

Creation of precise shoal elevation models from UAV-photogrammetry

Comparison of available (software) solutions using two surveys of the beach nourishment at the Oesterdam (The Netherlands)



Adriaan van Natijne <A.L.vanNatijne@student.tudelft.nl>

Supervisors: Dr. R.C. Lindenberg

Dr. ir. B.C. Van Prooijen

This report was written as Bachelor Thesis by Adriaan van Natijne at the Delft University of Technology. The process was supervised by Dr. R.C. Lindenberg and Dr. ir. B.C. van Prooijen.

I would like to thank Eric Geenen from Kragten for providing us with the UAV datasets and helping us with his experience.

Figure 1 (on the cover): Sand nourishment at the Oesterdam in Zeeland, The Netherlands, January 2014. (Photo: (Rijksoverheid, 2014a))

Abstract

Coastal sand nourishment require frequent, detailed and accurate monitoring to determine their morphological changes and effectiveness. Unmanned Aerial Vehicles could provide a flexible, cheaper alternative to conventional surveying techniques like GPS-RTK or total station measurements by employing stereophotogrammetry. The basic principles of (airborne) stereophotogrammetry are discussed. Four different software solutions (MicMac, IQmulus, Agisoft Photoscan and Pix4D Mapper) will be reviewed on the basis of a case study on a shoal near the Oesterdam (The Netherlands). The shoal is characterised by low texture terrain (sand). As most stereophotogrammetry techniques rely on texture based matching, processing problems and solutions are discussed. Comparison to GPS-RTK reference datasets shows that although promising the desired centimetre accuracy can not be met yet.

Contents

Abstract	2
1. Introduction	4
1.1 Problem description	5
2. Introduction to airborne stereophotogrammetry	7
2.1 Basic principles of stereophotogrammetry	7
2.2 Camera orientations and collinearity	8
2.3 Calculating real world coordinates	9
2.4 Feature detection and matching	9
2.5 Estimated accuracy	10
2.6 Summary of stereophotogrammetry	11
3. Available software	12
3.1 Ad hoc, auxiliary software	12
3.2 Software selected for testing	12
3.3 Summary on software	13
4. Available datasets	14
4.1 Actueel Hoogtebestand Nederland (AHN)	14
4.2 GPS/RTK (Rijkswaterstaat)	15
4.3 Images of the shoal near the Oesterdam, from Unmanned Aerial Vehicles	16
4.4 Overview of surveying techniques	19
4.5 Summary on datasets and surveying techniques	19
5. Workflow, results and processing difficulties when processing UAV photogrammetry	20
5.1 Workflow	20
5.2 Accuracy estimation	21
5.3 Results	22
5.4 Problems during processing	25
6. Possible applications and validation	30
6.1 Validation, comparison to other datasets	30
6.2 Possible applications	32
7. Conclusions and recommendations	33
7.1 Recommendations	34
List of figures	35
References	37
Appendix: list of photogrammetry software packages	40

1. Introduction

To monitor the result of the beach nourishment at the Oesterdam frequent surveys are required. Not only the morphology itself, but the changes of the morphology have to be monitored. Analysis of the movement of the sand will provide information on the long term effectiveness of the nourishment. Reliable Digital Surface Models (DSM) are required to monitor the erosion and settlement of sediment at the shoal. At least centimetre vertical accuracy and repeatability is required for this purpose.

The Oesterdam is a 10,5 kilometre long dam in the Dutch province of Zeeland (Figure 1.1). After completion in 1986 the dam connected Tholen to the Zuid-Beveland peninsula and separated the Oosterschelde from the Scheldt-Rhine Canal, creating a tidal free passageway for ships travelling between Rotterdam (Volkerak) and Antwerp.

In 1953 the Watersnoodramp (North Sea flood) killed nearly 2000 Dutch and made many more homeless. To prevent similar disasters in the future the Delta Plan was set-up, connecting many of the coastal islands by dams and strongly reducing the length of the sea-facing dikes. When the Oosterscheldekering (Eastern Scheldt storm surge barrier) was finished in 1986 the tidal amplitude in the water level in the Eastern Scheldt were reduced from 3,40 metre to 3,25 metre. As a consequence the waves are still eroding the shoals but the tidal fluctuations are not strong enough to leave new sediment on top (Natuurmonumenten, 2014a). This results in degradation of the shoals, that can erode up to 2 cm per year (Paree, 2014).

As the Oosterschelde became a national park in 2002 the Dutch committed themselves to the preservation of the unique environment. In the context of the 'coalitie Natuurlijke Klimaatbuffers' (Climate Buffer Coalition) seven environmental organisations, financed by the Dutch government, work on the realization of so called natural climate buffers: “nature areas specially designed to reduce the consequences of climate change.” (Climate Buffer Coalition, 2014) Beach nourishment and artificial oyster reefs at the basis of the Oesterdam compensate for the receding shoals elsewhere and will provide a safe retreat for many birds and other animals for many years to come (Natuurmonumenten, 2014b).

To effectively asses the effects of this human intervention surface maps are required. Differences between those maps will give insight into the movement, erosion and settlement of the sand. Airborne or terrestrial laser scanning (LiDAR), GPS-measurements, triangulations (theodolite, total station) and other surveying techniques require expensive equipment and/or costly manual labour. When working in coastal areas, walking over the sand is not only labour intensive, but also not instantaneous (enough) as the tides vary every six hours. Those techniques allow exact mapping but require considerable budget and effort, limiting the temporal frequency of surveys.

In recent years Unmanned Aerial Vehicles (UAVs) have taken a flight in detailed surveying. They provide a cheap, fast, simple alternative to conventional surveying. The UAV builder Parrot¹ has sold over 670 000 drones between 2010 and the end of the second quarter of 2014 (Parrot, 2014), to consumers alone. Instead of active techniques like laser scanning, stereophotogrammetry is used. Stereophotogrammetry is the reconstruction of 3D coordinates (e.g. elevations) based on photography. In combination with advances in computer vision, UAV based stereophotogrammetry

1 Parrot is only one of the major UAV builders and has different models for the consumer and professional market. Furthermore Parrot is a major investor in stereophotogrammetry software. (Melanson, 2012)

became a viable alternative to conventional surveying techniques (Barazzetti, 2010).

Given the fact that they are fast, freely available and relatively cheap, UAVs could provide a solution to those problems as the lower costs allow more frequent surveys and a fine grained temporal dimension.

The Dutch' firm Kragten experiments with UAVs as an alternative to conventional surveying. The two drones discussed in this report employed stereophotogrammetry to create a Digital Surface Model (DSM). Kragten did two flights on the Oesterdam and will do more in the future. They provided the raw data from the first flight for this report. Furthermore they provided the DSM they generated from their second flight.

1.1 Problem description

A lot of potential use cases of UAVs are well documented (e.g. (Bemis et al. 2014)), including accuracy estimations ((Barazzetti, 2010), (Bemis et al. 2014)). They all use the assumption that the surveyed terrain contains high texture and that reconstruction of low texture terrains is difficult to impossible (James & Robson, 2012). The shoals discussed here consist of sand that contains very few textural differences, only the artificial oyster reefs stand out in sandy plains. Although general accuracy rules should apply, software comparisons (like (Sona et al., 2014)) have to be rewritten for those complex, low texture, areas.

This report will asses UAV based stereophotogrammetry, especially on the low texture shoals at the Oesterdam. Can those surveys replace conventional surveying techniques and provide the preferred temporal dimension. This leads to the following research question:

Could UAV surveys replace conventional surveying?

- How is elevation data acquired from stereophotography (photogrammetry)?
- What accuracy can be expected from stereophotogrammetry?
- What processing software exists, and how does it perform in DSM generation?
- How do the results compare to other surveying methods?
- What applications do exist for the generated DSM's?

After this introduction of the why and how of this project, photogrammetry itself will be introduced in chapter 2. In chapter 3 a selection of the available software solutions will be described. This is followed by the datasets available to this study (chapter 4) and resulting models and the difficulties in obtaining them in chapter 5. The results will then be compared to the other datasets in chapter 6. Last a short overview of possible applications will be given.



Figure 1.1: Location of the Oesterdam, marked by a red star in the Dutch province of Zeeland. Three other major works of the Delta Plan can be seen on the sea-side, from north to south: Haringvlietdam; Brouwersdam and Oosterscheldekering. (Background: BRT Achtergrondkaart)

2. Introduction to airborne stereophotogrammetry

Jurists knew it all along: “*unus testis nullus testis*” - a single witness is no proof. As will be shown in this chapter two cameras are required to form a complete (3D) image of the situation. First the stereophotogrammetry and some of its products will be introduced. The mathematical principles and important coroutines will then be discussed. The expected accuracy will be mentioned last.

Stereophotogrammetry is the technique of using two or more cameras to reconstruct the 3D coordinates of an object. It is a non-intrusive mapping technique, independent of scale. Large areas can be recorded effectively using an aeroplane or satellite, smaller areas can be recorded by hand without damage to the object. Given the independence of scale this allows a wide range of applications, from small forensic specimens to worldwide mapping.

Stereophotogrammetry has been in use since 1840 for mapping purposes, the transition to digital processing started in the 1970s, and as of early 2000 acquisition and processing can be fully automated (Aerometrex, 2012). With the advances in computational power and computer vision new image processing algorithms became available. Those new algorithms, gathered under the name 'Structure from Motion' are able to extract 3D points from a sequence of 2D images (James & Robson, 2012).

The ASTER GDEM model, a digital terrain model of most of the earth, is a well known application of satellite based stereophotogrammetry (NASA, 2004).

Large parts of this chapter are inspired by the lecture slides of the MSc photogrammetry course GM1230, taught at the Delft University of Technology up till 2011. Information on feature detection is mostly taken from “Visual Object Recognition” (Grauman & Leibe, 2011).

2.1 Basic principles of stereophotogrammetry

We use the concept of stereo photography in our everyday life. If an object can be seen from two (slightly) different positions we can estimate the distance to the observer(s).

Our eyes function as two cameras, with a small horizontal separation. From experience rather than from a mathematical model we know how to interpret and combine the two images we see simultaneously. This method works especially well in close range: we can estimate the distance to our paper or screen with centimetre accuracy. Objects further away, the skyline of a distant city or the moon, require different techniques as those small differences require more precise measurements.

With a single camera the distance can only be estimated from known (surrounding) objects. This principle is illustrated by a red and green bar in Figure 2.1. When a second observer with a known position is added, the size of the object becomes clear. The system in Figure 2.2 is fixed, it is impossible to move the objects and/or cameras without disturbing the scene.




Figure 2.1: Single camera with two objects (red and green) that appear to have the same size.




Figure 2.2: Two camera's and a single object (red).

2.2 Camera orientations and collinearity

It is important to distinguish two worlds: the camera and the real world. The coordinates within the frame of the camera are known (pixel coordinates) and the position of the camera within the world is known (GPS, motion sensors). Bound together this information allows the calculation of the real world position of the detected object.

A light ray can be seen as passing through the centre of the lens. If an object is shown at p' in the image (at the focal length to the centre of the lens: f) a vector can be drawn through the centre of the lens. This principle is illustrated in Figure 2.3. As shown for a two dimensional image in the previous paragraph the object will be on this line, albeit unknown how far away. If a second observer with known position detects the same object at p'' the intersection of the rays equals the position of the object in real life, as illustrated in Figure 2.4.




Figure 2.3: Model of a lens. (Taken from ADAM Tech., 2010).




Figure 2.4: The basic principle of stereophotogrammetry. (Taken from Haggren, 2002).

Collinearity is the propriety of objects to lie on the same, single line. For both rays in Figure 2.1 this holds for any points on their line. The intersection of those lines/rays in Figure 2.2 equals the real world position of the object. The collinearity equation is about determining the position of the object by finding the intersection of those two lines.

In practice the cameras might not be parallel, as shown in Figure 2.4. These rotations, known in the aeronautical world as yaw, pitch and roll, have to be incorporated into the model. When the cameras rotate the rays rotate accordingly. The rotations, illustrated using aeroplanes in Table 1 apply

equally to cameras. All three rotations can be combined: $R_{\psi\theta\varphi} = R_{\psi} \cdot R_{\theta} \cdot R_{\varphi}$. (CH Robotics, 2014)




Axis		Rotation matrix	Example ²
Yaw	ψ	$R_{\psi} = \begin{bmatrix} \cos(\psi) & -\sin(\psi) & 0 \\ \sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix}$	
Pitch	θ	$R_{\theta} = \begin{bmatrix} \cos(\theta) & 0 & -\sin(\theta) \\ 0 & 1 & 0 \\ \sin(\theta) & 0 & \cos(\theta) \end{bmatrix}$	
Roll	φ	$R_{\varphi} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\varphi) & \sin(\varphi) \\ 0 & -\sin(\varphi) & \cos(\varphi) \end{bmatrix}$	

Table 1: Aeronautical angles of rotation.

The scale (λ) of the object is a function of the focal length (f) and distance to the object (d): $\lambda = d/f$. Combining this into a single formula results in the following equation for a single ray. Where X , Y and Z are the real world coordinates; x and y are the coordinates in the picture, relative to the centre; f is the focal length; X_0 , Y_0 and Z_0 are the coordinates of the camera in the real world. R and λ are rotation and scale, as discussed earlier.

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \lambda \cdot R_{\psi\theta\varphi} \begin{bmatrix} x \\ y \\ -f \end{bmatrix} + \begin{bmatrix} X_0 \\ Y_0 \\ Z_0 \end{bmatrix}$$

In practice the distortion of the lens has to be taken into account too. This will not be discussed here, as it exceeds the purpose of this report.

2.3 Calculating real world coordinates

After feature detection the images are matched with an increasing amount of other images. With each step the internal camera parameters (focal length, distortion parameters) and external camera parameters (position and orientation) are optimised.

If any of the variables is unknown the equations of two or more cameras can be combined into one, solving the equation for the unknowns. In practice extra variables are added to compensate for measurement errors. Those errors will then be kept to a minimum, for example using a Least-Squares method (Barazzetti, 2010).

The resulting matches form a sparse point cloud of points matched in two or more images. The calculated parameters now allow for a better, more precise search of matching points in the image, resulting in a dense, more accurate point cloud.

2.4 Feature detection and matching

To reconstruct the 3D coordinates for a point using the formulas mentioned above it is necessary to

² Images taken from the Toronto Soaring Club: http://toronto-soaring.ca/wordpress/?page_id=30 (last visited: 29-10-2014).

detect the same point in at least two images (with known camera positions). As we can see in Figures 2.2 and 2.4 we have to detect the same point on the red line in the images of multiple camera's.

In many cases this is done using the SIFT descriptor, described by D. Lowe in 2004. The algorithm is a combination of a *region of interest detector* and a *feature descriptor*. The detector decides which areas contain abrupt changes that may stand out in other images too, while the descriptor describes them in a way that is invariant to translation, rotation, scaling and most affine transformations. It is possible to use any combination of detector and descriptors. The algorithm published by Lowe in 2004 used his Difference of Gaussians detector (1999) and the (then new) SIFT descriptor. Both will be introduced briefly.

For the Difference of Gaussians (DoG) detector, a faster approximation of the Laplacian of Gaussians (LoG) detector, a Gaussian pyramid is created first. That is, a series of consecutive blurred images, as can be seen from left to right in the top row of Figure 2.5. The successive images are then subtracted from each other, creating the bottom row of the figure. Large differences indicate changes in the image at that scale space (level of blurring). The features (or key points) are extracted from those images by selecting the local minima and maxima (Sinha, 2014).




Figure 2.5: Example of the Difference of Gaussians. A Gaussian pyramid is shown in the top row, the matching Laplacian pyramid in the bottom row. (Taken from (Wallace, 2011).)




Figure 2.6: Illustration of the generation of a SIFT descriptor. (Taken from (Levi, 2013).)

The Scale Invariant Feature Transform (SIFT) descriptor is insensitive to changes in lightning and position. Even if photographed from different angles, similarities between features can be detected.

A 16×16 window of gradients is sampled around the detected feature. First the average direction of the gradients is extracted and compensated for to make the system insensitive to rotation. This window is then summarised as 4×4 matrix containing 8 histogram bins in each cell. The gradients of the 16×16 grid are aggregated based on intensity, direction (8 directions) and distance to the feature, creating a 128 wide vector ($4 \times 4 \times 8$). This procedure is illustrated in Figure 2.6. This vector can then be compared to features in other images. When comparing to possible matches a good match has minimal Euclidean distance.

2.5 Estimated accuracy

The accuracy of the survey is dependant on the ratio between the size of the pixel and the size of the pixel on the ground (ADAM Tech., 2010). In (James & Robson, 2012) this ratio is estimated as 1:1000, id est the accuracy is one thousandths of the distance between the camera and the object.

For best accuracy a pre-calibrated camera is used. The images are best taken at a slightly convergent angle (looking towards each other). When the lateral overlap is less than the overlap in the direction of flight, the survey should be repeated orthogonal to the first direction of surveying (Bemis et al. 2014).

Various formulas exist to estimate the accuracy. The main difference is the distinction between

stereo-pairs and convergent imagery. In any case the internal accuracy of the camera (σ_i) has to be determined first. Modern digital cameras are accurate to about half the size of a pixel (James & Robson, 2012) or better: $\frac{1}{3}$ pixel (ADAM Tech., 2010).

The first method to estimate the accuracy of the object coordinates is based on stereo pairs only, only two images are incorporated in the calculation. This formula is used by (ADAM Tech., 2010), the GM1230 photogrammetry course and (James & Robson, 2012). With σ_z the variance in the vertical direction; \bar{D} the average distance to the object; b the distance between cameras; f the focal length and σ_i the variance of the camera. For best results the ratio of b/\bar{D} is best kept between 0.5 and 1.0 (Compta Maps, 2011).

$$\sigma_z = \frac{\bar{D}^2}{b f} \sigma_i$$

The second methods from (James & Robson, 2012) involves two or more photos and is more often used in close range, convergent, photogrammetry. Using more photos will increase the accuracy (decrease the variance), but the effect will fade with every extra image, as can be seen in Figure 2.7. Here σ_c is the variance of the coordinates (x , y and z); \bar{D} the average distance to the object; k the amount of images; f the focal length and σ_i the variance of the camera. The variable q denotes the 'strength of the network'. This is an empirical value based on the number of images and their distribution, in practice values are in the range of 0.4 – 2 (Barazzetti, 2010).

$$\sigma_c = \frac{q \bar{D}}{\sqrt{k} f} \sigma_i$$

Ground control points allow the reconstructed scene to be placed in the real world, assigning the correct world coordinates and orientation. Nowadays many cameras record their position, using GPS, when taking an image. Those camera positions could serve as control points too, negating the necessity of ground control points for this purpose. Although possible the errors exceed 0.10 metre when using a GPS accurate to 10-20 cm (Turner, 2014).

2.6 Summary of stereophotogrammetry

If an object is seen by two or more observers with known position or has (ground) control references the 3D coordinates can be determined from the mutual dependencies in the images. The accuracy is dependant on the organisation of the images. Slightly convergent imagery, with large overlap and overlap in a direction orthogonal to the original surveying direction contribute to accurate results.

The ratio of 1:1000 for the ratio between accuracy (in any direction) and surveying distance serves as a rule of thumb. In practice higher accuracies are possible.




Figure 2.7: Accuracy in relation to increasing images. The average distance is taken as 100 metre, the focal length at 10 mm and pixel size as 5 μ m.

3. Available software

Various photogrammetry software packages were tested in this study. Even though they all implement similar algorithms large differences exist. The software suites will be introduced shortly, and some will be put to the test. The problems encountered will be discussed after the comparison of the results.

The following list of photogrammetry software is a compilation of the lists provided by (Bemis et al. 2014), (James & Robson, 2012) and (Sona et al., 2014), complemented with own findings. Many solutions consist of combinations of others. A detailed overview of the various packages is available in 'Appendix: list of photogrammetry software packages'.

Freely available: Bundler; SFMToolkit; Python Photogrammetry Toolbox (PPT); VisualSFM; 3DF Samantha; SURE; MicMac; CMVS.

(Online) services: Photosynth; Arc3D; CMP SfM Webservice; Autocad 123D Catch; My3DScanner.

Commercial, standalone software: Agisoft Photoscan; Acute3D; PhotoModeler; 3DF Zephyr Pro; Pix4D (former Pix4UAV); ERDAS IMAGINE, formerly Leica Photogrammetry Suite.

3.1 Ad hoc, auxiliary software

To assist in the evaluation of other packages, and to improve the understanding of the process a custom implementation of the orthophoto algorithm was made in Python. This software is not capable of matching on it's own. Various attempts were conducted using the OpenCV SIFT implementation. Due to the availability of various alternatives the benefits did not outweigh the (time) investment. Furthermore various auxiliary scripts exist to aid the sharing of data between various applications and formats.

3.2 Software selected for testing

Four software packages were selected for testing, three of the publicly or commercially available packages listed earlier in this chapter and one developed at the Delft University of Technology. A short introduction on those packages will follow.

Development on **Micmac** started in 2005 by the French IGN (Institut national de l'information géographique et forestière). It consists of many tools rather than a single binary. Its versatility comes at the price of extensive configuration requirements that are unsuitable for the novice user. Its manual is largely translated to English, but the more advanced chapters are only available in French (IGN, 2014). A simple graphical user interface is available (Cléry, 2013).

IQmulus is collaborative effort to “enable (optimized) use of large, heterogeneous geo-spatial data sets for better decision making through a high-volume fusion and analysis information management platform.” (IQmulus, 2014) In the context of this project a stereophotogrammetry system was developed at the Delft University of Technology.

The Russian **Agisoft** was founded in 2011. Their main product is Photoscan that is made for stereophotogrammetry, whether it be airborne or close range (Agisoft, 2014). Agisoft Photoscan is in use at the faculty, at the section of Applied Geology to analyse fracture patterns from UAV photography (Bisdorn, 2014).

Founded as a spin-off of the Ecole Polytechnique Fédérale de Lausanne, in Switzerland in 2011, **Pix4D Mapper** (then Pix4UAV) is made specially for UAV stereophotogrammetry (Pix4D, 2014). A major investment by UAV builder Parrot in 2012 confirmed this strategy (Melanson, 2012). The use of Pix4D was advised by Kragten as the 'low texture' option in newer versions had positive results.

3.3 Summary on software

A plethora of software solutions exists. Proprietary, open source and freeware. Most software is written for Microsoft Windows. With commercial packages prices range from less than € 100 to thousands of Euros in annual subscription fees. Furthermore webservices exist, that require little to none investments in hardware on the client side.

Four software packages are included in this report: MicMac, IQmulus, Agisoft Photoscan and Pix4D Mapper. Those packages were selected based on advise and/or experience from others.

4. Available datasets

Three different datasets were used in this report. The data gathered by the drones can be verified using those two datasets. The *Actueel Hoogtebestand Nederland* covering most of the dike and conventional GPS-measurements (RTK) done on the shoal. All three will be introduced in this chapter.

In 2012 the Dutch government started publishing their datasets, including geospatial data, as open data (Rijksoverheid, 2014b). The aerial background maps in many maps in this report, the AHN dataset and many more datasets are available through the 'Nationaal Georegister' (NGR, 2014). The number of available datasets is still growing and will be in the future (Rijksoverheid, 2013).

4.1 *Actueel Hoogtebestand Nederland* (AHN)

The AHN maps are based on aeroplane laser altimeter (LiDAR) dataset covering the whole of the continental Netherlands and is freely available. Laser scanning is an active technique. A laser (light) pulse is sent to the ground from an aeroplane. The light pulse bounces back on the ground and arrives back at the aeroplane with a slight delay, caused by the (two way) travel time between the aeroplane and the ground. The precise location of the aeroplane is known using GPS, an inertial measurement unit (IMU) tracks the orientation. The elevation of the ground below can be deduced from the measured travel time of the light, the known position and orientation of the aeroplane.

The campaign for AHN2 started in 2007 and was finished in 2012 (AHN, 2009); (AHN, 2012). The measurements on the Oesterdam were done in 2007 and 2009 (AHN, 2013), it is unknown which of them is in the published dataset. The systematic error is 5 centimetre at maximum, while the stochastic (random) error is 5 cm too. Combining those criteria 68,2% of the points falls within 10 cm vertical accuracy (AHN, 2013). A new survey was done in early 2014 but won't be published until November 2014 and is therefore not taken into account (AHN, 2014).

The coverage of the area is relatively bad, as can be seen in Figure 4.1. Laser light does scatter on water, there are therefore only few points outside of the fully covered dam. Objects on top of the surface, trees and cars for example, were filtered out before publication.

Although not having the desired accuracy, the AHN model allows a rough comparison of a model with the dam.




Figure 4.1: Actueel Hoogtebestand Nederland, a combination of tiles r49bz1 and r49dn1. (Background: PDOK Luchtfoto)

4.2 GPS/RTK (Rijkswaterstaat)

Regular GPS measurements are accurate enough for road navigation but their accuracy is far from the desired accuracy for this project. With regular GPS measurements 95% of the measurements error falls within 3,3 metre in the horizontal direction and 4,7 metre in the vertical direction (FAA, 2014). An error more than 400 times bigger than the desired centimetre accuracy.

Normal GPS units base their position on the different, delayed arrival times from the various satellites. As the position of the satellites and approximate speed of the signal is known the 'intersection' of the signals can be calculated based on the relative delays. This intersection equals the current position.

To achieve a higher accuracy, Real Time Kinematic (RTK) satellite navigation uses the wave form of the signal rather than the signal itself. Instead of using the code contained in the signal, the wave containing the signal is used. This allows much greater accuracy (up to sub-centimetre) but requires more expensive equipment as matching the raw signals is error prone.

In May 2014 Rijkswaterstaat completed an RTK survey on the shoal, a few days before the UAV survey. The layout of the RTK survey can be seen in Figure 4.2. This dataset, of 5355 points, is available to this study.



Figure 4.2: RTK survey of May 2014, consisting of 5355 points. (Rijkswaterstaat)
(Background: PDOK Luchtfoto)

4.3 Images of the shoal near the Oesterdam, from Unmanned Aerial Vehicles

Surveys performed using Unmanned Aerial Vehicles, sometimes called 'drones', have a lot in common with other airborne surveys. The used aeroplane or helicopter is a lot smaller and most notably controlled from the ground and/or flying autonomous. Various instruments can be attached to the drone, in this case a digital camera. With every image they record their position (GPS) and orientation (IMU). Other applications, like UAV mounted laser scanners are slowly introduced into the market (Riegl, 2014) and will not be discussed here.

The shoals west of the Oesterdam were surveyed twice by Kragten. The photo locations of both flights can be seen in Figure 4.3. Each survey with a new UAV: a Trimble X100 and a Trimble UX5. A comparison of those UAVs can be found in Table 2. A new survey is announced for the end of November 2014, but will not be incorporated in this report.

From the first flight the raw data was available to this study. The second survey was processed by Kragten using Pix4D (see chapter 5.3.d, page 24), the resulting point cloud was then shared with us.




Figure 4.3: Comparison of the two UAV datasets. Each point represents a picture taken. (Background: PDOK Luchtfoto)

4.3.a Kragten's UAVs

Kragten owns two *Unmanned Aerial Vehicles*, a Trimble UX5 and a Trimble X100 (Ven, 2014). With a wingspan of 1 metre and a cruising speed of 80 km/h the Trimble UX5 is capable of travelling 60 kilometre (~ 50 minutes) on one battery charge. Trimble X100, albeit smaller and lighter than the UX5, has a similar cruising speed but a slightly shorter range of 50 kilometre (~ 45 minutes) (Trimble, 2014).

Apart from technical, exterior differences, the main difference is the camera. Where the X100 is equipped with a Ricoh GR Digital III, the UX5 has a Sony NEX-5R attached. As the camera is one of the decisive factors in stereophotogrammetry this is an important difference.

An overview of the two drones and the surveys is given in Table 2.

UAV	Trimble X100 ³	Trimble UX5 ³
		
Cruising speed	80 km/h	80 km/h
Range	50 kilometre (~ 45 minutes)	60 kilometre (~ 50 minutes)
Camera	Ricoh GR Digital III ⁴	Sony NEX-5R ⁵
	Compact camera	Mirrorless interchangeable-lens camera
		
	Resolution	10 MP (3648×2736 pixels)
Sensor	1/1.7" (7,6×5,7 mm)	APS-C (23,4×15,5 mm ⁶)
Focal length (35 mm equiv. ⁷)	6 mm (27 mm)	15 mm (24 mm)
Survey	25 November 2013	May 2014
Images	523 (out of 876 ⁸)	1199
Time	10:54 – 11:20	
Shutter speed	1/800 second	
Aperture	f/4.0	
Average altitude	95 metre	124 metre
Area covered	~ 1 km ²	~ 3 km ² (Pix4D)
Georeferences (markers)	3	22

Table 2: Comparison of UAVs and flights.

The drones flew a pre-defined pattern at an approximate altitude of 95 and 124 metre respectively. The flightpath is constructed to have enough overlap (> 85% frontal overlap is advisable due to large uniform area's) and guarantee an equally spread coverage of the whole area (Ortega, 2013a).

3 Image taken from the Trimble website: <http://uas.trimble.com/media/galleries> . (29-10-2014)

4 Image taken from the Ricoh website: http://www.ricoh.com/r_dc/gr/gr_digital3/ . (29-10-2014)

5 Image taken from the Sony website: <http://www.sony.nl/support/nl/product/nex-5r> . (29-10-2014)

6 Officially APS-C is 25,1×16,7 mm, but variants exist among manufacturers.

7 The sensor of many digital camera's is smaller than the conventional 35 mm film, resulting in 'zoom' compared to the reference film size. Expressed here is the equivalent focal length when a 35 mm film would be used.

8 A total of 876 images was taken, but only 523 of them were available to this study.

All images were taken from a (close to) nadir perspective, keeping perspective distortion to a minimum. Shutter speed is fast to prevent motion blur, but as a result images are dark and have a low contrast ratio. Just like the aeroplane used in the AHN surveys the UAV is equipped with an inertial measurement unit (IMU) and GPS to track the orientation and position of the UAV and ensure the planned trajectory is followed as closely as possible.

4.4 Overview of surveying techniques

A self compiled schema of differences and similarities between different techniques is shown in Table 3. All listed techniques require a trained and skilled surveyor. According to Kragten a UAV survey may take up to three days (paperwork, flight and processing), while a RTK-survey of the same area may take up to eight days. Data acquisition using stereophotogrammetry is 80% faster than terrestrial laser scanning (James & Robson, 2012).

	Properties	Advantages	Disadvantages
Airborne LiDAR	Active laser pulse, from an aeroplane.	Densely covers large areas with high accuracy.	Expensive aircraft and sensors.
Terrestrial laser scanning	Active laser pulse, ground based.	Easy to deploy and accurate.	Slow compared to stereophotogrammetry, expensive sensor.
(RTK-)GPS	Manual GPS measurements.	Very accurate. Influence on coverage.	Requires access and extensive manual labour.
UAV Stereophotogrammetry	UAV mounted camera, computer processed to 3D information.	Fast, accurate when used in combination with GCPs. Dense cloud. Does not require human access.	Still under development. Possibly bad results on terrains with low texture or complex geometries (e.g. a see-through crashbarrier).
Total station	Manual measurements.	Very accurate. Influence on coverage.	Requires access and extensive manual labour.

Table 3: Comparison of various surveying techniques.

4.5 Summary on datasets and surveying techniques

Four datasets of three different types are available to this study: airborne LiDAR survey (AHN), a GPS-RTK survey and two stereophotogrammetry UAV datasets. Furthermore a comparison of the speed and costs of various surveying techniques is made in Table 3. We can conclude that although still under development, UAV based surveys may provide a fast and cheap alternative to conventional surveying techniques.

5. Workflow, results and processing difficulties when processing UAV photogrammetry

In this chapter the software introduced in chapter 3 will be run on the available datasets and results will be shown. Potential problems will be analysed and where possible addressed.

5.1 Workflow

Due to (practical) limitations the region of interest had to be limited. Pix4D, for example, recommends at least 16GiB of RAM for projects of around 500 images (Ortega, 2013b). As this was not available three main regions of interest were defined. Those regions were then fed to the various software packages.

5.1.a Regions of interest

Three distinct regions of interest were defined in the northern part of the first survey. All three show different obstacles for the used software.

- | | | |
|----------|---------------|---|
| Region 1 | The dam | With all distinct features on the dam, this area should be no problem for any of the used packages. |
| Region 2 | Shallow water | Shallow water with distinct traces of underwater structures. |
| Region 3 | Channels | The channels form distinctive features in a low texture environment. |




Figure 5.1: Regions of interest. As illustration the orthophoto from the second survey is shown as background. All three areas are fully covered by the first survey.

5.2 Accuracy estimation

Based on the flight information from Table 2 we can now estimate what accuracy to expect from the flights.

The sensor of the Ricoh GR Digital III camera, attached to the Trimble X100, measures 7,6×5,7 mm. Given the resolution of 3648×2736 pixels this results in square pixels with sides of $2,083 \times 10^{-6}$ metre. The average altitude during the flight was 95 metres. If we assume that the shoal is located at 0 metre elevation⁹ this results in an average camera to ground distance of around 100 metre, as the distance to points not centred below the camera is further than just the distance.

We can now calculate the size of the pixel on the ground. The scale factor is the ratio between the distance behind the lens (focal distance) and the distance in front of the lens (altitude).

$$\text{GroundPixelSize} = \frac{\text{Altitude} - \text{Elevation}}{\text{Focal distance}} \cdot \text{Pixel size} = \frac{94.7 - 0}{0.006} \cdot 2.083 \cdot 10^{-6} \approx 0.033 \text{ metre}$$

If we assume that the pixel is no more than half its width off (ADAM Tech., 2010). The (horizontal) accuracy right below the sensor is therefore: $\sigma_{\text{pixel}} \cdot \text{GroundPixelSize} \approx 0,5 \cdot 0,033 \approx 0,016$ metre.

The vertical distance is not only dependant on the horizontal accuracy but also on the distance to the next image. In this survey the average distance to the next image was 28 metres, from the ratio between the vertical and horizontal distance follows the vertical accuracy:

$$\frac{\text{Altitude}}{\text{Distance}} \cdot \text{HorizAccuracy} = \frac{94.7}{28.2} \cdot 0.016 \approx 0.055 \text{ metre}$$

This procedure is in fact equal to the formula for stereo-pairs given in chapter 2.5: $\sigma_z = \frac{\bar{D}^2}{b f} \sigma_i$.

In practice around 11 images are overlapping on average. The scene can be described as a convergent photogrammetry to benefit from this. Due to the lack of more detailed information we estimate the 'network strength' as $q=1$ (James & Robson, 2012). The σ_i is defined above as $0,5 \times 2,083 \times 10^{-6} = 1,042 \times 10^{-6}$. There are 11 overlapping images (k); the average distance to the ground (\bar{D}) is 100 metre. The focal length (f) is 6 mm.

$$\bar{\sigma}_c = \frac{q \bar{D}}{\sqrt{k} f} \sigma_i = \frac{1 \cdot 100}{\sqrt{11} \cdot 6} \cdot 1.042 \cdot 10^{-6} = 0.005 \text{ m} \quad (5 \text{ mm})$$

This results in a coordinate variance of only 5 mm! Adding more photos does not help much, as can be seen in Figure 5.2.

The found accuracies are much better than the 1:1000 (focal length:distance) introduced by (James & Robson, 2012): 1:1818 and 1:200 000 respectively. Please note that James & Robson described actual, realised, accuracies rather than estimations.




Figure 5.2: Variance in relation to increasing images. Given the parameters of the first survey from Table 2.

⁹ The average elevation found in the RTK measurements was -0,6 metre ($\sigma \approx 0,84$), justifies this assumption.

5.3 Results

Only Agisoft Photoscan and Pix4D returned usable results as IQmulus and MicMac had difficulties handling sections with few or no matches.

5.3.a Self developed

As described in section 3.1 the self developed software was not capable of processing datasets from images to 3D models. It was mostly used to check intermediary products from other packages. An example of an orthophoto based on coordinates rendered by Agisoft Photoscan is shown in Figure 5.3.




Figure 5.3: Orthophoto, rendered using own software. Based on coordinates Agisoft. Note the missing, unmatched images in the north.

5.3.b MicMac & IQmulus

Both MicMac and IQmulus had trouble processing the images. Although completely unrelated from a software point of view they produce comparable errors. The only region that produced a proper model was the dam. The software crashed when not enough matches were found between any of the image pairs involved. MicMac is highly flexible, and was build with mapping in mind. The fine-grained configuration files allow fine-tuning and still have as a lot of potential.




Figure 5.4: Dam, rendered using MicMac. Camera positions can be seen on top. (CloudCompare)




Figure 5.5: Only 39 points were rendered using IQmulus. Therefore only few 'dots' are visible. (CloudCompare)

5.3.c Agisoft Photoscan & Pix4D (first survey)

Processing the whole dataset from the first flight in Agisoft resulted in a model with many holes (Figure 5.6). In Figure 5.7 a 'valley' shape can be seen, as the east and west ends curl upwards. This dataset was processed without ground control points, relying on auto-calibration based on the on board GPS. Adding the ground control points to the model might solve the concave shape but will not repair any holes. Agisoft was run at 'normal' settings.




Figure 5.6: Point cloud rendered using Agisoft. (CloudCompare)




Figure 5.7: Digital Surface model as rendered by Agisoft. A concave distortion can be seen. (Background: PDOK Luchtfoto)

As Agisoft and Pix4D had demanding memory requirements (≥ 6 GB) the other experiments were performed on the regions defined in chapter 5.1.a. To further reduce computational costs only images that were estimated to contain more than 10% of the region of interest were used in the processing. The effects of this decision will be addressed in chapter 5.4.c.

Neither Agisoft Photoscan nor Pix4D returned results on the second, partly submerged, region. Both software packages returned results on region 1 (the dam) and region 3 (the channels). In Figures 5.8, 5.9, 5.10 and 5.11 the results are shown. The small black line corresponds to the region of interest as defined in Figure 5.1. Please note that, due to the lack of reference points in those small portions of the survey, the images are normalised on the mean (mean = 0 metre elevation).




Figure 5.8: Region 1, Agisoft Photoscan.




Figure 5.9: Region 3, Agisoft Photoscan.




Figure 5.10: Region 1, Pix4D.




Figure 5.11: Region 3, Pix4D.

The quality settings in Agisoft Photoscan were set to 'normal'. Pix4D was run using the 'alternative processing mode for planar surfaces with low texture'. At least four matches for the same point were required; the noise filtering radius was set to '10' and the smoothing radius to '20'; the resulting DSM was exported at $10 \text{ cm}/\text{pixel}$. Those settings match the ones used by Kragten in the second survey.

5.3.d Pix4D (second survey)

The model of the second survey was made by Kragten using Pix4D. Kragten did not supply a DSM, but only a point cloud. Pix4D was again run using the 'alternative processing mode for planar surfaces with low texture'. At least four matches for the same point were required; the noise filtering radius was set to '10' and the smoothing radius to '20'. While calculating Pix4D used up to 35 GB of RAM.

The point cloud made available by Kragten was converted to a DSM using LASTools' las2dem (Rapidlasso, 2012), again at $10 \text{ cm}/\text{pixel}$. The resulting DSM can be seen in Figure 5.12, as by-product of the dense point cloud a true colour orthophoto could be rendered (Figure 5.13). In Figures 5.14 and 5.15 other representations of the same DSM are given, for a more intuitive look at the results. The diagonal lines are the result of using an unlicensed version of LASTools.




Figure 5.12: DSM (10 cm/px) as rendered by Kragten using Pix4D.



Figure 5.13: A by-product of the process, a full colour (RGB) orthophoto of the area.




Figure 5.14: Slope-aspect plot of the DSM. All slopes $< 8.5^\circ$ are suppressed to reduce noise.




Figure 5.15: Hillshade of the DSM.

5.4 Problems during processing

While processing the datasets various problems surfaced. They will be discussed in this section and where available solutions or workarounds are provided.

5.4.a Lack of texture

The lack of texture hinders the detection of features. The most distinct features are described as only vague features are found the best matches are on small objects or even noise. As the difference in 'distinctiveness' is small those features might not be the ones described in neighbouring images. The vague gradient further complicates the matching process.




Figure 5.16: Graph generated by Pix4D illustrating which photos are best connected (darker) and which photos lack matching features (white). Due to the lack of matches two independent groups are formed. (Taken from the 1% overlap run, see 5.4.c. See Figure 5.3 for a visual reference of the area.)

As can be seen in the Figures 5.17 and 5.18 areas with distinct texture features (the dam) result in many matches between images. Areas with less texture (the sandplains) result in fewer matches, and only on points that stand out in their surroundings. The same can be seen in the network map Figure 5.16, showing strong links around the artificial oyster reefs and dam and almost no links on the sand (comparable to the image shown in Figure 5.17).

A possible work around is using higher resolution, low noise images that allow the description of smaller features. This can be achieved by, for example, flying at a lower altitude. More accurate GPS/IMU sensors could allow the matching algorithms to require less clear/distinct features for matches and start working with matching of low quality features from the start. This is now typically done in a second stage, after initial (self) calibration is complete.




Figure 5.17: MicMac matches between R0029210.JPG and other images, shown on the top and bottom rows.



Figure 5.18: MicMac matches between R0029468.JPG and other images, shown on the top and bottom rows.

5.4.b Inaccuracy of the sensors

The onboard GPS is used to follow the pre-defined course and to record the actual position when taking an image. Next to the GPS various other sensors (IMU) are available to record the rotation of the aircraft and its attached camera. The information from the GPS ensures that the photos taken along the planned trajectory, to ensure there is enough overlap between the images.

This information can be used to estimate the footprint of the images on the ground. Those footprints can then be used to estimate which images may contain matching features and allow greater computational efficiency.

A small comparison with the fitted camera parameters on 37 images¹⁰ with the model by Pix4D learned that the sensors on the Trimble X100 had doubtful accuracy. The roll is accurate within 10 degrees, the pitch has an average error of 10° but the yaw was more or less reversed. Having an error of 165° on average not only systematic (compared to 180°) but also at least 15° of random error. This inaccuracy is unfortunately linked to the UAV used.

This is a problem when estimating the footprints of the images for more efficient image matching as the estimation of the overlap might be wrong. This effect can be seen in Figure 5.19 where the footprints rotate and the distribution of shots is more randomly distributed than planned. As described by (Barazzetti, 2010) errors should not exceed 25° of rotation and 30% in scale for best results.

5.4.c Computational intensity

The computational load can be reduced by reducing the amount of images involved. In chapter 2.5 (page 10) this problem is seen from a mathematical approach. Here the problem is solved empirically: how many images do we need? More images requires more processing power and more importantly more working memory (RAM), unless heavy duty computers are available this forms the practical limit to analysing large areas using this method.

Pix4D was run multiple times on the same sandy area (shown in white). In Figure 5.19 the estimated footprints of the aerial photos are shown in the top row. The criteria for inclusion became increasingly strict from left to right. First all images that are estimated to be within 50 metres are used in the process, the same holds for a 25 meter buffer. Further restricting is the requirement of having at least 1% of the image overlap with the white box. This is comparable to having a zero metre buffer around the box. This is overlap requirement is further increased till the point that only the images perfectly overlying the area remain.

As can be seen in Figure 5.16 at least some of the surrounding images should be taken into account to allow points to be generated over the whole surface. Using more images allows connections to be made 'around the back', connecting the images using strong links elsewhere. For example images with matches on the dam or oyster reefs form a stiff frame to match weakly linked image in between.

¹⁰ The images used were from the dam (region 1), to ensure enough matches for proper matching.




Figure 5.19: Comparison of region sizes. Top row: estimated images footprints; centre row: footprints after correction; bottom row: resulting DSM from Pix4D. (Background: PDOK Luchtfoto)

5.4.d Classification

Attempts were made to automatically detect if a pixel covered water, sand or the dam. This could help reducing the computational load further by not processing images lacking information (photos containing only water for example).

With the data from the Trimble X100 survey the problem was the Automatic (Daylight) White Balance from the camera. This setting made the image grey on average, as shown in Figure 5.20: red, green and blue are equally present in all images. This is a reasonable assumption for your holiday pictures, that include a bright sky that is grey (white). Absolute information is lost, it is therefore impossible to check if an image is more red, green or blue than its neighbour. It is therefore impossible to classify the images based on their colour distribution alone.




Figure 5.20: Ternary diagram of the average colours of the images from the first survey. All images (shown as many small dots) have an equal distribution of red, green and blue.

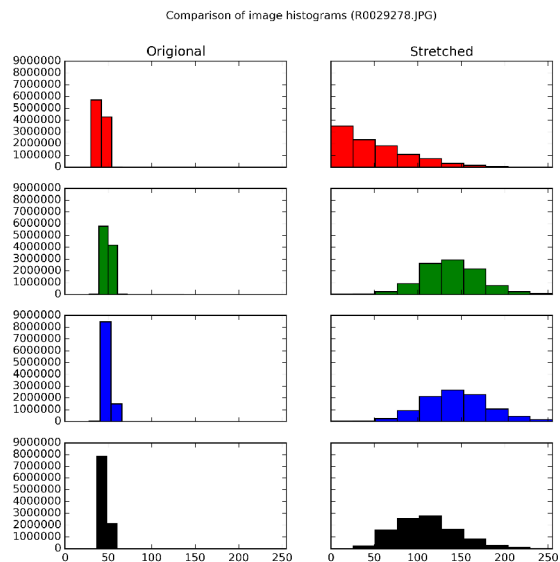


Figure 5.21: Histogram differences between the original and stretched histogram of R0029278.JPG (see Figure 5.22).

5.4.e Histogram stretching

Histogram stretching was thought to be useful as a method of contrast enhancement. Increasing the contrast would allow better detection of features in images.

The images from the camera use only very little of the available spectrum (around $\frac{1}{5}$ of the available range). When stretching the histogram the small range is mapped to the full extent, this process is visible in Figure 5.21. Examples of resulting images are shown below.

The approach was not effective as the feature detection and description mechanisms in the used software were not sensitive to this approach, as they work on the differences between pixels. Those inter pixel differences are stretched accordingly.




Figure 5.22: Original image on the left, image with stretched histogram on the right. Underwater plants become visible through the water. (R0029278.JPG, see Figure 5.21.)




Figure 5.23: The effect is less dominant in this image that includes part of the (bright) dam (original left, stretched histogram on the right).

6. Possible applications and validation

In this chapter the generated Digital Surface Models from chapter 5.3 will be compared to the reference datasets (LiDAR and GPS-RTK measurements, see chapter 4). Furthermore possible applications of the results will be discussed.

6.1 Validation, comparison to other datasets

Although not tested here systematic errors may cause subsequent surveys to have a relative offset, in particular when relying on auto-calibration based on the UAVs position. The geometry itself will be preserved (K. Bisdom, personal communication), but the relative offsets will complicate comparison between surveys. In the first survey this problem is apparent due to the lack of GCPs, and will have to be dealt with.

6.1.a First survey

Due to the lack of reference data (RTK) for the first flight those results can only be compared to the dam using AHN data.

In Figure 6.1 a comparison with AHN of the attempt at reconstructing the whole survey is shown. Not only are the absolute Z-coordinates off, possibly due to the lack of GCPs, but the elevation difference (± 40 metre within the DSM) is far from realistic (expected: -2 to $+2$ metre). Furthermore a 'valley' shape can be seen in Figure 6.1, as the east and west ends curl upwards. Nonetheless the difference on the dam is constant. This result can not be used for any morphological comparison, though the orthophoto might be usable.

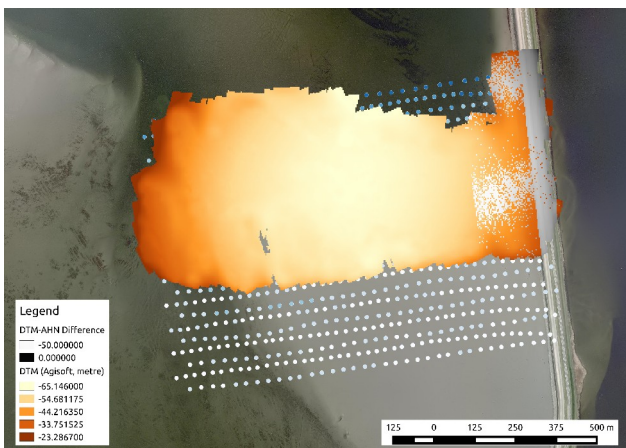


Figure 6.1: Difference (in grey) between the DSM rendered by Agisoft Photoscan and AHN. A concave distortion can be seen. (Background: PDOK Luchtfoto)

The three regions can best be compared to the results of the other software package (Agisoft Photoscan versus Pix4D), as shown in Figures 6.2 and 6.3 for Region 1 and 3 respectively. No results were returned for Region 2, this is remarkable as (Paree, 2014) showed that it is possible to have (accurate) results in shallow water.

Please note that the morphology between the two software packages is very different. The Agisoft

model of the dam shows a concave shape, similar to the one in Figure 6.1. Furthermore the results on Region 3 show different morphologies in the results of Agisoft and Pix4D: an upwards slope to the north in Agisoft, to the south in Pix4D.




Figure 6.2: Region 1




Figure 6.3: Region 3

In all cases no ground control points were used. The auto-calibration, based on the GPS measurements, as described by (Turner, 2014), did not provide stratifying results.

6.1.b Second survey

For the second survey a RTK survey is available from a few days earlier (see chapter 4.2). The generated DSM can be compared to this model. As can be seen in Figure 6.4 the error is largely concentrated around zero. In fact, the average error is only -0.012 metre, with a standard deviation of 0.12. In other words: 68% of the differences is in the range of -0.012 ± 0.12 . This is not as accurate as required, but the distribution of differences is important. As can be seen in Figure 6.5 the changes concentrate in three areas: near the edges of the model (north and south); parallel to the dam; a braided pattern from north-west to south-east. As we can see in Figures 5.12 and 5.13 the line parallel to the dam corresponds with a steep slope. Erosion may have driven the coast back in the days between the surveys. The braided pattern could be the result of (lateral) movement of the tidal channels. When those three areas are ignored the accuracy is still only -0.28 ± 0.06 metre ($\mu \pm \sigma$).




Figure 6.4: Histogram of differences between the DSM and RTK survey.




Figure 6.5: Difference between the DSM and RTK survey.

Even when cherry picking the best matches between the DSM and RTK survey the model does not

meet our criteria of centimetre accuracy. Still the results are still more accurate than predicted based on stereo-pairs only. A variance of 0.055 was estimated in chapter 5.2, in practice a variance of 0.014 was realised, slightly worse than the variance estimated for convergent photogrammetry networks (0.005).

6.2 Possible applications

As shown in the previous paragraph current models do not meet the requirement of centimetre accuracy. Larger features, like channels or (sand) ridges may still be extracted from the generated DSM. Their size allows a more robust handling of the measurement errors and their presence does not depend centimetre accuracy.

Apart from monitoring morphological changes the true colour images obtained in this survey could be used for other purposes. For example morphological patterns could be detectable, even if the data is less detailed than the ripples themselves (Visser, 2014). The series of (raw) images could for example be used in detecting animals or plants on the shoal, as shown by (Sirmacek, 2014).

Equipped with (slightly) different (measuring) devices UAVs can be used for a variety of other applications. Monitoring of crops using infrared cameras; monitoring wildlife; measurements on weather conditions at different altitudes or even parcel delivering.

7. Conclusions and recommendations

UAV based stereophotogrammetry has proven to be a reliable surveying technique in terrains with high texture. The shoals at the Oesterdam provide a challenge, as they contain low texture, complicating feature detection and matching – the basis of stereophotogrammetry.

The research questions can now be addressed:

How is elevation data acquired from stereophotography (photogrammetry)?

A series of images is linked together by detecting similar objects across images. This is done by detecting key features (Difference of Gaussians) and describing them (Scale Invariant Feature Transform: SIFT). Similar features have maximal similarities between their descriptions. The extracting of key features can be challenging in areas with low texture, like shoals. After feature selection and matching a model of features, ground control points and cameras is compiled. This model is then iteratively optimised. The result is a sparse point cloud and detailed information on the camera, scale, world position and orientation of the model. This information is then used as the bases for the creating of a dense point.

What accuracy can be expected from stereophotogrammetry?

Theoretically sub-centimetre accuracies are possible in UAV stereophotogrammetry, as the addition of multiple images over stereo-pairs increases the accuracy. In practice only decimetre accuracy can be reached on the low texture shoals. Higher accuracies might be possible on different terrain.

What processing software exists, and how does it perform in DSM generation?

Various software solutions exist, only four of them were put to the test in this report. Performance varies widely across packages and reports (literature). The new 'low texture' mode in Pix4D provides promising results. Agisoft Photoscan, praised by many papers, failed to generate a adequate DSM for the shoal.

How do the results compare to other surveying methods?

The results are very fine grained, only $10 \text{ cm}/\text{pixel}$, compared to many meters between RTK measurements. Laser scanning may produce more fine grained and accurate results at the cost of speed. Limited results are returned on water, where shallow water may still be surveyed using RTK measurements.

What applications do exists for those DSM's?

Apart from monitoring morphological changes various other applications exist, for example the detection of morphological patterns. Furthermore the hundreds to thousands of photos taken in a single survey contain all sorts of information.

We can now answer the main question:

Could drone data replace conventional surveying?

Not now, not completely. Although promising the accuracy did not meet the required centimetre precision. The availability by-products like true colour orthophotos could still provide substantial benefit over conventional surveys.

7.1 Recommendations

To provide the best coverage, full datasets should be processed rather than excerpts should be processed were possible. This requires significant processing power, most notably more RAM (up to 30 GB).

In theory higher accuracy could be possible by lowering the flight height, therefore taking pictures with a smaller ground pixel size. Though this requires the camera to take pictures at a higher rate, and introduces the risk of losing images as the camera is not ready yet.

Further applications of the true colour images could be explored. A possible application could be the monitoring of the flora and fauna on the shoal. Although the accuracy is not sufficient to monitor morphological changes might be possible to extract larger features like channels and ridges.

Most software packages used provided limited to no configuration options. More customisable software could be better tailored to low texture areas.

When the second model becomes available it will be possible to compare the two DSMs. It will then be possible to verify if it is possible to detect morphological changes from subsequent DSMs.

List of figures

Most maps in this thesis were produced using QGIS 2.4.0 – Chugiak¹¹. Images of 3D models and point clouds were made using CloudCompare v2¹².

Many maps refer to the 'PDOK Luchtfoto' or 'BRT Achtergrondkaart' as background image. This set of aerial imagery¹³ ('PDOK Luchtfoto') and topographic maps¹⁴ ('BRT Achtergrondkaart') is published by the Dutch government under the Creative Commons Attribution Non-Commerical license¹⁵.

Figure 1 (on the cover): Sand nourishment at the Oesterdam in Zeeland, The Netherlands, January 2014. (Photo: (Rijksoverheid, 2014a))	1
Figure 1.1: Location of the Oesterdam, marked by a red star in the Dutch province of Zeeland. Three other major works of the Delta Plan can be seen on the sea-side, from north to south: Haringvlietdam; Brouwersdam and Oosterscheldekering. (Background: BRT Achtergrondkaart) ...	6
Figure 2.1: Single camera with two objects (red and green) that appear to have the same size.	8
Figure 2.2: Two camera's and a single object (red).	8
Figure 2.3: Model of a lens. (Taken from (ADAM Tech., 2010)).	8
Figure 2.4: The basic principle of stereophotogrammetry. (Taken from (Haggrén, 2002)).	8
Figure 2.5: Example of the Difference of Gaussians. A Gaussian pyramid is shown in the top row, the matching Laplacian pyramid in the bottom row. (Taken from (Wallace, 2011).)	10
Figure 2.6: Illustration of the generation of a SIFT descriptor. (Taken from (Levi, 2013).)	10
Figure 2.7: Accuracy in relation to increasing images. The average distance is taken as 100 metre, the focal length at 10 mm and pixel size as 5 µm.	11
Figure 4.1: Actueel Hoogtebestand Nederland, a combination of tiles r49bz1 and r49dn1. (Background: PDOK Luchtfoto)	15
Figure 4.2: RTK survey of May 2014, consisting of 5355 points. (Rijkswaterstaat) (Background: PDOK Luchtfoto)	16
Figure 4.3: Comparison of the two UAV datasets. Each point represents a picture taken. (Background: PDOK Luchtfoto)	17
Figure 5.1: Regions of interest. As illustration the orthophoto from the second survey is shown as background. All three areas are fully covered by the first survey.	20
Figure 5.2: Variance in relation to increasing images. Given the parameters of the first survey from Table 2.	21
Figure 5.3: Orthophoto, rendered using own software. Based on coordinates Agisoft. Note the missing, unmatched images in the north.	22
Figure 5.4: Dam, rendered using MicMac. Camera positions can be seen on top. (CloudCompare) ..	23
Figure 5.5: Only 39 points were rendered using IQmulus. Therefore only few 'dots' are visible. (CloudCompare)	23
Figure 5.6: Point cloud rendered using Agisoft. (CloudCompare)	23
Figure 5.7: Digital Surface model as rendered by Agisoft. A concave distortion can be seen.	

11 <http://www.qgis.org/en/site/> (29-10-2014)

12 <http://www.danielgm.net/cc/> (29-10-2014)

13 <http://www.nationaalgeoregister.nl/geonetwork/srv/dut/search?#|5bda5b66-ae22-4a12-9236-13ad2c4d0730> (last visited: 29-10-2014)

14 <http://www.nationaalgeoregister.nl/geonetwork/srv/dut/search#|3373be8c-8539-4763-bc22-eba23ac1898fa> (last visited: 29-10-2014)

15 <http://creativecommons.org/licenses/by-nc/2.0/> (29-10-2014)

(Background: PDOK Luchtfoto)	23
Figure 5.8: Region 1, Agisoft Photoscan.	24
Figure 5.9: Region 3, Agisoft Photoscan.	24
Figure 5.10: Region 1, Pix4D.	24
Figure 5.11: Region 3, Pix4D.	24
Figure 5.12: DSM (10 cm/px) as rendered by Kragten using Pix4D.	25
Figure 5.13: A by-product of the process, a full colour (RGB) orthophoto of the area.	25
Figure 5.14: Slope-aspect plot of the DSM. All slopes < 8.5° are suppressed to reduce noise.	25
Figure 5.15: Hillshade of the DSM.	25
Figure 5.16: Graph generated by Pix4D illustrating which photos are best connected (darker) and which photos lack matching features (white). Due to the lack of matches two independent groups are formed. (Taken from the 1% overlap run, see 5.4.c. See Figure 5.3 for a visual reference of the area.)	26
Figure 5.17: MicMac matches between R0029210.JPG and other images, shown on the top and bottom rows.	26
Figure 5.18: MicMac matches between R0029468.JPG and other images, shown on the top and bottom rows.	26
Figure 5.19: Comparison of region sizes. Top row: estimated images footprints; centre row: footprints after correction; bottom row: resulting DSM from Pix4D. (Background: PDOK Luchtfoto)	28
Figure 5.20: Ternary diagram of the average colours of the images from the first survey. All images (shown as many small dots) have an equal distribution of red, green and blue.	29
Figure 5.21: Histogram differences between the original and stretched histogram of R0029278.JPG (see Figure 5.22).	29
Figure 5.22: Original image on the left, image with stretched histogram on the right. Underwater plants become visible through the water. (R0029278.JPG, see Figure 5.21.)	29
Figure 5.23: The effect is less dominant in this image that includes part of the (bright) dam (original left, stretched histogram on the right).	29
Figure 6.1: Difference (in grey) between the DSM rendered by Agisoft Photoscan and AHN. A concave distortion can be seen. (Background: PDOK Luchtfoto)	30
Figure 6.2: Region 1	31
Figure 6.3: Region 3	31
Figure 6.4: Histogram of differences between the DSM and RTK survey.	31
Figure 6.5: Difference between the DSM and RTK survey.	31

References

- ADAM Tech., 2010:** ADAM Technology, "How Accurate is Photogrammetry?", 2010, <http://www.adamtech.com.au/Blog/?p=167> (last visited: 2014-10-29)
- Aerometrex, 2012:** Aerometrex, "A brief history of aerial photogrammetry", 2012, <http://aerometrex.com.au/blog/?p=203> (last visited: 2014-11-14)
- Agisoft, 2014:** Agisoft, "About Agisoft", 2014, <http://www.agisoft.com/about/> (last visited: 2014-10-29)
- AHN, 2009:** Actueel Hoogtebestand Nederland, "AHN-2 beschikbaar", 2009, <http://www.ahn.nl/nieuws/2009/10/ahn-2-beschikbaar.html> (last visited: 2014-10-29)
- AHN, 2012:** Actueel Hoogtebestand Nederland, "Vluchten AHN2 afgerond", 2012, <http://www.ahn.nl/nieuws/2012/04/vluchten-ahn2-afgerond.html> (last visited: 2014-10-29)
- AHN, 2013:** Actueel Hoogtebestand Nederland, "Kwaliteitsdocument AHN2", 2013, http://www.ahn.nl/binaries/content/assets/hwh---ahn/common/wat+is+het+ahn/kwaliteitsdocument_ahn_versie_1_3.pdf#page=8 (last visited: 2014-10-29)
- AHN, 2014:** Actueel Hoogtebestand Nederland, "Inwinning AHN2014 op de helft", 2014, <http://www.ahn.nl/nieuws/2014/03/inwinning-ahn2014-op-de-helft.html> (last visited: 2014-10-29)
- Barazzetti, 2010:** Barazzetti, L., Scaioni, M., Remondino, F., "Orientation and 3D modelling from markerless terrestrial images: combining accuracy with automation", 2010 (doi : 10.1111/j.1477-9730.2010.00599.x)
- Bemis et al. 2014:** Bemis, S., Micklethwaite, S., Turner, D., Jamers, M.R., Akciz, S., Thiele, S., Bangash, H.A., "Ground-based and UAV-based photogrammetry: A multi-scale, high-resolution mapping tool for Structural Geology and Paleoseismology", 2014 (doi : 10.1016/j.jsg.2014.10.007)
- Bisdom, 2014:** Bisdom, K., "About GeoBee", 2014, <http://homepage.tudelft.nl/f0b7w> (last visited: 2014-11-12, click on 'GeoBee')
- CH Robotics, 2014:** CH Robotics, "Understanding Euler Angles", 2014, <http://www.chrobotics.com/library/understanding-euler-angles> (last visited: 2014-10-29)
- Cléry, 2013:** Cléry, I., "Interface MicMac, user documentation", 2013, <http://logiciels.ign.fr/IMG/pdf/docinterface.en.pdf> (last visited: 2014-11-14)
- Climate Buffer Coalition, 2014:** Climate Buffer Coalition, "The Netherlands naturally safe", 2014, <http://www.klimaatbuffers.nl/english-homepage-2> (last visited: 2014-10-28)
- Computa Maps, 2011:** Computa Maps, "Stereo Photogrammetry – How We Create 3D Models ", 2011, <http://www.computamaps.com/newsletter/2-2/newsletter2-2.html> (last visited: 2014-11-15)
- FAA, 2014:** William J. Hughes Technical Center, NSTB/WAAS T&E Team, "Global Positioning System (GPS); Standard Positioning Service (SPS) Performance Analysis Report", 2014, http://www.nstb.tc.faa.gov/reports/PAN86_0714.pdf#page=22 (last visited: 2014-10-29)
- Grauman & Leibe, 2011:** Grauman, K., Leibe, B., "Visual Object Recognition", 2011 (ISBN : 9781598299687)
- Haggrén, 2002:** Haggrén, H., "Seven pictures for stereography", 2002, http://foto.hut.fi/publications/paperit/hhaggren/bildteknik_2002/stereo_photography.htm (last visited: 2014-11-15)
- IGN, 2014:** Unknown, "MicMac, Aperol, Pastis and Other Beverages in a Nutshell!", 2014, <http://logiciels.ign.fr/IMG/pdf/docmicmac.pdf> (last visited: 2014-10-29)
- IQmulus, 2014:** IQmulus, "Project overview", 2014, <http://www.iqmulus.eu/index.php/project-overview>

(last visited: 2014-10-29)

James & Robson, 2012: James, M.R. and Robson, S., "Straightforward reconstruction of 3D surfaces and topography with a camera: Accuracy and geoscience application.", 2012 (doi: 10.1029/2011JF002289)

Levi, 2013: Levi, G., "A Short introduction to descriptors", 2013, <http://gilscvblog.wordpress.com/2013/08/18/a-short-introduction-to-descriptors/> (last visited: 2014-11-17)

Melanson, 2012: Donald Melanson, "Drone-maker Parrot invests \$7.5 million in two EPFL spin-offs, sets sights beyond toys", 2012, <http://www.engadget.com/2012/07/26/drone-maker-parrot-invests-7-5-million-in-two-epfl-spin-offs-s/> (last visited: 29-10-2014)

NASA, 2004: NASA, "ASTER Global Digital Elevation Map Announcement", 2004, <http://asterweb.jpl.nasa.gov/gdem.asp> (last visited: 2014-11-15)

Natuurmonumenten, 2014a: Natuurmonumenten, "Nationaal Park Oosterschelde, Zandhonger", 2014, <https://www.natuurmonumenten.nl/zandhonger> (last visited: 2014-10-29)

Natuurmonumenten, 2014b: Natuurmonumenten, "Oesterdam: innovatief project Oosterschelde", 2014, <https://www.natuurmonumenten.nl/oesterdam-innovatief-project-oosterschelde> (last visited: 2014-10-29)

NGR, 2014: Nationaal Georegister, "NGR: Nationaal Georegister", 2014, <https://www.pdok.nl/nl/ngr/ngr-home> (last visited: 2014-10-29)

Ortega, 2013a: Ortega, M., "Selecting the images acquisition plan type", 2013, <https://support.pix4d.com/entries/27002457-Step-1-Before-Starting-a-Project-1-Designing-the-Images-Acquisition-Plan-a-Selecting-the-Images-Acqu#label5> (last visited: 2014-10-29)

Ortega, 2013b: Ortega, M., "System Requirements: Minimum and Recommended Hardware and Software Specifications", 2013, <https://support.pix4d.com/entries/26799527-System-Requirements-Minimum-and-Recommended-Hardware-and-Software-Specifications> (last visited: 2014-10-29)

Paree, 2014: Paree, E., "STOWA, Casus Sophiastand (Rijkswaterstaat)", 2014, <https://www.youtube.com/watch?v=10riYZcmUwg> (last visited: 2014-11-16)

Parrot, 2014: Parrot, "Parrot Q2 2014 earnings, investors presentation", 2014, <http://www.parrotcorp.com/en/financialpublications/q22014earningspresentationroadshow> (last visited: 2014-10-29)

Pix4D, 2014: Pix4D, "About us", 2014, <http://pix4d.com/about-us/> (last visited: 2014-10-29)

Rapidlasso, 2012: Rapidlasso, "las2dem", 2012, <http://rapidlasso.com/lastools/las2dem/> (last visited: 2014-10-29)

Riegl, 2014: Riegl, "UAS/UAV Laser Scanning", 2014, <http://www.riegl.com/products/uasuv-scanning/> (last visited: 2014-11-16)

Rijksoverheid, 2013: Rijksoverheid, "Open Overheid", 2014, <https://data.overheid.nl/openoverheid> (last visited: 2014-10-29)

Rijksoverheid, 2014a: Rijksoverheid, "Zandsuppletie Oesterdam goed voor natuur en veiligheid", 2014, <http://www.rijksoverheid.nl/nieuws/2014/01/29/zandsuppletie-oesterdam-goed-voor-natuur-en-veiligheid.html> (last visited: 2014-10-29)

Rijksoverheid, 2014b: Rijksoverheid, "Data.overheid.nl: het opendataportaal van de Nederlandse overheid", 2014, <https://data.overheid.nl/> (last visited: 2014-10-29)

Sinha, 2014: Sinha, U., "SIFT: Finding keypoints ", 2014, <http://www.aishack.in/tutorials/sift-scale-invariant-feature-transform-keypoints/> (last visited: 2014-11-17)

Sirmacek, 2014: Sirmacek, B., "Animal observation case study: Automatic population counts of birds from low-cost terrestrial images", 2014

- Sona et al., 2014:** Sona, G., Pinto, L., Pagliari, D., Passoni, D., Gini, R., "Experimental analysis of different software packages for orientation and digital surface modelling from UAV images", 2014 (doi: 10.1007/s12145-013-0142-2)
- Trimble, 2014:** Trimble, "Specifications", 2014, <http://uas.trimble.com/specifications> (last visited: 2014-10-29)
- Turner, 2014:** Turner, D., Lucieer, A.; Wallace, L., "Direct Georeferencing of Ultrahigh-Resolution UAV Imagery", 2014 (doi: 10.1109/TGRS.2013.2265295)
- Ven, 2014:** Van, P. van de, "STOWA, Bedrijfspresentatie Kragten", 2014, http://www.youtube.com/watch?v=x_lwKZbPPhk (last visited: 2014-11-15)
- Visser, 2014:** Visser, A.H., "Detecting morphology changes due to oyster reefs in a tidal basin using Terrestrial Laser Scanning", 2014
(<http://repository.tudelft.nl/assets/uuid:929a9c1a-2b9e-4703-9f75-c6142d4de162/BScThesisAHVisser.pdf>)
- Wallace, 2011:** Wallace, E., "CS 143: Hybrid Images", 2011, <http://cs.brown.edu/courses/cs143/2011/results/proj1/edwallac/> (last visited: 2014-11-17)

Appendix: list of photogrammetry software packages

Name	Website	Developer	Platform	Price	Notes
Bundler	http://www.cs.cornell.edu/~snavey/bundler/	N. Snavely	Windows, Linux	Open Source (GPL)	
SFMToolkit	http://www.visual-experiments.com/demos/sfmtoolkit/	H. Astre	Windows	Open Source (MIT)	Combination of Bundler and CMVS/PVMS ¹⁶ .
Python Photogrammetry Toolbox (PPT)	http://184.106.205.13/arcteam/ppt.php	Arc-Team	Windows, Linux	Open Source	A Python toolkit, with own GUI.
VisualSFM	http://ccwu.me/vsfm/	C. Wu	Windows, Linux, Mac	Free for personal, academic use.	Frontend for PVMS/CMVS ¹⁶ , preparation for CMP-MVS (CTU Prague) and SURE.
3DF Samantha	http://www.3dflow.net/technology/samanta-structure-from-motion/	3DFlow	Windows	Free for non-commercial use.	Only a SfM implementation.
SURE	http://www.ifp.uni-stuttgart.de/publications/software/sure/index.en.html	Institut für Photogrammetrie, Universität Stuttgart	Windows, Linux	Free for non-commercial use. Commercial product: 'nFrames'.	
MicMac	http://logiciels.ign.fr/?Micmac	IGN	Windows, Linux, Mac	Open Source (CeCILL-B)	
CMVS ¹⁶	http://www.di.ens.fr/cmvs/	Y. Furukawa	Linux, (Windows)	Open Source (GPL v3)	Intermediate step between Bundler and PMVS2 ¹⁶ .
Photosynth	https://photosynth.net/	Microsoft	Windows, Windows Phone, iOS	Free	

¹⁶ Clustering Views for Multi-view Stereo (CMVS)
Patch-based Multi-view Stereo Software (PVMS)

Arc3D	http://www.arc3d.be/	KU Leuven	N/A	Unknown	Webservice
CMP SfM Webservice	http://ptak.felk.cvut.cz/sfm-service/	CTU Prague	N/A	Unknown	Webservice
Autocad 123D Catch	http://www.123dapp.com/catch	Autodesk	Windows, Android, iOS	Free for non-commercial use, 'premium' or commercial models require a \$10 monthly subscription.	Webservice
My3DScanner	http://www.my3dscanner.com/	My3DScanner	N/A	Free	Webservice
Agisoft Photoscan	http://www.agisoft.com/	Agisoft	Windows	\$ 180 - \$ 3500, educational licenses \$ 60 - \$ 550	
Acute3D	http://www.acute3d.com/	Acute3D	Windows	Free for small models, starting at € 2500 annually for larger models.	
PhotoModeler	http://www.photomodeler.com/	PhotoModeler	Windows	\$ 1150 - \$ 3500	
3DF Zephyr Pro	http://www.3dflow.net/3df-zephyr-pro-3d-models-from-photos/	3DFlow	Windows	\$ 3200	Complete solution. (Compared to 3DF Samantha)
Pix4D Mapper (former Pix4UAV)	http://pix4d.com/	Pix4D	Windows	Basic free version, € 2600 annually. Educational license € 1500 once. Non-commercial license € 3900 once.	Standalone and webservice
ERDAS IMAGINE (Photogrammetry)	http://www.hexagongeospatial.com/products/photogrammetry/imagine-photogrammetry-%28formerly-lps%29/overview	Hexagon	Windows	Unknown	Formerly Leica Photogrammetry Suite.