

Photogrammetric Procedure for Modeling Castles and Ceramics

Aaron Pattee, M.A.

Institute for European Art History, Heidelberg University

25 October 2016

Institut für europäische Kunstgeschichte
Seminarstraße 4
69117 Heidelberg

Table of Contents

Introduction	2
Step 1: Location and Environment	2
Step 2: Equipment	2
Software	2
System Requirements for <i>Agisoft Photoscan Pro</i>	3
Cameras.....	3
Drones.....	3
Equipment Checklist—in addition to a camera, <i>Agisoft</i> , and potentially an UAV.....	4
Step 3: Taking the Photos	4
Light Settings	6
Step 4: Entering the Photos into <i>Agisoft</i>	6
How to Save the Photos	6
Creating Chunks in <i>Agisoft</i>	7
Step 5: Generating the Alignment.....	7
Preparation for Alignment.....	7
Creating a Mask	8
Aligning the Photos.....	8
Alignment for Large Datasets	9
Digital Markers	9
Optimizing the Point Cloud.....	11
Combining Large Datasets	11
Aligning Chunks	12
Merging Chunks.....	12
Step 6: Generating the Dense Cloud	12
Dense Cloud Generation for Large Datasets	13
Step 7: Generating the Mesh.....	14
Referencing Duplicate Chunks with a Python Script.....	14
Step 8: Generating the Texture	15
Step 9: Scaling the Model.....	16
Step 10: Exporting the Model.....	17
Step 11: Decimating the Model	18
Decimating the Mesh in <i>Agisoft</i>	18
Decimating the Mesh in <i>Meshlab</i>	18
Uploading to <i>Sketchfab</i>	19

Introduction

This is a complete modeling procedure from beginning to end for generating models of both large sites and small objects. The primary example for large sites is the site of *Burg Hohenecken* (a castle in southwest Germany). Examples regarding smaller objects are provided by a project concerning a ceramics collection excavated from a former cistern on the campus of the University of Nebraska Lincoln. Both examples follow essentially the same procedure, though they differ at various stages, and will be described parallel to one another throughout this procedure. The full project concerning *Burg Hohenecken* can be found in the University of Nebraska Digital Commons under the title: *Integrative 3D Recording Methods of Historic Architecture—Burg Hohenecken from Southwest Germany*. The cistern project is currently under development, though a publication soon to come.

Step 1: Location and Environment

The most important thing to determine before taking any photos, is to analyze the area in which the photos are to be taken. Photography is a light based technology and is therefore sensitive to light. Photogrammetry will not work well at night for a large object (e.g. a castle) or when photos are taken over the course of many hours due to the changing of the position of the sun. When the shadows move, the calculations for aligning the photos are affected. Taking all of the photos at an outdoor site requires haste and efficiency! Indoor settings are a different situation as long as the light can be properly controlled. Due to the somewhat hurried nature of outdoor photography for photogrammetry, it is best to determine a path along which the photos can be taken in the most efficient manner in order to prevent unnecessary backtracking. Once the area has been analyzed and a path has been determined, continue to step 2. For indoor settings, identify whether the object is glossy or matte. In the case of a glossy object, it is best to limit direct light as much as possible, allowing only ambient light throughout a darkened room. For matte objects, direct light can be applied, though a light box to fade the light is preferential.

Step 2: Equipment

Software

Before beginning the process, be sure to have a version of the updated *Agisoft Photoscan Pro*, not the standard version. *Agisoft Photoscan Pro* uses Scale Invariant Feature Transform (SIFT), Structure from Motion (SfM), and Multi-View Stereo (MVS) in order to create 3D models.ⁱ Other software programs include *PhotoModeler*, which was used in a 2008 study of *Burg Landenberg* in Switzerland.ⁱⁱ *PhotoModeler* is an older platform that uses close-range photography for the photo-sets. Another photogrammetric program being developed at the Institute for Photogrammetry at the University of Stuttgart, called *SURE* (**S**urface **R**econstruction from **I**magery), is a specialized software for dense image matching. It uses the MVS technique based upon a modified version of the *Semi Global Matching* algorithm (called the *tSGM*). The modification leads to a more efficient dense cloud generation, reducing RAM requirements up to 70% while accelerating the computation by 30%—crucial savings for large datasets. In contrast to *Agisoft Photoscan Pro*, it is not designed to conduct SfM. Instead, orientation data are imported from a SfM software (like *Agisoft*) into *SURE*, which then generates the dense clouds (up to one 3D point per pixel) as well as the mesh and texture. *SURE* presents an excellent method to increase efficiency and decrease computation time and memory.ⁱⁱⁱ *Agisoft* was used for this procedure because it can complete the entire workflow from alignment to texture. In contrast to other modeling systems, a high resolution camera and photogrammetric software is far cheaper than a laser-scanner with its associated software. Photogrammetry is more cost effective and is not nearly as logistically limited as laser-scanning because of the relatively small size of a camera compared to a laser-scanner.

System Requirements for Agisoft Photoscan Pro

Figure 1 is taken directly from Agisoft Photoscan Pro's website as of the 24th of August 2016: <http://www.agisoft.com/downloads/system-requirements/>

Basic configuration up to 32 GB RAM	Advanced configuration up to 64 GB RAM	Extreme configuration more than 64 GB RAM
CPU: Quad-core Intel Core i7 CPU, Socket LGA 1150 or 1155 (Haswell, Ivy Bridge or Sandy Bridge)	CPU: Six-core Intel Core i7 CPU, Socket LGA 2011-v3 or 2011 (Haswell-E, Ivy Bridge-E or Sandy Bridge-E)	For processing of extremely large data sets a dual socket Intel Xeon Workstation can be used.
Motherboard: Any LGA 1150 or 1155 model with 4 DDR3 slots and at least 1 PCI Express x16 slot	Motherboard: Any LGA 2011-v3 or 2011 model with 8 DDR4 or DDR3 slots and at least 1 PCI Express x16 slot	
RAM: DDR3-1600, 4 x 4 GB (16 GB total) or 4 x 8 GB (32 GB total)	RAM: DDR4-2133 or DDR3-1600, 8 x 4 GB (32 GB total) or 8 x 8 GB (64 GB total)	
GPU: Nvidia GeForce GTX 780 or GeForce GTX 980 (optional)	GPU: Nvidia GeForce GTX 780 Ti, GeForce GTX 980 or GeForce GTX TITAN X	

Figure 1: Agisoft Computer Configurations.

Cameras

A DSLR is the most preferred camera for photogrammetric purposes. The Nikon D3300 has worked incredibly well and was the camera used for all examples in this procedure. For large objects (e.g. castles), a lens with a focal length of 55mm is sufficient, though for smaller objects, a longer lens (55-200mm) may be necessary depending upon how small the object is. Maintaining the same focal length is a necessity in order to insure proper alignment. A substantial memory card for the camera is also a necessity—64 GB cards seem to be more than enough even for projects involving thousands of photos. Additional accessories are described in the following paragraphs.

Drones

The size of the site being recorded plays an essential role in how the photos are taken. Larger sites that have standing walls cannot be completely recorded at ground level. In some cases, full coverage can only be accomplished by using aerial photography. UAVs (Unmanned Aerial Vehicles) for recording tall structures are incredibly useful, though are also very costly. UAVs present key benefits including fast realization of images in-flight, instant quality control on the field, more stability and flexibility than a helicopter, and are low cost compared to manned aerial vehicles (planes, helicopters, etc.).^{iv} They are described by a variety of names including Remotely Piloted Vehicle (RPV), Remotely Operated Aircraft (ROA), Remote Controlled (RC), Unmanned Vehicle System (UVS), and more colloquially, drones. The device used at *Burg Hohenecken* was a UAS (Unmanned Aerial System). Additionally, UAVs are separated into three general categories based upon size, weight, endurance, range, and flying altitude. The smallest are Tactical UAVs which are short-range, low altitude (a few hundred meters to five kilometers), generally low endurance, and low carrying capacity (fifty-one thousand kilometers). Strategic UAVs, are long-range, high altitude (higher than 20,000 meters), and have a much longer endurance. Special Tasks UAVs are the combat vehicles used by militaries.^v


Tactical UAVs are the most common forms for archaeological purposes, featuring rotary wings and do not require a run-way. The capabilities of these small vehicles ranges from more common kits, carrying relatively small cameras (e.g. *GoPros*) to more complex kits which can be equipped with DSLR cameras or laser-scanners. The UAS used at *Burg Hohenecken* had eight wooden rotors, an aluminum frame with carbon fiber joints to reduce the overall weight, and was equipped with a Sony NEX full-frame camera. UAVs with laser-scanners are much heavier and generally more expensive, requiring more powerful engines than those carrying only cameras. These systems allow for close-range and long-range photos to be taken, encompassing the entire site in one shot—which typically cannot be done in the absence of an aerial device.^{vi} However, UAV cameras are often auto-focusing, requiring all photos to be individually calibrated to targets or other reference points before performing any sort of metric reconstruction from the images, though this dramatically lowers the accuracy.^{vii}

Equipment Checklist—in addition to a camera, *Agisoft*, and potentially an UAV

- Camera accessories:
 - A protective, waterproof, case for carrying the camera and all accessories
 - Lens protector
 - Camera strap
 - Extra batteries (preferably 2-3 extra battery packs)
 - Centimeter scale to be placed near the object
- For indoor:
 - Tripod
 - Contraption for revolving camera around the object
 - Light boxes and LED lights for creating ambient light
 - A scale bar to be placed near the object

Step 3: Taking the Photos

Taking the photos generally follows two requirements: a fixed focal length and a substantial amount of overlap between photos. All photos should be taken perpendicular to the object (shown in Figure 2). The more photos the better, though too many photos will only serve to extend the computation time. Each photo should have at least 50% (upwards of 70%) overlap relative to neighboring photos. Given the focal length (which is entirely dependent upon the site), an average distance between camera positions can be approximated. Photos should also not reach above 20° orthogonal to the axis of the camera. In the case of corners, edges, columns, and other circular objects, photos should be taken hemispherically to the object. Every aspect of the site or object should be recorded, even the bottom of small objects (which can be merged with the top later on).



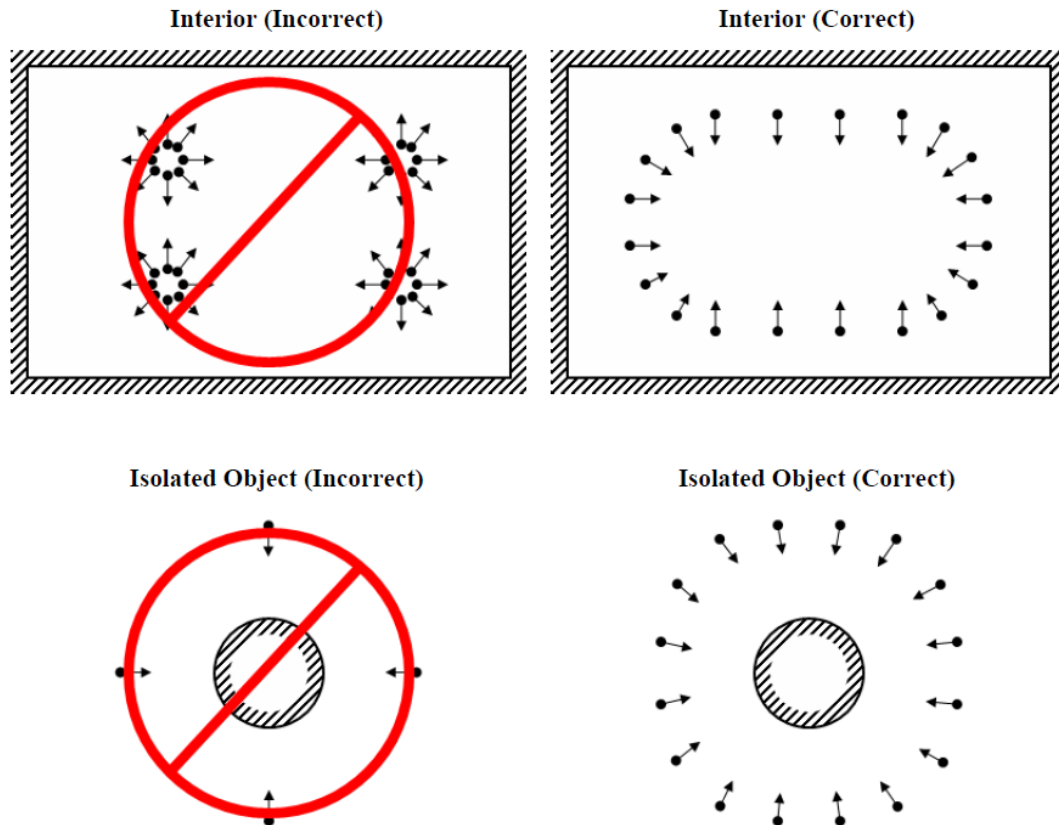


Figure 2: Correct method of photography for photogrammetry.^{viii}

For small objects, higher resolution models for glossy or glass objects are achieved by rotating the camera around the object while the object remains still. This serves to prevent light from shifting, causing a subsequent change in shadows. Having the camera attached to a tripod atop a platform also ensures that the focal length will remain unchanged. A simple set up seen in Figure 3 can be constructed from wood and PVC at a cost of less than \$100. For non-glossy/glassy objects, a simple turn-table placed within a light box with the camera in a fixed position facing the object will suffice. In fact, this method is the most used and works very well for matte objects.



Figure 3: Photogrammetric Contraption for glossy and glass objects.

Light Settings

In outdoor settings, the photos should be taken as quickly as possible, though not at the expense of quality. Glossy small objects (like ceramics) should be placed in a darkened room, with only two light sources. The light sources should be placed within light boxes, whose opaque bottoms should be placed facing the object. The other sides of the boxes will provide indirect ambient light throughout the room, eliminating glares and providing just enough light for the photographs to be taken. Given the settings of the room, the ISO of the camera must be adjusted accordingly. To make this procedure more exact, it would be worthwhile to calculate the exact light exposure of the room used for the small ceramics experiment to ensure a proper replication.

Step 4: Entering the Photos into Agisoft

How to Save the Photos

Once all of the photos have been taken, whether aerially by an UAV or terrestrially (both at large sites and for small objects), the photos should be immediately separated into two general groups. For large sites, separate the photos into aerial and terrestrial. For small objects, separate the photos into top and bottom. Organizing the folders efficiently on the computer will make the whole process much easier. Be sure to save all of the files and photos directly to the internal storage of the computer. It may be easier to allocate the storage among multiple internal hard drives. In the case of the models for the castle and the ceramics, *Agisoft* ran from the internal OS hard drive (a 250 GB SSD) whereas all of the folders and photos were saved to an internal 4 TB HDD. Saving outside of the internal storage can significantly increase the risk of file corruption. Within the project folder, save the *Agisoft* model as a .psx file rather than a .psd file and remember to save often! Each time a process has been run, save a copy of the project prior to running a new process in case *Agisoft* crashes (which may happen if the computer does not have the proper requirements).

Creating Chunks in *Agisoft*

Within *Agisoft*, locate the left side panel labeled “Workspace” shown in Figure 4. Notice that at the bottom of the panel there is another tab labeled “Reference.” Each of the tabs will be used throughout the project, so be sure to become familiar with how they appear. To make a chunk, right click in the white space of the workspace panel and select “Add Chunk.” For large sites, in which aerial photos were made, add only the aerial set for now. For small objects, create the chunks for both the top and bottom photos.

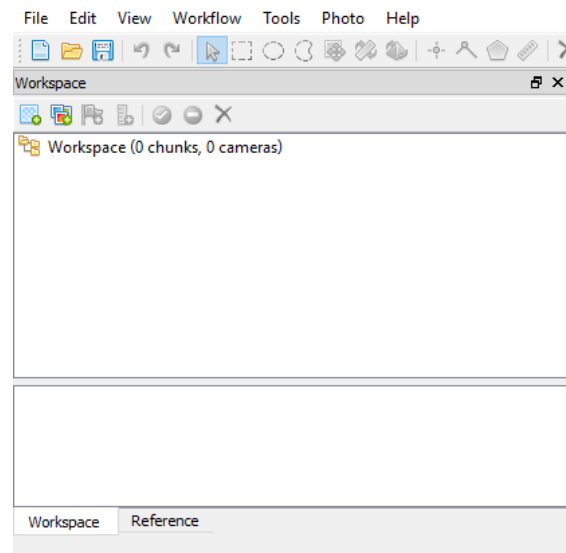


Figure 4: Workspace panel and tabs. Note: the reference tab is at the bottom left.

Step 5: Generating the Alignment

Preparation for Alignment

This is the orientation phase of the project in which the SfM (Surface from Motion) is conducted, though in *Agisoft* it is called “Alignment.” If pre-calibrated markers (shown in Figure 5) were set at the site, then select the chunk of interest and select Tools→Markers→Detect Markers. The accuracy of a photogrammetric model can be significantly increased if coded targets are used on-site. These targets (shown in Figure 5) can be printed via the software and are already calibrated.^{ix} They increase the measurement accuracy compared to manually placed digital markers or automatically generated tie-points. However, the lack of physical markers increases the image calibration error by 45%. This is typically the case in models in which a variable lens with autofocus is used. Despite this short-coming, variable lenses with autofocus can be adjusted in *Agisoft Photoscan Pro* by individually calibrating each photo with different foci, rather than having them bundled together.^x




Figure 5: Pre-calibrated marker from *Agisoft*.

Markers on-site should be used whenever possible, because they significantly reduce the time necessary for the generation of the 3D models. Additionally, the combination of a fixed lens, a full-frame sensor, and a significant number of targets that can be seen in every image, the precision of a photogrammetric model can reach a ratio of at least 1:50,000.^{xi} Regardless if markers were set or not, proceed to Tools→Camera Calibration (Figure 6). This is to check the focal lengths of the individual photos—be sure that the camera positions are all within either one or two groups of focal lengths. An example of the camera groups from one of the ceramics is shown below:

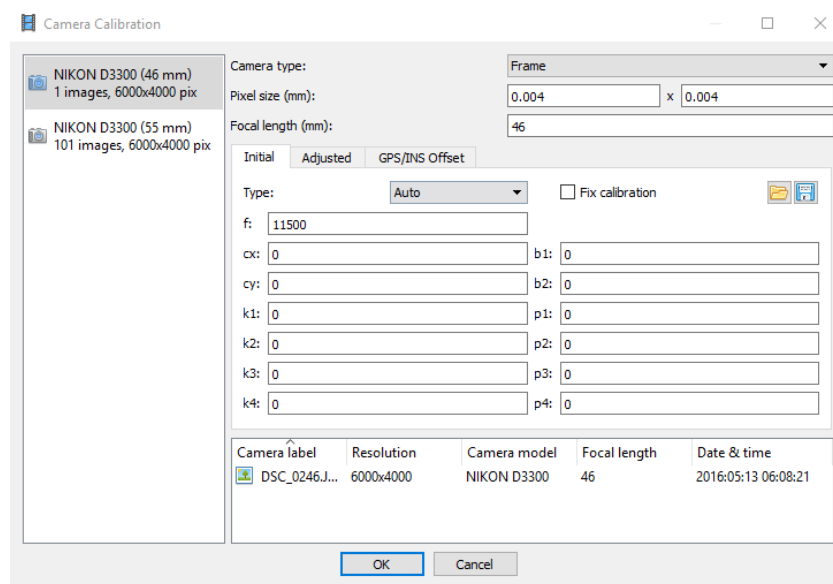


Figure 6: Camera calibration box.

Creating a Mask

Masks essentially tell the software where to not pay attention whilst aligning the photos and can be time intensive if each photo is individually masked, however this is not advisable. It is much more efficient to create a mask by first creating a low resolution model from the alignment stage to mesh stage and importing the model as a mask. In order to do this, run the model on low settings through all three steps. Look ahead in this procedure to see how to do this, though all of the steps are seen in the “Workflow” drop-down at the top of the program. Be sure that the model actually looks somewhat like it is supposed to after the alignment stage. It is advisable to either save the project as a new project or to duplicate the chunk of photos in the workspace. After running everything on low and receiving a fairly poor model, use the selection to cut out all of the background points that may yet be in the model. Then go to Tools→Import→Import Mask. This will open a box for creating the mask. For “Method” select “From Model,” for “Operation” choose “Replacement,” and for “Apply to” select all cameras. Selecting all cameras will only affect the cameras in that chunk. To apply the mask to all cameras in each chunk, select “Entire Workspace.” The software will then have automatically masked all of the photos. Now continue with the regular procedure of the high quality alignment.

Aligning the Photos

To align a chunk, select it and then select Workflow→Align Photos. Once this has been selected, a box will pop up shown in Figure 7.

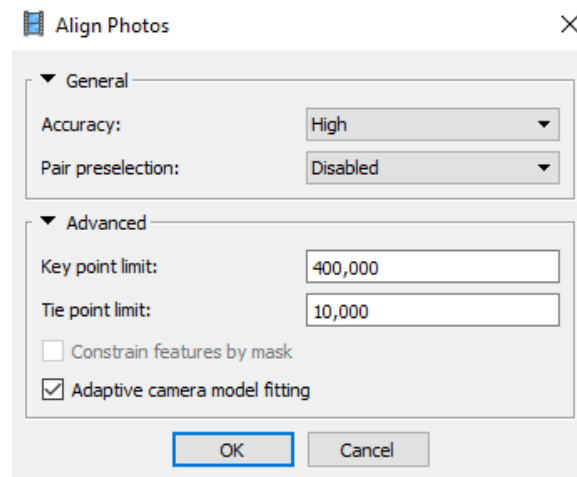


Figure 7: Alignment settings.

Selecting for high accuracy is a key point in the alignment stage, although it is really only time efficient for chunks of 400 photos or less. Adjusting the accuracy of the alignment from low to medium, or from medium to high yields a four-fold increase in accuracy. For much larger data sets, proceed to the next sub section titled “Alignment for Large Datasets.” Disable “Pair Selection” because it is an unnecessary option to select if the photos were taken with significant overlap as discussed before. It is only useful, if there are areas in which there is no clear overlap between neighboring positions. Selecting for a high “Key Point Limit” will increase the amount of points analyzed in each photo, consequently increasing the probability that the software will recognize points common among multiple photos. These points are called “Tie Points” and are what essentially patch the photos to one another. Selecting a high “Tie Point Limit” will ensure that each photo has many tie points. This number can vary, depending upon the dataset. It is important to note that this does not mean that each photo will find the exact number of tie points selected as the limit in the “Tie Point Limit” function. Constraining features by masks is only necessary if a mask has been made (as described in a previous section). Then select “OK” and prepare to wait for anywhere from 30 minutes to 12 hours, depending upon the dataset. Typically 100 photos will require three hours on high alignment settings given the computer requirements of “Basic Configuration” found above.

Alignment for Large Datasets

For large datasets, particularly for sets in which aerial and terrestrial photos were taken, time can become a serious issue. If markers were used, then the process will be much quicker because the software will already have recognized calculated reference points throughout the dataset. If no markers were placed, and the site is especially complicated (i.e. dozens of walls and different rooms), then run the aerial photos (which are typically fewer in number) independently of the terrestrial photos and select the options for the alignment as shown above. A set of 391 aerial photos from *Burg Hohenecken* required about 12 hours for processing. These will be combine with the terrestrial photos later on.

Digital Markers

Digital markers are only effective after the alignment has been done because the software will attempt to locate the positions of the markers in other photos. The software cannot do this if the photos have not already been aligned. Simply aligning the photos does not guarantee an accurate alignment! Once a marker has been placed, a red line will appear in photos were the software believes the position to be. The markers must be manually selected in at least three different photos in order to determine that the markers are in the correct places throughout the

model. Once the software has located the markers, a green check will appear the photos on the right side panel as well as a shadowed person (shown in Figure 8), which means that the marker has been found in other photos besides the photos in which the markers were manually placed.




Figure 8: Identification of digital makers.

If digital markers are chosen to be incorporated into the model, then they can be adjusted to GNSS coordinates or local coordinates. GNSS is the umbrella term for all positioning systems, including the American GPS (the largest such system), the Russian GLONASS, the Chinese BeiDou, and the upcoming systems Galileo and DORIS (from the EU and France respectively). This is accessed in the reference panel by selecting the “Reference Settings” tool—which looks like a wrench and a hammer crossing over one another shown in Figure 9.




Figure 9: Reference Settings tool.

Once this is selected, a box will open allowing the option to choose between Global coordinates or local coordinates (shown in Figure 10). If digital markers were placed in the absence of GNSS coordinates, as is the case with small objects, then select the “Local Coordinates (m)” option. There are other options within the box relating to camera and marker accuracy which allows users to tell the software what the margin of error should be while it looks for matching positions for markers.

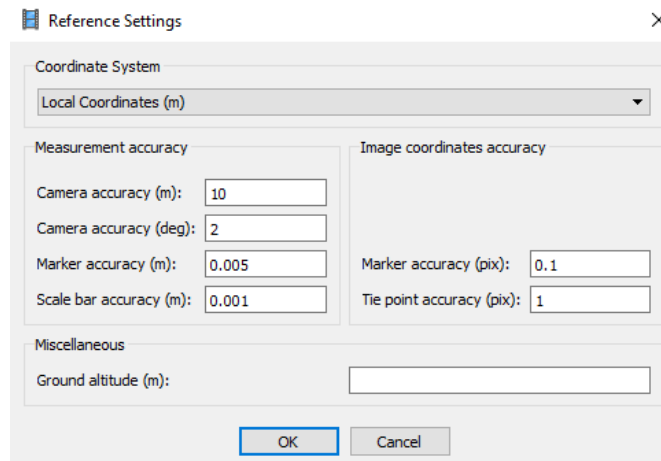


Figure 10: Settings selection in Reference Settings.

If this is done and the accuracies are adjusted, then be sure to select the “Optimize Cameras” tool in the reference tab to optimize the alignment to the adjusted accuracies (shown in Figure 11).



Figure 11: Optimize Cameras tool.

Optimizing the Point Cloud

Spend time with the point cloud deleting miscellaneous points that are unrelated to the object in question. Extra points serve only to slow the whole process down and make the model “messier.” To do this, select the chunk that has been aligned and then select the “Free-form Selection” tool (Figure 12) to highlight extraneous point located outside the realm of the immediate model. Once these points have been selected, delete them.




Figure 12: Free-form Selection tool.

Combining Large Datasets

Previously in this step, a separation was made for aerial and terrestrial datasets. Thousands of photos on high accuracy in *Agisoft* with the basic configuration of computers (the most common configuration—and cheapest) will require hundreds of hours to complete. As test, all of the 1782 photos taken at *Burg Hohenecken* (aerially and terrestrially) were placed in the same chunk and set to align on high accuracy. The time estimation was over 250 hours, which is neither efficient nor realistic for any research purposes. For this reason, the groups of photos were separated in order to be combined at a later stage. Now we have come to the later stage in which the photo groups will be combined. For this procedure, select the aerial photos chunk that has been aligned, go to Tools→Export→Export Cameras, and save it as a txt. file with a name like “Exported Camera Positions” within the same project folder. Due to the many exports done in *Agisoft*, make a folder within the project folder named “Exports.” To combine the aerial and terrestrial photos, make a new chunk in which all of the photos are placed. Then go to the reference tab and select the “Import” button shown in Figure 13.



Figure 13: Import tool in reference panel.

Select the exported camera positions from before and click OK. This will import the previously calculated positions of the aerial photos and is an important step in order to reduce computation time. The previously aligned photos were orientated on “High” which means that they are already very accurate. Entering the exported alignment into the combined chunk allows the software to recognize the positions of hundreds of the cameras already, in addition to the relative shape of the object being modeled. In the case of *Burg Hohenecken*, upon importing the aerial photos via the same procedure, a new alignment of all the photos was run on medium. Due to the aforementioned accuracy of the aerial photo alignment, re-running everything on high again would be unnecessary, though it would have increased the accuracy—a stronger computer configuration would have been useful in that case. The time allotted for the combined alignment was roughly 13 hours—still a substantial decrease from the ~250 hours of everything at once without imported camera positions.

Aligning Chunks

In many cases, it is difficult to model every side of an object whilst the object is in one position. Therefore it is necessary to photograph as much as possible with the object in one position and the switch it to another position in order to photograph the parts unseen by the first round (i.e. the tops and bottoms of smaller objects). Simply proceed with the alignment and optimization as stated before for both chunks individually, and then use the selection tool to remove the extraneous points and backgrounds (e.g. the table top upon which the object is resting). Then go to Workflow→Align chunks and enter in the values for a regular alignment. It will be relatively quick compared to the individual alignments because the software already knows where the camera positions are. Once this is complete, continue on to merge the chunks.

Merging Chunks

In order to merge any chunks, they must first be aligned, as described immediately before this. Simply go to Workflow→Merge chunks and select the chunks to be merged that have been aligned to one another. Be sure to select the options “Merge models” and “Merge markers” if any markers had been placed. Merging the dense clouds will not be necessary, unless the model was partitioned into separate chunks.

Step 6: Generating the Dense Cloud

Generating the dense cloud is where the model begins to appear more as the object than a phantom alignment of dots. Essentially the software is creating a more complex point cloud in which more information is attached to the alignment, making it denser. This is often times the longest procedure, but after it is complete, it will be clear whether or not the model is going to work. This all depends upon the accuracy of the alignment. To begin, simply go to Workspace→Build Dense Cloud and the following box will pop up shown in Figure 14.

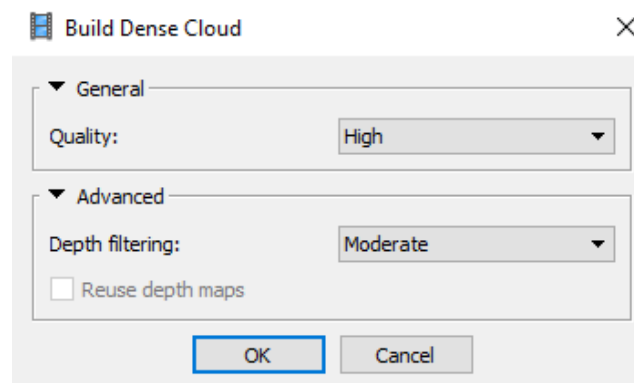


Figure 14: Dense Cloud settings.

Depending upon the size of the dataset and the configuration of the computer, select medium, high, or ultra-high—never select below medium unless a mask is being made. However, ultra-high really only works for advanced or extreme computer configurations. For smaller objects like ceramics or anything in which the data set consisted of 200-300 photos or less, then absolutely select “high” for the dense cloud. For larger datasets, consider selecting medium for the sake of time. This will only negatively affect the model if the alignment had not been done on high—once again stressing the importance of an accurate alignment. In the “Build Dense Cloud” box, select Mild, Aggressive, or Moderate for depth filtering. For complex objects, choose “Mild” whereas for objects with few small details, choose “Aggressive.” The option “Moderate” is somewhere between “Mild” and “Aggressive” which is the option that is usually best, but again, it depends upon the object. Regardless of which option is chosen, the process is long. A ceramics model of roughly 100 photos on high quality and moderate depth filtering required two hours. Larger projects like *Burg Hohenecken* required around 30 hours total on medium.

Dense Cloud Generation for Large Datasets

Due to the sheer size of the sparse cloud (alignment) of large models, it is best to generate the dense cloud in small portions at a time. This partitioning can be done by re-sizing the region of the model by clicking on the “Resize Region” button shown in Figure 15.




Figure 15: Re-size Region tool.

The region can also be rotated, which is most useful in order to place the object within the region so that rectangular smaller regions can be made. The rotating region button is shown in Figure 16.




Figure 16: Rotate Region tool.

Before actually resizing the region of the model, be sure to duplicate it by right-clicking on the chunk and selecting “Duplicate.” Do this prior to making new regions in order to avoid missing parts of the model later on. Most likely there will be a lot of duplicate chunks. Be sure to rename the chunks so that it is clear what is in each chunk. An example of the duplicate chunks from the *Burg Hohenecken* project is shown in Figure 17.

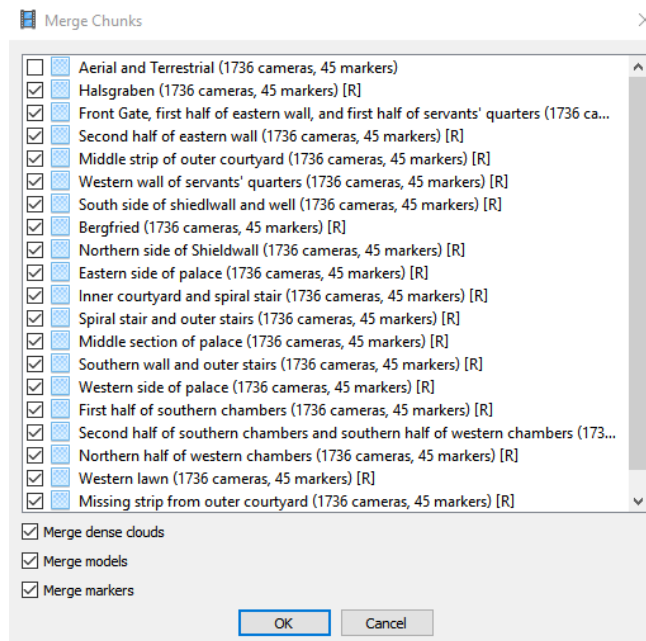


Figure 17: Merged chunks that have been referenced. Note the [R] on the right side of each name.

The problem that will arise in this process of duplicating and then generating the dense cloud separately for each chunk is twofold. First, each chunk on medium to high quality will require two to three hours—the more chunks the longer the process. Second, they will no longer be combined to one another. If the site has set Ground Control Points (GCPs) that have been geo-referenced or calibrated markers, then all that has to be done is to merge the chunks after the dense cloud, mesh, and texture has been done for each. When merging the chunks later on, be sure to check the boxes for “Merge dense clouds,” “Merge models,” and “Merge markers.” If no GCPs or calibrated markers were used, then a python script must be inserted into each chunk via the following procedure.

Referencing Duplicate Chunks with a Python Script

This step actually should not be done prior to generating the textures, but it should be kept in mind during this part of the procedure. To reference duplicate chunks, go to Tools→Run Script. Then enter in the following python script:

```
PhotoScan.app.document.chunk.transform.matrix = PhotoScan.app.document.chunk.transform.matrix
```

The script will first have to be saved as a txt. file in order to open it in the “Run script” option. Upon entering this script individually for each chunk, a “[R]” will appear within the name of each of the chunks on the right side (shown in Figure 17). This indicates that each chunk has been referenced to one another, which was made possible because everything had been aligned successfully prior to duplication. After entering the script into each of the desired chunks, then select the “Merge Chunks” option and merge the referenced chunks back into a single model. Do not re-align the chunks via “Align Chunks” because they were already re-aligned by the python script.

Step 7: Generating the Mesh

Generating the mesh is arguably the most important stage of the modeling procedure because it allows the texture to be applied to the object and also allows the object to be compatible for alignment with laser-scan models. The mesh does not take nearly as long as the dense cloud or the alignment. To begin the meshing, go to Workflow→Build Mesh and the following box will appear (shown in Figure 18).

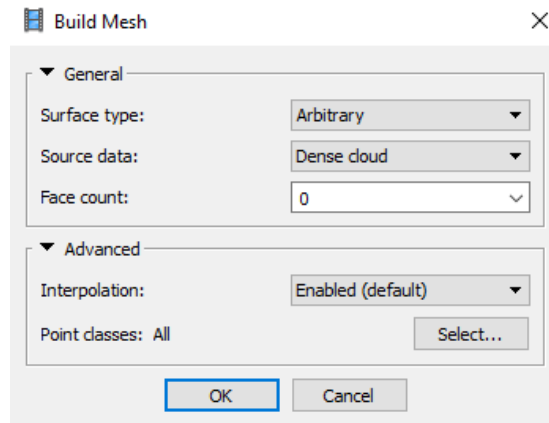


Figure 18: Mesh settings.

Unless the object is a flat surface, always select the surface-type to “Arbitrary” because it is designed for more complex objects that are not simply flat. The source data will always be the dense cloud from the same project. Essentially the mesh triangulates the data so that the dense cloud becomes a network of triangles with faces and vertices rather than points. Imagine having a table with objects all over it and then casting a net atop all of the things on the table. The result is a more fluid depiction of the surface of the table plus the objects. The mesh can be made more detailed by increasing the face count. By selecting “0,” *Agisoft* will build the full detailed size of the mesh using the data from the dense cloud. The face count also has a drop-down selection so that varying levels can be chosen for the size of the mesh. Typically the software will recognize the capabilities of the computer being used and will adjust accordingly if “0” is typed in the box which makes things much easier. Entering “0” can create incredibly large meshes. For example, a ceramic pot with 135,000 points yielded a mesh of over 55 million faces after “0” was entered. The model looked great, but is not entirely practical for transfer. There is a balance between quality and practicality that has to be taken into consideration for every model made.

Note: the failure or success of the mesh can be assumed whether or not the progress bar passes the ten percent mark. It will typically linger around the six percent mark for a while before reaching the ten percent mark, if it fails to pass ten after 30 minutes or more, start the mesh over at a lower face count setting.

Step 8: Generating the Texture

The texture is the simplest portion of the entire process and the quickest. It is also the form of the model that is most appealing to audiences and for surface analyses. It essentially drapes the photos over the mesh that was created in the previous step. Go to Workflow→Build Texture and the following box will appear (shown in Figure 19).

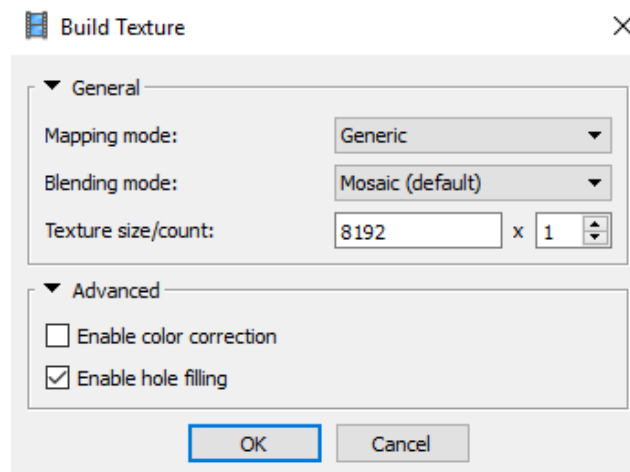


Figure 19: Texture settings.

Keep the settings for “Mapping mode” on “Generic” and the “Blending mode” on “Mosaic.” A texture is essentially a high resolution mosaic that adds detail to the faces of the mesh. Selecting higher texture sizes per count means selecting a higher resolution to the exported texture file. However, the larger the number, the more important the computer configuration tends to be. The “Texture size/count” of 8192 x 1 works well with 32 GB of RAM but less RAM may lead to difficulty with a texture size of 8192 per count. If the computer does not have 32 GB of RAM or more, enter 4096 for the texture size.

Upon finishing the texture, go through the model and delete any extraneous points. Note that any points deleted on the mesh/texture will not be deleted from the alignment or the dense cloud, but will be deleted from the exported model. It is much easier to see which point are part of the model and which are not when viewed as a texture as opposed to a dense cloud.

Step 9: Scaling the Model

Scaling the model is an essential part of the photogrammetric process. If calibrated Agisoft markers were used in the study, then it is not necessary to re-scale it. If no markers were used, then the model can be scaled quite easily, though not quite to the accuracy of calibrated markers. Click on one of the photos in the set that clearly displays the centimeter scale. Create two digital markers on the corners of one of the centimeter boxes, by right-clicking and selecting “Place marker.” This is essentially the same process as placing digital markers discussed earlier. Rename them “0 cm” and “1 cm” to keep them separated from other digital markers that have perhaps been placed, but more importantly to make them intuitive. Go to the reference tab, select both scale makers, right-click and select “Create scale bar.” A new scale bar will appear directly below where the markers are located in the reference tab (Figure 20). Next to the newly created scale bar, double-click in the “Distance (m)” column to insert the correct scale.

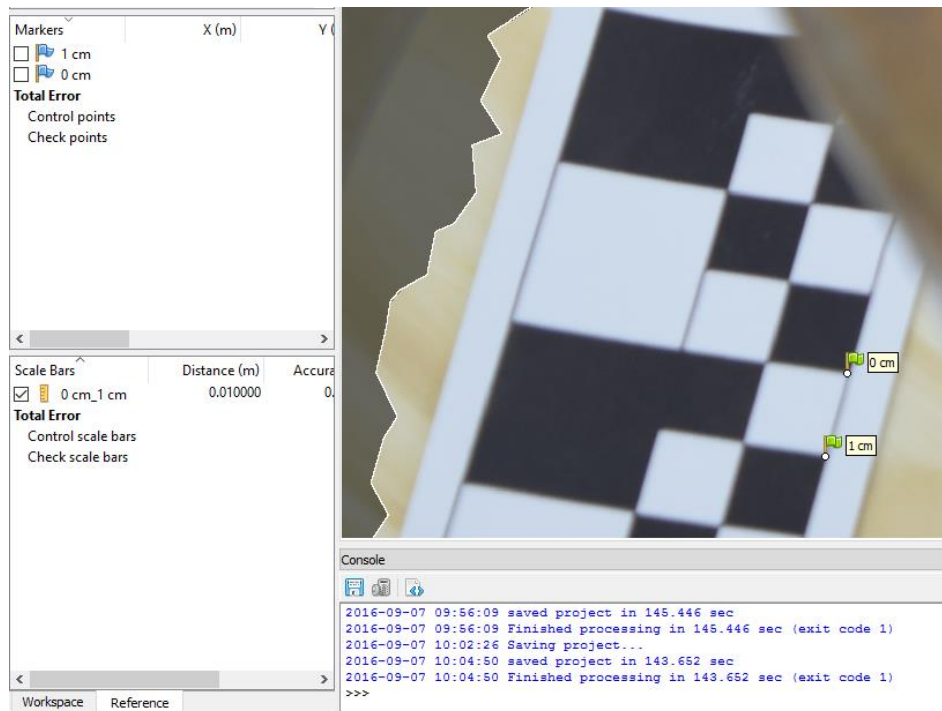


Figure 20: Scaling the model in the reference tab.

Note that this is in meters! Thus, 1 centimeter (as in the case of this example) is 0.01. This process will apply the scale bar to the entire model after selecting the update button shown in Figure 21.



Figure 21: Update the photoset and model.

Step 10: Exporting the Model

Most platforms for demonstrating models use the obj. file format. In order to export the model from Agisoft as an obj. file, go to File→Export Model, name the file something intuitive and press OK. A popup will appear referring to the textures shown in Figure 22.

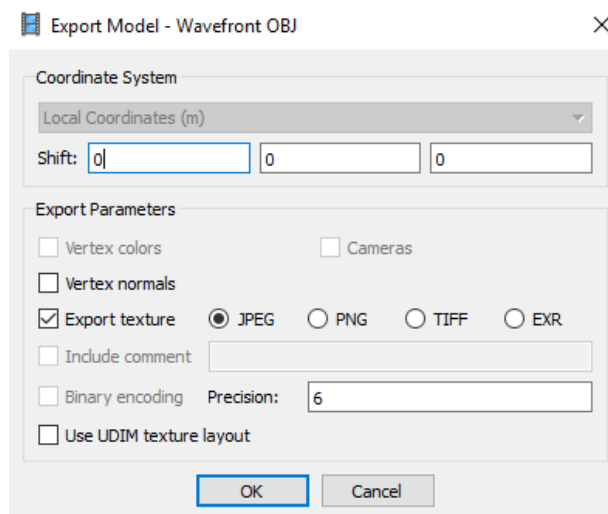


Figure 22: Export Model settings.

The exported texture files will be saved separately from the obj. file which means that if they are not transferred along with the obj. file, then the model will lack color and will resemble a gray mesh. Depending on the size of the model, there can be anywhere from 1-20 texture files. Save these as either a jpeg. or png. file as they are more compatible with most platforms. Exporting the textures as png. files will result in much larger files than jpgs, making it unlikely to be able to upload the model to Sketchfab.

Step 11: Decimating the Model

Decimation is done in order to reduce the size of a model so that it can be uploaded. The various upload platforms have various file size limits. The platform discussed here is the website *Sketchfab*. *Sketchfab* is very particular about the size of a model. They limit the upload size to 50MB, which can be substantially smaller than the size of the model created. For example, the model of *Burg Hohenecken* created using a high quality alignment and high resolution dense clouds, meshes and textures yielded a 1.4GB file.

Decimating the Mesh in Agisoft

There are two ways to decimate the model. The first method is in *Agisoft*, though once decimated, the textures will have to be run again to adjust for the lower face count in the decimated mesh. To do this via *Agisoft*, duplicate the chunk in which the model is located and then go to Tools→Mesh→Decimate mesh. The following box will appear (shown in Figure 23).




Figure 23: An example for the Decimate Mesh settings.

This can be a frustrating process due to the varying sizes of the models. The best way to estimate the decimation of the mesh is to first save the model as an obj. and then look to see what the size of the file is. For example, a model of a glass bottle with over 8 million faces in its mesh yielded a model of 603,000MB. To estimate the face count in order to yield a model of only 50MB, perform a proportion to solve for the number faces. Keep in mind that the texture files also have to be uploaded, but they cannot be uploaded separately and are included in the *Sketchfab* 50MB limit. Therefore, create an obj. around 45MB to allow room for the jpeg. or png. texture files.

Decimating the Mesh in Meshlab

Another, more effective way to decimate the mesh without having to re-run the textures is to open the originally exported model into a program called *Meshlab*—a free downloadable program designed for post-processing of meshes. To decimate the model exported from Agisoft, open Meshlab and drag only the obj. file into the program. Depending upon the size of the model, it will take some time for it load. Once it has loaded, go to Filters→Remeshing, Simplification and Reconstruction→Quadratic Edge Collapse (with texture). This will then open a box in Figure 24. For the purpose of uploading to Sketchfab, type in 350,000 (roughly equivalent to 45 MB in size) as the new face count in the box at the top, click “Apply,” and wait again for a bit. Save the project and export the model under File→ Save Project and Export Mesh As, respectively.

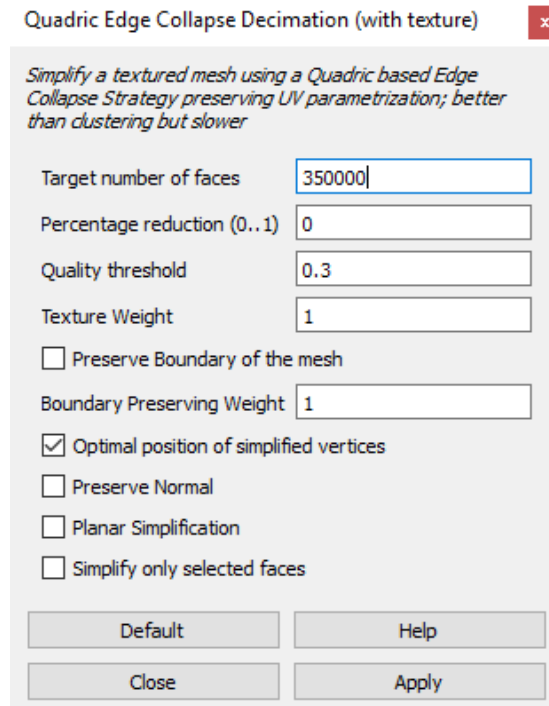


Figure 24: Decimation of a mesh with textures in *Meshlab*.

Uploading to *Sketchfab*

First make an account in *Sketchfab* and then go to the folder in which everything has been saved for the model. To upload everything together, select the obj., mtl., and mlp. files that were created after saving the project and exporting the model from *Meshlab*, and select the texture files generated from exporting the model from Agisoft. Save these files together as an archive using *Winzip*, and then upload this archived model to *Sketchfab*. If everything is under 50MB, then it should upload without a problem.

ⁱ Sapirstein, Philip. 2016 Accurate Measurement with Photogrammetry at Large Sites. *Journal of Archaeological Science* 66:137.

ⁱⁱ Püschel, Hannes, Martin Sauerbier, and Henri Eisenbeiss. 2012 A 3D Model of Castle Landenberg (Ch) from Combined Photogrammetric Processing of Terrestrial and UAV-Based Images. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. 37:93-98.

ⁱⁱⁱ Wenzel, Konrad, Mathias Rothermel, Norbert Haala, and Dieter Fritsch. 2013 SURE—The ifp Software for Dense Image Matching. *In Photogrammetric Week '13*. Dieter Fritsch, ed. Pp.59-70. Berlin: Wichmann Verlag.

^{iv} Püschel, Hannes, Martin Sauerbier, and Henri Eisenbeiss. 2012 A 3D Model of Castle Landenberg (Ch) from Combined Photogrammetric Processing of Terrestrial and UAV-Based Images. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. 37:94.

^{v v} Remondino, F., L. Barazzetti, F. Nex, M. Scaioni, and D. Sarazzi. 2011 UAV Photogrammetry for Mapping and 3D Modeling—Current Status and Future Perspectives. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*. 33:25-31.

^{vi} Püschel, Hannes, Martin Sauerbier, and Henri Eisenbeiss. 2012 A 3D Model of Castle Landenberg (Ch) from Combined Photogrammetric Processing of Terrestrial and UAV-Based Images. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. 37:94.

^{vii} Sapirstein, Philip. 2016 Accurate Measurement with Photogrammetry at Large Sites. *Journal of Archaeological Science* 66:140.

^{viii} Agisoft LLC. 2016 Agisoft Photoscan User Manual: Professional Edition, Version 1.2

^{ix} Sapirstein, Philip. 2016 Accurate Measurement with Photogrammetry at Large Sites. *Journal of Archaeological Science* 66:138.

^x Sapirstein, Philip. 2016 Accurate Measurement with Photogrammetry at Large Sites. *Journal of Archaeological Science* 66:141.

^{xi} Sapirstein, Philip. 2016 Accurate Measurement with Photogrammetry at Large Sites. *Journal of Archaeological Science* 66:144.