Chapter 8

Results of the volumetric method

In this chapter, we describe the hardware used to acquire the range data and how we treat the range scanner's lines of sight (Section 8.1). In Section 8.2, we explain our method for addressing the problem of aligning range images. Finally, we demonstrate the effectiveness of the volumetric method on several models (Section 8.3).

8.1 Hardware Implementation

The data used in this chapter was acquired using the Cyberware 3030 MS laser stripe optical triangulation scanner, as described in Chapter 4. We collected data both using the Cyberware's traditional triangulation method and using the spacetime analysis (by way of the JPEG compressor board as shown in Figure 4.1).

When using traditional triangulation analysis implemented in hardware in our Cyberware scanner, the uncertainty in triangulation for our system follows the lines of sight of the expanding laser beam as described in Section 5.2. When using the spacetime analysis, however, the uncertainty roughly follows the lines of sight of the CCD camera. For each method of triangulation analysis, we adhere to the appropriate lines of sight when laying down signed distance and weight functions. The space carving, however, only follows the lines of sight of the laser. Owing to the geometry of the Cyberware scanner, the lines of sight of the laser are significantly simpler to follow than those of the CCD camera. The decision to carve from the point of view of the laser is not a fundamental limitation of the algorithm, merely a matter a convenience for the implementation. To maximize the amount of possible space carving, it remains an area for future work to follow the lines of sight of the CCD camera.

8.2 Aligning range images

As described in Chapter 5, proceeding with pairwise alignment while walking around an object may yield good results at each stage, but the errors will accumulate. The general solution is to minimize the alignment errors among all of the range images simultaneously. Only recently are solutions to this total alignment problem emerging [Bergevin et al. 1996] [Stoddart & Hilton 1996].

For the results obtained in this chapter, we devised and implemented a practical solution that utilizes information from the scanner's motion control system. The Cyberware scanner described in the previous section has a rotational degree of freedom about the vertical axis. By calibrating the position and orientation of this rotational axis, we can perform a set of scans, with a specified rotation between each scan, and begin with good alignment information for this set of scans. After merging these scans, the resulting isosurface will usually have good coverage of the object so that it may be used as an anchor for future alignments (using the method of [Turk & Levoy 1994]). As scans are aligned and merged, new isosurfaces can be extracted to serve as better anchors - better in the sense of more surface coverage and in the sense of improved surface estimation. Usually, two or more additional scans are immediately required after the initial set of rotations to capture the top and bottom of the object. If these scans are well chosen, they will overlap the existing isosurface well enough to perform pairwise alignment while providing enough new data to offer an anchor for future alignments against the top and bottom of the object.

We do not claim that this practical procedure is optimal. In fact, the alignment results will depend on the order in which scans are merged. Nevertheless, the measurable quality of the results indicate that the procedure is viable and accurate for our scanning system.

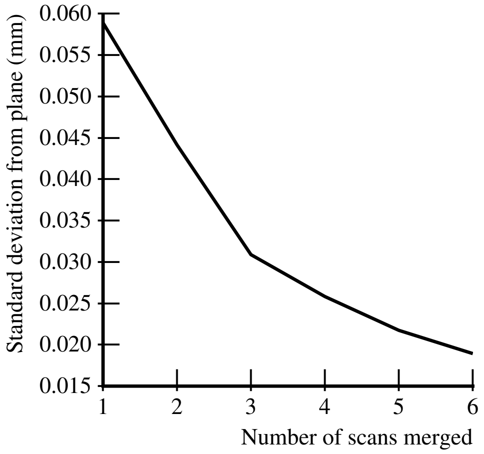

Figure 8.1: Noise reduction by merging multiple scans. A planar target was scanned 6 times with a 15^o rotation of the viewpoint after each scan. Standard deviation from planarity for the reconstruction was determined after merging each scan volumetrically. Note how each scan improves the estimate of the shape of the target.

8.3 Results

We show results for a number of objects designed to explore noise reduction, robustness of our algorithm, its ability to fill gaps in the reconstruction, and its attainable level of detail.

To explore the noise reduction properties of our algorithm, we scanned a planar target from 6 different viewpoints. After merging each scan into the volume, we extracted an isosurface, fit a plane through this surface, and computed the standard deviation of the reconstructed vertices from the planar fit. Figure 8.1 shows the results of this procedure. Clearly, after each scan is merged, the planar fit improves, indicating a reduction in noise.

To explore robustness, we scanned a thin drill bit (about the thickness of the laser sheet) from 12 orientations at 30 degree spacings using the traditional method of optical triangulation. Due to the false edge extensions inherent in data from triangulation scanners using traditional analysis (see Figure 4.11b), this particular object poses a formidable challenge. The rendering of the drill bit reconstructed by zippering range images [Turk & Levoy 1994] shows the catastrophic outcome typical of using a polygon-based approach. The model generated by the volumetric method, on the other hand, is without holes and preserves some of the helical structure of the original object, even though this structure is near the resolution limit of the scanner.

To demonstrate the effectiveness of carving space for hole filling, we show the stages of our algorithm when reconstructing the dragon model in Figures 8.3 and 8.4. In the first stage, we confine the signed distance ramps and weight functions to the vicinity of the observed surface. The result is a reconstruction with holes where we failed to see the surface. In the second stage, we follow the laser lines of sight away from the observed surface and "empty out" much of the space around the dragon. The resulting isosurface indicates some success in filling holes, but a number of undesirable protrusions have been generated. These excess surfaces are the result of insufficient carving of space. In the third stage, we follow up with some additional scans using a backdrop to carve through sections of the model. This results in the removal of the extraneous regions, but the hole fill regions are still jagged due to aliasing in the transitions between unseen and empty. In the final stage, we filter the mesh to smooth the hole fill regions. This filtering applies only to the hole fill regions while preserving the detail in the observed portions of the surface.

The "Happy Buddha" in Figures 8.5 and 8.6 shows that our method can be used to generate very detailed, hole-free models. From a set of 48 range images, we generated the model shown in Figure 8.6a. This model has gaps which we were able to fill with space carving and an additional 10 scans against a backdrop. The final model contains over 2.5 million triangles, which we decimate to 800,000 triangles without appreciable loss of detail using the method of [Schroeder & Lorensen 1992]. The rendering in Figure 8.6b, made using RenderMan and the decimated mesh, incorporates a number of rendering algorithms including accessibility shading and simple simulation of layered surface scattering. Because the computer model is free of holes and self-intersecting surfaces, we successfully manufactured a hardcopy using stereolithography (Figure 8.6c). For a more detailed description of the stereolithography process, refer to Appendix B. Figure 8.7 shows both wireframe and shaded renderings of a single range surface, the reconstructed surface, and the decimated mesh. The reconstruction from multiple range surfaces is less noisy and more detailed than a single range surface. The number of triangles generated, however, is more than necessary. In the figure, a 3:1 decimation approximates the reconstruction without appreciable loss of detail.

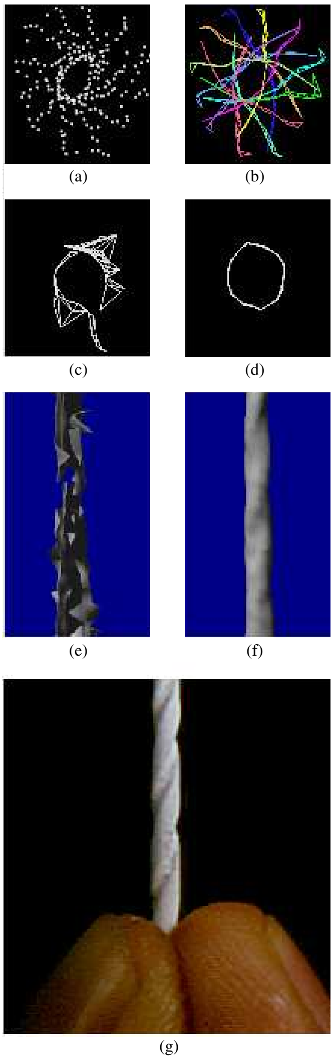

Figure 8.2: Merging range images of a drill bit. We scanned a 1.6 mm drill bit from 12 orientations at a 30 degree spacing using traditional optical triangulation methods. Illustrations (a) - (d) each show a plan (top) view of a slice taken through the range data and two reconstructions. (a) The range data shown as unorganized points: algorithms that operate on this form of data would likely have difficulty deriving the correct surface. (b) The range data shown as a set of wire frame tessellations of the range data: the false edge extensions pose a challenge to both polygon and volumetric methods. (c) A slice through the reconstructed surface generated by a polygon method: the zippering algorithm of [Turk & Levoy 1994]. (d) A slice through the reconstructed surface generated by the volumetric method described in this thesis. (e) A rendering of the zippered surface. (f) A rendering of the volumetrically generated surface. Note the catastrophic failure of the zippering algorithm. The volumetric method, however, produces a watertight model. (g) A photograph of the original drill bit. The drill bit was painted white for scanning.

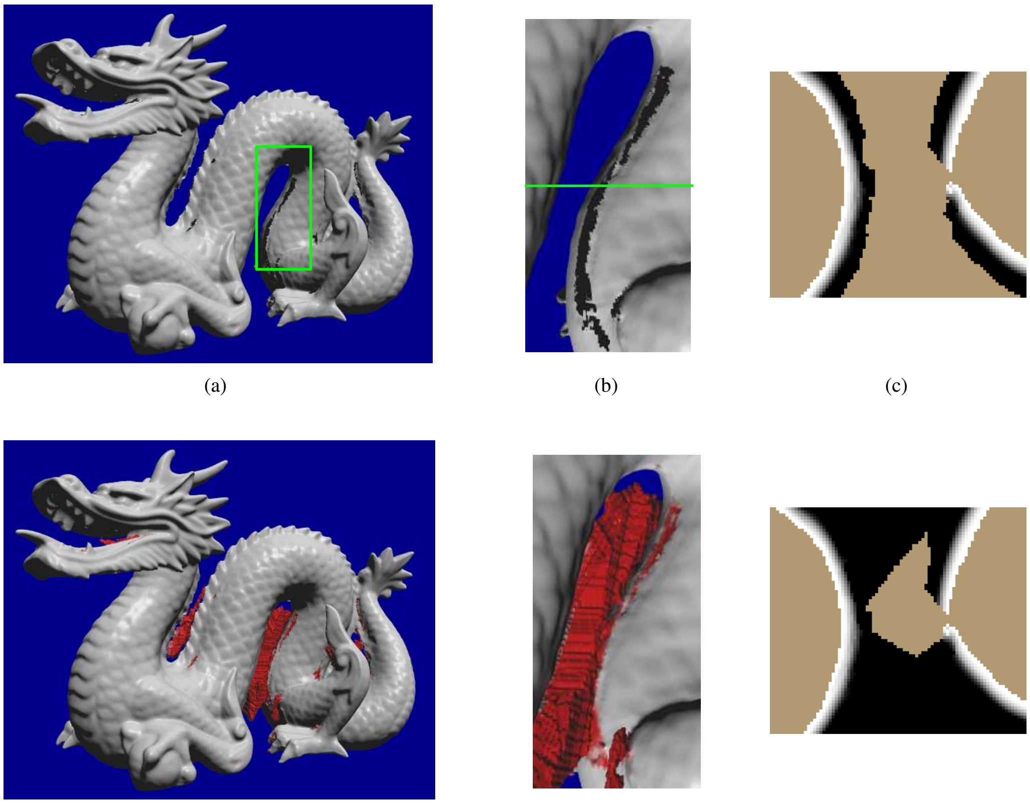

(a) (b) (c)

(d) (e) (f)

Figure 8.3: Reconstruction of a dragon – Part I of II. Illustrations (a) and (d) are full views of the dragon. Illustrations (b) and (e) are magnified views of the section highlighted by the green box in (a). Regions shown in red correspond to hole fill triangles. Illustrations (c) and (f) are slices through the corresponding volumetric grids at the level indicated by the green line in (b). (a)-(c) Reconstruction from 61 range images without space carving and hole filling. The magnified rendering highlights the holes in the belly. The slice through the volumetric grid shows how the signed distance ramps are maintained close to the surface. The gap in the ramps leads to a hole in the reconstruction. (d)-(f) Reconstruction with space carving and hole filling using the same data as in (a). While some holes are filled in a reasonable manner, some large regions of space are left untouched and create extraneous tessellations. The slice through the volumetric grid reveals that the isosurface between the unseen (brown) and empty (black) regions will be connected to the isosurface extracted from the distance ramps, making it part of the connected component of the dragon body and leaving us with a substantial number of false surfaces.

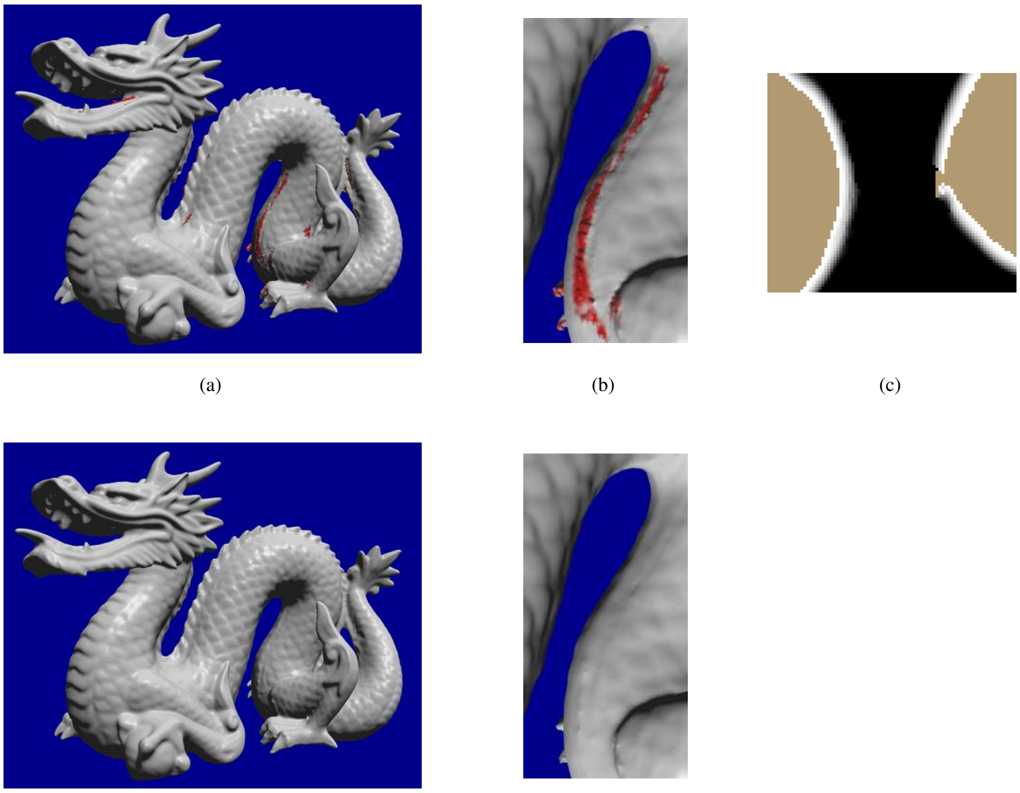

(a) (b) (c)

(d) (e)

Figure 8.4: Reconstruction of a dragon – Part II of II. Following Figure 8.3, (a) and (d) are full views of the dragon, (b) and (e) are magnified views of the belly, and (c) is a slice through the volumetric grid. (a)-(c) Reconstruction with 10 additional range images using "backdrop" surfaces to effect more carving. Notice how the extraneous hole fill triangles nearly vanish. The volumetric slice shows how we have managed to empty out the space near the belly. The bumpiness along the hole fill regions of the belly in (b) corresponds to aliasing artifacts from tessellating over the discontinuous transition between unseen and empty regions. (d) and (e) Reconstruction as in (a) and (b) with filtering of the hole fill portions of the mesh. The filtering operation blurs out the aliasing artifacts in the hole fill regions while preserving the detail in the rest of the model. Careful examination of (e) reveals a faint ridge in the vicinity of the smoothed hole fill. This ridge is actual geometry present in all of the shaded renderings, in this and the previous. The final model contains 1.8 million polygons and is watertight.

(a) (b) (c)

Figure 8.5: From the original to a 3D hardcopy of the "Happy Buddha" – Part I of II. (a) The original is a plastic and rosewood statuette that stands 20 cm tall. (b) Photograph of the original after spray painting it matte gray to simplify scanning. (c) Gouraud-shaded rendering of one range image of the statuette. Scans were acquired using a Cyberware scanner, modified to permit spacetime triangulation. This figure illustrates the limited and fragmentary nature of the information available from a single range image.

(a) (b) (c)

Figure 8.6: From the original to a 3D hardcopy of the "Happy Buddha" – Part II of II. (a) Gouraud-shaded rendering of the 2.4 million polygon mesh after merging 48 scans, but before hole-filling. Notice that the reconstructed mesh has at least as much detail as the single range image, but is less noisy; this is most apparent around the belly. The hole in the base of the model corresponds to regions that were not observed directly by the range sensor. (b) RenderMan rendering of an 800,000 polygon decimated version of the hole-filled and filtered mesh built from 58 scans. By placing a backdrop behind the model and taking 10 additional scans, we were able to see through the space between the base and the Buddha's garments, allowing us to carve space and fill the holes in the base. (c) Photograph of a hardcopy of the 3D model, manufactured by 3D Systems, Inc., using stereolithography. The computer model was sliced into 500 layers, 150 microns apart, and the hardcopy was built up layer by layer by selectively hardening a liquid resin. The process took about 10 hours. Afterwards, the model was sanded and bead-blasted to remove the stair-step artifacts that arise during layered manufacturing.

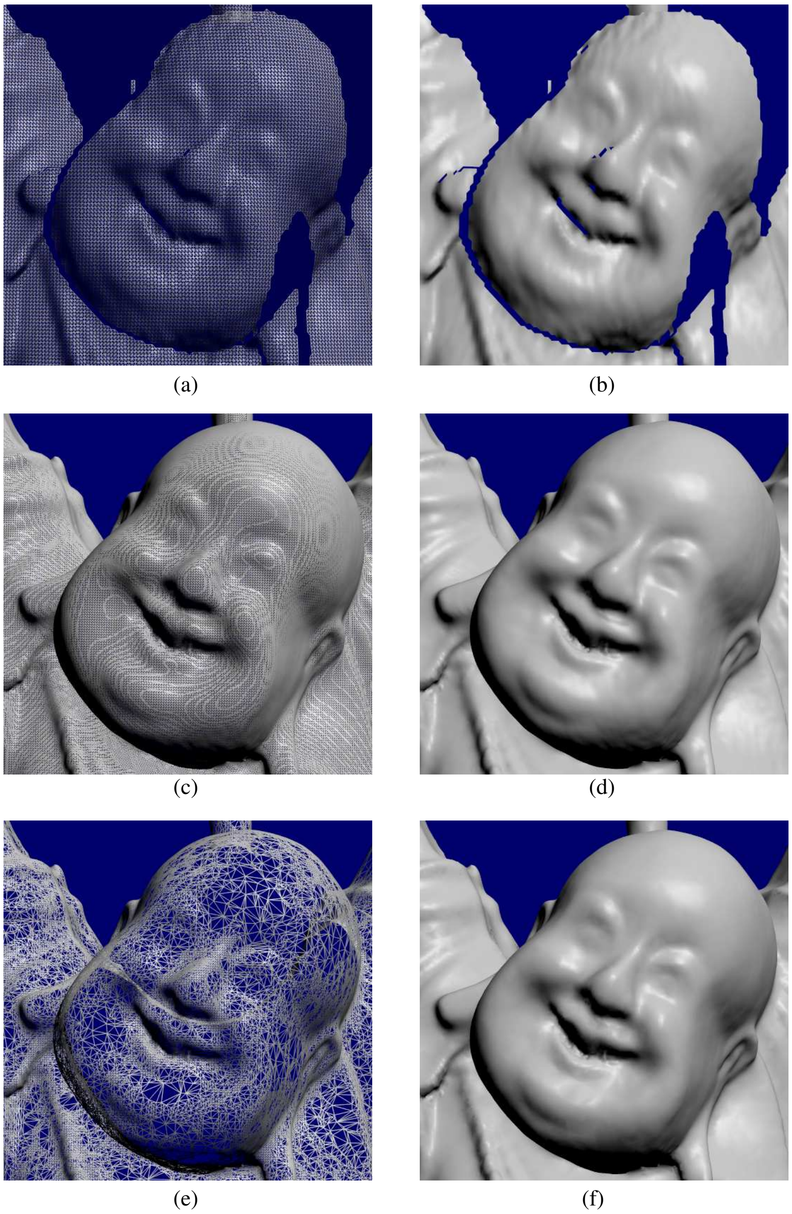

Figure 8.7: Wireframe and shaded renderings of the Happy Buddha model. (a) and (b) Wireframe and shaded rendering of a single range surface. (c) and (d) Wireframe and shaded rendering of the 2.6 million triangle reconstruction. The contours in the wireframe correspond to small triangles created when the isosurface clips the edges and corners of voxel cubes. The shaded rendering demonstrates noise reduction and increased detail after multiple scans are combined. (e) and (f) Wireframe and shaded rendering of the 800,000 triangle decimated mesh. The decimation process effectively collapses nearly coplanar triangles into larger triangles following the method of [Schroeder & Lorensen 1992].

Model	Scans	Input triangles	Voxel size (mm)	Volume dimensions	Exec. time (min.)	Output triangles	Holes
Dragon	61	15 M	0.35	712x501x322	56	1.7 M	324
Dragon + fill	71	24 M	0.35	712x501x322	257	1.8 M	0
Buddha	48	5 M	0.25	407x957x407	47	2.4 M	670
Buddha + fill	58	9 M	0.25	407x957x407	197	2.6 M	0

Table 8.1: Statistics for the reconstruction of the dragon and Buddha models, with and without space carving.

Statistics for the reconstruction of the dragon and Buddha models appear in Table 8.1. With the optimizations described in the previous section, we were able to reconstruct the observed portions of the surfaces in under an hour on a 250 MHz MIPS R4400 processor. The space carving and hole filling algorithm is not completely optimized, but the execution times are still in the range of 3-5 hours, less than the time spent acquiring and registering the range images. For both models, the RMS distance between points in the original range images and points on the reconstructed surfaces is approximately 0.1 mm. This figure is roughly the same as the accuracy of the scanning technology, indicating excellent alignment and a nearly optimal surface reconstruction.