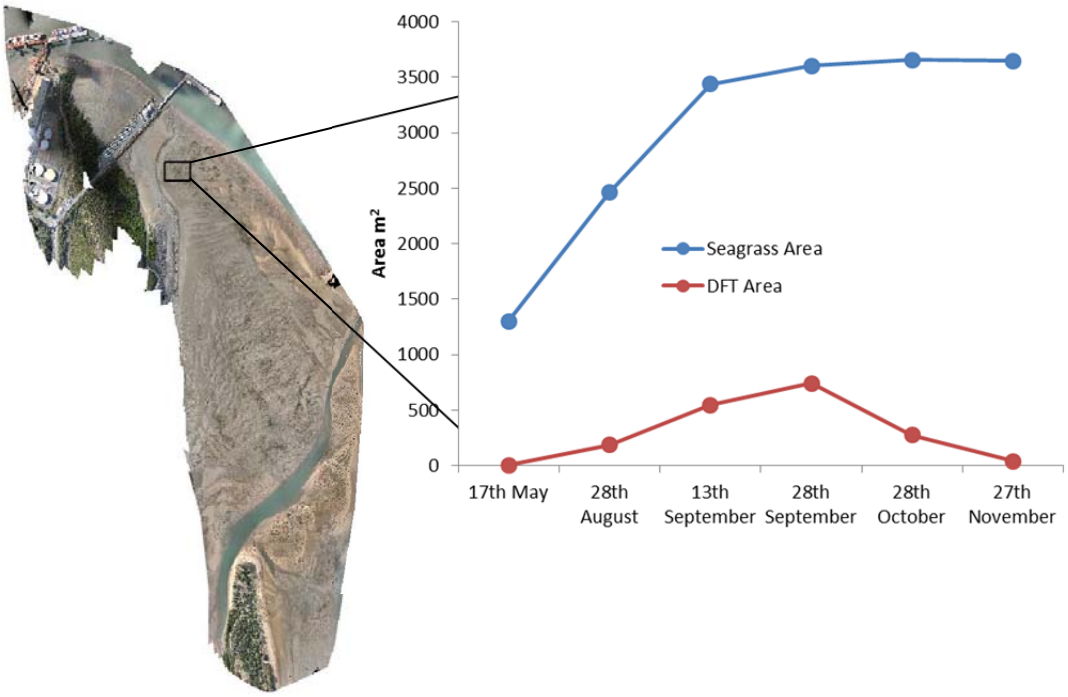

Figure 15: Change in seagrass area and area of Dugong Feeding Trails for the interim assessment area extracted from 2015 orthomosaic

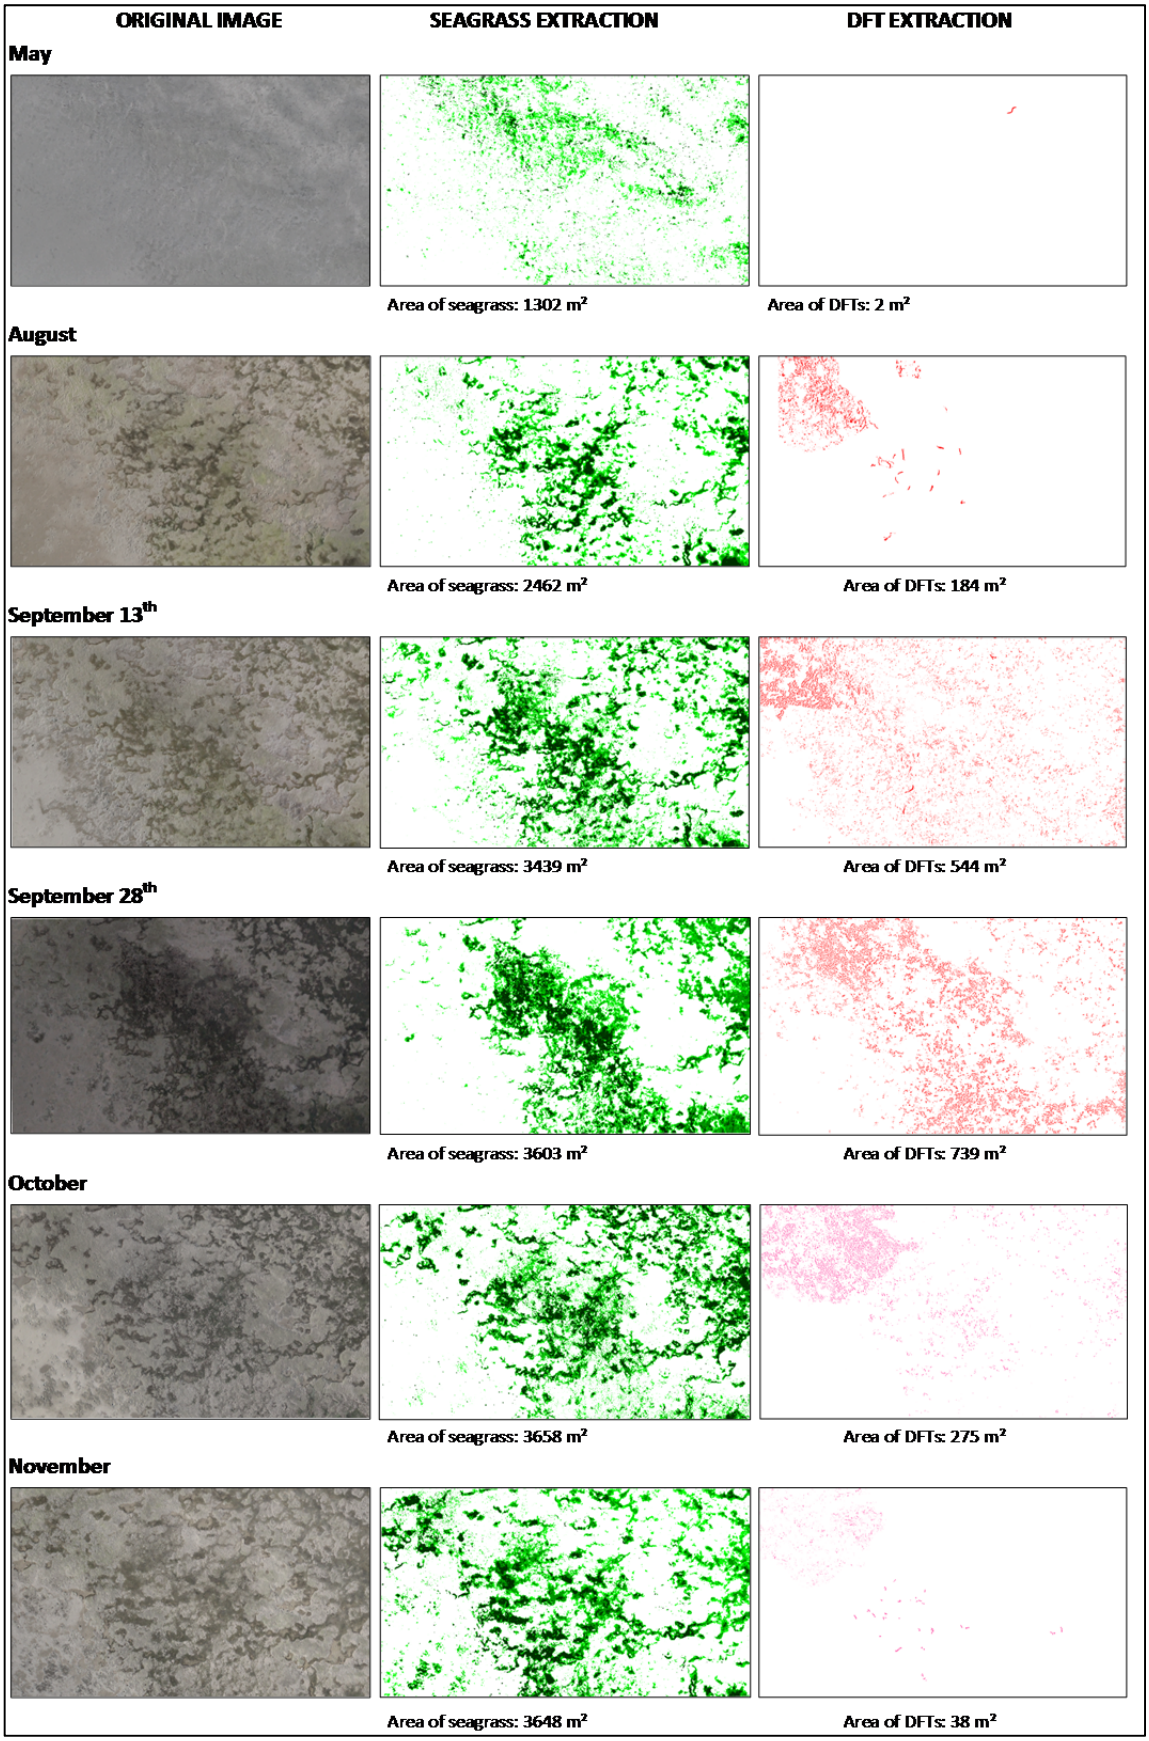

Figure 16: South Trees investigation area original images from the orthomosaics and extractions of seagrass and dugong feeding trails (DFT) using the interim algorithms for surveys between May and November 2015.

Longevity of feeding trails between surveys

Qualitative observations of a random selection of DFT's identified in August indicated that DFT's within seagrass were generally unable to be seen in the imagery from between 2 to 8 weeks later (Figure 17). DFT's that were found in unvegetated areas appeared to have a longer residence time with some trails evident 2 to 3 months after first being recorded.

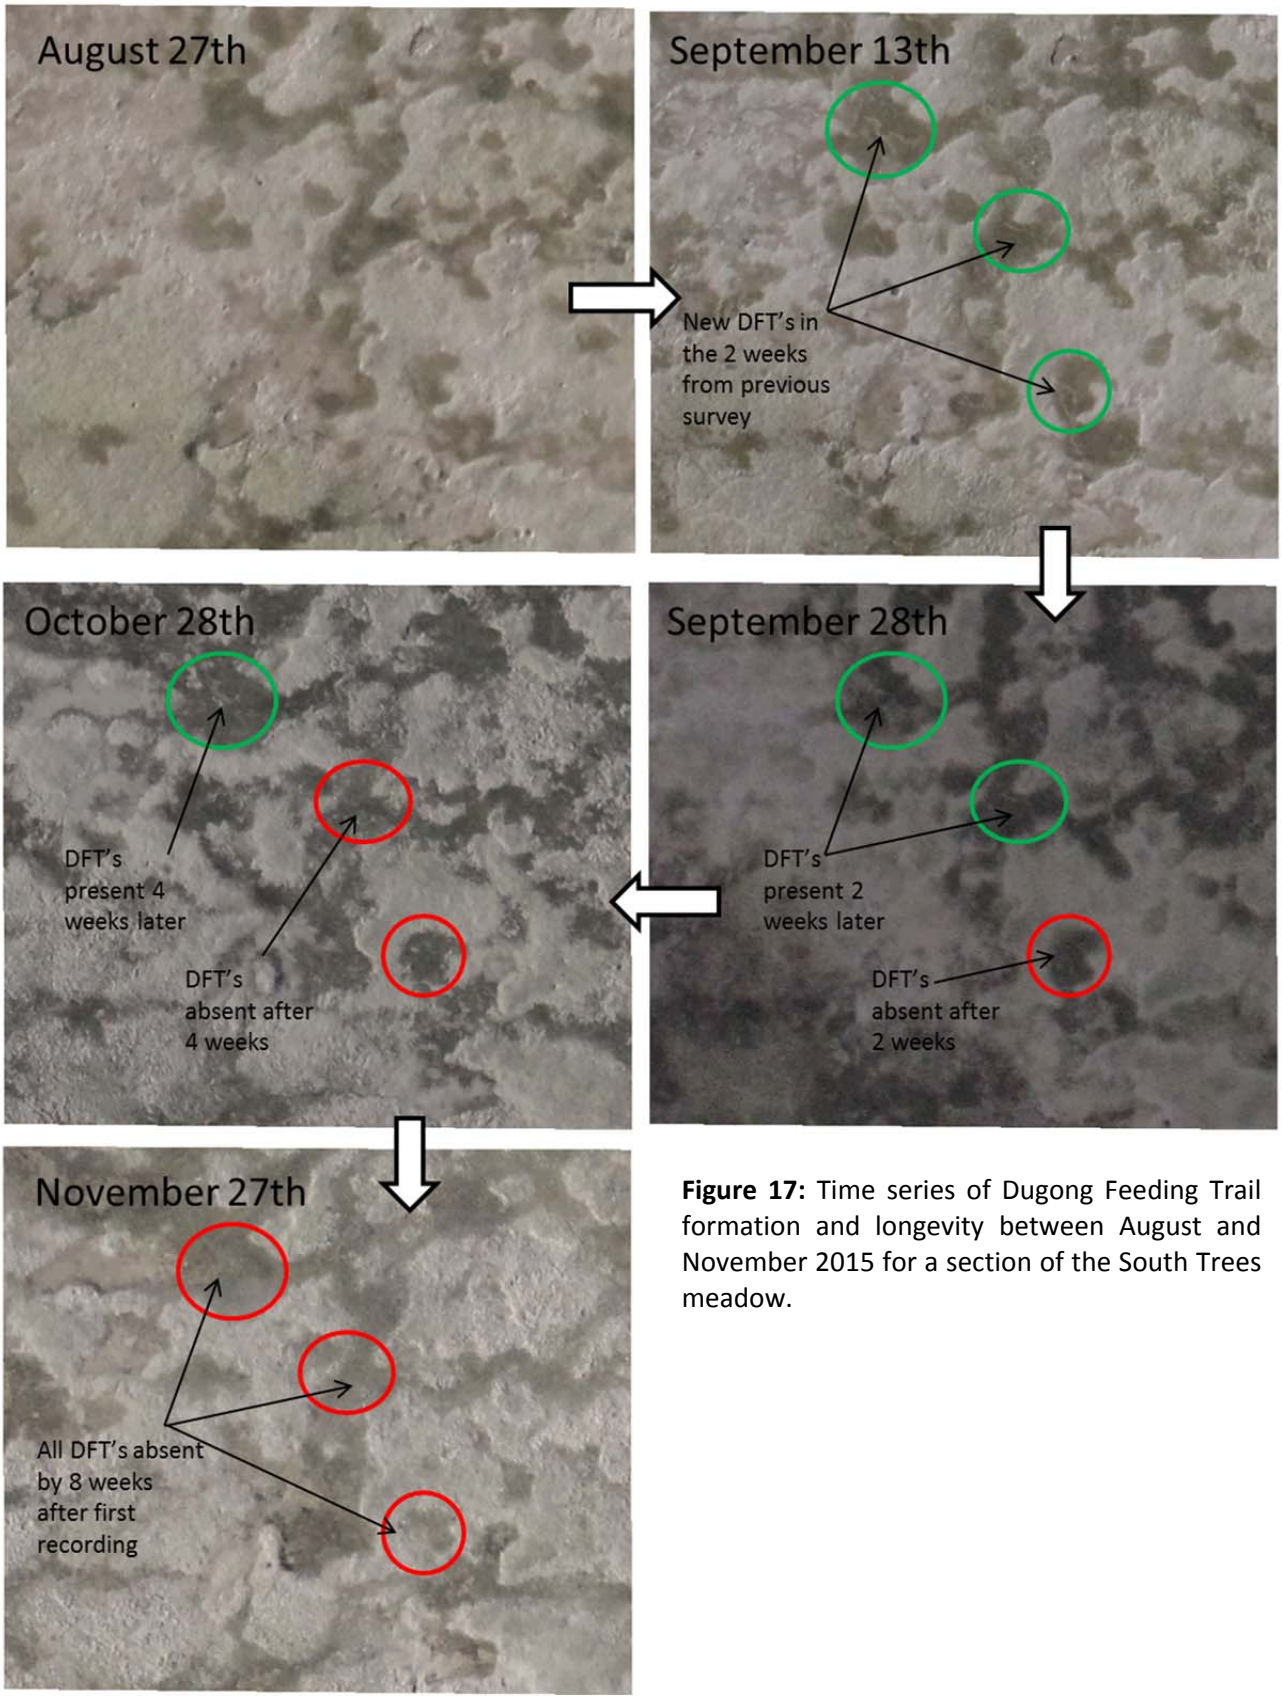

Figure 17: Time series of Dugong Feeding Trail formation and longevity between August and November 2015 for a section of the South Trees meadow.

Comparison with lower altitude flight

Imagery from the lower level flights (90m) of South Trees (from the 13th September onwards) has also been processed with orthomosaics generated. These flights led to images with a much reduced pixel size compared with the standard 175m altitude imagery, with pixels reduced from ~5cm to <2.5cm. Figure 18 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, as expected, formal analysis over the next 12 months will reveal if it confers an improvement in identification and extraction of DFT features.

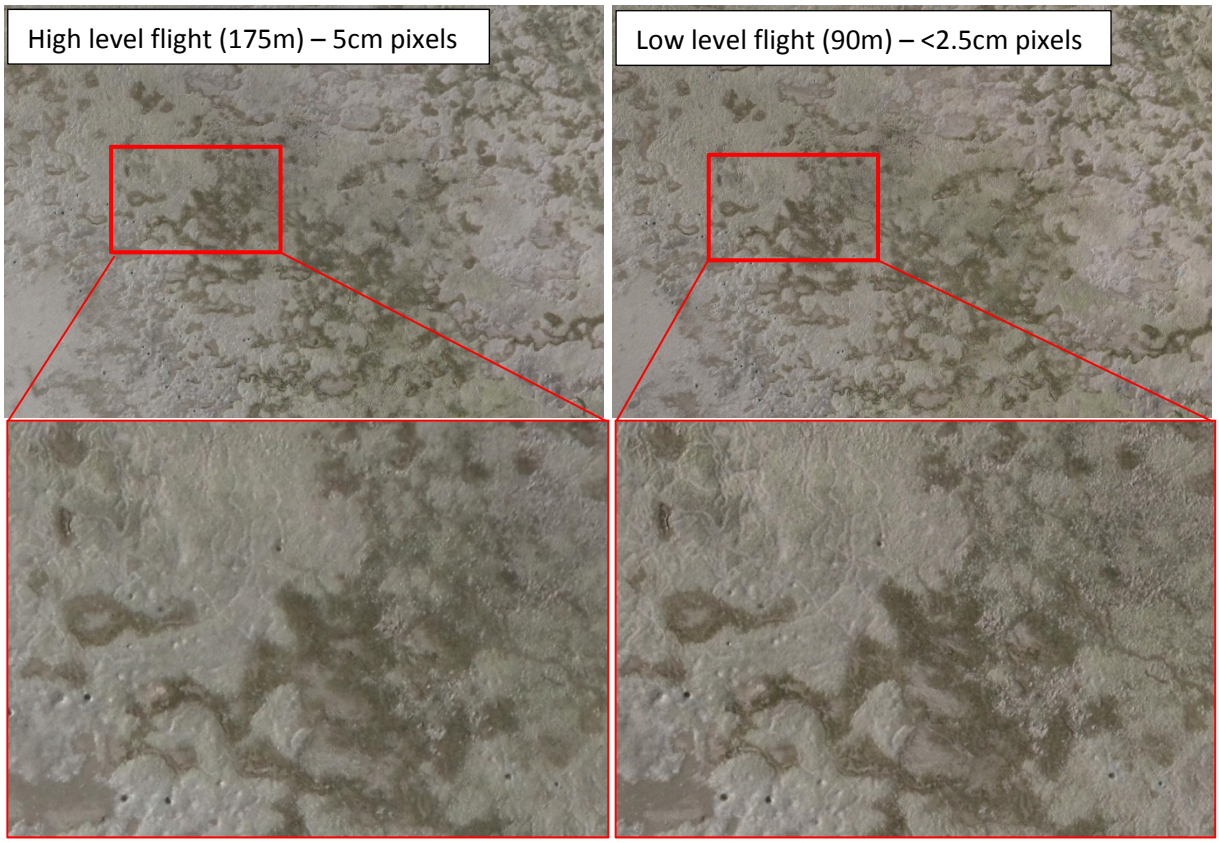

Figure 18: Comparison of high level (175m) and low level (90m) flights and resultant orthomosaic generation from the September 13th survey.

4 DISCUSSION

Interim Results

The first 8 months of this project have seen significant progress in developing a novel assessment technique for dugong feeding activity in intertidal seagrasses. As this is the first time this has been attempted it has required substantial investment of time into trailing and refining techniques both in field data acquisition and in post field processing and analysis. In the first year we have developed a successful low-cost technique to capture imagery of sufficient resolution to identify dugong feeding activity and seagrass changes at a relatively large scale (meadow). Refinements of the image analysis and data extraction algorithm generation will continue in 2016 but the interim assessments have revealed a positive demonstration of the value of the technique as well as some preliminary information on temporal and spatial change of dugong feeding activity in the intertidal areas of Gladstone.

A strong seasonal change in seagrass cover was evident in the interim investigation area of the South Trees meadow. This was in line with the expected seasonal cycle for seagrass in the region (Bryant 2014 a & b). The level of dugong feeding activity increased initially with seagrass increases between May and September but the subsequent drop off in dugong feeding, despite seagrass cover being maintained, suggests that dugong may have moved elsewhere to feed in the harbour. Substantial new areas of seagrass had expanded in the Wiggins Island region during October & November and there were large areas of dense feeding trails obvious in these areas from the field and imagery that has yet to be formally analysed. Clearly dugong had targeted this new area of seagrass to feed on at the time decreased feeding activity was observed at the South Trees site.

As we process and analyse more of the datasets for the five locations these spatial and temporal changes in feeding activity will become more apparent. Our initial qualitative scans of the imagery sets indicate that there are further broad-scale examples of the patterns formally analysed at the small scale interim investigation area.

From our initial assessment of dugong feeding trail longevity it appears that DFT's recorded in each quarterly survey are likely to be unique to those surveys especially in areas where coverage of seagrass is maintained. There may be some temporal/seasonal variation to this however, 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) and our assessments of longevity were performed during the high season for seagrass growth (August-November). Residence times during the senescent season, when seagrass asexual colonisation is slower, are likely to be longer but quantifying this would require a repeat of the longevity study during this time of year. For the regions that we have examined in more detail, there were no obvious trails from the May survey evident 3 months later in August (when they occurred in areas with seagrass present) so even in the senescent season it seems likely that DFT's between quarterly surveys are likely to be unique. Where DFT's have been recorded in areas where seagrass coverage was not maintained, their residence time appears to be substantially longer, as would be expected as trail longevity is determined by seagrass re-growth. It is possible that these trails remain in place between the quarterly sampling, and we will further quantify this during the second year of the study.

Algorithm development and improvements

The algorithm for seagrass and DFT feature extraction will continue to be developed and refined over the next year. However this is not a simple process with major challenges in training the algorithm requiring many re-runs as an iterative process.

The segmentation of DFT's may be seen as a classic feature extraction problem that requires the characterisation of DFT's in terms of texture, colour and shape. Spectral bands were limited to red, green and blue, and although some benefit may be gained using NIR to better differentiate photosynthetically active seagrass from other dark matter, the frequent covering of regions of the study areas with shallow water would likely render the NIR variable and unreliable. The nature of the environment in areas not

covered with seagrass does not provide much spectral variation, and recognisable features are generally identified by their texture rather than their colour. This makes the classification problem akin to the task faced in medical imagery in the segregation of tissues of different density in ultrasound, x-ray or magnetic resonance imagery, or in radar imagery.

Reliance on spectral characteristics are further confounded by the fact that DFT's run through different backgrounds or substrates. These might be different hues of seagrass (a function of species or density), or they may be found as scars in bare sand, visible by internal shadow or external highlights, caused by ridges of excavated material to the side of the troughs (and very sensitive to solar angle). It is no surprise, therefore, that the algorithm being developed for the broad-scale automated classification of DFT's relies most heavily on the shape and texture of the trails to identify them. This method also presents challenges. DFT's come in a range of lengths and widths. Furthermore, their length is often broken by intermittent sections of seagrass regrowth or shallower foraging. Our human brains are able to easily determine that a series of channels is likely to belong to the same broken trail, but this is a challenging undertaking for a computer, involving many parameters which need to be trained into the algorithm.

Perhaps the biggest challenge arises from the experience we have gained with DFT's this far. As we observe and map DFTs in seagrass, and see how the trails widen, disappear and form areas of prominent regrowth once again, it is easy to come to the conclusion that the trails are far more prolific than it would first appear. Certainly, in observing the time series of data, there are cases where we can clearly observe a trail widening (possibly by a retreat of the seagrass at its boundaries), and then being merged with other feeding activity which renders the cluster of trails into a patch of bare sand within the seagrass. In such a case, we can correctly label the bare patch as due to dugong feeding, but the question as to what we may assume with other similar patches of bare substrate requires more testing. The most important task in the classification of DFT's is therefore their definition, and this will only be determined by careful analysis of the progress of individual trails across the full data set. The work so far has therefore addressed the most conservative description, defined by a single trail of reasonably consistent width (around 200 mm) and with sufficient length to rule out other causes. Having established the metrics, the bulk of this work is the training of the algorithm, which involves finding as many examples of all of the different scenarios in which DFTs may be found, and labelling them, and then doing the same with "false trails" - identifying channels which are not DFTs, but which are likely to be confused as such. Only in this way can the thresholds and boundary conditions be clearly determined. The training process is underway, and with more aerial surveys yet to be carried out, we have the opportunity to validate of the resultant classification by observation, which will provide valuable feedback into the process.

As a trail becomes covered by seagrass regrowth, its shape becomes fragmented and narrows. The sensitivity of the DFT mapping process is therefore determined in part by the smallest measurable unit or feature discernible, which is, in turn, a function of spatial resolution or pixel size and definition. The chosen height of flight has been shown to be fit for purpose, and experimentation with camera settings has resulted in sub-optimal (though usable) results for one survey (September 28th), which has informed settings for all subsequent flights. Modifications to the camera bracketry and trigger mechanism settings, as well as the recent construction of a new housing for the camera to be positioned where vibrations may be minimised, are expected to help ensure that the quality of imagery produced so far will be at least maintained, and likely improved upon.

During 2016 in addition to the work on the algorithm we will be:

- Conducting 4 additional quarterly surveys
- Processing the back catalogue of imagery
- Undertaking formal analysis of dugong trail longevity
- Assessing cost/benefits of lower altitude higher resolution imagery
- Testing the applicability of the technique using a UAV mapping drone platform (customised 3DR X8-M – pictured)

Detailed analysis and full reporting will be presented on the project in the final report at completion of 2016 sampling (report due in February 2017).

5 REFERENCES

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

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

Celik, T (2009), Unsupervised Change Detection in Satellite Images Using Principal Component Analysis and k-Means Clustering, IEEE Geoscience and Remote Sensing Letters, vol. 6, no. 4, pp. 772-776.

Haralick, R. M., Shanmugam, K., and Dinstein, I. (1973). Textural features for image classification. IEEE Transactions on Systems Man and Cybernetics, SMC-3(6):610–621.

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

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

Rasheed, M. A. (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, M. A. (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, M. A., McKenna, S. A., AB, C. and Coles, R. G. (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.

Searle, K. R., Hobbs, N. T. and Shipley, L. A. (2005) Should I stay or should I go? Patch departure decisions by herbivores at multiple scales. Oikos 111: 417–424.

Sheppard, J.K., Marsh, H., Jones, R.E. and Lawler, I.R. (2010) Dugong habitat use in relation to seagrass nutrients, tides, and diel cycles. Marine Mammal Science, 26: 855-879.

Sobtzick, S., Grech, A., Coles, R., Cagnazzi, D. and Marsh, H. (2013) Status of the dugong population in the Gladstone area. A Report for Gladstone Ports Corporation Limited for Project CA 120017: Monitoring of Dugongs, Centre for Tropical Water & Aquatic Ecosystem Research (TropWATER) Publication, James Cook University, Townsville, 37 pp.

Westoby, M., Brasington, J., Glasser, N., Hambrey, M., and Reynolds, J. (2012). "Structure-from-Motion" photogrammetry: A low-cost, effective tool for geoscience applications. Geomorphology, 179:300–314.

6 APPENDIX – Summary of the algorithm development process

6.1 Segmentation of DFTs

Overview of how it works

The extraction of DFTs from the imagery involves two stages. Firstly, areas where DFTs are present must be located. Secondly, the trails themselves must be delineated and extracted from within those areas. This second stage is a classification process which is automated using a machine-learning algorithm. This itself also has two stages: training and prediction. Training machine-learning classifiers to identify clusters based on multi-band imagery often simply involves an operator digitising polygons over a sample image. This is problematic in our case, as the DFTs, by their very nature, are long and thin, and make for an extremely difficult and error-prone exercise. This coupled with the need for as much training data as possible, leads to the necessity to help automate even the training process, to some extent. This was done by carrying out preprocessing of the imagery to first extract all "long thin" features of the right scale which may possibly be DFTs. These features were then converted to vector objects. This simplifies the training process considerably, whereby the operator needs just to click on an object and indicate whether the object is, or is not, a DFT.

The second stage involves the derivation of spatial layers representing all of the parameters that the classification algorithm can use to identify DFTs. These can be values for each of the red, green and blue image bands, as well as textural measures characterising the variation of those values within a shape. These characteristics vary with scale, and the relevant parameters are therefore calculated at many scales. Once this is done, all of these parameters are associated with each of the objects, DFTs and non-DFTs. These form the basis on which the algorithm tries to predict whether or not a shape is a DFT.

Development of multiscale layers

The finest scale with which to work is the resolution of the imagery resulting from the bundle-adjustment photogrammetry software. This was typically 3.5 cm in our case. Creating the multiple scaled surfaces involved progressively choosing coarser scales, and resampling the raster values by an aggregate.

6.2 Simple aggregates

The simple aggregates used were mean and standard deviation. The scales chosen where multiples of the base scale, and choice of the number of scales was governed by processing time and the relative contribution of the resultant layers to the success of the classification process. This was determined by trial and error of sampled regions.

6.3 Textural measures

Apart from the mean and standard deviation, other aggregate, or textural measures were tried. These measures indicate spatial dependencies in the variation of cell or pixel values in a neighbourhood. They are all aggregate functions of several cell values, which are taken from cells surrounding the cell under study, and the resultant value is given to the central cell. Raster surfaces were calculated for each of these measures, at each of the chosen scales.

6.4 Wavelet decomposition

One method of analysis of imagery at multiple scales is Wavelet Decomposition. This process decomposes values in an image into components representing different scales using harmonic functions known as *wavelets.* The values of each of the components may be added together to produce the original image. The benefit of the process is that features that have a consistent dimension, such as the width of DFTs, tend to fall within only one or two components. Taking these components on their own and disregarding the other components, therefore, tends to filter out any features that are of a scale in which we are not interested, and results in a concentrated representation of possible candidates for DFTs, based on scale. The red, green and blue image bands were combined using a Principal Components transform (described later), and decomposed to the maximum possible number of scales (limited by the number of

times you can double the pixel size before reaching the extents of the image = $\log_2(x)$, where x is the number of pixels across, or down, whichever is least). The resultant raster surfaces where included in the list of predictors whose values were to be fed into the training algorithm.

6.5 Creation of the training dataset

6.5.1 Masking

Once the image data has been processed to full orthomosaics using Pix4D, some filtering is done to reduce large amounts of unnecessary processing. This is a manual, but reasonably fast, process of digitising areas of the image which are definitely NOT worth searching for DFTs. These will include areas of distortion caused by the presence of water, piers, jetties, rock walls and other areas away from the beach. These latter areas played their part in the imagery to assist in the georectification process, but are now no longer required. The masked regions are then inverted to create masks governing where training and prediction will occur.

6.5.2 Treating pixels as members of a single entity

Training the algorithm involves identifying what is and what is not a dugong feeding trail, as many times as possible, in as many different background environments as possible. The visual identification of a DFT is reasonably straightforward to the trained eye, such as the ecologists that have worked on the seagrass meadows under study for many years. What is more difficult is to define the boundaries of the DFT. This is necessary at the training stage, as we must be able to identify all of the pixels which make up the shape of the DFT. Shape and texture (for example how the transition of DFT to non-DFT values at the edges) will play a part in the segmentation process, and there are other features around the seagrass meadows that have similar shapes to the DFTs, which must also be selected as objects and positively identified as NOT a DFT.

6.5.3 Principal Components Analysis

Possible DFTs must therefore be identified and converted to objects. The first step is to decide which image band or bands to use. When working with remote sensing data, we consider each spectral band separately. In this case we have Red, Green and Blue (RGB) bands. One well known way of extracting the most important information from multidimensional data is to use Principal Components Analysis (PCA). This involves a transformation of n layers of data into n components, the first component of which results in a new layer whose values represent those values of all of the layers that are most highly correlated.

6.5.4 Coefficient of Variation

Image: Two examples of DFT clusters with very different relative pixel values to their backgrounds.

The next difficulty to overcome is the fact that the backgrounds of the DFTs in different regions are not the same. In fact, we often get opposites occurring, where DFTs appear as light trails on dark backgrounds in some areas, and dark trails on light backgrounds in others. To overcome this, we "normalise" the relationship of DFT pixel values with those of the background. We do this by calculating a new surface using the Coefficient of Variation (CV), such that

$$C_v = \frac{PCA_1 - \mu}{\sigma}$$

where C_v is the coefficient of variation, PCA_1 is the first PCA component, and μ and σ refer to the mean and standard deviation, respectively. The values from which μ and σ are sampled are those pixels from a window surrounding the pixel under study. The size of that window must be decided by trial and error, but should ideally cause the window to straddle the entire DFT.

Once the CV has been calculated we apply a threshold to its absolute value, based on observation. This threshold on abs(CV) is used to produce a binary map, where contiguous values are treated as one potential DFT. These areas are then converted to vector polygons. This is a useful step, as now shape attributes may be evaluated to each polygon to help in the classification process.

6.5.5 Vector shapes analysis

- Sinuosity

Our sinuosity measure is calculated as the ratio of the length of the centreline of the shape with its area. The centreline is found by "thinning" the shape along its shortest axis.

- Fractal Dimension

In a similar way, the Fractal Dimension is calculated as

$$FD = 2 \times \frac{\log(perimeter)}{\log(area)}$$

- Compactness

The compactness is calculated as the ratio between the perimeter and the area of each shape.

In order to filter out shapes that are unlikely to be DFTs, we calculate a range of shape metrics, which are then each added to the attribute table of the vector shapes.

6.5.6 Training

Training involves an operator going through as many as the shapes as possible, giving each a "class" attribute value of either 1 for a DFT, or 0 for a non-DFT. This is made much more practical by the following method. A new point shape file is created called "DFT" on a GIS which is displaying the previously-created vector shape file. They simply click a new point over every shape that they believe to be a DFT, choosing as many as possible. Then a new point file is created called "non-DFT". The operator then goes through identifying shapes that are definitely not DFTs, aiming for approximately the same number of non-DFT hits as DFT hits. Care is taken to try to select regions that might seem to easily confuse the segmentation algorithm. An example might be linear features such as dune edges or water ripples and wave fronts.

6.6 Store region

The algorithm works within the confines of a geographical region set at the beginning. This region defines not only the geographical bounds (the corner coordinates), but also the spatial resolution at which the algorithm starts. This is the base resolution, usually the Ground Sampling Distance (GSD) for the original imagery.

6.7 Create multiscale layers

The first spectral layer is taken, and is resampled to the next highest spatial scale. At this scale, two further digital surfaces are calculated, which represent the mean and the standard deviation of the lower scale values, aggregated to the new scale.
This process is repeated for each spectral layer, and for each scale.

6.8 Link to training layer database

The attribute table of the training vector file is stored as a database file (DBF format). This data is tapped and brought in to the algorithm.

6.9 Extract training data

For each polygon in the training dataset, the underlying values for each of the predictors (the multiscale layers) is extracted and fed into the training vector attribute table.

6.10 Random Forest

From this point, the Random Forest Algorithm is used.
The algorithm uses the training data to best predict the class (whether a polygon is or is not a dugong feeding trail) using the statistics gathered for each multiscale layer as predictors. The process results in a set of decision trees, which are saved as an object for the prediction stage.

6.11 Tiling

Whilst the training was done on a subset of the total remote sensing data, the prediction process must be carried out on the whole dataset, in order to identify the location of DFTs. This means that the value for each layer, at each scale, must be assessed for each 4 cm pixel over tens of square kilometres, which amounts to a lot of processing.

For this reason, the processing script tiles the data into manageable sized files. A tile size of 500 × 500 pixels was chosen, as this took available memory to close to maximum capacity. Each tile formed the extents of the data being passed to the RF classifier. The RF classifier operates on each individual cell, one at a time. This makes it possible to divide the process into any number of batches. In our case, the workload was divided into batches of 100,000 cells, meaning that each 500 × 500 tile was processed in 3 batches.

6.12 Prediction

The RF prediction is run for the tile. The output values, rather than being a binary decision as to whether or not a cell lies within a DFT, are values between 0 and 1, representing the probability that the cell lies within a DFT. This allows for further analysis and thresholding later. It allows for decisions to be made about whether to be conservative in the prediction to reduce the number of errors of commission. In this case, a high probability threshold can be used to determine membership of a DFT. For a more liberal prediction, where reduced errors of omission are sought, a lower probability threshold can be chosen.

6.13 Mosaicking

The DFT probability tiles are combined for broad-scale output.