



From Scan to Print: 3D Printing as a Means for Replication

Susanne Klein, Michael Avery, Guy Adams, Stephen Pollard, Steven Simske

HP Laboratories
HPL-2014-30

Keyword(s):

3D printing; scanning; software; fidelity

Abstract:

Replication, or making exact copies with consistent results, is at the heart of manufacturing. It is used in mass production of all kinds of items, from foodstuff to cars, from houses to books. But it is also used to reproduce already existing objects. In the 18th and 19th centuries plaster casting was used to bring the wonders of the world to private collections and museums. In the cast court of the Victoria and Albert Museum in London, a life sized replica of Trajan's column [1] can be admired. The combination of a 3D scanner and printer offers the possibility of a new way to make a three dimensional copy of an existing object. Whereas in a plaster cast, where high fidelity is achieved by creating a physical mould from the original object, scanning does not require physical contact to the original. This can be an advantage when the object is fragile, but can lead to loss of fidelity during the reproduction process. We discuss the difficulties in achieving a truly high fidelity copy of even simple objects when a scanner and 3D printer are used for object replication.

External Posting Date: September 21, 2014 [Fulltext]
Internal Posting Date: September 21, 2014 [Fulltext]

Approved for External Publication

From Scan to Print: 3D Printing as a Means for Replication

Susanne Klein¹, Michael Avery², Guy Adams¹, Stephen Pollard¹, Steve Simske³; ¹HP Labs Bristol, UK, ²Bristol Centre for Functional Nanomaterials, University of Bristol, UK, ³HP Labs Fort Collins, USA

Abstract

Replication, or making exact copies with consistent results, is at the heart of manufacturing. It is used in mass production of all kinds of items, from foodstuff to cars, from houses to books. But it is also used to reproduce already existing objects. In the 18th and 19th centuries plaster casting was used to bring the wonders of the world to private collections and museums. In the cast court of the Victoria and Albert Museum in London, a life sized replica of Trajan's column [1] can be admired. The combination of a 3D scanner and printer offers the possibility of a new way to make a three dimensional copy of an existing object. Whereas in a plaster cast, where high fidelity is achieved by creating a physical mould from the original object, scanning does not require physical contact to the original. This can be an advantage when the object is fragile, but can lead to loss of fidelity during the reproduction process. We discuss the difficulties in achieving a truly high fidelity copy of even simple objects when a scanner and 3D printer are used for object replication.

Introduction

3D printing is seen as a threat to copyright. The doomsday scenario predicts the market being swamped by pirated copies of all kinds of commodities. It is certainly true that 3D printing offers the opportunity to generate objects without the need for intricate and expensive tooling. But, just how good are the current generation of moderately priced 3D scanners and printers at faithfully reproducing existing 3D objects? In the experiments we report here we use the combination of a DAVID Structured Light Scanner (SLS-2) and associated software [2] and an HP DesignJet 3D printer manufactured by Stratasys [3] as a proof point of the state of the art performance.

Typically when a 3D scanner is used to capture a 3D object, a calibrated combination of light source and high resolution camera are used to triangulate points on the object's surface (a few systems use time of flight imaging devices to measure 3D data directly but these tend to be noisy and low resolution). The light source can be either a laser, as used in LiDAR (for example see [4] for an application), or a structured white light source, as used in the SLS-2 Scanner [2], which is faster as multiple stripes, at varying spatial frequencies, can be projected together. High fidelity reproduction of a complete 3D object requires the acquisition and combination of multiple high fidelity digital sets, each recorded with the object oriented in different positions to adequately cover the viewing sphere of the object. Each subset starts life as a point cloud, i.e. points in an X,Y,Z coordinate system representing the external surface of the scanned object. In the case of the DAVID software used here the geometry can only be viewed as a triangular mesh or output as such in the Alias Wavefront OBJ file format [5] this helps reduce the overall data size with the coarseness of the mesh selectable by the user. The separate views have to be stitched

together to create a more complete surface description of the 3D object. While useful for visualisation (including rendering with image texture extracted from the real object) and shape manipulation, the combined surface description cannot yet be printed on a 3D printer. It has to be converted into a solid watertight 3D mesh model. The density of mesh vertices defines the upper limit on the resolution of the printed object. Too many vertices make the OBJ files large and cumbersome; too few blur detail. In fact, the achievable resolution is limited by the physical printing process. Details smaller than the bead size, for extrusion printing, or the particle size, for powder bed printing, cannot be reproduced. So, in practice, there is little point using a mesh density beyond the resolution of the print.

Step 1: Preparation of the original object for scanning

Almost all objects require some form of modification before they can be scanned. Figure 1 shows examples of two simple objects that required modification for scanning.

Problem 1: Moving parts

The bell of the cow in Figure 1(a) is a moving part. It had to be immobilized to avoid different locations of the bell during re-orientation of the cow for scanning. To avoid damage to the object, we did not want to glue the bell into place. Instead, we used a reusable putty to fix it in place. The putty had to be carefully repositioned during the scanning process to avoid including it in the scanned data.

Problem 2: Holes

The bell in Figure 1(a) is hollow and deep enough to cause problems during scanning and data manipulation. Concavities of this kind require too many additional viewpoints to obtain complete surface descriptions. The software that comes with the scanner will close the 'hole' during fusion of the different stitched mesh files, i.e. make a solid object out of a hollow bell.

Problem 3: Specular reflections

Interestingly a specular reflection creates the same problem as a hole: missing data points. The ring in Figure 1(b) could not be scanned successfully without being sprayed with a matte chalk layer. This is standard practice for shiny or transparent objects and recommended by scanner manufacturers, but can damage the original object since chalk residues stay in indentations. The solvent in the chalk spray can damage the surface of a more fragile object.

Problem 4: High Contrast

The dark collar in Figure 1(a) caused the cow's head to become detached from its body. The gap in the data between main body and head was large enough for the software to define the

head and body as two different objects. Stitching and fusion led to a headless cow with a bodiless head.



Figure 1: a) wooden toy cow, the white haze is chalk, used to improve the scan quality, which did not wash off again; b) skull ring.

Step 2: Scanning

Accessing the whole surface of a three dimensional object can be done in three ways. Firstly, the object is stationary and the scanner is moved around the object. Secondly, the scanner is stationary and the object is rotated. Thirdly, the object is



Figure 2: set up of scanner and stage

surrounded by multiple stationary scanners. In our case, we use a stationary scanner and a rotating object as shown in Figure 2. To rotate the object in a controlled way, the object was mounted on a rotation stage. The scanner needed to be calibrated as a function of the size of the object, i.e. when the dimensions of the object changed by one or more orders of magnitude, the scanner had to be recalibrated. By covering the stage and the background with black velvet we reduced reflections not belonging to the scanned object and minimized scanning artefacts.

The stitching software supplied with the scanner could not handle scans more than 60 degrees apart which posed a problem for the underside or the top of the object. Laying the object on its side to scan the underside, for example the belly of the crocodile, means rotating the object by approximately 90 degrees. Even though this procedure is recommended by the manufacturer, we often had to prop the object up at a shallower angle to stitch the underside to the rest of the scans successfully. This may well be a

specific problem of that particular software but the collection of data from the entire surface is not trivial and is often not achieved. For example during scanning for a lifelike figurine the underside of the feet is not scanned. It is not perceived as a fault since a complete replica is not expected, but it is still an incomplete data set.

Step 3: Stitching and converting

On average 8 to 16 scans have to be combined to create a complete 3D object. The software for the scanner completes this process in 2 stages. First all the viewpoints are stitched together transforming all the mesh data into a single coordinate frame. If there are too few scans or the visual angle between them is too large the mesh data will not be successfully stitched. The data is then fused and the various overlapping surfaces that come from each viewpoint combined to create a single object surface and a solid model that can be printed. This approach has 3 limitations: (1) the fusion can fail resulting in multiple ghost surfaces, especially if lots of dense viewpoints are combined that include small calibration errors; (2) extraneous data, which comes from parts of the setup which do not belong to the object, needs to be cropped by hand from the individual scans before stitching; and (3) to (re)create hollow objects the individual scans need to be manipulated by hand using external software packages (e.g. we used Mathematica to transform the individual surfaces into 3D layers and combined them to generate a hollow object; a process that is both time-consuming and fiddly). While these problems are specific to the software we used, they are also indicative of the results obtained using more expensive software packages that we have also tested. Counterintuitively we found that fused files with a higher density mesh data tend to generate more uneven surfaces than those with lower mesh densities when the surface of the original is smooth, as seen in Figure 3. This suggests more sophisticated data fusion is needed to combine the data and deal with inevitable calibration errors

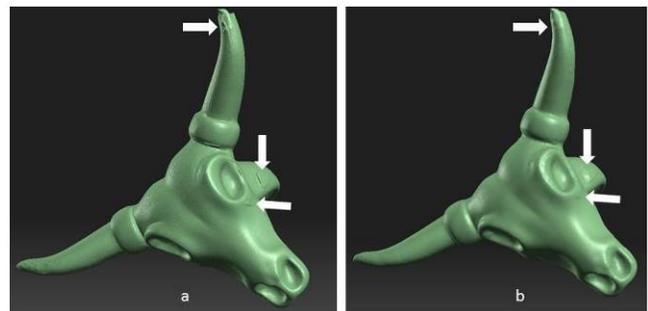


Figure 3: a) scans fused with a vertex spacing of 0.0251mm which led to a file size of 110 MB; b) scans fused with a vertex spacing of 0.401mm which led to a file size of 3.07 MB. The arrows point at artefacts in the high resolution file which partly disappear in the low resolution file.

After the scans have been fused together into one three-dimensional mesh file, the file has to be converted into a format that can be read by the printer. We chose the Standard Tessellation Language (STL).

Step 4: Printing

Before the STL file can be printed, it has to be checked for misoriented surfaces, self-intersecting or open curves, which lead to missing surfaces. If the file is 'watertight', i.e. the surface is complete without holes or inverted triangles, then the object is ready to be sliced. At this point, the final dimensions of the printout can be selected. The object can be scaled arbitrarily to fit the print volume. We are interested here in printing small objects at actual size so a scaling factor of 1 is employed. The slicing software then accounts for the printer resolution and divides the file into printable layers. This defines the resolution of the final print. If the printer resolution is higher than the resolution in the sliced file, all detail in the slice will be reproduced but the opportunity to print more detail will have been lost. If the resolution is lower, only coarser features will appear in the print.

Example: Toy crocodile

To avoid all issues highlighted in the previous sections, we chose a plastic toy crocodile as the object we would like to replicate. It has neither moving nor metallic parts and is made out of a single material. It is reasonably low contrast, so that while the eyes and teeth are white, black is only used in tiny specks. The object has an overall length of 152 mm and a maximal width of 66.55 mm. The smallest details, i.e. the iris of the crocodile's eye and the small teeth, are roughly 0.5 mm in their smallest dimension (see Figure 4).

The crocodile was put on the rotation stage as shown in figure 2. The projector and camera were angled 20 degrees downwards.



Figure 4: Plastic toy crocodile, the original.

We took scans every 30 degrees, rotating the object about its vertical axis. We took another 3 scans of the bottom of the object and one scan from the top since the back of the crocodile includes artefacts after stitching without the top scan. Figure 5 shows all the

scans used by the software that comes with the scanner to produce the 3-dimensional representations in Figure 6.

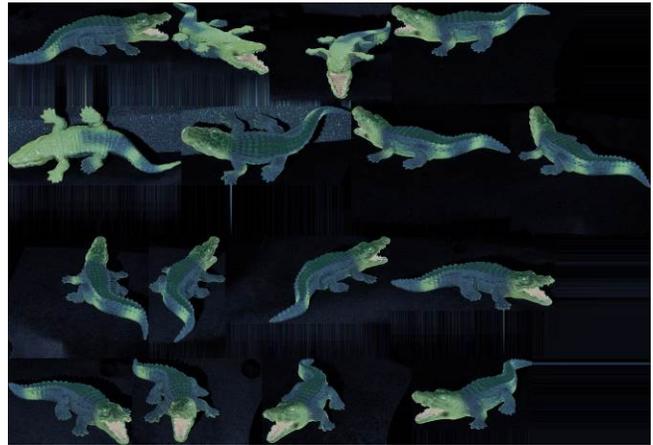


Figure 5: Collection of images representing the orientations of the scans used to generate the 3 dimensional files shown in figure 6.

As discussed above, the file with the higher density mesh tends to have more obvious artefacts (the artefacts in low density meshes are smoothed out and less easy to see – they may still be geometrically incorrect but at least they are smooth) as shown in Figure 6. However there is almost no difference in the resolution of the final print (see Figure 8) as both the low and high density meshes produce very similar prints on the HP DesignJet 3D printer. Figure 7 shows the ready to print file including the support material.

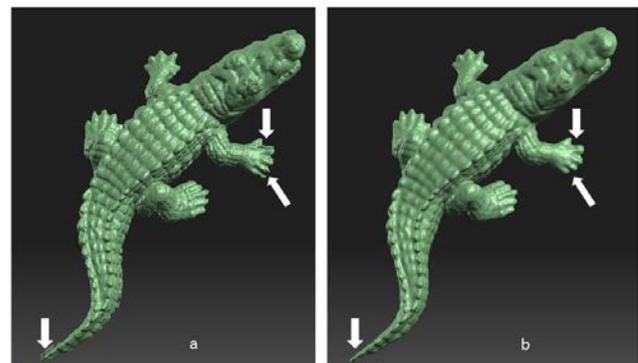


Figure 6: The scan in a) contains four times as many vertices as the scan in b). Artefacts are pointed out by arrows. The detail in b) blurred compared to a).

Figure 8 shows the 3D prints. It is difficult to tell which print comes from which file. In both files the detail was finer than the printer resolution; therefore the detail reproduced in the two prints is ostensibly the same. The dimensions of the two prints are the same, but not identical to the original. The overall length of the original is 152 mm. The overall length of the print is 150 mm. The mouth opening of the original is 20.20 mm, that of the print 19.40 mm. The widest part of the torso of the original is 27.80 mm, that of the print 28.15 mm. The change in the proportions of the crocodile's body suggest that there is either a calibration issue or

errors have been introduced during the stitching and fusion of the completed 3D model.

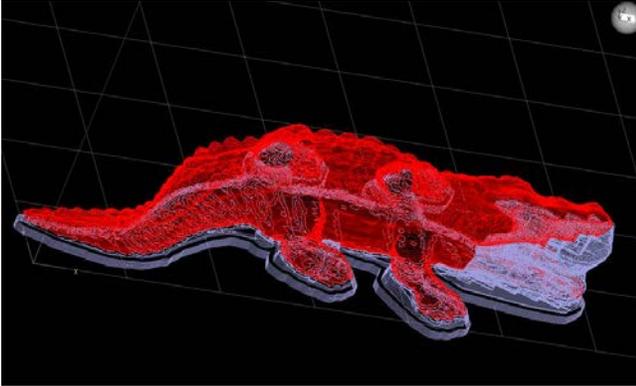


Figure 7: Ready to print files. The file has been sliced taking the resolution of the printer into account, the tool path has been calculated and support material, in light purple, has been added.

To make the prints more similar to the original, their surfaces need to be post processed. From figure 8 it is clear that the surface roughness of the print is much higher than that of the original. The bead of the machine is very visible. Stratasys (who manufactured the HP DesignJet 3D) offers a smoothing station where the print is exposed to solvent vapours and then dried [6]. Acetone vapour is widely used in the hobbyist community. Both professional and hobbyist methods require the part to be exposed to the vapour and checked several times before an acceptable surface finish is reached. Overexposure to the solvent vapour leads to loss, or melting, of finer detail.

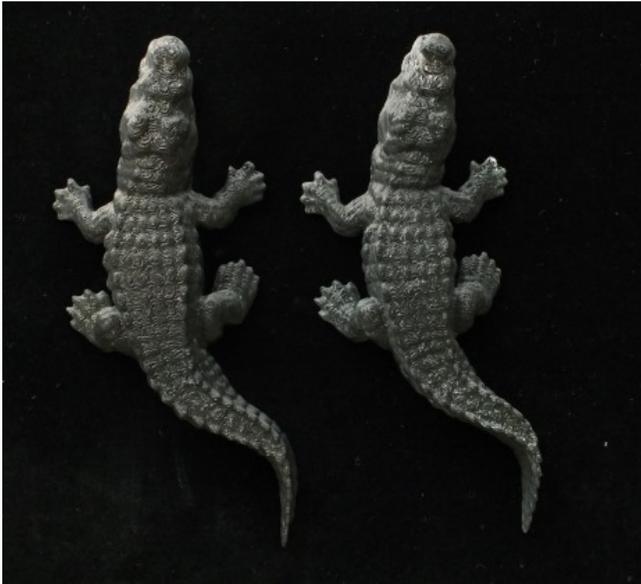


Figure 8: Prints of the files in figure 6

After the surface has been smoothed, the crocodile needs to be painted. The original was hand-painted as stated by the

manufacturer (see [7]). Painting by hand is probably the most cost and time effective method to generate a replica.

Post processing is one of the most important steps to generate an acceptable replica. Depending on the complexity of the surface colouring and structure it can be extremely time-consuming, especially when a complicated surface finish is desired, for example turning a plastic into fake marble.

Fidelity

Quality control is at the heart of manufacturing. One way to assure quality is the so called 'First Article Inspection' [8]. A sample is tested against agreed specifications, for example drawings, stiffness, colour, reflectance, surface finish, dimensions, etc.

Several difficulties are encountered when replicating an object via scanning and 3D printing.

Without the original design file, it is not possible to assess the accuracy of the digital file generated by scanning. Deviations can be generated by a defective original, by the scanning process and by the stitching of the different files into one 3D point cloud.

When the original is next to the printed replica, a physical inspection is possible. Stiffness, colour, reflectance, surface finish, density and dimensions can be measured and compared.

We did not post process the crocodile in the example discussed earlier. But measurements of the bigger dimensions showed deviations from the original. Even if we wanted we could not correct this mismatch. The length of the print is smaller than that of the original but the width is bigger, i.e. it cannot be corrected by simple rescaling.

Conclusions

Scanning and printing allow reproducing an object similar in dimensions without tool making. We could make a model of the original crocodile within 8 hours. Scanning takes about 1 hour, generating a STL file about 3 hours and printing again 3.5 to 4 hours. With a more efficient setup and more efficient software scanning and generating the printing file will be faster. We estimate that post processing, i.e. smoothing of the surface and painting will take another 2 hours. Here of course the appropriate tools are necessary. At the end we will have a similar object to but not really a replica of the original. Weight and dimensions will be similar but not identical, since the densities of the materials of print and original differ. Even without the original next to it, the print could be identified as print since it differs in material stiffness from the original. Unless the same materials can be used during printing, this will be always a major problem. Without expansion of the materials portfolio for 3D printing, it will be very hard to make pleasing replicas.

References

- [1] <http://www.vam.ac.uk/content/articles/t/trajans-column/>
- [2] http://shop.david-vision-systems.de/product_info.php/info/p124_DAVID-Structured-Light-Scanner.html
- [3] <http://www.stratasys.com/3d-printers/production-series/fortus-250mc>
- [4] <http://www.auto-map.co.uk/overview/data-collection/>
- [5] <http://www.fileformat.info/format/wavefrontobj/egff.htm>
- [6] <http://www.stratasys.com/3d-printers/finishing-touch-smoothing-station>

- [7] <http://www.schleich-s.com/en/GB/consumer-service/faqs/>
- [8] Pennella, C. Robert, *Managing Contract Quality Requirements*, American Society for Quality, pp. 48–49 (2006).

Author Biography

Susanne Klein holds a Diploma in Physics and a PhD in Medical Physics from the University of Saarland. Since 1995 she has worked and lived in the UK, first as a Royal Society Research Assistant at the University of Bristol and then as a member of HP Labs. Her main research interests lie in new materials for displays and 3D printing