Projeto Redentor

White Paper

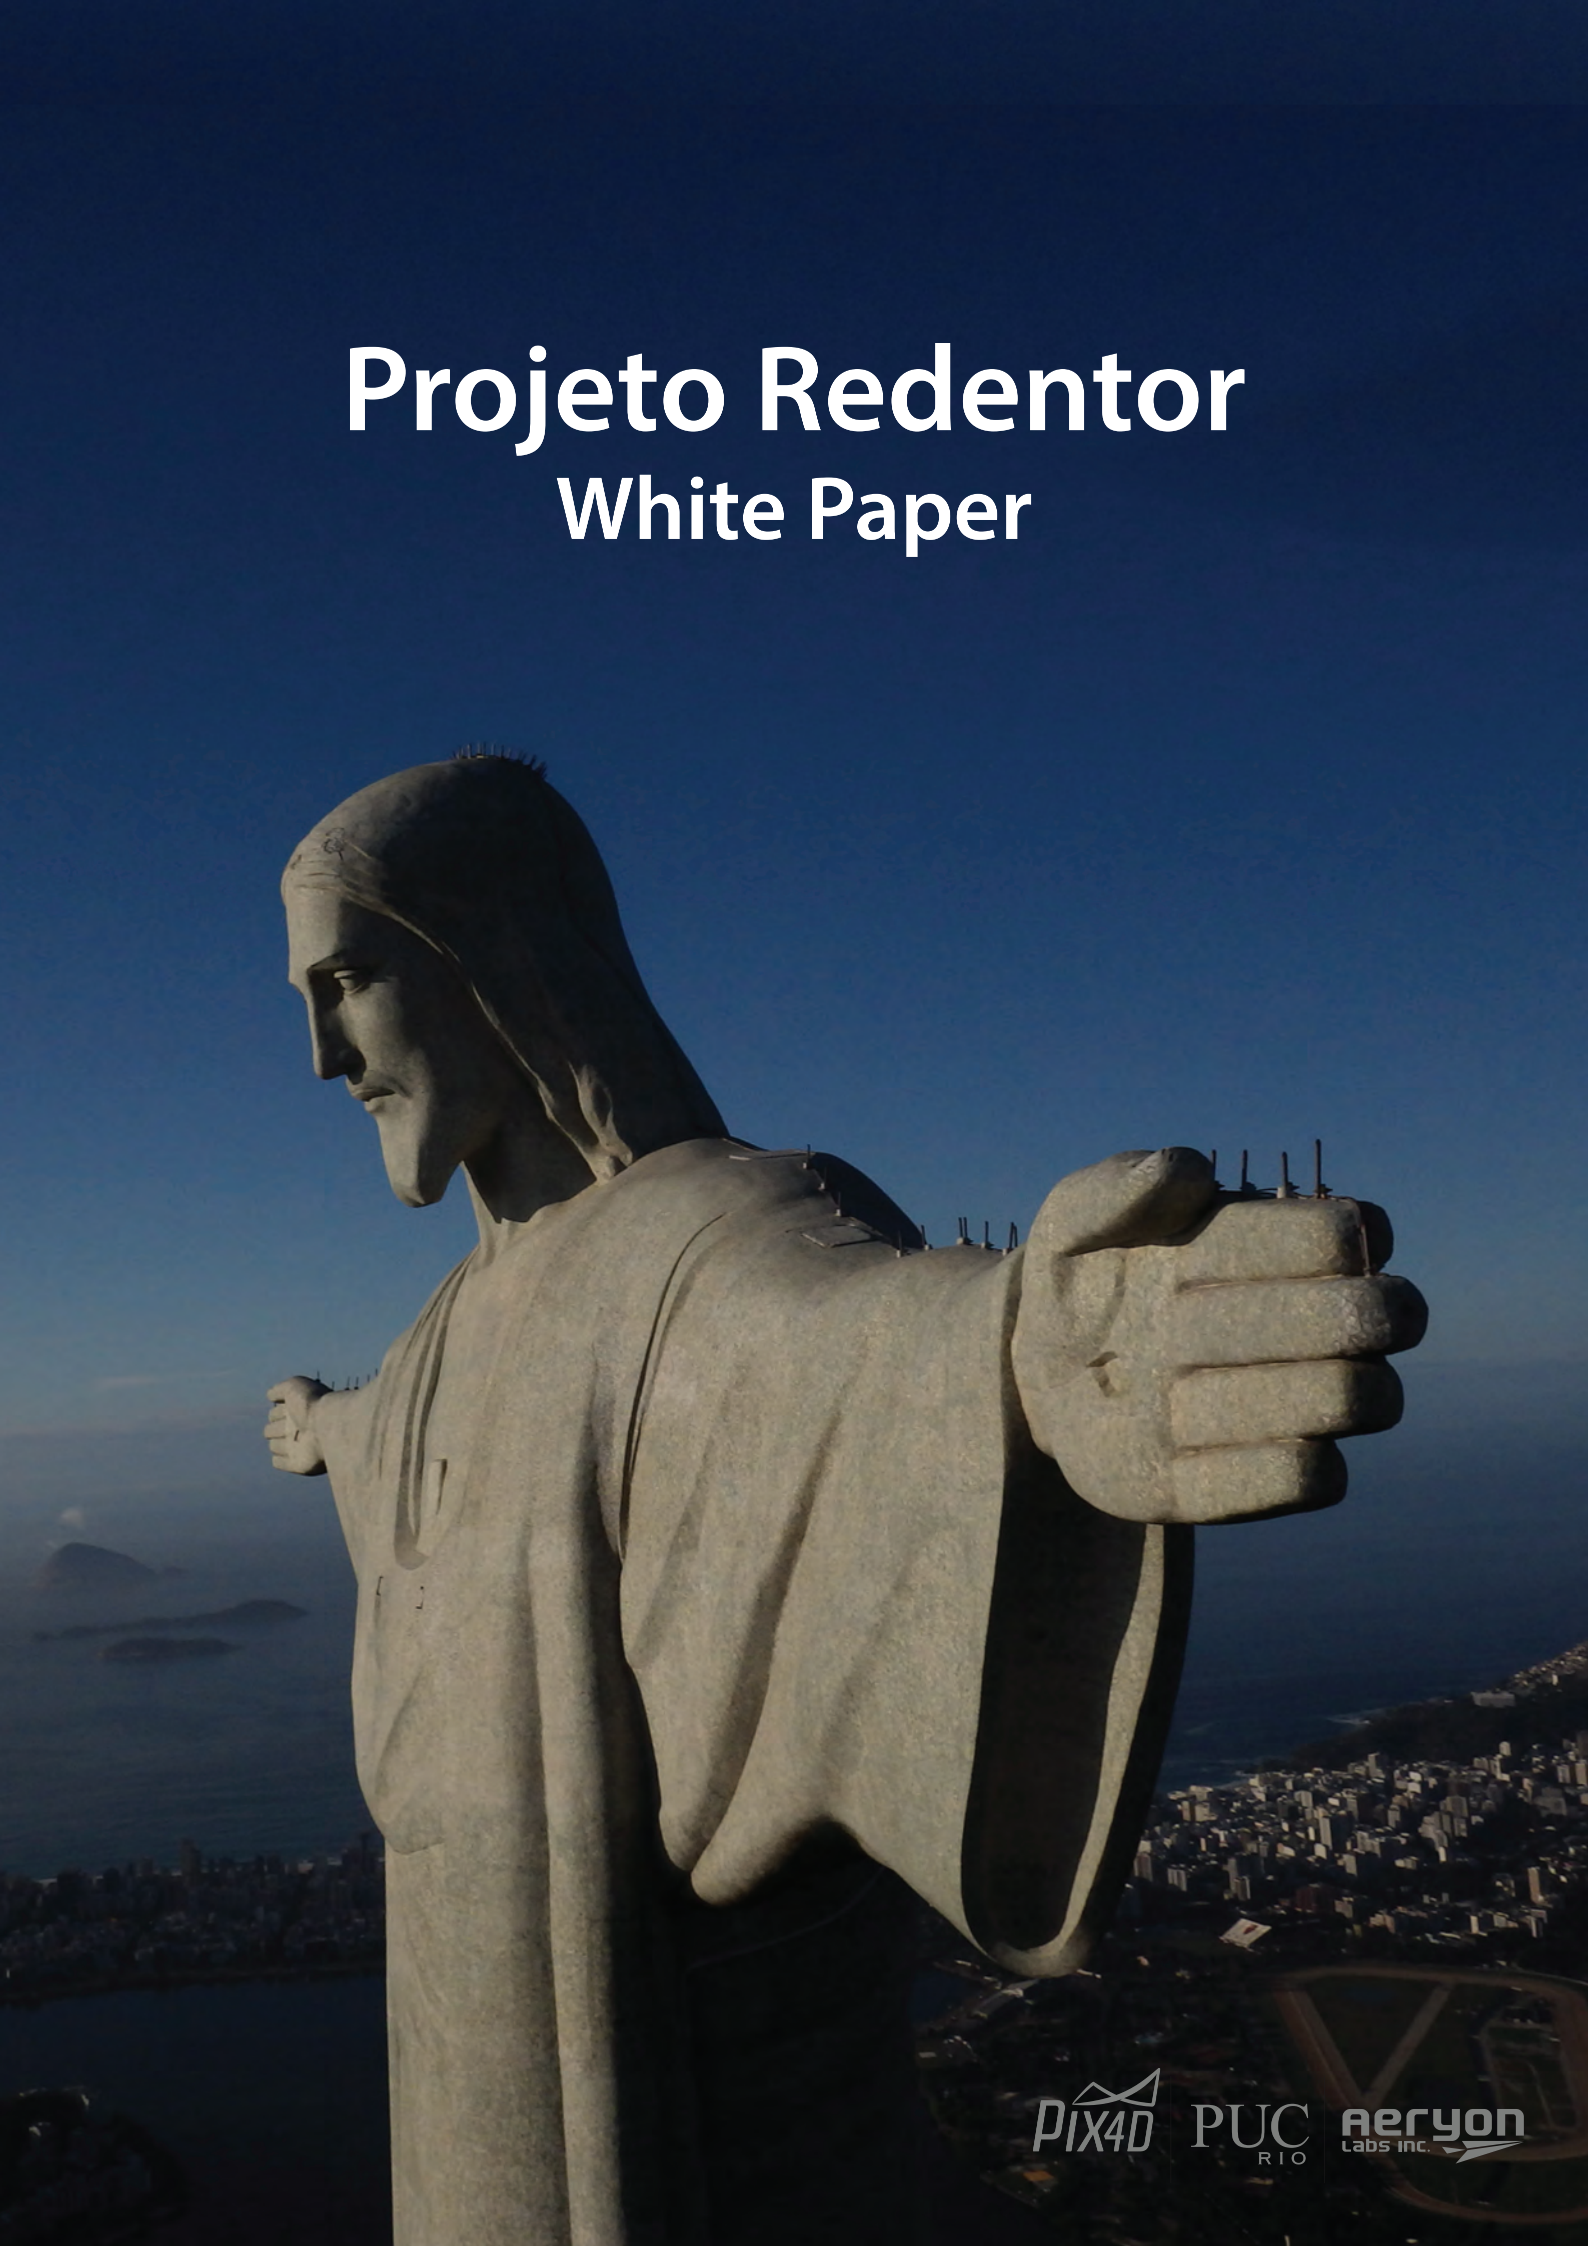

Projeto Redentor: High-resolution 3D modelling of large, hard-to-reach objects

How can one create a highly accurate 3D model of large objects, in hard-to-reach locations, for which traditional scanning methods cannot be applied? Is it possible to reconstruct both the object and its surroundings (for example, the Christ the Redeemer statue of Rio de Janeiro) and combine everything into one model?

This project, Projeto Redentor, was planned and executed by Pix4D in partnership with Aeryon Labs Inc. and PUC University of Rio de Janeiro to create a use case that shows the feasibility, accuracy and efficiency of image-based 3D modelling for large, extremely detailed objects in complex surroundings.

Authors: *Lisa Chen, Sonja Betschart (Pix4D) / April Blaylock (Aeryon Labs Inc.)*

1. INTRODUCTION

The Christ the Redeemer (Cristo Redentor in Portuguese) statue is the icon of the city of Rio de Janeiro and the most important statue in Brazil. It was designed by the Brazilian engineer and architect Heitor da Silva Costa, in collaboration with artist Carlos Oswald, sculptor Paul Landowski and French engineer Albert Caquot between 1922 and 1931. The monument was opened on October 12, 1931.

The statue measures a total height of 38 meters (30 meters for the statue itself and 8 meters for the base) and its arms stretch 28 meters wide. It is located at the peak of the 700 meter high Corcovado Mountain in the Tijuca Forest National Park, surrounded by a very narrow terrace that allows visitors from around the world to admire not only the statue, but also the view over the city of Rio.

Until today, all the 3D models and replicas of the statue have been designed by hand. Accurate 3D reconstruction has not been possible since technologies, such as LiDAR, have not been able to scan the complete statue, due to its size, location, difficult access and challenging weather conditions.

The Núcleo de Experimentação Tridimensional (NEXT) department of PUC University in Rio de Janeiro contacted Pix4D to determine if a 3D reconstruction could be achieved using image processing technology and Unmanned Aerial Vehicles (UAVs) for data acquisition. Pix4D took on the challenge, in collaboration with the Canadian UAV manufacturer Aeryon Labs Inc. for the data acquisition and PUC's NEXT Lab for the organization of all project logistics, including special permission to fly UAVs at the Christ the Redeemer heritage site. Safe operation was an extremely important aspect of this project, for both visitors as well as the statue itself. For this reason, Aeryon Labs Inc. and their highly reliable and enduring quad rotor aerial platforms were chosen as the UAV partner for this project.

The success of this project was dependent on the project team overcoming a number of challenges, including: the size of the statue and its remote location on the peak of the mountain, changing weather and wind conditions, restricted hours for data acquisition (flights could only take place before and after visiting hours), inconsistent lighting conditions and shadows in early morning and late afternoon resulting in the inhomogeneous color balancing of imagery, just to mention the most important.

This White Paper describes the challenges of this project and how they were overcome. It also highlights the workflow – from data acquisition to image processing and the generation of the 3D model (point cloud and textured mesh) as the final output result.

Watch the project video on YouTube: http://youtu.be/-ucLIckILT4
Discover the 3D model on Sketchfab:

  Christ Statue only: https://sketchfab.com/models/fa8c4b24898c40ec86a40449cec47474
  Corcovado and Christ Statue: https://sketchfab.com/models/41b2a83c75ab4040b490a7d61fdb08c6

2. THE EQUIPMENT

2.1 Aeryon Scout™

The Aeryon Scout™ is a Vertical Take-Off and Landing (VTOL), quad rotor small UAV – ideal for commercial and industrial users. The Aeryon Scout can fly lower and slower than manned aircraft, as well as in confined environments where fixed-wing UAVs cannot operate safely. The Scout can be deployed quickly, easily and frequently – collecting imagery (with an integrated camera payload) that can be imported into GIS databases, stitched together and used to generate 3D maps and reconstruction models. Control of the UAV navigation and camera is through an intuitive, point-and-click touchscreen interface, on a handheld tablet.

2.2 Integrated camera payload

The interface between the Aeryon Scout and payload, used in this project, has a single pitch control servo (azimuth) with yaw (heading) provided by the aerial vehicle orientation. The images that were captured by the Aeryon Scout were stored on the system, as well as streamed real-time over Wi-Fi back to the tablet display and the team on the ground. Basic camera controls like taking snapshots, changing zoom levels, and recording video were supported by the Mission Control System (MCS) software on the tablet computer. User commands, like taking a snapshot, were sent to the Aeryon Scout over Wi-Fi from the tablet interface and then turned into the corresponding IR pulse train sequences to trigger the camera to take the appropriate action.

2.3 Pix4Dmapper Pro software

Images were processed using Pix4D's desktop software Pix4Dmapper Pro. Rapid Check processing mode was used on a Lenovo laptop computer while still on site to verify the completeness and quality of the acquired images. Full processing of the project was done on a desktop computer with 32GB of RAM, i7 CPU and Windows 8. Details of the Full Processing workflow are described in Chapters 5 and 6.

2.4 Ground Truth

The width of the upper base was measured by the project team on site and used to refine the scale of the whole project. Section 6.2 describes how this manual measurement was used to provide a correct scale.

Figure 1: Layout of statue and base with approximate measurements

3. THE SITE

Christ the Redeemer statue stands atop the Corcovado Mountain, which is a 710 meter granite peak shaped like a dome overlooking the city of Rio de Janeiro. The complete monument is comprised of two sections - the 30 meter tall statue of Christ, that is composed of reinforced concrete and covered with thousands of soapstone mosaic tiles that measure 2,7 by 2,7 by 3,7 cm (each), and the marble base that is 8 meters high.

3.1 Weather conditions

The complex hill shape, as well as micro weather and wind patterns encountered at the top of Corcovado Mountain, created unique and challenging flying conditions. In October (time of data acquisition) the prevailing wind direction for the region is from the south west, with an average wind speed of approximately 6 knots (11 km/h)[1]. The Aeryon Scout can operate reliably in sustained wind speeds (up to) 50 km/h, which enabled the team to have confidence that the system would operate safely.

Wind that hits a peak or dome like Corcovado creates higher pressure on the wind facing side and lower pressure on the shadow side.This, in turn, can create a high velocity upward air stream to the front side and chaotic flow patterns that pull down on the shadow side[2]. It can also result in a discontinuous boundary at a certain elevation above the peak where the transition between the low and high wind speed is quite sudden.

Figure 2 below shows the discontinuous effect of elevation on wind velocity on the Corcovado Mountain. To capture this data, the Aeryon Scout was launched from the foot of the statue and ascended to 13 meters. The Aeryon Scout flight control system performs position corrections at a hundred times a second. This allowed the UAV to stay in place while transiting the discontinuity. By comparison, a UAV flown manually with a remote control, would not operate as reliably under these forces and could potentially drift off course 10 meters to the north-east in 1 second, before the operator could respond appropriately.

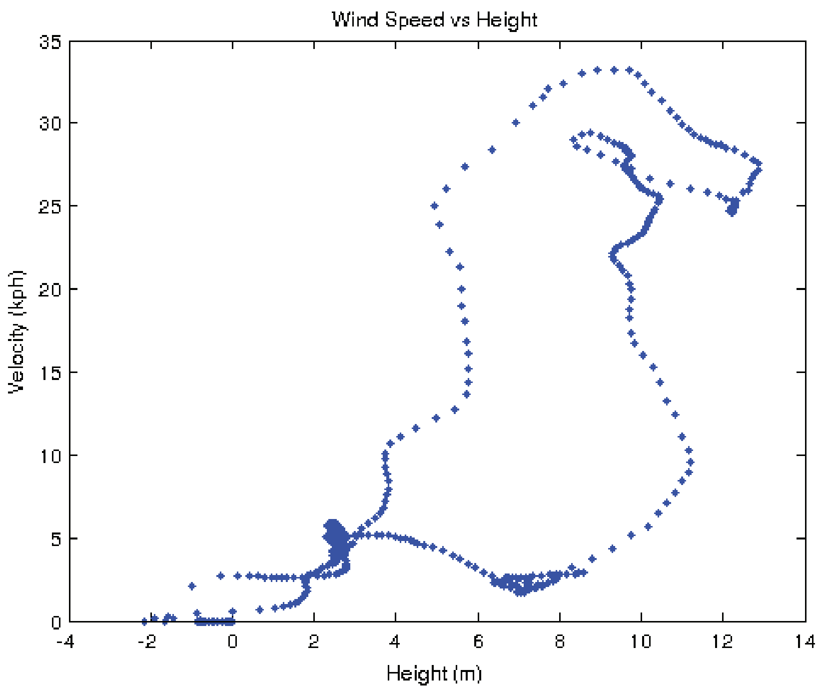

Figure 2: Wind velocity vs height above ground on Corcovado Mountain

3.2 Flight and safety considerations

For commercial applications, most UAV regulations throughout the world require that the operator maintain line of sight with the vehicle at all times. As well, other air traffic must also be monitored to ensure safety. During this project, line of sight was maintained at all times between the operator and the Aeryon Scout. Poor visibility, from fog and cloud cover, was a challenge that the team needed to consider and adapt to each day during flight planning.

Generally, the clouds were low and often eclipsed the statue's view from the city below early in the morning or after the frequent early morning rain showers. As a result, the ability to see the statue was reduced to between 5 and 50 meters for several days during the data acquisition time period.

Christ the Redeemer is the most visited monument in Brazil. To ensure the safety of visitors, the project team was granted special permission from various administrations to enter the site before and after official visiting hours.

3.3 Take-off and landing position

All flights were launched approximately 10 meters from the foot of the statue. However, these take-off locations were limited by the effect of the electromagnetic disturbances around the statue. The disturbances were caused by the large amount of iron in the terrace support structures that are designed to support the weight of the statue, as well as the thousands of visitors. The take-off area was also restricted to a 1m x 1m area due to the many electrical cables that were used to light the monument and surrounding area each night.

The orientation of the Aeryon Scout is controlled by an electronic compass. Knowing which direction the aerial vehicle faces is important when applying the correct direction of thrust when maintaining or moving to a desired position. The Aeryon Scout was capable of taking off under the strong magnetic interference at the site by using a forced heading (informing it which way it is actually facing before it takes off).

Another consideration, when capturing images and data along the back of the statue, was how the vehicle should handle a possible loss of communication between the vehicle and the control tablet. Aeryon UAVs have several built-in fail safes for, such as communications loss, that can be set up before launching the system: "Home and land", "Land in place" and "Home and hover". All options carried risk when the statue was located between the home position and the Aeryon Scout in flight, or when the vehicle was hovering along the side of the mountain. Landing outside the perimeter of the terrace would have resulted in an unrecoverable tumble down the side of the Corcovado. To minimize these risks, and whenever flights were conducted on the far side of the monument, the automated command "Home and land" was set up. Additionally, the return home altitude was set to 50 meters to avoid the tip of the statue and the low battery margin was increased to 300 seconds to double the time needed to ascend and descend the 50 meter height.

Accounting for all of these factors left little time for flights and called for well-defined flight paths planned ahead of time.

Figure 3: Take off and landing position at the base and in front of the statue

4. DATA ACQUISITION

4.1 Preparation and flying time

The operational flight time was shorter than expected because of the weight of the custom mapping payload. The usual 20 minute flight time for Scout was reduced to 15 minutes due to the higher rpm motor demands that were required to maintain lift. Additionally, a 5 minute safety margin reduced flight time further; resulting in an actual flight time of 10 minutes for each battery. Before and after each session, a 5 minute pre- and post-flight preparation time was needed to download data, change the batteries, estimate flying and lighting conditions and launch again.

Limited by the weather conditions and data capture flights restricted to 1.5 hours each morning (7:00am to 8:30am), not many changes could be made to compensate for the light conditions that were very different from one day to another - ranging from extremely low light (in cloudy and foggy weather) to bright sunshine reflecting off the white surface of the soapstone tiles covering the statue. Additionally, flights could not take place at the end of the day (after the site was closed to visitors) due to the very strong winds that occurred each evening.

In total, 3'584 images were captured in a cumulated flight-preparation and flying time of approximately 8 hours, distributed over 7 days. See Section 4.3 for more information about the acquired datasets.

4.2 Image acquisition and overlap

The camera payload was able to capture and process a full resolution image in approximately 2 seconds. The required image overlap for 3D mapping an object, like the Christ Redeemer statue, is 80 percent. These two factors determined the maximum coverage per flight. Image optical zoom was maintained at 1:1 to avoid a reduction in focus depth and poor image quality.

Given focal length (f) and sensor size (w and h), one can calculate the UAV's horizontal and vertical translation rate to generate images with an 80 percent overlap at a given standoff distance (d).

Field of view[3]: $fov = \arctan\left(\frac{w}{2f}\right)$

Given the capture frequency (r) and the effective coverage per image ratio (c) we can solve for:

Translation rate (m/s): $2drc \cdot \tan(fov)$

This resulted in a minimum of 1.1 m/s horizontal and 0.9 m/s vertical velocity settings for the Aeryon Scout with a standoff distance of 10 meters. Since images were captured within 5 - 10 meters, the maximum velocity rate was set to 0.5 m/s for both translations while mapping.

4.3 Acquired datasets

Datasets were collected over 7 days on site:

Date	Weather Condition	Number of Flights	Acquired Data
Wed, October 22	Heavy fog, light wind	0	None due to zero visibility and statue covered in fog
Thu, October 23	Heavy clouds, strong wind	1	Front of statue body
Fri, October 24	Sun, light wind	2	• Front of statue body, including arms and head • Base
Sat, October 25	Light clouds, medium wind	4	• Statue body back and sides • Oblique flights of the terrace • Top Corcovado Mountain
Mon, October 27	Heavy fog, light wind	1	Left hand, side and back (The images were unusable because of heavy fog and visibility of less than 5 meters)
Tue, October 28	Sun, medium wind	6	Chin, head, hands, body sides
Wed, October 29	Sun, medium wind	5	• Back of statue body • Inside of arms, chin and base • NADIR flights of the terrace and top part of Corcovado

5. DATA PROCESSING

Images were captured at various view angles; not all of them contained accurate geotags, as GPS signals close to structures and buildings are attenuated due to signal reflectance. All sub-projects were processed[4] without geotags as GPS information was too inaccurate to bring any benefits to the reconstruction and would have misled the camera calibration. Sufficient manual tie points were added to merge all sub-projects and a linear measurement was used to provide the correct scale to the final result.

Each set of images and data collected during a flight is called a project. For this series of flights, there are three main projects (statue, base and surrounding area) and nine sub-projects. A total of 2'090 images were used for the reconstruction.

5.1 Statue of Christ the Redeemer

The following sections describe the processing of the nine sub-projects that were created to ensure that the reconstruction included the intricate details of the statue with clarity and accuracy.

5.1.1 Head

The head section of the statue provides the most detailed view of the tiny triangular mosaic pieces on the statue. We decided not to process all of the head images as a single project since most of the images of the back of the head would have been un-calibrated and would be filtered out from the whole dataset.

To successfully calibrate all images of the head, three sub-projects were created by separating the images (front, back and top). By adding three manual tie points, the sub-projects were then merged together to create one project of the complete head.

The three datasets consisted of 105 images for the front of the head, 42 images for the back and 142 images for the top.

Figures 4, 5 and 6: Head reconstruction results of the front, back and top projects

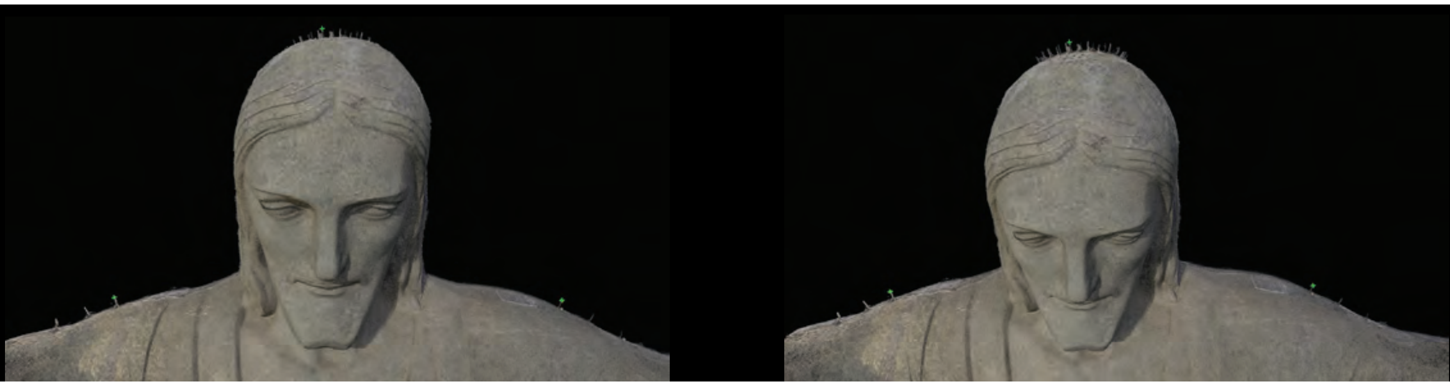

Figure 7: Merged project of the head, showing the three manual tie points (green crosses)

The high resolution image created from the merged sub-projects of the head clearly shows the detail of the mosaic tiles on the face.

Figures 8 and 9: Clear, accurate images of forehead and cheek

5.1.2 Left Hand

The hands of the statue were also processed separately because the images were collected at different locations resulting in different resolutions and view angles. Figures 10 and 11 show the contrasting images for the same object (in this case the left hand). The dataset of the left hand contains 143 images taken from all angles (360 degrees), including the fingers and the sleeve.

Figures 10 and 11: In Figure 10 the left hand is farther away from the camera and thus has much lower resolution than Figure 11

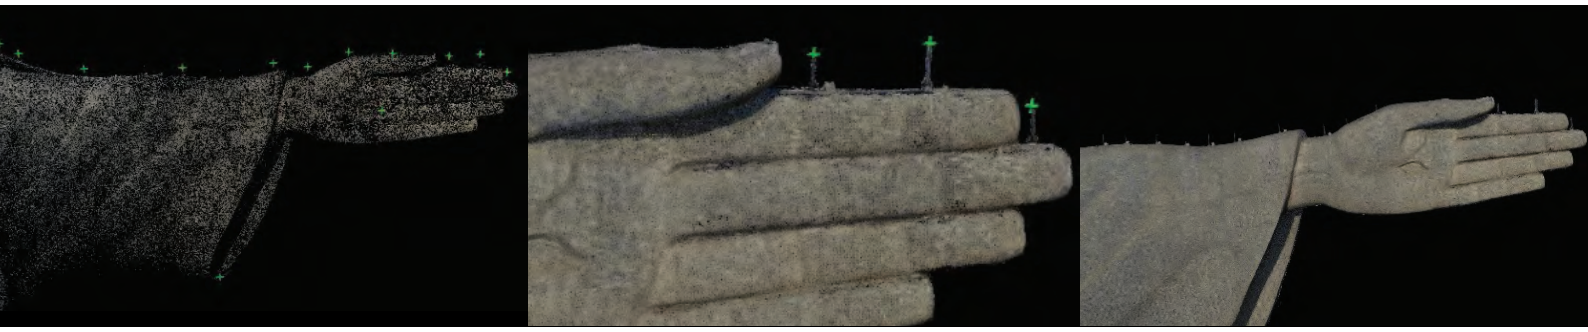

Figure 12: A *B* *C*

Figure 12A: Manual tie points were added to the tips of various spikes on the hand to assist the automatic calibration process on these extremities. Two additional manual tie points were added on the lower sleeves of the arm to evenly distribute all tie points.

Figure 12B, 12C: Generated point cloud of the left hand and the manual tie points on top.

5.1.3 Right Hand

Images of the right hand were collected in very early mornings (with the rising sun) as well as later mornings with full sunshine. As a result, the shading and light saturation of the images prevented the calibration of a large number of the images. Therefore, these datasets were also separated and processed individually.

There were 124 images processed of the back of the hand and 32 images of the front. As with the left hand dataset, manual tie points were added along the tips of the spikes.

Figure 13: A *B* *C*

Figure 13A, B: Example images show the shading and saturation differences for the right hand data set
Figure 13C: Generated point cloud of the right hand dataset

5.1.4 Body

The images for the body part of the statue were taken from a slightly further distance and thus have less detail. They were all processed as one project. The reconstruction of the body section used 626 images. The point cloud of the chest area was processed along with the head section of the statue. Figure 14 below shows the resulting calibration of the body section of the statue.

Figure 14: Calibration of the body

5.2 Base of the statue

It was determined that three different datasets were required to successfully create the base of the statue. The images in each dataset were taken at similar positions but different view angles. To avoid confusion between the datasets they are labelled as "Lower base", "Middle base" and "Upper base".

5.2.1 Lower Base

The Lower base dataset contains 88 images and was taken at a 90 degree (forward/up-looking) angle from the terrace. It has the highest and the most consistent resolution; however, it was highly influenced by the reflective surface of the base.

Figure 15: Sample image of dataset

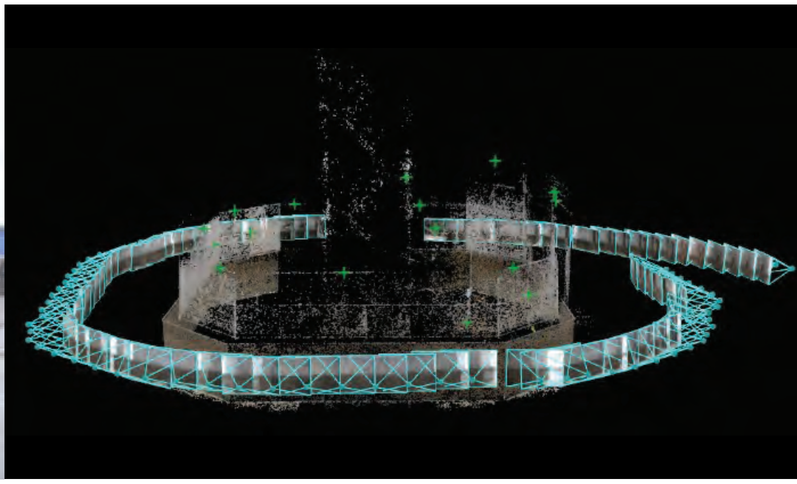

Figure 16: Reconstructed 3D point cloud of lower base

The sample image shows the reflected balustrade which causes ambiguity for Pix4Dmapper as the software will recognize the reflection as real features and use it for reconstruction.

5.2.2 Middle Base

The dataset of the middle base was taken at approximately a 110 degree (forward/up-looking) angle from the terrace and includes 81 images. It serves as the overlap region between the upper and the lower base datasets. This overlap enabled the manual tie points to accurately join the three datasets (lower, middle and upper base) to each other.

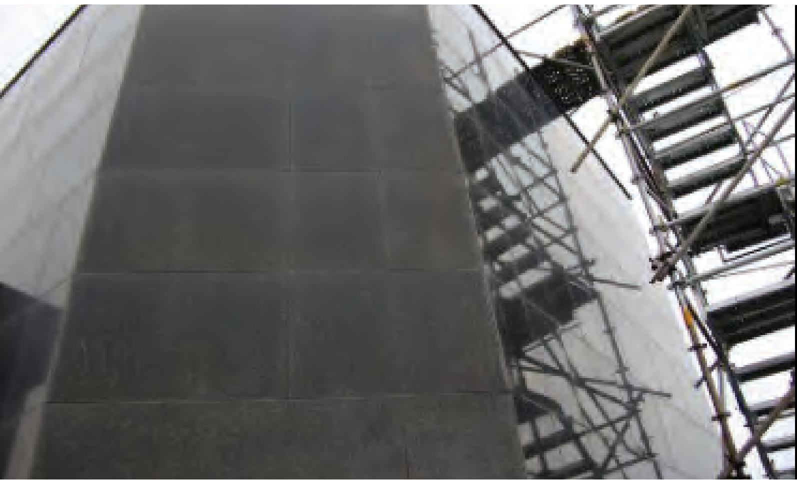

Figure 17: Sample image of dataset

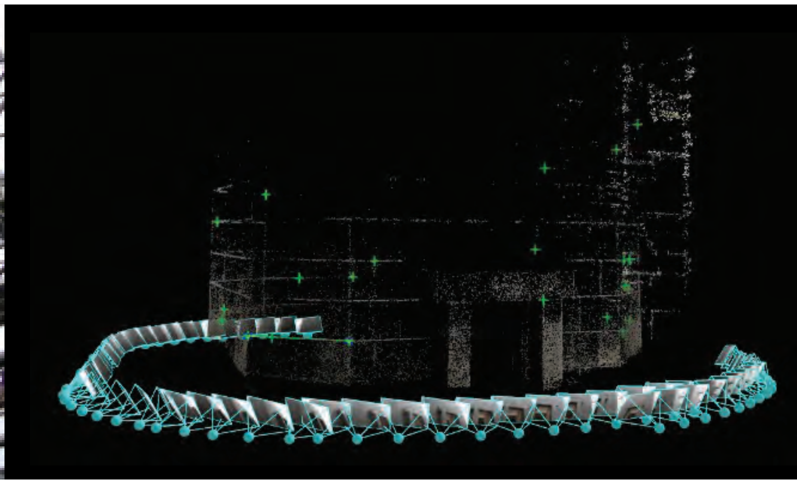

Figure 18: Reconstructed 3D point cloud of middle base

5.2.3 Upper Base

The upper base dataset was captured at approximately a 135 degree (forward/up-looking) angle and includes 84 images. This portion of the reconstruction had the fewest points (of the three datasets) because of the angle (looking towards the sky) that was required, as well as the fact the lower section of the statue was included in the images. This additional information made it more difficult to process this dataset.

Figure 19: Sample image of dataset

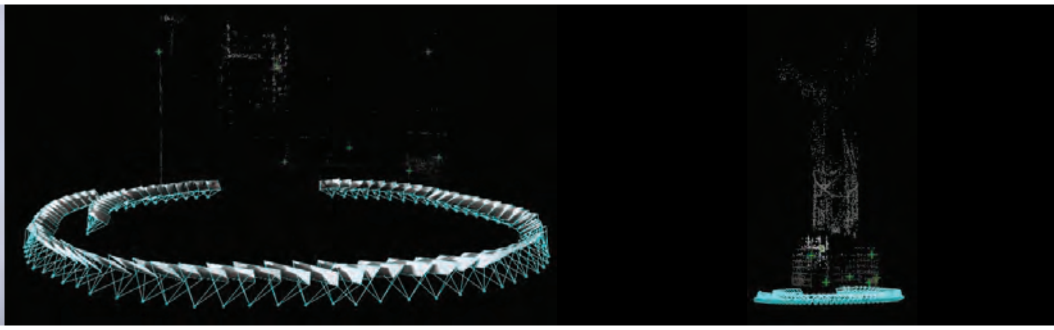

Figure 20: Reconstructed 3D point cloud of upper base

5.2.4 Merging of the base

Selecting the correct manual tie points was the most time consuming part for merging the three datasets (of the base) because of the variation in the view angles. The reflection of the surface and images which did not contain the whole base or contained only one façade also contributed to the difficulty in determining where the tie points should be made. At least three manual tie points were selected for every façade of the base.

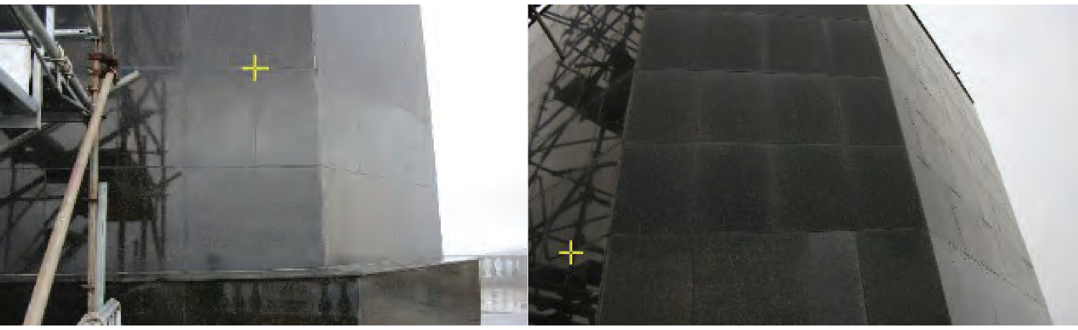

Figure 21: Importance of picking appropriate manual tie points

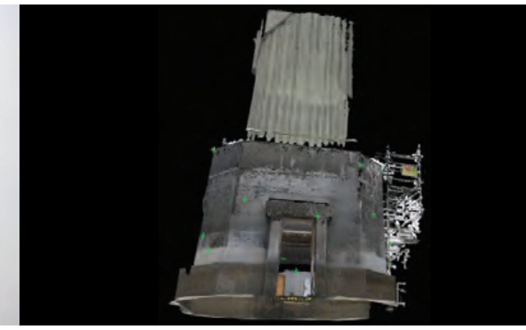

Figure 22: Selected manual tie points used and 3D point cloud of the merged base

5.3 Surrounding area

5.3.1 Cliff

The project team flew in a helicopter to collect the first data set of the cliff. This data set includes 165 images which were captured using a hand-held Canon 6D. The processing of this data set was relatively straight forward, but resulted in a comparatively low Ground Sample Distance (GSD). This dataset provides an overview of the statue and surrounding area. In Figure 23 below, only a rough outline of the statue can be seen because of the low resolution.

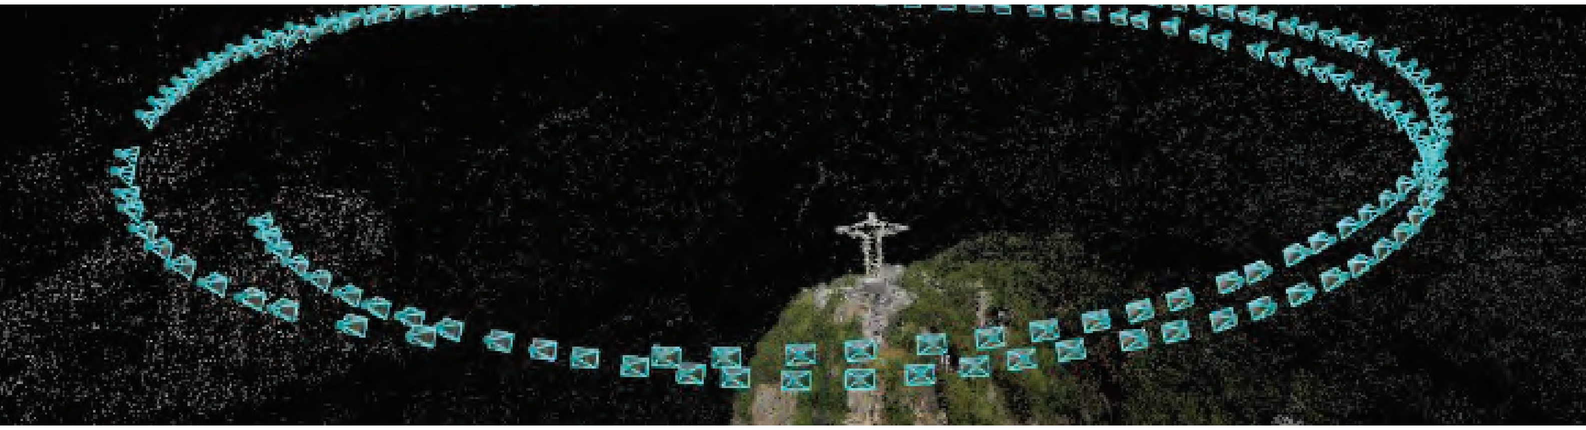

Figure 23 : 3D point cloud generated from the cliff data set

5.3.2 Terrace

The project for the terrace includes two different data sets: one data set was captured at an oblique angle to provide the side view and facades. The other dataset was captured at a Nadir angle of view to provide information on the top surface.

5.3.2.1 Terrace Oblique

During the two UAV flights, 271 oblique images were collected and processed automatically and without any manual tie points.

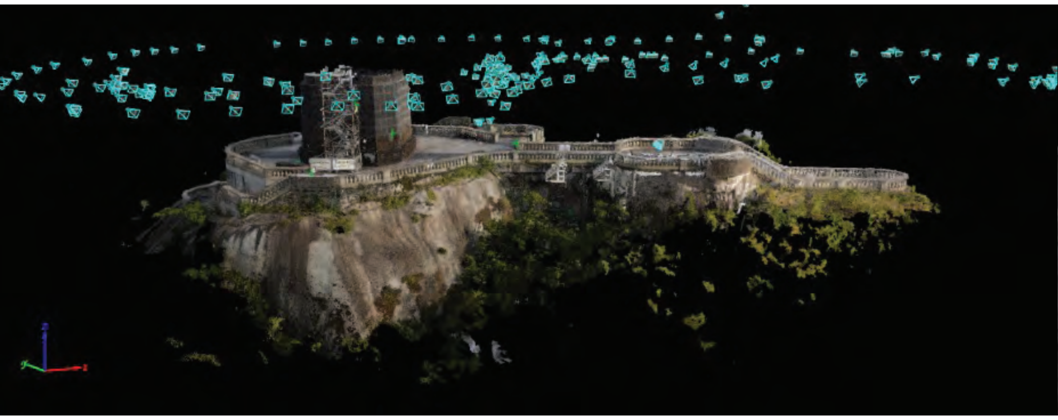

Figure 24: Flight plan of the terrace oblique data set

Figure 25: Point cloud of oblique dataset, including low-resolution of base

5.3.2.2 Terrace Nadir

The Nadir terrace data set includes 187 images, also collected using the Aeryon Scout. This data set provides a complete texture of the top surface which cannot be achieved by oblique flights. It is crucial to have information from this view point for the final output.

Figure 26: Flight plan terrace Nadir data set

Figure 27: Point cloud from terrace oblique data set with the manual removal of the statue

6. OVERALL RECONSTRUCTION

6.1 Project merging

All of the projects were merged simultaneously with a total of 82 common manual clicked tie points. After the densification, Pix4Dmapper's point editing tools were used to crop out the unwanted parts and the annotation feature was used to filter out the points from the background (such as the sky).

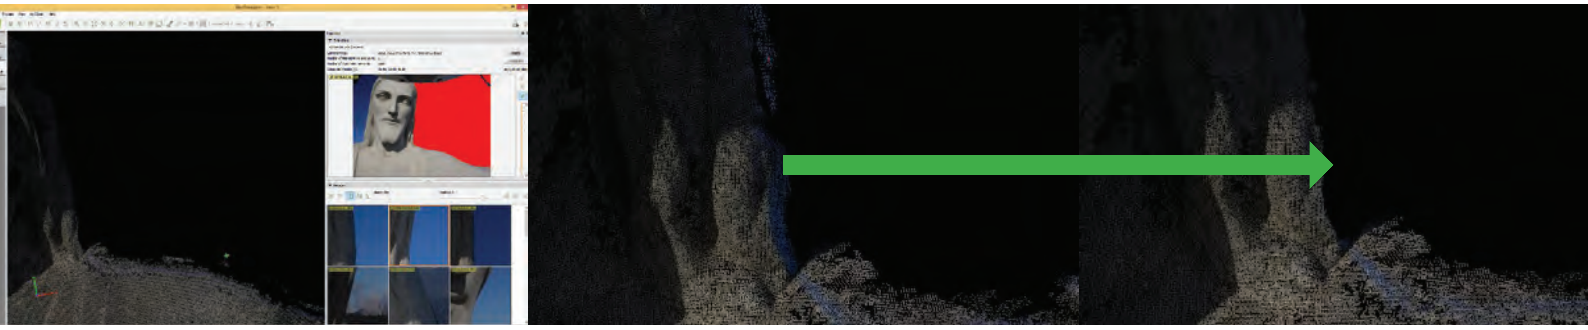

Figure 28: Annotation example and the difference when applying it

The terrace Nadir data set created a point cloud of the entire top surface. To ensure the unity of the surface and remove any influence from the different orientation of shadows, the top surface generated from the terrace oblique data set was cropped out of the merged dataset. Therefore, the facades were collected from the oblique images and the top surface from the nadir images.

Figure 29: Result of cropped surface, filled with the nadir-view point clouds, and a zoomed-in view of the combined outcome

6.2 Providing the correct scale

All images used in this project are without accurate geotags, and thus the point cloud generated is relatively correct in position. This final step demonstrates how one linear measurement was applied to fit the entire project to the correct scale. The base serves as a comparative object for the measurement since it has clear boundaries with sharp corners.

While on site, the project team measured the width of the upper base with a measuring tape (at 5.3 meters). This measurement was used for the scale refinement in Pix4Dmapper.

Figure 30: Manual measurement of the base

Ideally, at least three control points with accurate three dimensional coordinates are needed to place the full project in real position. In this case, only scale has been corrected; absolute position has not been considered.

Two control points are sufficient to provide the scale. In Pix4Dmapper, two pseudo ground control points were added with coordinates: (0,0,z) and (5.3,0,z). They were set as 3D GCPs in the local coordinate system. In this case, z can be an arbitrary value and the value was set at 20.

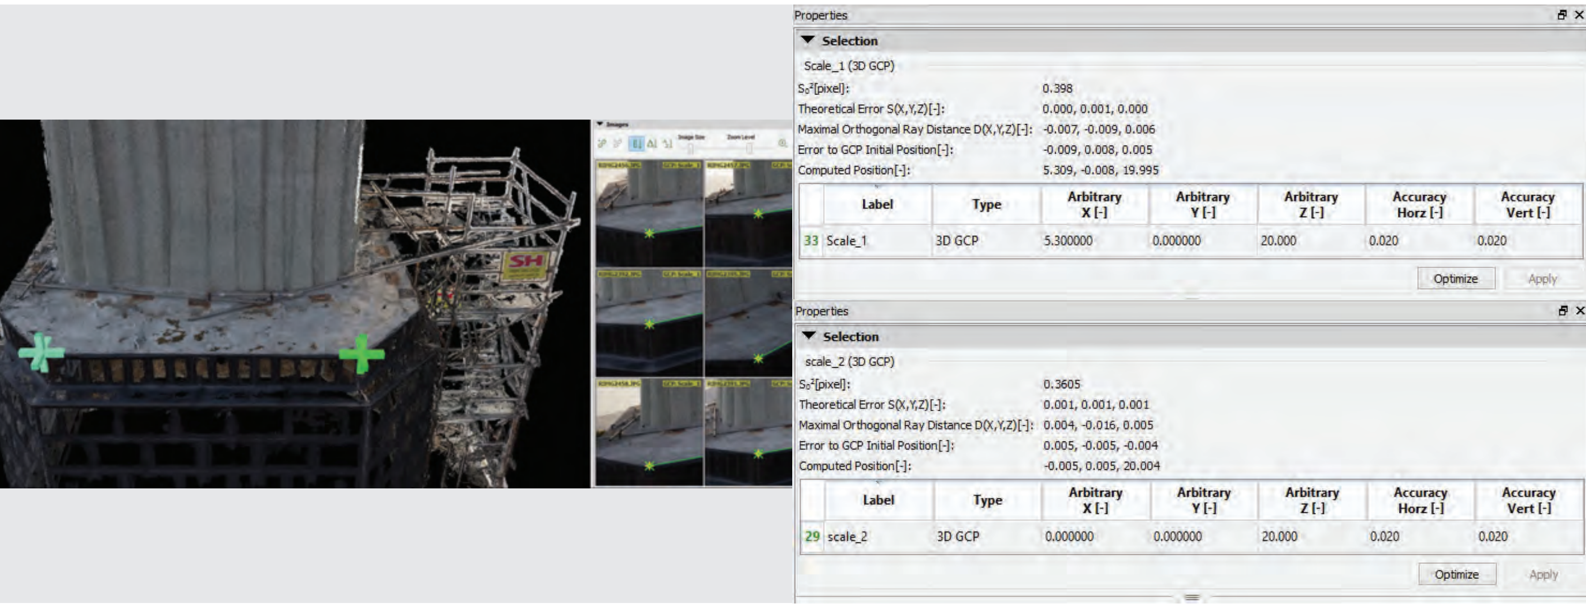

Figure 31: Pseudo 3D GCP marked in Pix4Dmapper's rayCloud editor

After re-optimizing the project using the two pseudo GCPs, the distances were measured directly in the point cloud for verification purposes. The accuracy measured was within centimeters. A certain difference in result was to be expected since only the point to point distance could be measured rather than a true vertical distance of the base in the software.

The vertical distances measured in Pix4Dmapper were as follows:

- Height of statue: 30.00 meters
- Height of base: 8.09 meters
- Horizontal distance from fingertip to fingertip: 29.42 meters

The claimed value (the only measurements available are in the original drawings of the statue, dated 1924) for these three measurements is 30.03 meters, 8.00 meters and 29.40 meters respectively.

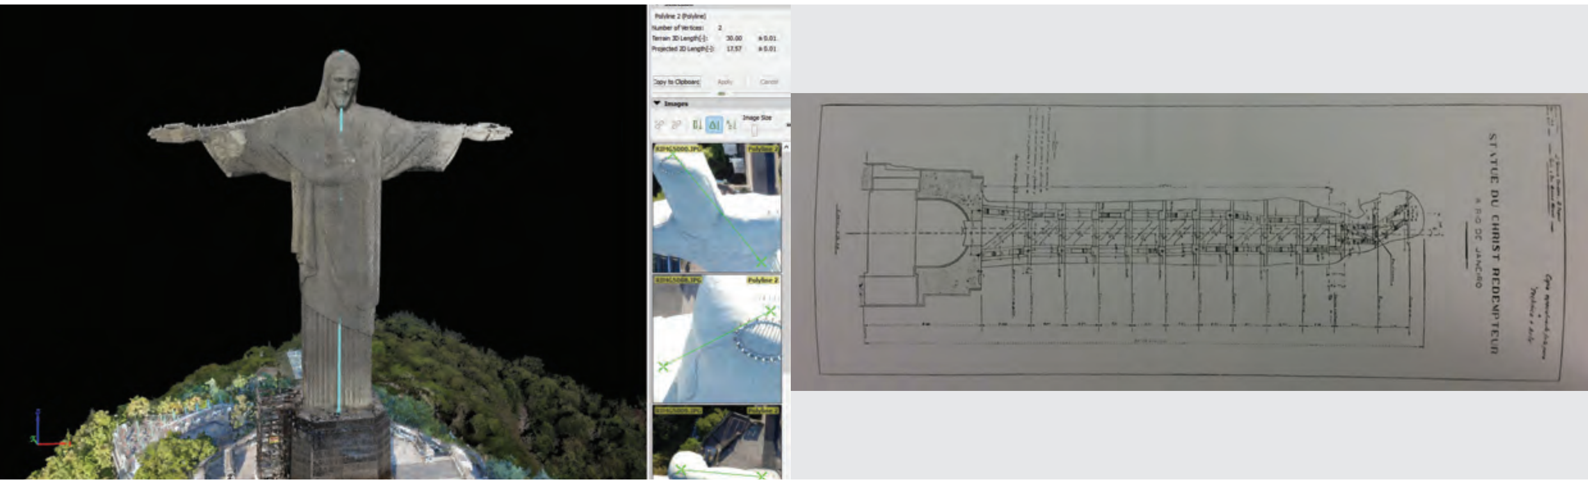

Figure 32: The height of the statue documented in the official architectural design

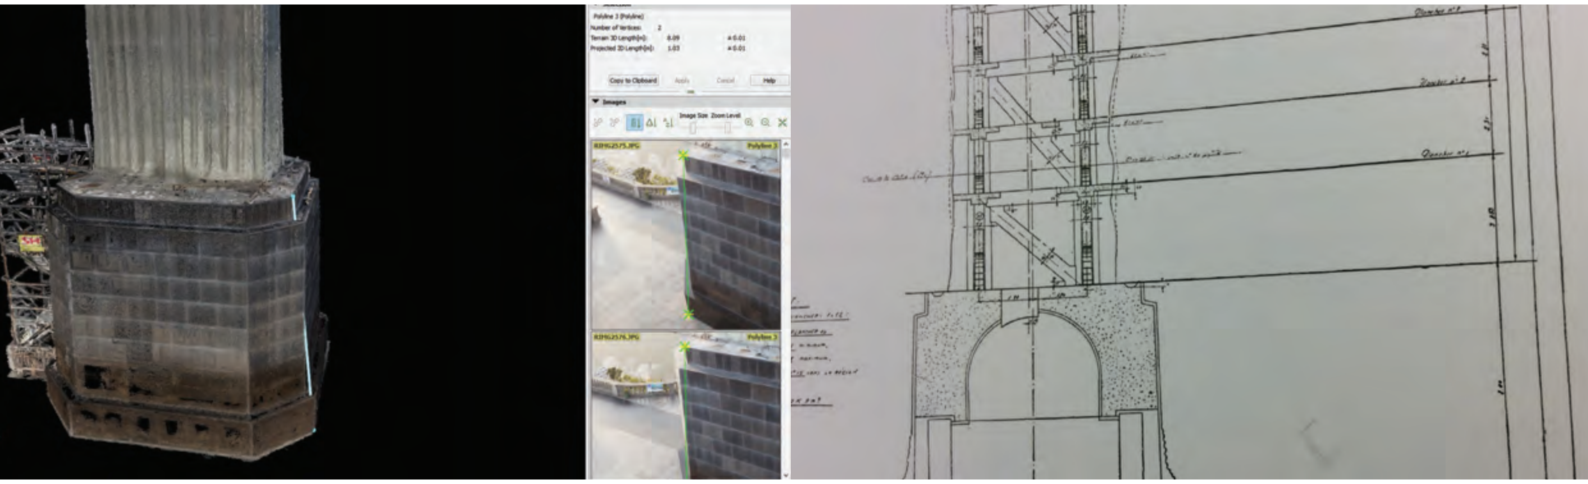

Figure 33: The base height as documented in the official architectural design

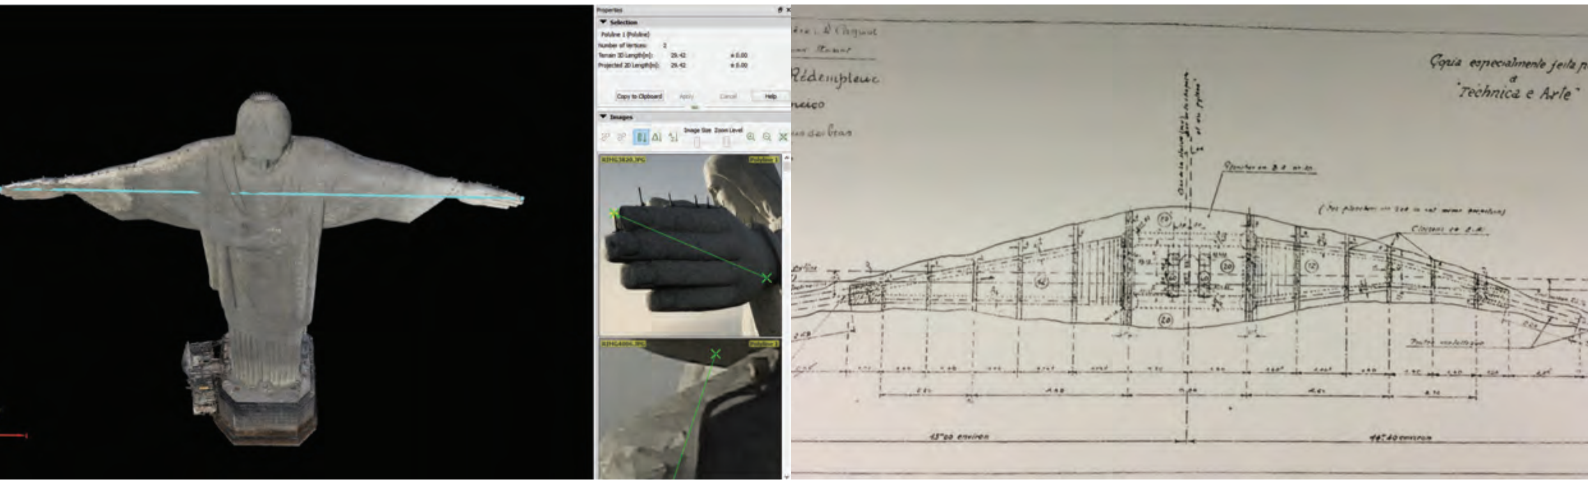

Figure 34: The arm span of the statue documented in the official architectural design

7. FINAL RESULTS

Point Cloud of statue and surrounding area

Merged point cloud of all the subprojects, including the statue, the terrace (both nadir and oblique), and the surroundings taken by Canon 6D on a helicopter. Point cloud with a total of 134.4 million points (97.6 million for the statue, 30.7 million for the terrace and 6.1 million for the surroundings). The 3D textured mesh is the new feature of Pix4Dmapper software version 1.3, the final 3D textured mesh consists of 2.5 million triangles with a texture of 16384px x 16384px.

Orthographic view of 3D textured mesh of Christ the Redeemer

3D textured mesh, upper body of statue

3D textured mesh, zoom of head

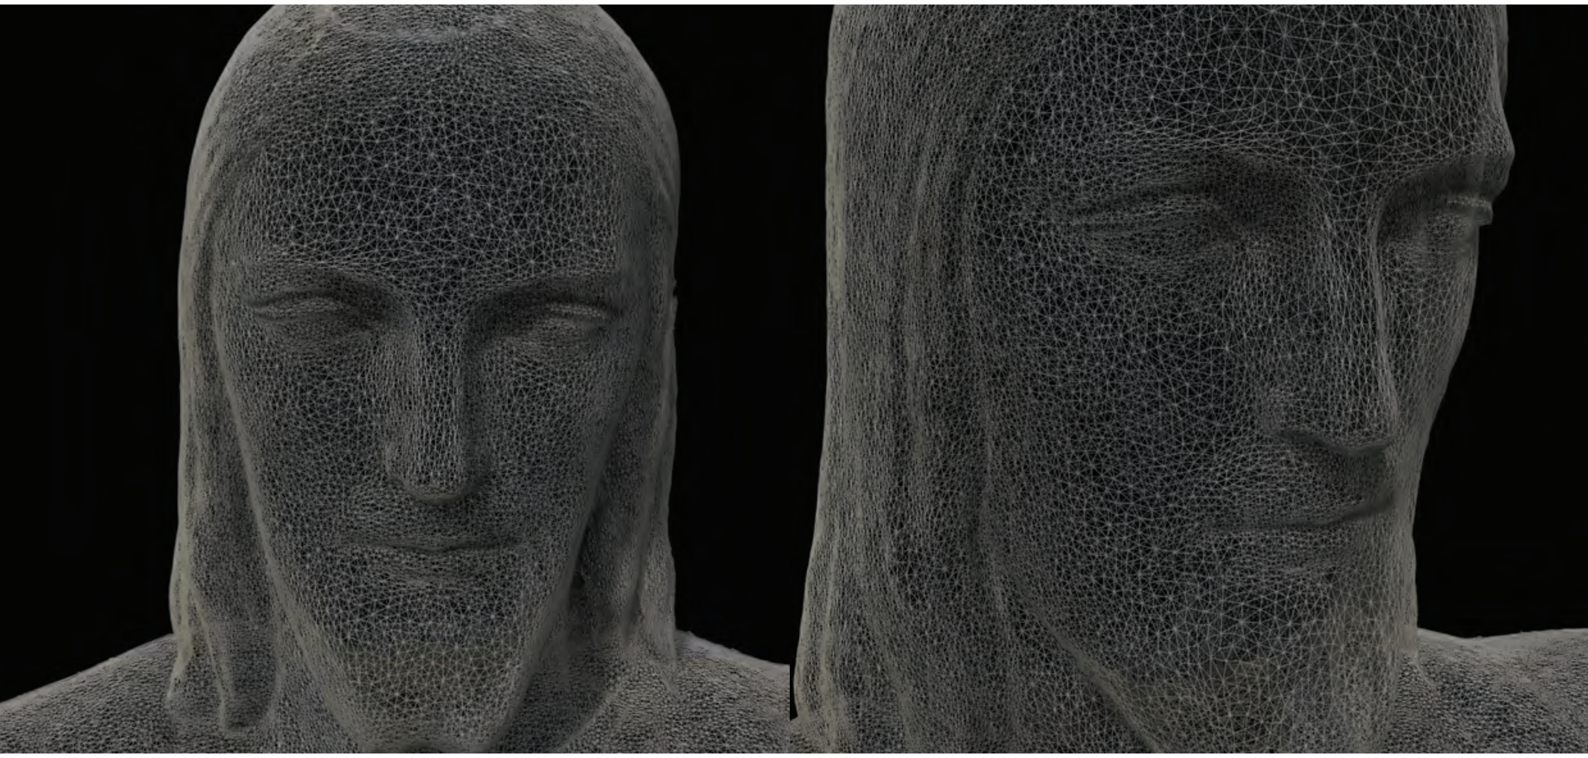

Triangle mesh of head

ACKNOWLEDGEMENT

Special thanks to Celso Santos, Jorge Lopes and Cláudio Magalhães of PUC-Rio / NEXT Lab, and Priest Omar of The Archdiocese of Rio de Janeiro for making this project possible.

REFERENCES

1. WindFinder.com GmbH & Co. KG . 2014. Weather and Wind Statistics. [ONLINE] Available at: http://www.windfinder.com/windstatistics/barra_da_tijuca_rio_de_janeiro. [Accessed 06 January 15].
2. Antenna Systems & Solutions, Inc.. 2014. Antenna Systems: Wind Turbine Towers. [ONLINE] Available at: https://antennasystems.com/Merchant2/merchant.mvc?Screen=wind-turbine-towers. [Accessed 06 January 15].
3. Wikimedia Foundation, Inc.. 2014. Angle of View. [ONLINE] Available at: http://en.wikipedia.org/wiki/Angle_of_view. [Accessed 06 January 15].
4. Pix4Dmapper processing steps "Initial processing" and "Point Cloud Densification".