MicroPasts Technical Note 3a (2014)

3D Models of Artefacts via Structure-from-Motion (PhotoScan)

Andrew Bevan

1. Introduction

This set of working notes follows on from those concerned with photographic strategy and photo-masking for making 3D models, and deals with the Structure-from-Motion/Multi-view stereo (hereafter SfM) process itself. There are online SfM solutions for building 3D models from photographs, such as 123DCatch (http://www.123dapp.com/catch) or Photosynth (http://photosynth.net/) and you could certainly take raw photos from the MicroPasts site and explore these online possibilities yourself. Typically, however, they do not provide as much flexibility or as detailed a final result as more user-guided offline alternatives, so this tutorial focuses on the latter, and takes you through the entire workflow in a commercial software package called PhotoScan Pro (note that several steps will not work with the 'Standard' version of this software, only the 'Pro' version, but you can always try out a demo version of the latter, albeit without being able to save your work: http://agisoft.ru/products/photoscan/professional). Although MicroPasts is almost entirely an open source software project, we do currently use PhotoScan Pro for our model-building, because it supports (a) photo-masking and (b) control points that currently produce the best results we have encountered. However, in a parallel set of notes to these, we will also take a look at the equivalent workflow achieved in the open source packages VisualSfM and Meshlab.

You can find the datasets used here and any further resources on our website. Note also that the example photographs used below are not meant to be considered 'optimal'. They have been captured quite quickly in reasonably poor lighting conditions with automatic camera settings and some out-of-focus areas on the object. We have also not colour-balanced (aka 'white-balanced') the individual photographs. The 44 photographs used here are adequate for a good model, but on the low side. Using more photographs would increase processing time but typically improves model output (and many people prefer 100+). So there are plenty of ways to further perfect capture of a single object if this your goal.

2. Steps in Agisoft PhotoScan Pro

2.1 Launch PhotoScan Pro and your screen should look like the one below, with normal menu bars at the top, a model viewer window middle-right, a photo thumbnail viewer at the bottom and a Workspace on the left.

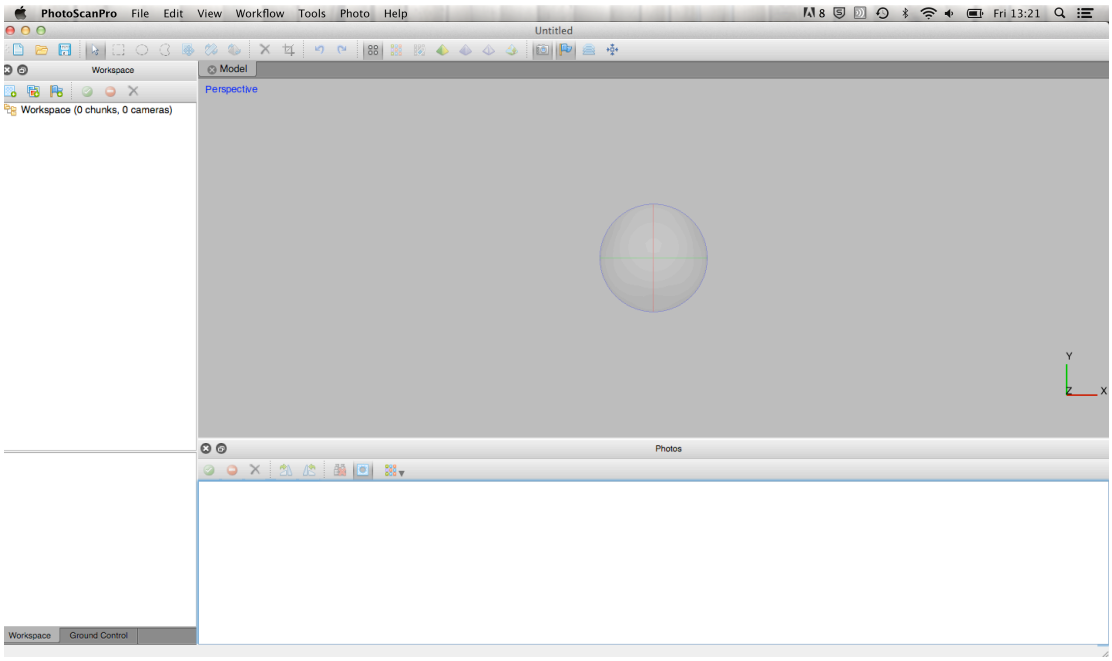

2.2 If you are lucky in terms of the object you are working with, then you might be able to capture it entirely without ever turning the object over. However, for most artefacts you will probably need to rest them on a flat surface and then flip them over to photograph the other side. This, then, is also a modelling task best tackled in at least two steps, and with PhotoScan these steps are handily split into what are called 'Chunks'. We will look at modelling just one Chunk to begin with, covering all the photographs taken with the object in a single position (in the example used here the photographer usefully put a paper label in the frame with an 'A' for the first object position and a 'B' for a second one when the object was flipped over). Let us start by adding the photographs for a chunk 'A' then. First click on the Add Photos button at the top of the workspace. 

2.3 Browse to where you have stored the tutorial dataset and make sure to add only the original JPG images, not their corresponding masks (which have a '_mask' ending and to which we will come back in a moment, see also tutorial 2), and choosing all the photos in the sequence meant to belong to this chunk (in this example, they are IMG_0441.JPG to IMG_0465.JPG). They should appear on your screen as below:

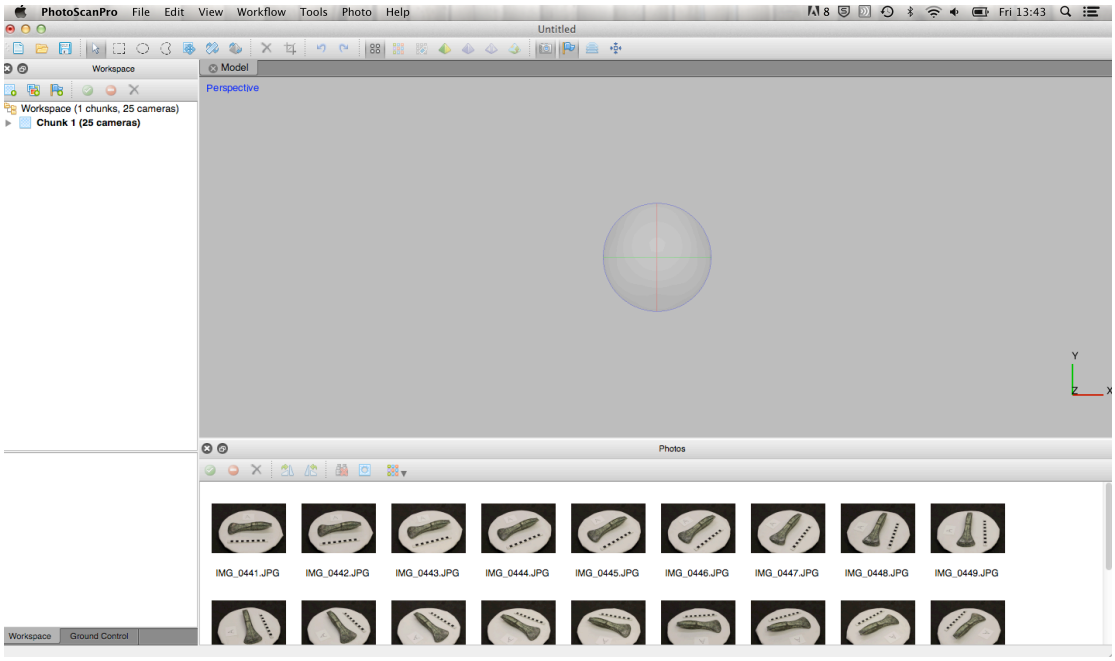

2.4 There is some useful information that is brought in with these images, and resides in the EXIF metadata embedded in the image header. For example, right click on the first image thumbnail, and choose "Show Info…". In the example below, we get information indicating that the photo was taken by a Canon camera with a focal length of 43 (mm), amongst other things. Focal length in particular provides an important starting point for the SfM method and most model-building works much better when such information is available in the image metadata. Most cameras today create such data as a matter of course, but image manipulation programs sometimes strip it out when saving a file afresh. This, along with the fact that any cropping of the image leads to an inaccurate assessment of focal length, means that it is best to work with the direct output from the camera (e.g. files produced directly by the camera or directly from camera RAW format) and in the case of PhotoScan and VisualSfM, JPGs are perhaps the most straightforward despite the potential loss of detail they risk due to their compression method.

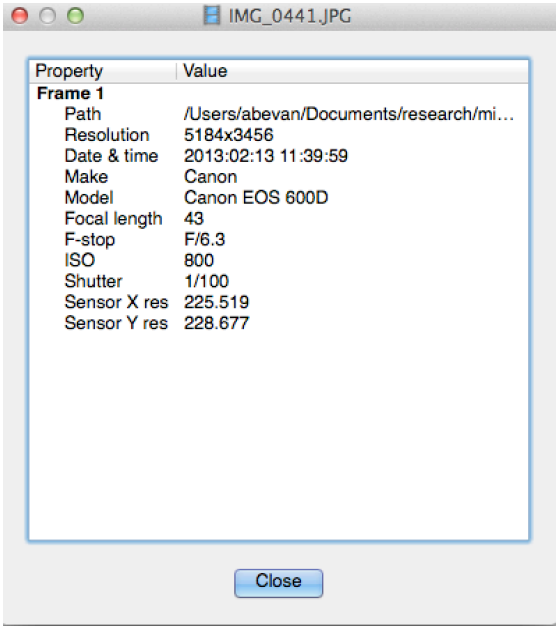

2.5 Now right click on the first image thumbnail again and choose "Import Mask". From the resulting dialogue (see below), choose to 'Import masks for' "Active chunk", and the 'Method' should be "From File". Make sure the 'Mask file names' correspond with your mask images names, i.e. in this case change the default .png format to .jpg format, and then click OK.

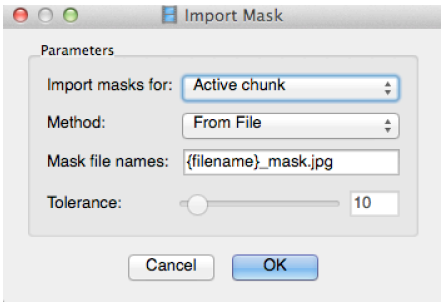

You will be asked to choose a folder with the masks in it. Choose your tutorial folder with both the original photos and their accompanying masks (the mask files might be ones that you learnt to create by hand in image processing software in a previous set of MicroPasts technical notes or that you downloaded from the MicroPasts site). The software will now go through and look for .jpg files with a '_mask' ending and endeavour to match them up with the original photos that you have already loaded.

2.6 Once it has imported them (you should see a progress bar momentarily letting you know that something is being imported), then there will be no observable change on screen. However, you can look more closely at a given image by double-clicking on its thumbnail at the bottom of the screen, and it appears in the main model viewer window. Double click on the first photo and your screen should look like this:

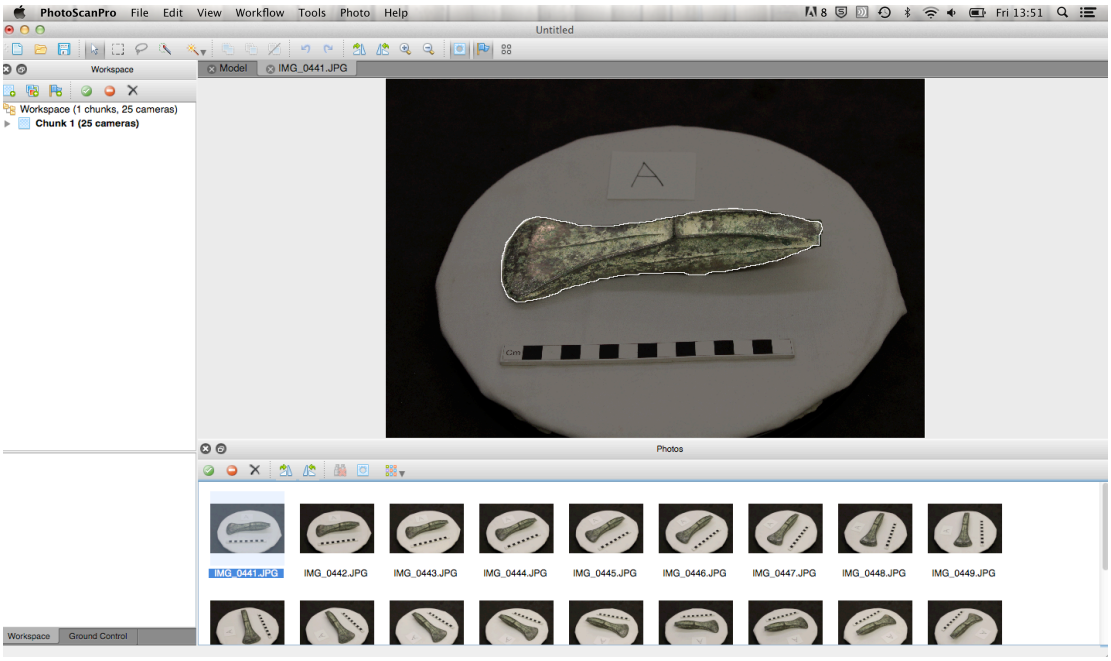

The white outline indicates the mask border (with the unmasked part of the object in the middle shown as a more saturated set of colours). Once you have checked this, close the view of this image by clicking on the X at the top of its window.

2.7 To see whether the batch import routine for the masks has been successful in all cases, first close the image you have opened by clicking on the X at the top of the model viewer window. Then click on the 'Show Masks' button on the thumbnail viewer. Any photos that have masks will now appear in mask form, and in the case of our tutorial example, we should see that all of the images have corresponding masks, as per below:

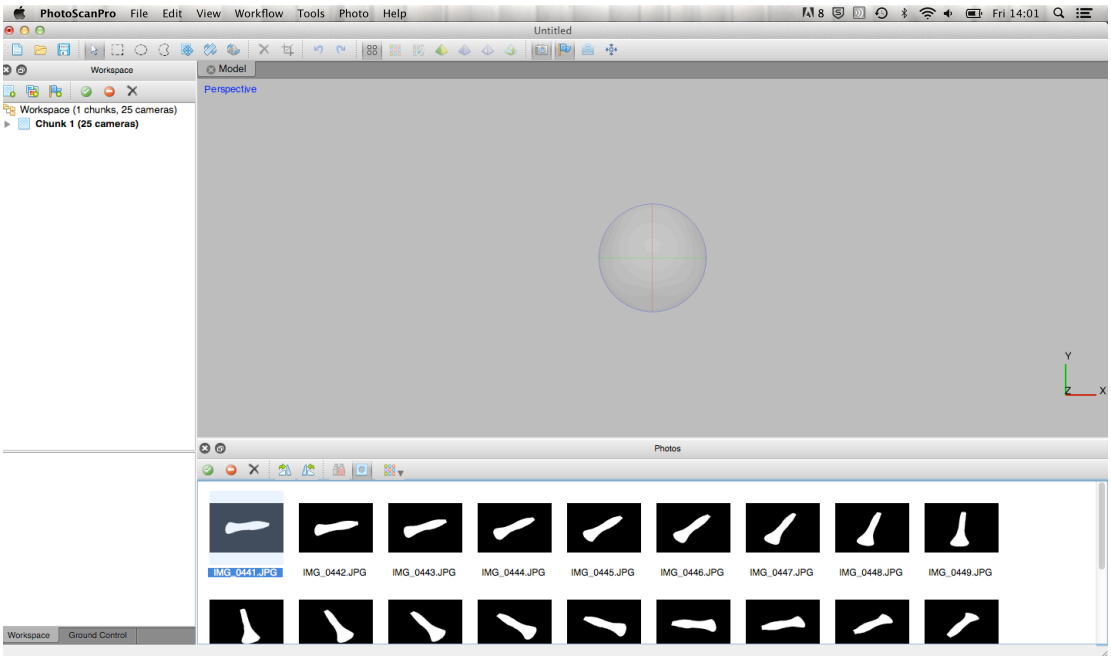

2.8 On the left hand side of your screen, you should now see that you have 'Chunk 1' with 25 cameras (i.e. 25 photos from which we wish to guess at the physical positions of 25 cameras). Right click on Chunk 1, choose "Rename…" and rename it 'Chunk A'.

2.9 Now save your PhotoScan project by clicking on File→Save and save a project file called P1991_3-3_5.psz into your tutorial folder. Note that this file does not include all the model data, but only some of it, and is otherwise simply a pointer to where the raw photos are held and how they should be treated. Warning: if you move this file or move the data itself at all, then it will not longer work properly!

2.10 We are now in a position to run the first two stages of a typical SfM method, which involve the software seeking to go through each photograph and identify distinct groups of pixels (e.g. with similar colour gradients) that constitute 'features' that are likely to be discernible in several images. After these features have been described for each image, they are matched across multiple images to produce a network of spatial relationships from which individual camera position for each photograph can be reconstructed. The end result is a sparse cloud of 3D point locations that mark the successfully matched features. To achieve this in PhotoScan, we choose Workflow→Align Photos. Note that in the resulting dialogue you have several options. Higher 'Accuracy' takes more time to process but produces better estimates of camera positions (we will accept the default of 'High'). 'Pair preselection' can be used to speed up the processing time for this step, by reducing the number of image pairs that need to be compared to one another, either by running a low accuracy matching routine first ('Generic') or by taking advantage of camera locations that are already known in some other way ('Ground Control'). We will accept the default to keep this 'Disabled'. If you expand the 'Advanced' tab, you will also see that you can also choose to alter the 'Point limit' which again seeks to save on processing time, by capping the maximum number of features that can be identified in any one photo and thus be compared across photos (we will accept the default of 40,000). The final setting 'Constrain features by mask' is important as this allows us to take advantage of the mask files we have just loaded (so leave it ticked as is the default). Now click OK and the resulting process may take some time, but when it is finished, you should have a screen roughly like the one below (note that these notes were made with a black background to the Model pane… Yours may be different and we will look at how to change this setting later).

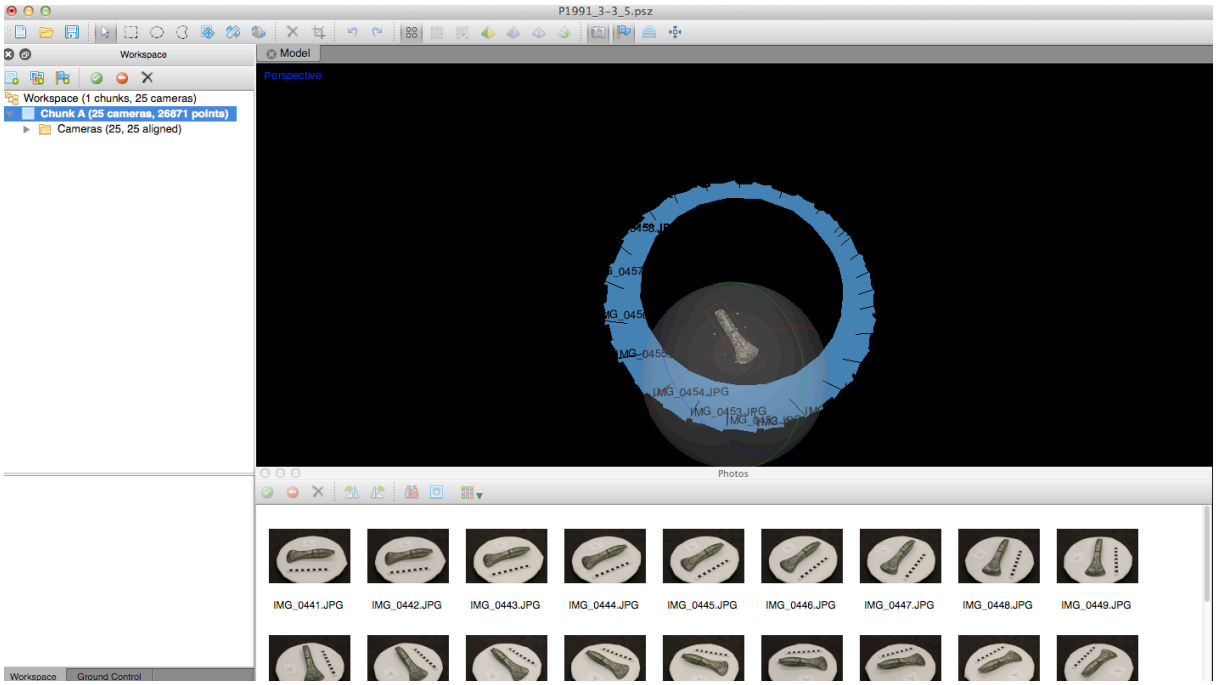

In the main Model pane, you now have a series of reconstructed camera positions shown in blue and labeled according to the photograph to which they refer. On the left-hand side in the Workspace pane, it also now says that 25 out 25 cameras were successfully aligned. Try navigating in the Model pane: the navigation controls vary by computer platform and hardware settings a bit, but a left mouse click usually rotates the view, a right mouse click pans, and you can zoom with the mouse scroll wheel (or by pushing up or down on a trackpad).

Clicking on any given photograph in the Photos pane causes the corresponding camera placement in the Model pane to be highlighted in pink, as below.

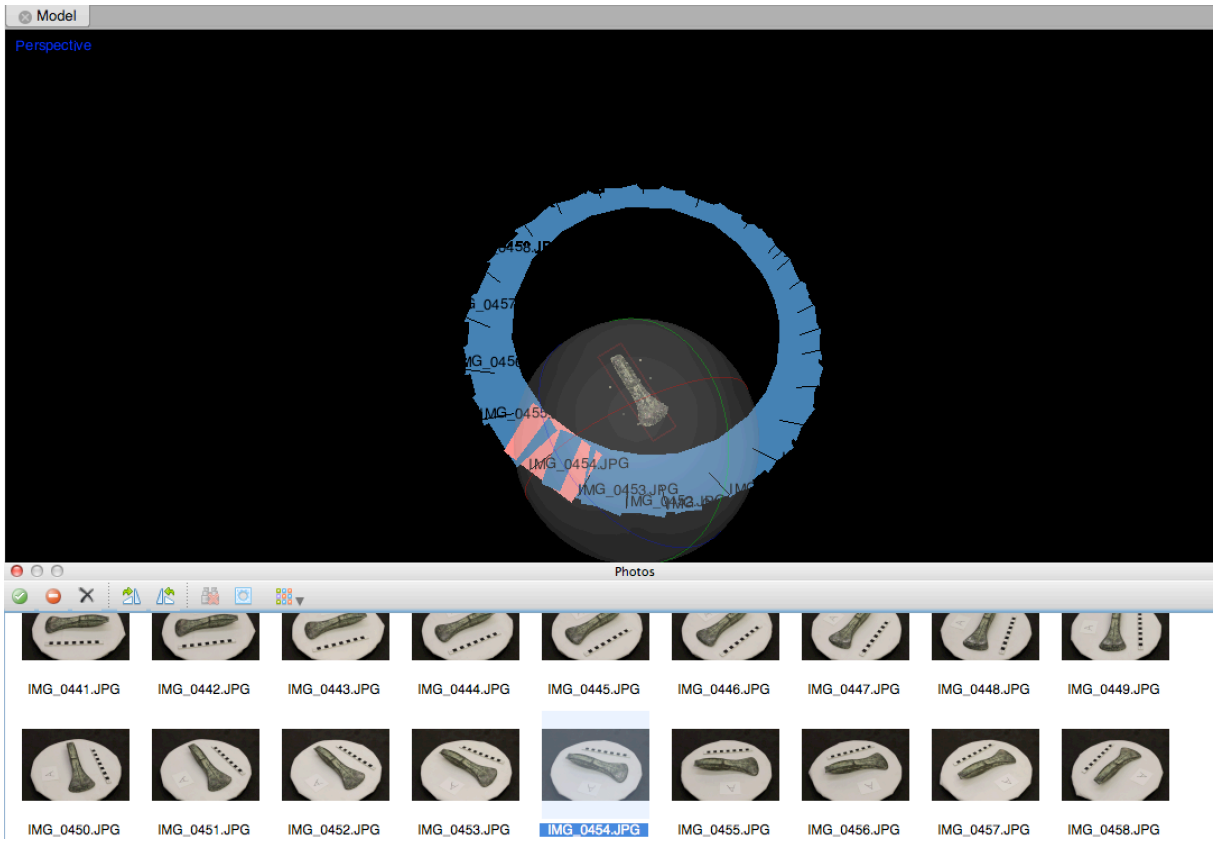

The camera positions in this example make it obvious that the object was photographed fairly systematically, but it is also a little bit of a Galilean illusion because, in this particular case, the camera was set up on a stationary tripod, and it was the object that was 'moved' by rotating it in a series of small steps on a turntable! Regardless, the reconstructed relative camera positions are fairly accurate ones and, it is clear from the 'sparse' point cloud in the middle of the scene that the reconstruction is promising. Zoom in to the sparse point cloud so that your screen looks a bit like the one below.

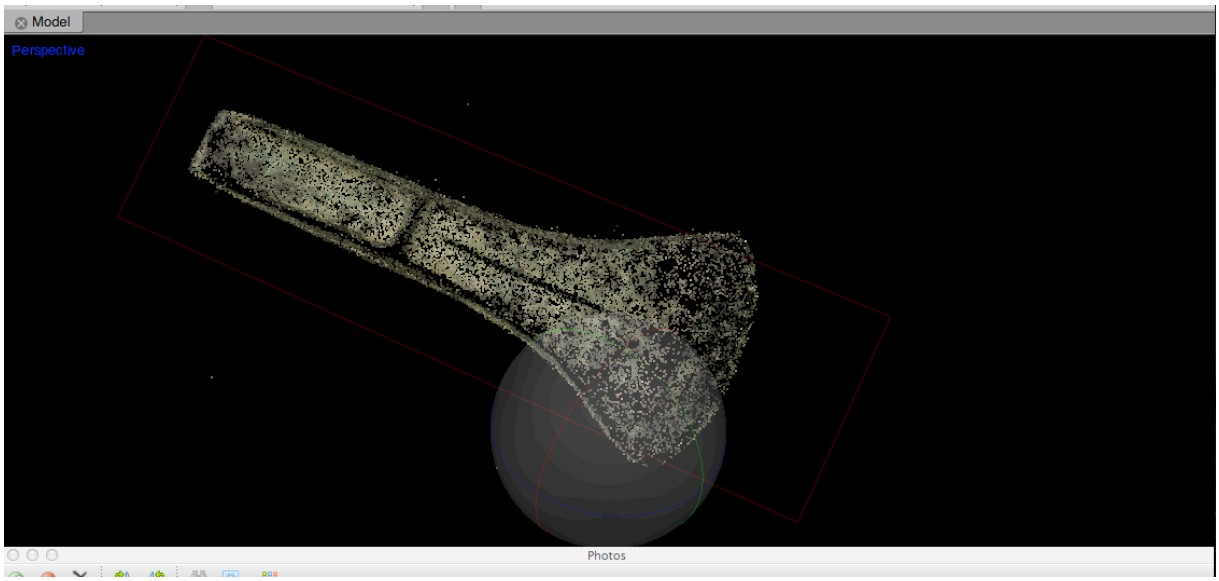

2.11 Having gone through these steps, now is another good time to save your work by clicking on File→Save.

2.12 Currently this model has been reconstructed in a largely 'scale-free' way, in the sense that while the shape of the object is well-modelled, there is no sense as yet of its absolute dimensions (e.g. the axes' actual length). Before we model the rest of the object, it is therefore useful to adjust this sparse point cloud for approximate scale. We can do this by creating some 'Markers' on the raw photographs with reference to the scale bar shown in the photographs (although you could do something similar for any object of known length in the photograph: it does not matter if that portion is inside or outside of the image mask). In the Photos pane, double-click on the first image (IMG_0441.JPG) and navigate on this image until you have a good zoomed-in view of the whole scale-bar in the photograph. Now right-click on the top-left corner of the scale-bar (i.e. the thin black outline, not the edge of the cardboard) and choose 'Create marker'. This will place a marker at that location called 'point 1'. If you want to be more precise then you can zoom in a bit more and drag the marker into the exact place you wish.

2.13 Now do the same for the top-right hand corner of the scale-bar, adding 'point 2'. Your screen should now look something like this:

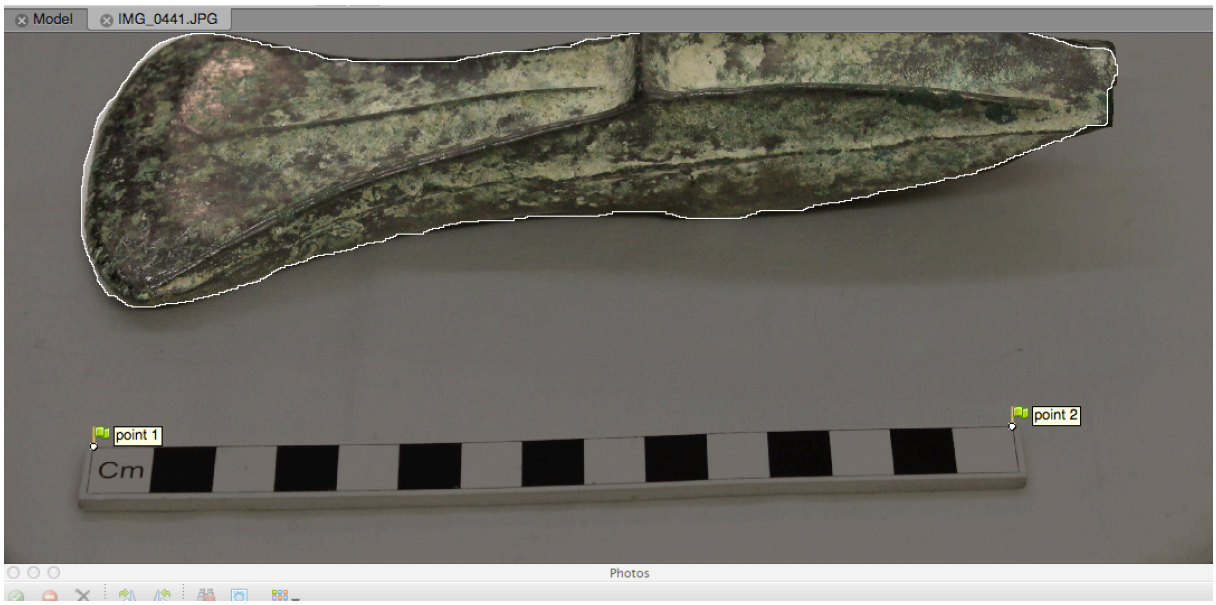

2.14 Now we will specify the same two markers in a second photograph. In the Photos pane, double-click on the second image (IMG_0442.JPG) and navigate until you have a good zoomed-in view of the top-left corner of the scale-bar in this photograph. Right click on this top-left corner and choose Place marker→point 1. The location of the marker you originally placed in the first image is now also placed in the second. Go and do the same for the top-right corner of the photograph, choosing to place 'point 2' there.

2.15 Having now placed both point 1 and point 2 on two separate photographs, you should notice that there are two green flags showing on the corner of the two

photographs in the Photos pane . There are also now a series of grey-colored 'rubber-sheet' icons on the corners of the other photographs in this pane . The software is seeking to identify (via its knowledge of camera positions) where the same two markers might be on the other photos, with the grey icons indicating that this has been done mathematically and without your direct intervention as a user. Now double-click on the third image (IMG_0443.JPG). You will see the two marker positions shown in grey in roughly the correct positions on the photo (N.B. Note also some funny-looking portions of the photo where the image mask left a few background pixels unmasked and appearing as white-outlined little speckles – a bit untidy, but not something to worry about as these cannot be found on multiple images). It is quite possible that you will not be able to improve upon the current position of the markers, but just to demonstrate that you are allowed to, please zoom in and click on point 1. Moving it fractionally if you wish or leave it where it is. You have now made a user edit to this marker and it has been turned into a green flag, both on the photo and in the Photos pane.

2.16 Normally with this scaling procedure, you merely need to pick two points on a scale-bar or length of known dimension, repeat for a second photo and then do some spot checks to ensure that these positions have been correctly calculated for the other photos in your chunk. If the camera reconstruction you did earlier was successful, there should be little need for a lot of fiddling around with re-positioning the markers manually (if you find that you are doing a lot of this, reconsider your photo strategy and/or the way you are chunking your photographs sets). You are now in a position to tell the software exactly what the distance between point 1 and point 2 is in real world units so that it can scale the entire model accordingly. At the bottom-left of your screen are two tabs: 'Workspace' and 'Ground Control' with the former currently selected. Click on the latter to switch to it and you should see that the panes on the left-hand side of your screen switch to ones like those below:

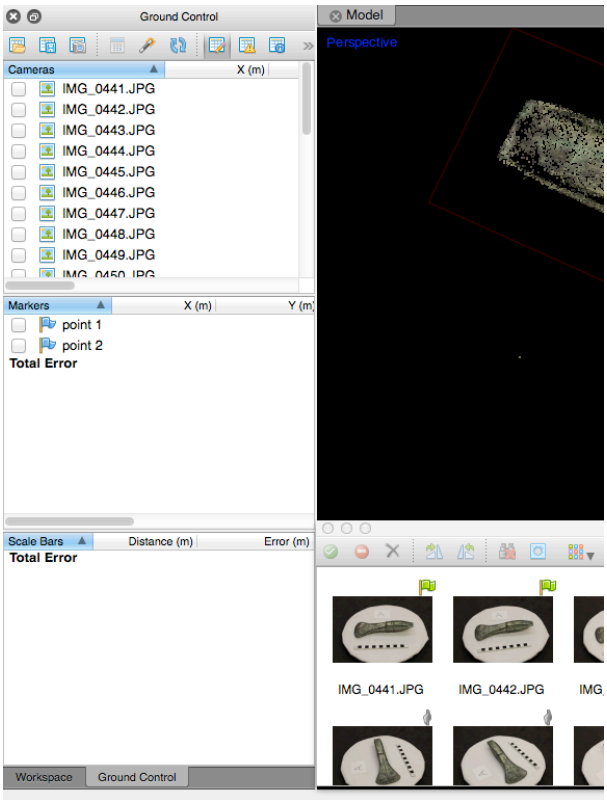

2.17 In the middle pane on the left-hand side, tick the boxes next to both point 1 and point 2 and also use the SHIFT key and click on both point 1 and point 2 to select them (in blue). Now right click on this selection and choose the option 'Create Scale Bar'. A yellow line will appear between these points in your Model pan, whilst 'scale 1' will also be added to the bottom pane on the left-hand side. Depending on where you are zoomed in to, your screen should look something like the one below.

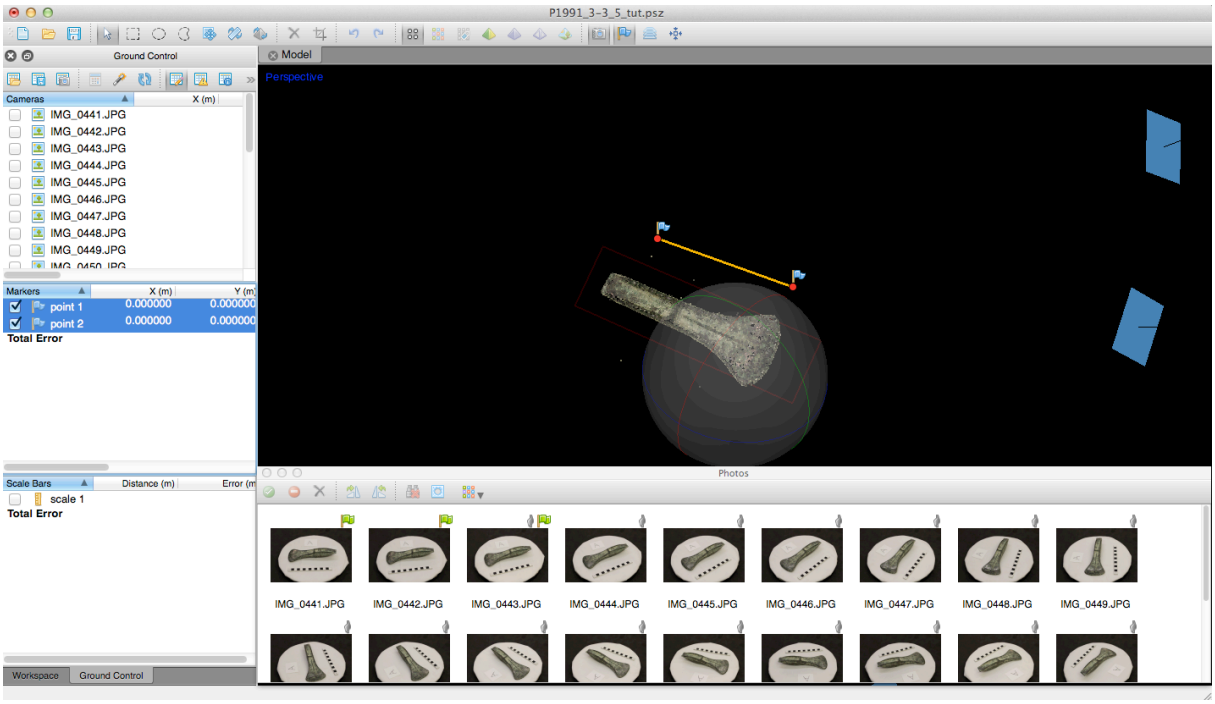

2.18 Now tick the box next to 'scale 1' at the bottom left of your screen and the distance measurement on this scale should appear as '0.0000'. Double-click on this measurement and change it to 0.15 (the scale bar is 15cm long; for consistency between object models and larger models of buildings and landscapes we may as well work with units in metres). Now at the top of the 'Ground Control' pane on the left-hand side, click on the 'Update' button  .

You have now resized the entire model to reflect the specified dimensions. Overall, our experience is that this scaling method produces absolute dimensions that are generally pretty accurate, but we have not yet compared them exhaustively to real world measurements, and it does all depend on your objectives. Note that the key issue is really not whether an absolute length measurement is perfect at the end of an SfM modelling process, but whether the technique reconstructs *relative* dimensions correctly (i.e. length to width ratio, and more generally overall shape), since absolute XYZ dimensions can always be slightly re-scaled later on if need be. Our experience is that with a sufficient quantity of well-taken photographs, this technique produces models with very reliable overall shape properties that are consistent with those, for example, obtained via laser scanner (though the very best, well-calibrated laser-scanning equipment still probably can boast mildly better results… not enough comparisons out there yet to really say however!).

2.19 Now is a good time to save your overall work by going to File→Save.

2.20 Having created a first 'Chunk A' with reconstructed camera positions, a sparse point cloud model and the correct dimensions, we are now in a position to add a second 'Chunk B' covering the other side of the object, which was photographed in the same way

as the first, after the object had been flipped over on the turntable. First switch back to the Workspace pan, by clicking on the Workspace tab at the bottom-left of your screen. Then at the top of the Workspace pane, click on the 'Add Chunk' button. This will add a new chunk to the pane. Right click on this chunk, choose Rename… and call this 'Chunk B'. Your workspace pane should now look similar to the one below.

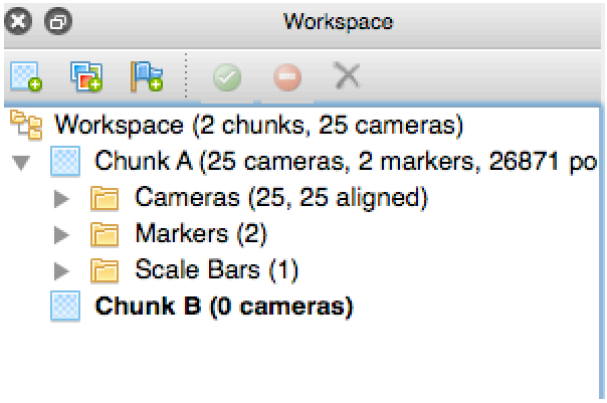

2.21 To add a new set of photos to this chunk, right-click on the name 'Chunk B' in the Workspace pane and choose 'Add Photos'. Browse to where you have saved the example set of photos and add those labeled 'B' to your chunk (IMG_1688.JPG to IMG_1714.JPG), excluding the mask files associated with them for now.

2.22 Now batch import the relevant mask files in the same manner as you did for Chunk A above (see 2.5-2.7).

2.23 Now choose Workflow→Align Photos, accepting the same defaults as you did for Chunk A (see 2.10).

2.24 You should now have a second model (of camera positions and sparse point cloud) showing in your Model pane, similar to the one below:

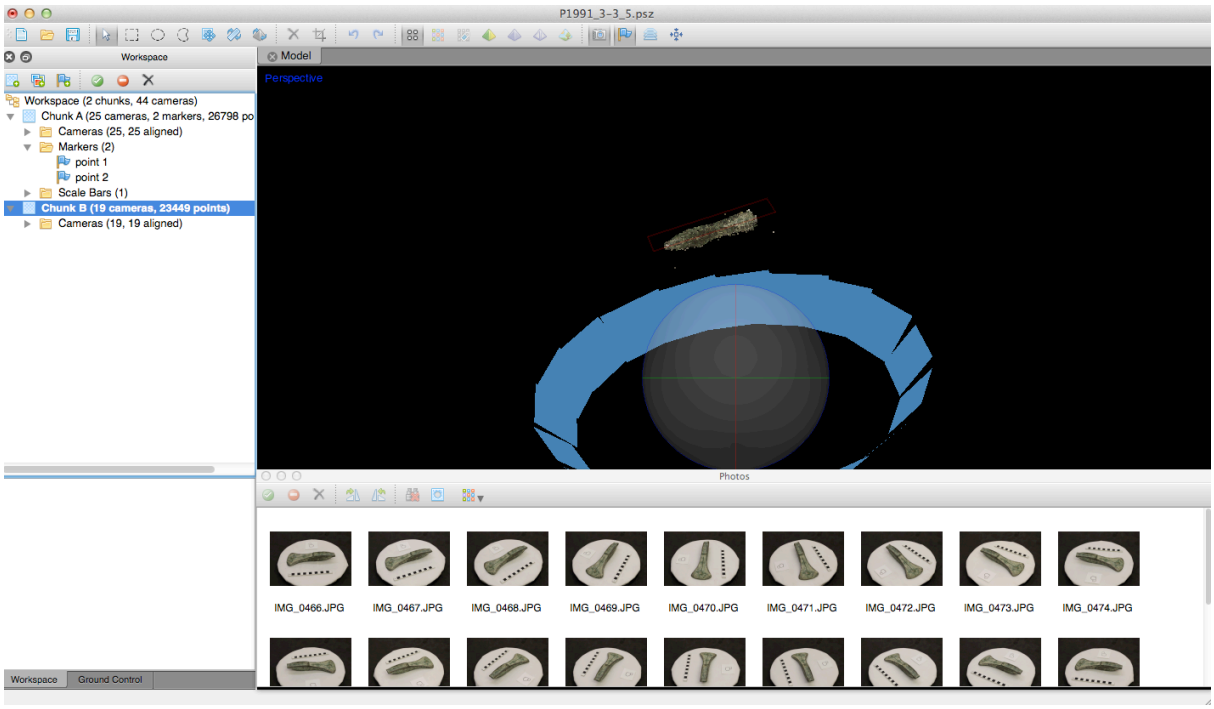

2.25 The model of Chunk A focuses on one side of the Bronze Age axe, whilst Chunk B reflects what was visible in the photographs once the axe was flipped over. We obviously want to merge these two into one model and then proceed to create a final consolidated version. There are a variety of ways of correctly aligning more than one chunk in PhotoScan, but here we will use a marker-based approach that works well for objects such as this. As with the construction of the scale bar earlier, it involves some manual ground control, but this time, we cannot use the scale bar in the photos since its position relative to the object has changed when the object was flipped over. We need to identify features on the object visually that we can hope to find in BOTH chunks. The most obvious areas showing in both chunks are the corners of the blade and features on the narrow sides of the object. Unfortunately, there is no easy way to go step-by-step through this selection of markers in the notes below, but you need to add a series of markers in Chunk A at points on the object that you can also see in Chunk B. Then you must add the corresponding points in Chunk B, making sure to label/rename each point *so it matches* the name of the equivalent point in Chunk A. As with the scale-bar, you need to find the location in at least two photographs in one chunk before the software will auto-suggest this locations in the rest of the photographs in that chunk (see 2.13-2.15). The best strategy on an object such as this axe is to find small shape changes (e.g. corners, divots etc.) or easily recognisable corrosion patches and use a landmark on those to place you markers. For example, here, you can see point 4 has been placed at the end of a distinctive dark patch on the seam of the axe's side.

Note that ALL points appear in all photographs in a chunk even if they are notionally on the other side of the object from what can be seen in a given photo (because the points are being fitted in 3D relative to the sparse point model, not just on the 2D photo). If you wish to adjust the markers that are being auto-suggested, you certainly can and should where there are possible errors, but only do so on photos where you have a good view of the location in question. Finding about 6-8 markers in both chunks leads to good results, for example with the points shown below:

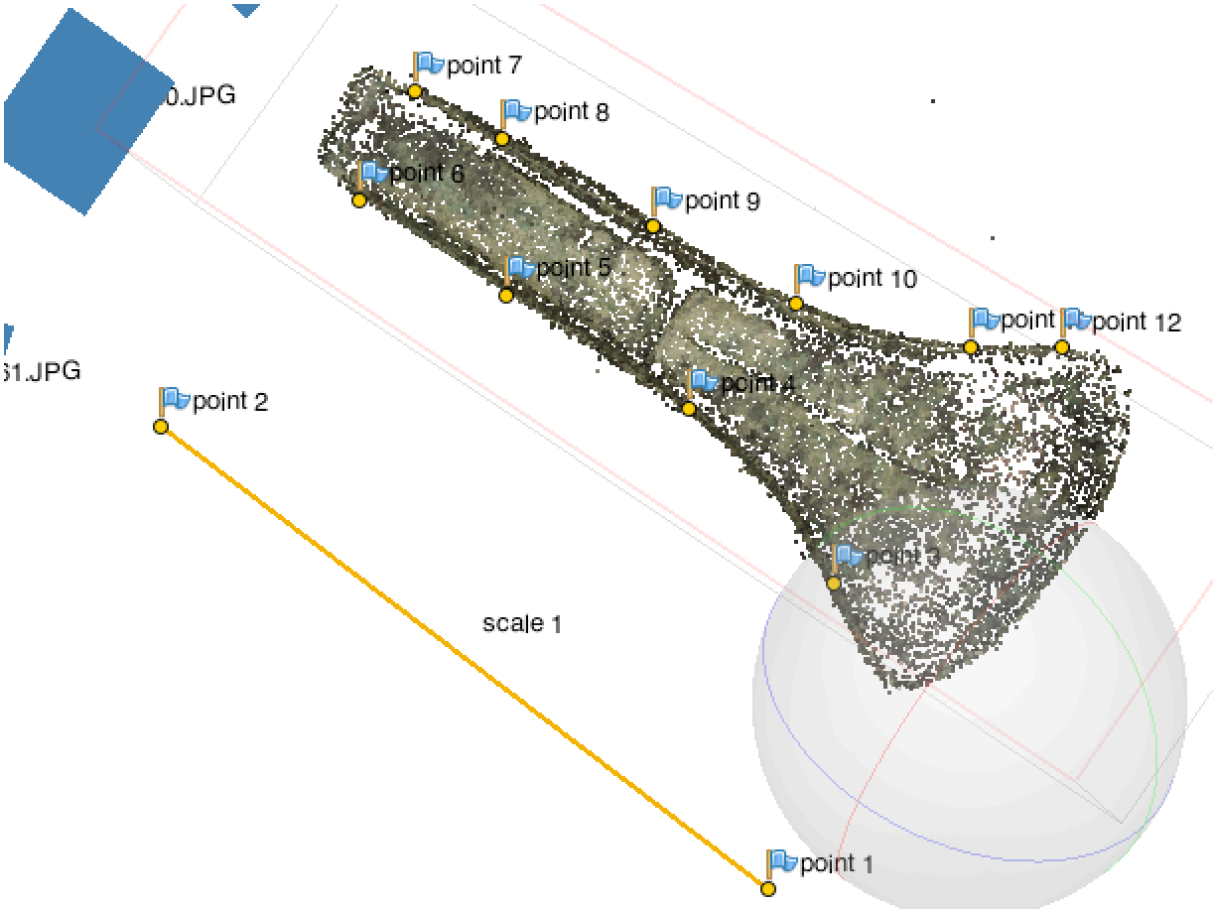

2.26 Marker choice for chunk alignment is perhaps the fiddliest part of the SfM process for object-based models, but a very important one. Once you have a set of markers you are happy with, then you can choose Workflow→Align Chunks. This brings up a dialogue and you want to make sure that Chunk A is shown in bold, as this is the one that you have already scaled correctly via the scale bar, so its dimensions should be treated as the dominant ones. Under Parameters, choose the 'Marker based' method, accept the other defaults (e.g. leave 'Fix scale' unticked) and click OK.

2.27 Your chunks should now be aligned with one another and can now be merged. Choose Workflow→Merge Chunks, and in the result dialogue tick to 'Combine models' and 'Merge markers'. Double click on the resulting merged model in the Workspace pane to make it active and your screen should look similar to the one below:

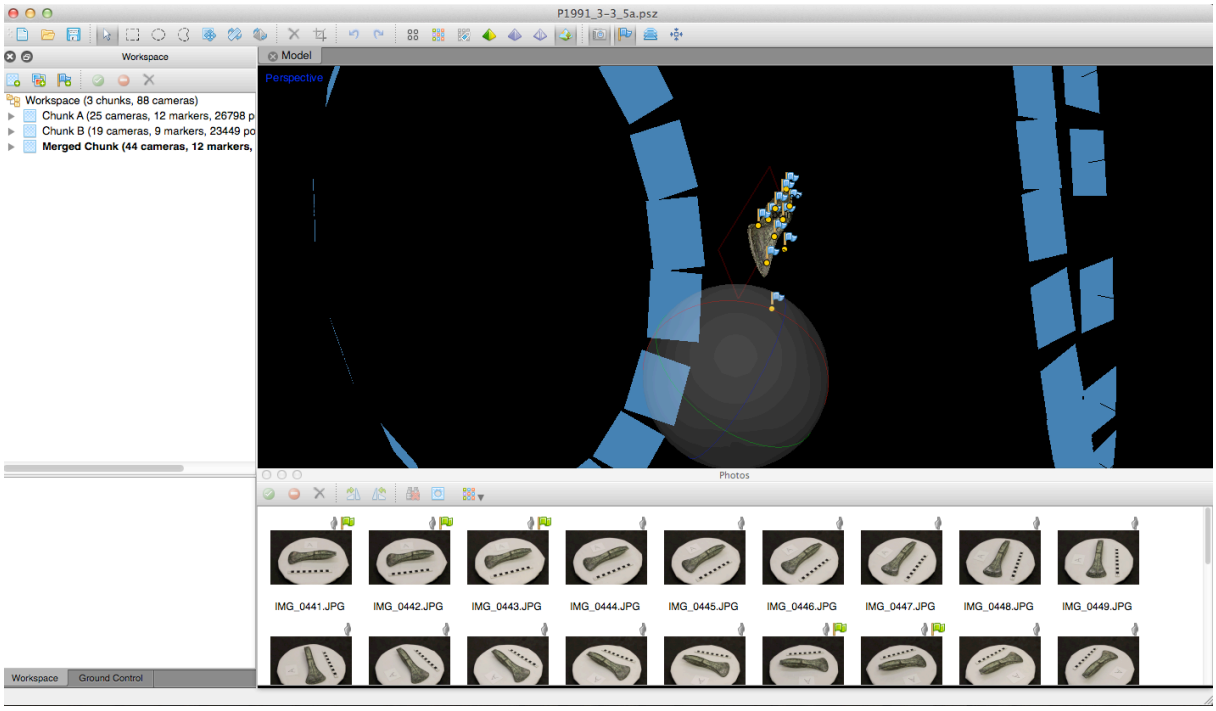

2.28 The sparse point cloud of the merged model has been used by the software to establish a default 'bounding box' around the area with meaningful data. The background of your Model pane may not be the same as the one shown above, but we can set it consistently by going to PhotoScan Pro→Preferences and setting the 'Background color to' grey (e.g. 'Silver' from the crayon colour picker). You should now be able to see the bounding bow similar to the one below:

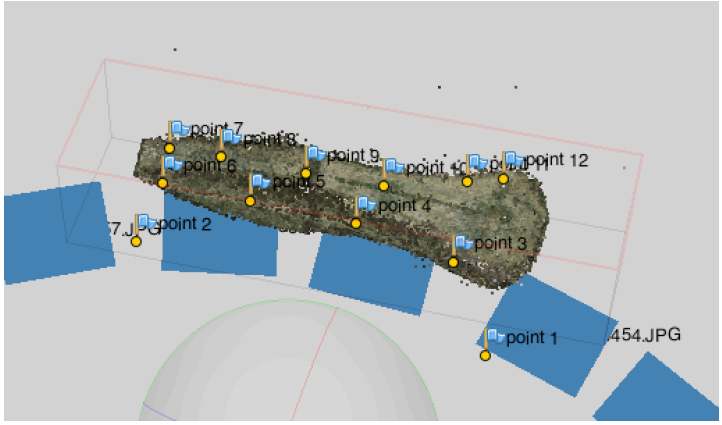

2.29 Sometimes this initial box is a bit too tight and risks excluding some small parts of the object from subsequent processing. This is the case here, for example, with the flared edges of the axe-blade. We can however slightly resize the bounding box by first orienting the model so it is roughly face on, then clicking on the Resize Region button and dragging the corners of the box a little to create a slightly larger box. Then click back on the navigation arrow to accept this change.

2.30 We now have a merged model and are ready to construct a dense point cloud. The nature of this step varies slightly form one SfM software implementation to another, but the goal is to return to the individual photos, and in the light of their reconstructed camera positions to revisit the identified features and produce a denser mapping of individual features. In PhotoScan, this involves creating individual 'depth maps' for each image. We do this by choosing Workflow→Build Dense Cloud. In the resulting dialogue, the 'Quality' parameter produces more or less dense point clouds, with High and Ultra High requiring significant amount of extra processing time (accept the default of Medium here for now). Under the 'Advanced' parameters, it is possible to re-use existing depth maps (e.g. if you have run this process before) and 'Depth filtering' controls whether the software smooths out perceived outlier points. A Mild setting (the one to go with in this example) means that small-scale changes in the surface topography of the model are kept, rather than smoothed out (so better where the surface object has real surface topographic detail to retain). Click OK. This process is often the most computationally-intensive of all the SfM steps, but when it finishes, you should not initially see any difference in the model, but by clicking on the 'Dense Cloud' button you should be able to see a denser modelled surface such as the one below:

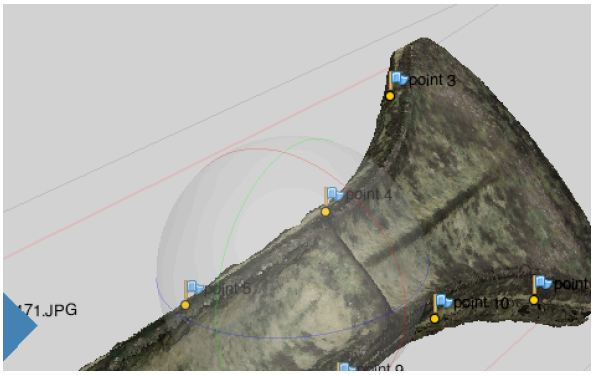

2.31 This dense point cloud has a colour value (RGB) per point, derived as a weighed average of values from the closest photographs. For certain purposes it is this dense point cloud that would be most useful as a final model output. If there are any stray points that should be deleted from the model, we can select these with the select tools  and then hit Delete.

2.32 We can run a script that will realign the bounding box so that it runs along the XYZ planes. This is not only useful as a starting point for comparing objects, but makes it much easier to pan, zoom etc. in most 3D viewers. The Python script for this is courtesy of Agisoft (also thanks to Lee Perry-Smith). To run the Python script, go to Tools→Run Scripts… and browse for coordinates_to_bounding_box_rotate.py in the resources accompanying this technical note. Then click OK. All of the points in your model have now been transformed so that the object is roughly flush with the XYZ planes and centred on the origin.

2.33 Once you are happy with a final model, then you could now save it, by choosing File→Save. To export it to a standard format used in other software and viewers, such as PLY, OBJ or STL, choose File→Export Points.

2.33 For many presentational purposes, however, a triangular mesh with accompanying photo-texture is a sharper, more attractive end result than a dense point cloud. The latter is created in two steps. First go to Workflow→Build Mesh… In the resulting dialogue, you have two options for 'Surface type': 'Height field' is appropriate when your model is of a landscape orientation because, in that case, the modeled topographic features should never wrap back around on themselves (i.e. they are what some people call only 2.5D). However, for most other purposes (as here), Arbitrary is more appropriate. For creating the mesh, you can use either the 'Sparse cloud' data or the 'Dense cloud' data (here we will choose the latter). Under 'Polygon count', increase the target number of polygons to 300,000. In the advanced settings, it is possible to determine the degree to which the mesh seeks to fill wholes in the point cloud or not (Interpolation) and which sub-types of points to include or exclude from the mesh generation (if they have been coded by type, which they have not in our case). Accept these defaults, and click OK.

2.34 The mesh model is then generated but does not appear instantly. You can view, with averaged colour on the mesh faces, with monochrome colour or as a wireframe only, via the following buttons  .

2.35 To create a photo texture, go to Workflow→Build Texture… In the resulting dialogue, there is a range of available parameters. Please refer to the PhotoScan manual for what these parameters each seek to do. For our purposes in building an object-scale model, we will accept the default options that create a medium resolution texture that takes a crisp texture mapping (rather than averaged colour value) from the nearest camera position. Click OK. After it finished processing you will not see the texture immediately, but can make it visible by clicking on the Textured button  . Your screen should look very roughly like the one below.

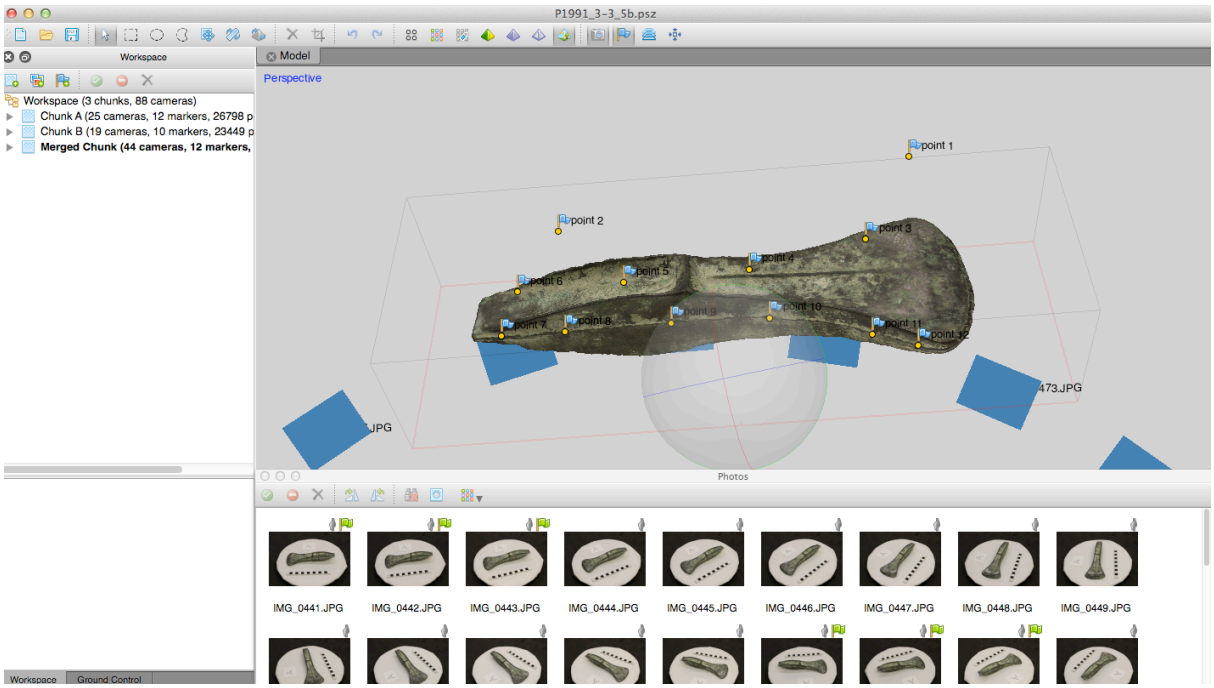

2.36 You can now save the project again, by choosing File→Save. You could also export the mesh and texture (e.g. in PLY or OBJ formats), by choosing File→Export Model.