

Introduction

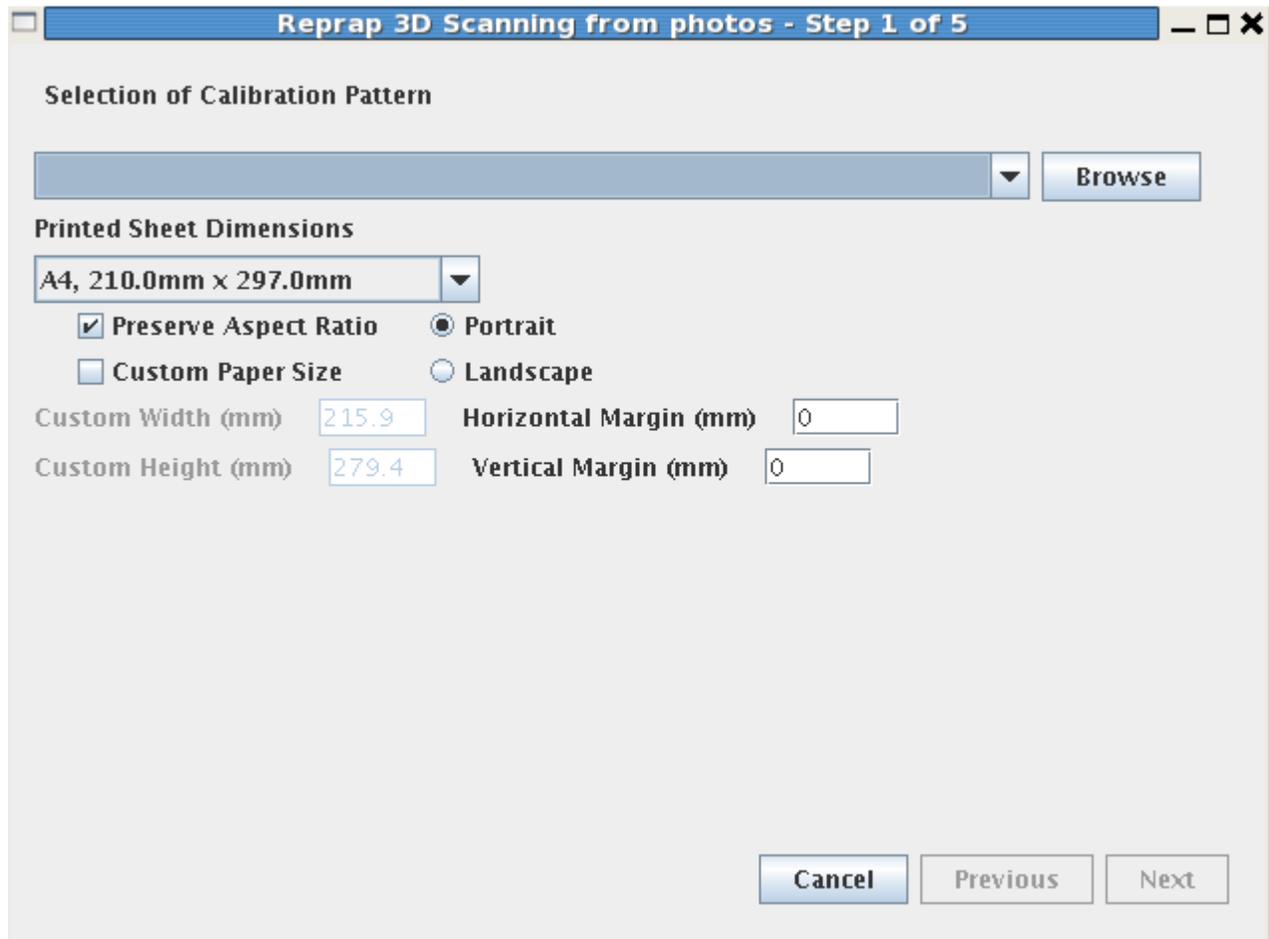
This tutorial will guide you through the process of turning digital photos into a basic 3d model using the example calibration sheet and some photos of a juggling ball on this calibration sheet that are packaged with the software.

The program is a 5 step wizard with the first two steps for user input and the last 3 being simply automated processes that are logically split into three stages.

Once you have used the program once it will remember all your settings so if you want to re-run the wizard with the same settings all you have to do is select an output file.

For more detailed information about how the program works, see the Software Technical Report, included with this file.

Step 1



Upon running the program for the first time the above screen will be displayed. At the top is a drop down list of previously selected calibration sheet patterns. As this is the first time running the program this list is empty so you need to click browse, navigate to the example calibration sheet and select it.

The paper size relates to the size of the paper the selected calibration sheet is printed out on in the images to be given to the computer. It is needed so that the computer has an accurate sense of scale and can relate pictures to real world measurements. Currently the only paper sizes loaded into this list by default are A4 or Letter. If you use a different size you can specify it as a Custom paper size. This list is to be expanded for the 1.0 release to include all the standard sizes.

As the margins are only a small fraction of the paper size they can normally be left at zero if the calibration sheet was printed out without user specified margins or unknown margins. If they are known the scale of the resulting object is better estimated but at the present time the accuracy of the model is not good enough to make this worth worrying about. The horizontal margin is the sum of the left and the right margins and the vertical is the sum of the top and bottom.

An inaccurate measurement of the margins will make the dimensions of the object slightly wrong but that can be fixed by importing into a 3d modeling tool and re-scaling afterwards if need be.

Note that in most cases, when printing out the calibration sheet, the aspect ratio will normally be preserved by default unless the user specifies otherwise but when it isn't the circles on the calibration sheet are printed out as ellipses. If the paper size and margins are accurate then the ellipse dimensions can be accurately estimated by the computer.

In this example the calibration sheet was printed out on an A4 sheet, with the aspect ratio preserved, and an unknown margin. This came about due to the basic image viewing program being used to print out the calibration sheet not having the option to set page margins. Therefore, if you wish truly accurate scaling for the object make sure to print out the calibration sheet using a program with margin setting options as part of the print process. This functionality is available in most image editing programs but does not seem to be common in most basic image viewing programs.

Click the Next button to go on to Step 2.

Step 2

The screenshot shows a software dialog box titled "Reprap 3D Scanning from photos - Step 2 of 5". The dialog is divided into several sections:

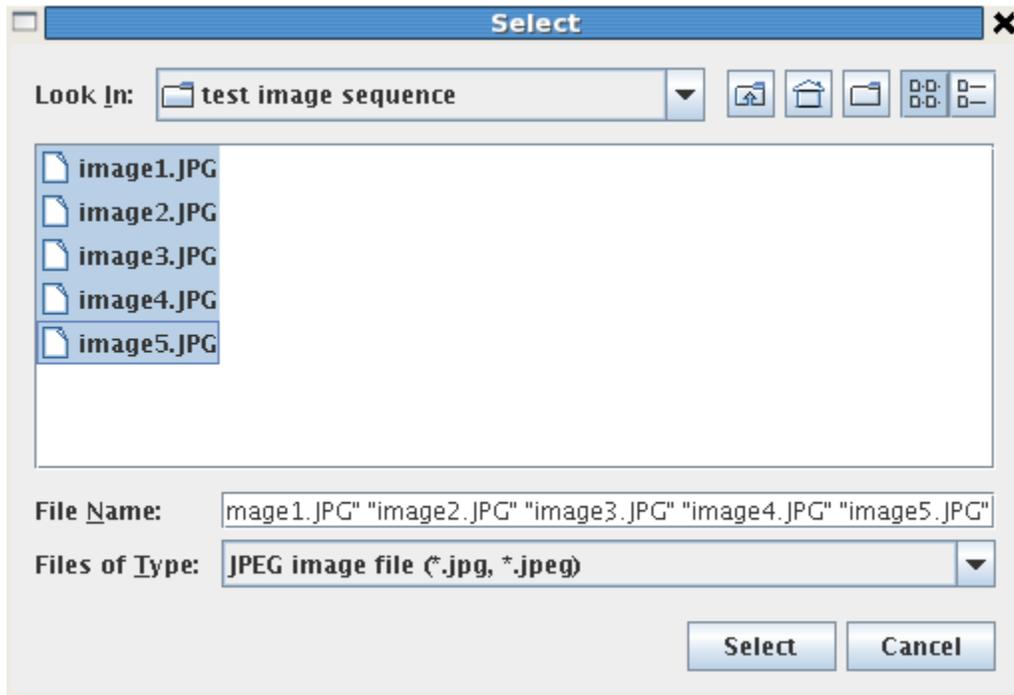
- Selection of Images:** A large empty rectangular area for displaying a list of images. To its right are two buttons: "Add" and "Remove".
- Selection of output file:** A text input field for specifying the output file name, followed by a "Browse" button.
- Optional Internal Object Name:** A text input field for providing an internal name for the object.
- Navigation:** At the bottom right, there are three buttons: "Cancel", "Previous", and "Next".

In this step the next button will not work until some images are selected and an output file has been chosen. In most cases you can leave the object name blank but it may be of use to help identify the object if you are planning on importing multiple objects into a 3d modeling tool that supports this identification.

For this example, click on the Add button and navigate to the folder containing the test image sequence and select all the files in this folder to add.

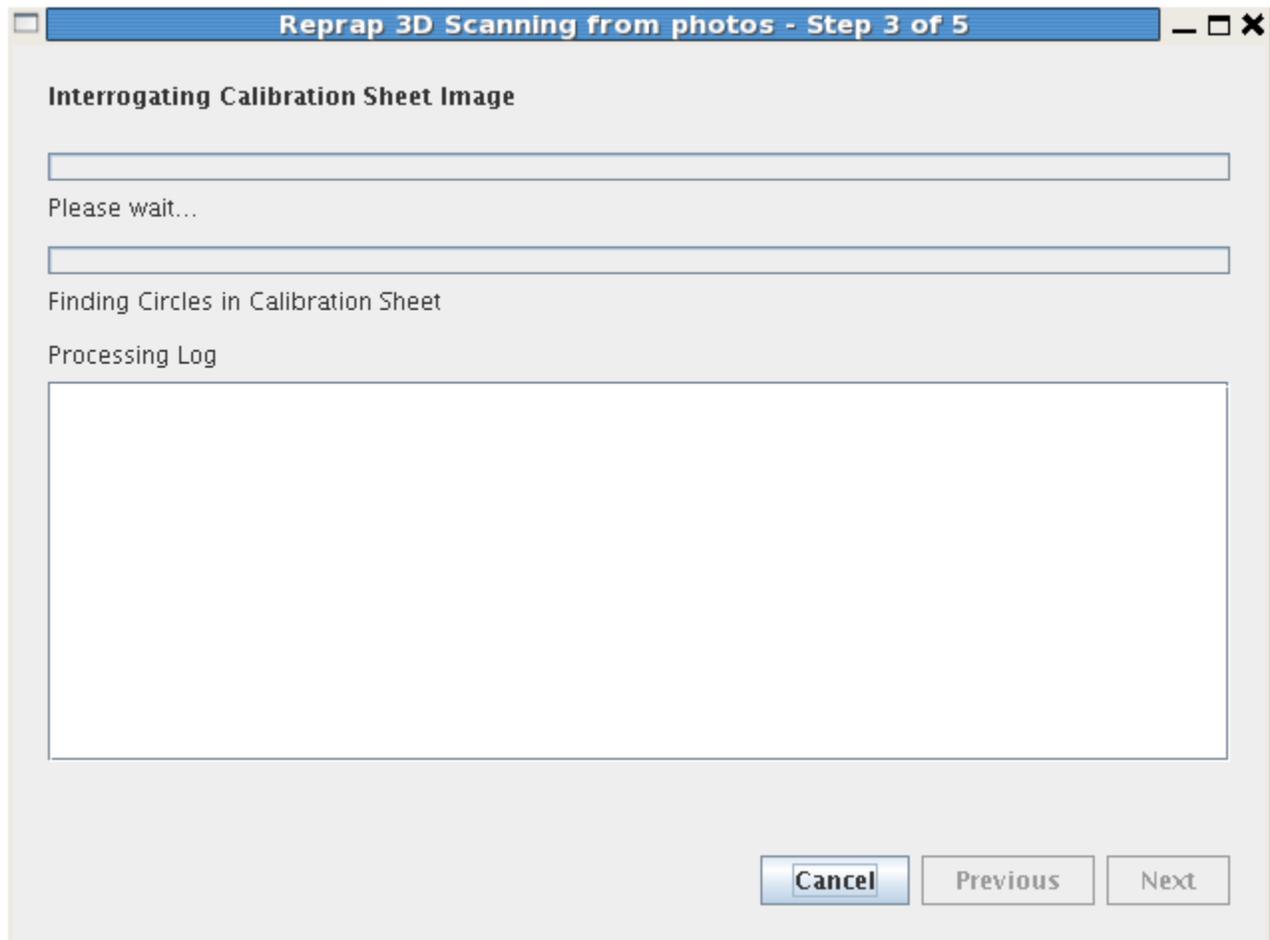
Once files are in the list, if you wish to remove them, select the filename and click the remove button.

Note that when selecting files in you can select multiple files in a list as in the example below. The exact specifics on how to do this in all situations vary between operating systems but one constant is that you can select a sequence of files by clicking on the top one, hold the Shift key down, and click on the bottom one.



When you have selected the images you want to process and also chosen an output file the Next button will be able to be clicked. Once you click the Next button the automated processes take over and all you need to do is wait until they are finished.

Step 3



This is the first of the automated stages. In this stage the computer is looking at the calibration sheet specified in the first step and trying to find the circles in it. Whilst it knows that all the circles will be the same size, it does not know what this size is. The progress bar will be quite sluggish to start with as it is trying to find the first circle but once it knows what size circles to look for it speeds up significantly.

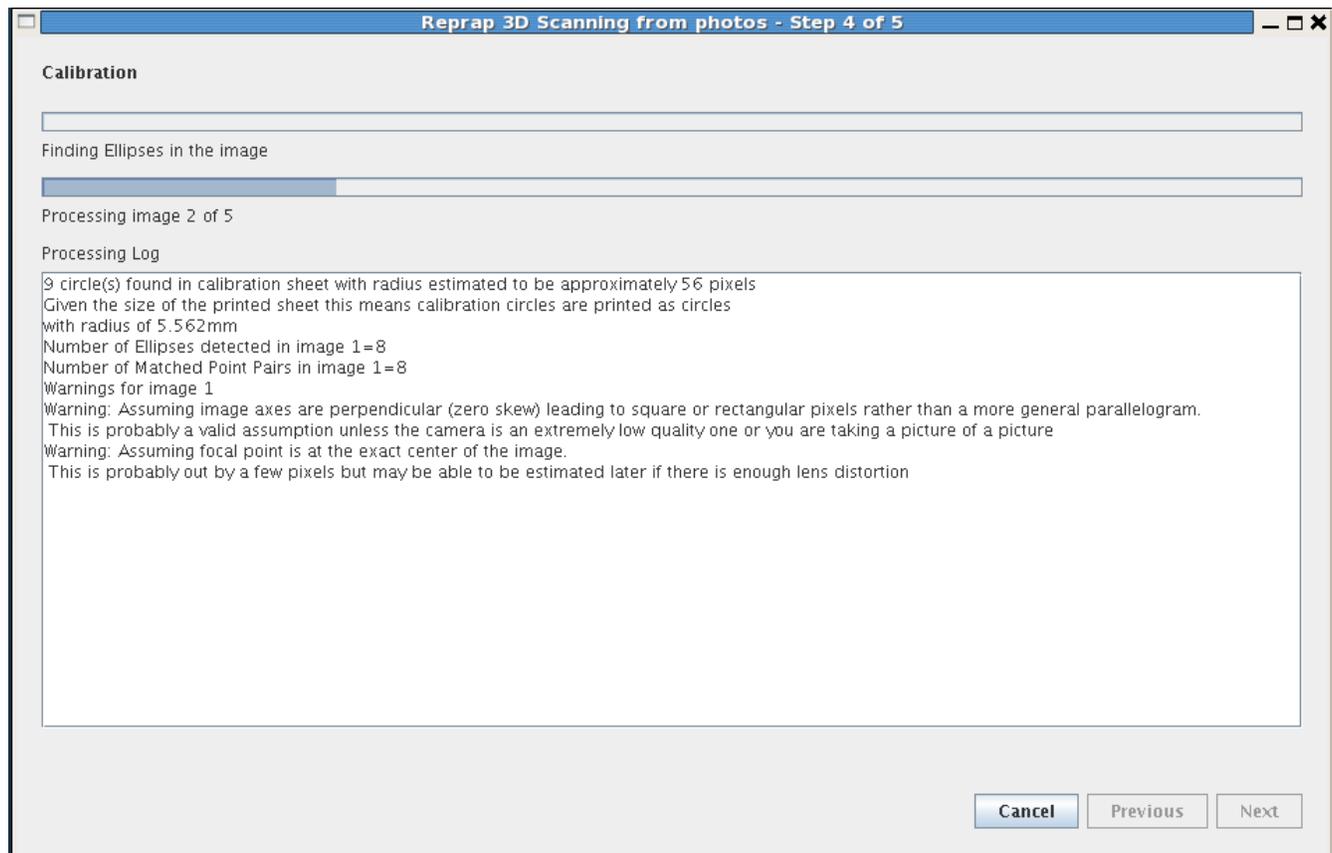
Once it has found the circles it takes the other information given in the first step to estimate how big the circles on the printed calibration sheet are. The log output will look similar to the below:

```
9 circle(s) found in calibration sheet with radius estimated to be approximately 56 pixels
Given the size of the printed sheet this means calibration circles are printed as circles
with radius of 5.562mm
```

Note that if the calibration sheet image was printed without keeping the aspect ratio this message will say that a calibration circle is printed as an ellipse and it will specify the size and orientation of the semi-axis lengths. Orientation is the direction of the major semi-axis length and is specified as being in the direction of either X or Y, i.e. left-right or up-down.

The program will then automatically proceed to Step 4.

Step 4

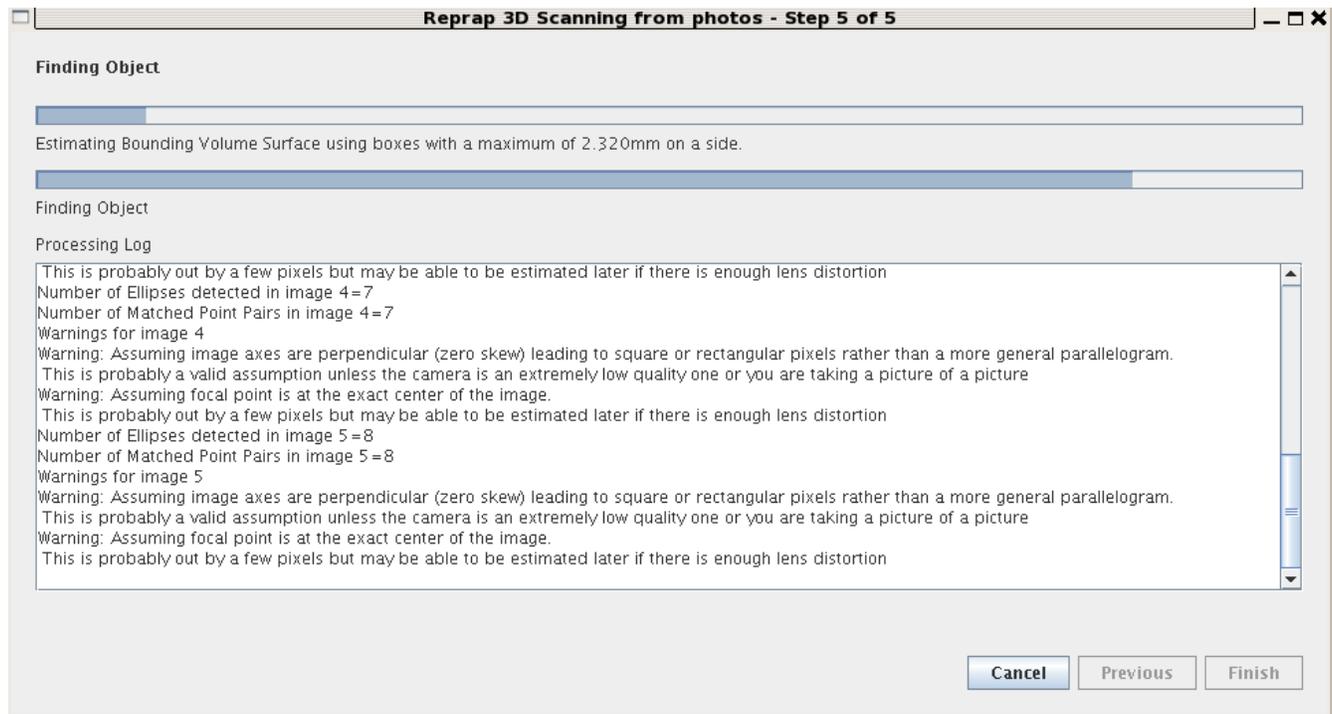


This is the processing of the images given in step 2. This processing has 5 sub-steps for each image:

- Find the candidate ellipses in the image that are potentially part of the calibration sheet viewed at an angle. This is similar to the finding of the circles in the calibration sheet but there is a known maximum and minimum ellipse size to limit the search based on the assumption that the whole calibration sheet is in the image.
- Match these ellipses up with the calibration sheet circles so that the camera parameters can be estimated. These camera parameters include the 3d position and orientation of the camera in space as well as estimating parameters to do with the structure of the camera: the size and shape of the pixels, the focal point and the amount of lens distortion. Given that we can't get all this information from one image some assumptions are initially made as can be seen by the warnings above. Note that currently the functionality to estimate the focal point when there is large distortion is currently untested and so commented out. Although this step is complicated it can be done very fast and in normal circumstances the progress bar will appear to flash past.
- Once these parameters have been estimated they are further adjusted to see if there is a better fit. This can take a while as the parameters are minutely adjusted and refined again and again.
- When the best estimate of lens distortion has been made, this distortion is corrected for.
- Finally the image is segmented into distinctly labeled parts so that the known calibration sheet is labeled with the unknown object able to be perceived as a silhouette against this calibration sheet.

Once all images have been processed the next step is automatically started.

Step 5



This is the final stage. It consists of 3 sub-steps:

- The 3d maximal outline of the object is estimated by dividing the volume of space above the calibration sheet into small boxes and categorizing them based on whether or not the calibration sheet can be seen in any image through them. If this is the case then the box is not part of the object, otherwise it is. Note that the maximum dimension of these boxes is displayed between the progress bars during this sub-step.
- The surface boxes are identified and the rectangular faces are subdivided into triangles. The log is updated to show the number of boxes and triangles used.
- These triangles are then written to a file in the industry standard STL format.

The volume of space that is of interest is based on the size of the calibration sheet. The distance to which the volume of interest extends above the calibration sheet is the larger of the width or height of the calibration sheet pattern. This volume of interest is subdivided into, by default, 128x128x128 sub-volume boxes, or voxels, to be processed.

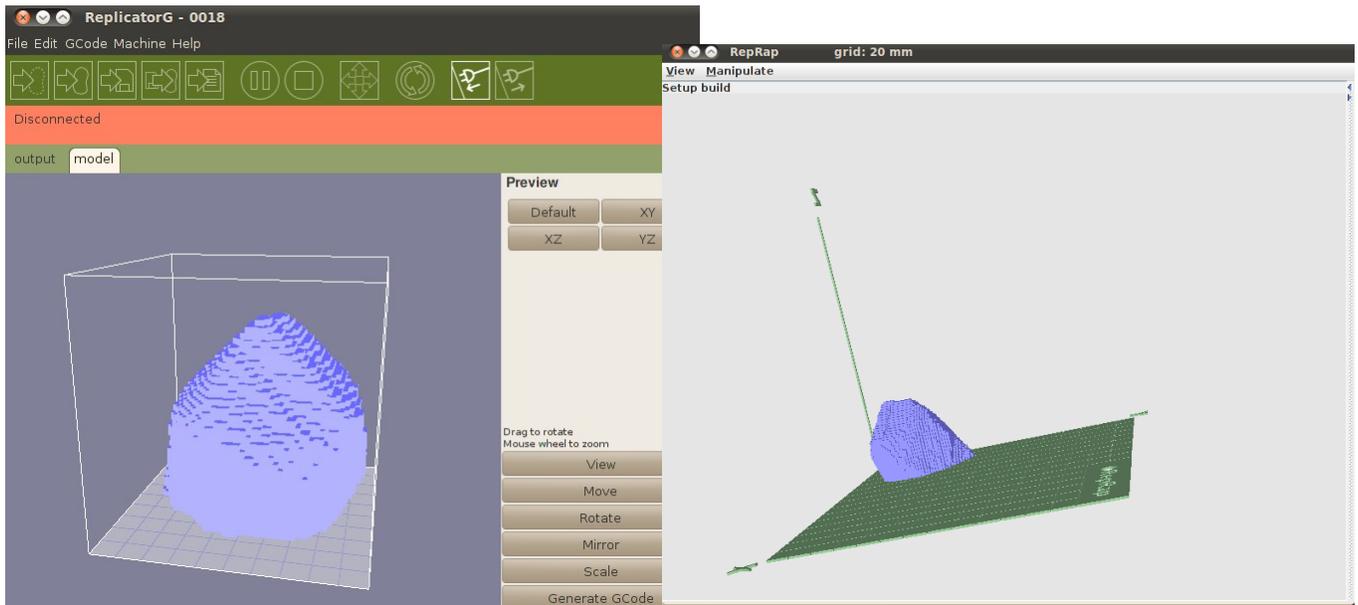
As each sub-step is finished it will put some basic information in the log as below.

```
Object bounding volume surface estimated using 8283 boxes.  
Object bounding volume surface estimated using 17136 triangles.  
Processing complete. Click Finish to exit.
```

When these steps are finished the Finish button will be available and can just be pressed to exit the program.

Conclusion

If you wish to view the output file you will need to have installed a separate program capable of reading STL files. Two examples are the RepRap host software, available from <http://sourceforge.net/projects/reprap/>, and ReplicatorG, available from <http://replicat.org/download>



Screen shots of the estimated object as viewed from these programs are above. Note that although the estimated object is vaguely ball like it has a few problems.

The overall scale is approximately correct as the juggling ball is in reality approximately 7cm in diameter and the highest point on the model is a little over 8cm above the surface but this is part of a spike artifact. This problem can be minimised by intelligently thinking about what photos to take as this artifact is caused by all the photos having been taken from approximately the same height above the object. Simply varying the height and angle from which photos are taken can minimize this sort of issue.

This is not the only problem however as there is currently an issue with the image segmentation routine which has some trouble with shadows being perceived as part of the object leading to odd protrusions being added to the object. It is intended that this be fixed before the 1.0 release but until then the best that can be done is to try and have as diffuse a light source as possible.

This maximal object outline is really just an intermediate step for further refinement in the future whereby it is hoped that additional information such as colour and texture of the object can be matched over multiple images to give a better and more fine-grained object estimation.