Locating the source of topological error in reconstructed 3D models

Eric Firestone, Craig Povey, and Zoë J. Wood

California Polytechnic State University, Computer Science Department, San Luis Obispo, CA;

ABSTRACT

Although range scanning technology has offered great improvements to digital model creation in recent years, it has also introduced some new concerns. Specifically, recent work shows that topological errors such as tiny handles can significantly lower the overall quality of range-scanned models for down-stream applications (such as simplification and parameterization). In this paper we present our investigation into the source of this topological error in the range scanning process, and our methods to alleviate the error. We concentrated our investigation of the scanning process on: (1) signal noise or calibration error in the laser scanner (resulting in bad data points) and (2) error during the model reconstruction phase. We found that by modifying the surface reconstruction phase of the range scanning process, we were able to reduce the amount of topological noise in the resulting 3D model by up to 60 percent.

Keywords: range scans, surface reconstruction, topological noise

1. INTRODUCTION

Creating digital models of physical objects is essential for use in a vast array of different applications including: computer animation, video game development, rapid prototyping, medical research, archaeology, art history, and fashion design.[1,2] One very popular method of model acquisition is range scanning. This method yields a point cloud of 3D data, called a range image, that can then be aligned and combined with other range images captured from different view angles and then processed to produce a digital surface model. The final step in the scanning process is generally referred to as surface reconstruction.[3–5] The most popular approach to surface reconstruction merges each individual scan into a volumetric grid.[3] A triangle mesh is then extracted using the marching cubes algorithm.[6]

Although range scanning technology has contributed greatly to model creation, it has also introduced some concerns. Specifically, recent work shows that topological errors such as tiny handles can significantly lower the overall quality of range-scanned models.[7,8] A handle can be defined as a portion of the surface that is topologically equivalent to a genus one surface, like the handle on a coffee cup. In most cases, the handles found on the final reconstructed models are extremely small. They degrade the overall quality of a model indirectly by complicating subsequent geometry processing procedures, such as model simplification, smoothing, and parameterization (including texture mapping and remeshing).[9–11] Finally, some applications, such as the fitting of organ templates to medical MRI data, strictly require topologically correct models.[12,13] Reasonable progress has been made toward removing the topological defects from the constructed model, however, much of this work offers little to no insight regarding the fundamental cause of those handles, much less how to alleviate or even prevent their creation. We therefore aim to identify the source of these excess handles in the surface scanning process and suggest methods to alleviate their creation.

The digital model generation pipeline from scanned data typically includes the following stages: data acquisition, alignment and merging data, and surface reconstruction. We focused our investigations primarily on the first and last stages of the reconstruction pipeline as we hypothesized these to be the most likely stages for error to be introduced. In the first stage, any error introduced would be due chiefly to either signal noise or calibration error in the scanner device itself, subsequently resulting in the generation of outlier data points, and

therefore extraneous handles in the final model. The final stage in the generation of a digital surface model is surface reconstruction. In general terms, this stage is responsible for transforming a raw or basic representation of the model (i.e. cloud of data points or volumetric representation) into a closed, usually manifold mesh. In this stage, extraneous handles could be created in the construction of the surface from the model data. One of the main challenges to model reconstruction at this final stage is **hole filling**, where a **hole** is caused by gaps in the input data, which left untouched would result in a surface with numerous jagged boundaries.

One common reconstruction tool widely used in the research community is the VRIP application.[3] We base our exploration of the source of excess topology on this well known surface reconstruction tool. We found that by examining the acquisition and reconstruction phases, we were able to reduce the amount of erroneous topology by over 60 percent. In particular, we found that the surface reconstruction phase contributed strongly to the presence of excess topology.

2. RELATED WORK

Surface reconstruction has been tackled by many researchers,[3–5, 14, 15] with the work of Curless and Levoy[3] being used to create many of the most popular digital surface models. VRIP[3] is widely accepted and utilized in the research community and as a result, the excess topological handles in the models produced by VRIP became more evident to the researchers working with them. Recent research has focused on the problem of topological noise removal.[7, 8, 16, 20] This work focuses on removing the excess topology in either the volumetric or surface setting. Additional work from the domain of human brain cortex modeling addresses excess topology as well.[13, 17–19] Our work is not focused on identifying or removing this excess topology but on identifying its source in the reconstruction pipeline.

The methods we used to analyze the role of surface reconstruction in creating excess topology is broadly related to the wide variety of research that has focused on improving model representation and on hole filling in particular. Again these methods are either typically applied to the surface or volumetric setting.[24–26] For example, a recent paper by Tao Ju[21] proposes repairing polygonal models by taking a polygon soup as input, then constructing an intermediate volume grid, and finally generating a surface by dual contouring this grid. For an excellent survey of many of these methods, see the 2007 SIGGRAPH course on geometric modeling.[15]

One of the challenges of surface reconstruction is hole filling. Our work, led us to examine this stage of the surface reconstruction pipeline in particular. Curless and Levoy[3] incorporate a hole-filling approach called space carving into VRIP which uses scanner information to carve out volumetric regions and create surfaces over holes. These added surfaces serve to produce watertight models. Davis et al.[22] propose an improvement to this method called volumetric diffusion. This method is intended to address situations in which holes are too geometrically complex to fill using traditional triangulation algorithms. This technique consists of converting a surface into a volumetric representation with a given signed distance function. The function is initially defined only near observed regions of the surface, and then through alternating blurring and compositing steps, the function is diffused through the volume until its zero set covers all the existing holes. Another hole-filling approach proposed by Wang and Oliveira[23] is based on a moving least squares (MLS) algorithm and is intended to recover both geometry and shading information for the hole by using an interpolation procedure based on the context of the surrounding surface. For each hole, an MLS algorithm is used to repetitively resample and refit a surface to the hole until a reasonable fit is reached. Our work focuses on exploring the source of excess topology during the surface construction pipeline and our hole filling approaches are intended to expose and explore this issue. Our volumetric hole filling exploration follows the work of Nooruddin and Turk,[24] with some exceptions.

3. APPROACH

Our investigations include an analysis of the raw range data involved in the data acquisition stage, and an examination of hole-filling as a main component of the model reconstruction process. By exploring both of these paths, we discovered some potential sources for the excess topology created using standard processes for generating digital models from scanned data sets.

3.1 Raw Range Image Analysis

Our initial investigation involves taking a closer look into the raw range image data that serves as the input to the entire model creation process. We conjectured that it was possible that the topological handles observed were caused by erroneous points in the input data set generated by the laser scanning device. We focused on identifying and removing error in terms of outlier data points; that is, areas where the distance from one or more points to their neighboring points is greater than the average neighbor-to-neighbor distance. Furthermore, we hypothesized that the most likely candidates for regions such as this were areas of overlap between adjacent range images. Our reasoning behind this was twofold; first, these overlapping regions correspond to the edges of individual range images, which are generally known to be the least accurate portions of any given range image and because these regions are the most susceptible to inaccuracies in the alignment process. Although these outlier data points are typically weighted less in the final reconstruction of a surface,[3] we wanted to explore removing these points completely from the reconstruction pipeline.

Since we cannot gain any sense of overlapping regions or range image adjacency without first aligning all the range images, we actually perform our data point analysis and removal algorithm after the range image registration phase. The basic idea of our algorithm was to remove any range data points whose distance to its neighbors was beyond a given threshold value. Our tests consisted of the following steps: (1) run our algorithm to identify likely erroneous points, (2) remove the flagged points and write out a filtered version of the input data set, (3) input the filtered data set into the reconstruction phase (in place of the original data), and (4) compare the genus (handle count) of the model produced from our filtered data to that of the model produced from the original, unmodified data.

Our filtering algorithm may be broken into the following high-level steps:

1. Find all regions of overlap between each pair of range images

2. For each point P in an overlap region, find the closest point on the closest neighboring range surface to P

3. If the distance to this closest point is greater than a given threshold (DT), remove P as a potential erroneous data point

In order to determine the overlapping regions between neighboring range images, we used spatial partitions to efficiently gain information regarding both whether or not a given region in the volume contains overlapping range images, as well as to narrow the search for closest neighbors on adjacent scans. We chose to set the value of DT for a model to be twice the average distance between neighboring vertices in a given range image from that model. In our analysis, we use this filtered range data as input into the reconstruction stage of the model creation pipeline in order to determine whether removing these flagged vertices improves the genus of the generated model. Not surprisingly, most of the outlier points we identified were indeed at the edges of each range image, however, we found mixed results in terms of the effect of removing this data before constructing a surface. For results see Section 4.

3.2 Hole-Filling Analysis

In addition to looking at the raw range data as a possible source of topological error, we also speculated that the process of filling holes in the mesh during model reconstruction could contribute to topological noise in the final model. Various hole filling methods exist. Due to the popularity of VRIP, we choose to examine a comparison between the VRIP hole-filler and two alternative hole fillers. We compared VRIP output against both our own simple mesh-based post-process hole filler and our own combined volumetric-surface hole filler. Both hole-fillers are somewhat simplistic in that they do not aim to produce aesthetically pleasing result; our goal was to simply close open boundaries (i.e. holes) in the most straight-forward manner with the intent of reducing the introduction of excess topology.

The following summarizes the functionality of VRIPs hole-filling mechanism. Their hole-filler takes a volumetric approach.

1. Set all voxel points as unseen

Figure 1. An example of an invalid hole-spanning face. The mesh on the left is a valid manifold-with-boundaries mesh. The mesh on the right, however, is non-manifold since it contains invalid vertices (marked in red) that are neither disks nor half-disks. These vertices represent an invalid hole-spanning face or isle.

2. Update voxels near to or containing the surface with continuous signed distance and weight values

3. Follow the lines of sight back from the observed surface and mark the corresponding voxels as empty (this step is known as space carving)

4. Extract an isosurface made up of: (1) the zero-crossing of the signed distance function (i.e. the scanned surface) and (2) the surfaces corresponding to boundaries between unseen and empty voxel points (i.e. filled holes)

The basic idea behind our mesh based hole-filler was to start with extracting a surface from the volume which contained boundaries for any holes (i.e. that only extracted a surface for the zero-crossing of the signed distance function) and then to simply fill each hole in the mesh with a triangle fan. While this will not produce a visually realistic result in most cases, it is a simple solution that is not likely to add any extraneous topology. Despite the simplicity of this approach, there are some important considerations for the method to generate a manifold mesh.

Intuitively, the first step in the hole-filling process would be to find the holes in the mesh. However, before doing this, it is necessary to ensure that our mesh is manifold-with-boundaries; that is, that every vertex in the mesh is topologically equivalent to either a disk or a half-disk. This is necessary because we ultimately want to output a manifold and our process of filling in holes is to cap each boundary with a disk and in order for this to produce a manifold, we must start with a manifold-with-boundaries. In some cases, the holes in the mesh will already conform to this standard; however, we found that many meshes contain cases that do not. For example, it is not uncommon to encounter a face or group of faces which span a hole, touching the holes boundary only at vertices (we refer to this case as containing hole-spanning faces, which are sometimes called isles)(see Figure 1). In order to identify the hole-spanning faces in a mesh, we take a simple approach of positively identifying all faces that are known to be non-hole-spanning faces, implicitly identifying any unmarked faces as hole-spanning faces. We start with any random face in the mesh and mark it as seen, we then add any face sharing an edge with it to the stack and repeat this process, marking all faces reachable across an edge. Any faces *not marked* in this traversal are considered isles and removed.

By finding and removing all the hole-spanning faces from the mesh, we now have a manifold-with-boundaries. This greatly simplifies the process of identifying holes because once a boundary edge is identified, we need only recursively trace along neighboring boundary edges until we arrive back at the starting edge to find a hole. Once we have traversed all the boundary edges and identified boundary loops (sets of boundary edges), the only remaining step is to fill each hole by inserting a surface such that all boundary edges for a given hole are adjacent to the inserted surface. Specifically, we fill each hole using a triangle fan located at the center of the hole. Although this technique only makes little attempt to reconstruct the original geometry of the physical model, it does guarantee that the hole-filled regions on the reconstructed model will be manifold and reduces the introduction of extraneous topological handles, by filling holes as simply as possible.

3.3 Hybrid Volumetric-Surface Hole Filler

We also explored the use of a hybrid volumetric and surface hole filler. This exploration was motivated by the fact that our surface hole filler did reduce the introduction of extraneous topological handles, however, it caused

some topological side-effects. See Section 4, specifically, Figure 2. Our hybrid hole filler worked first on the volume data directly by using morphological operators similar to those in Noorudin and Turk,[24] however the expansion and contraction are only performed on boundary voxels and in a direction orthogonal to the surface normal at the boundary. This volume hole filling was then followed by our own surface hole filler. See Section 4.2 for a discussion of this method.

4. RESULTS

We present results from our analysis of both the raw input range data as well as the hole-filling stage of model reconstruction. The range image data used in our experimental tests was obtained from Stanfords 3D Scanning Repository website and includes the following models: a Buddha statuette, a sculpture of a Chinese dragon, and an armadillo action figure (see Figures 2 and 3). Each of these models was scanned using a Cyberware 3030 MS laser scanner device, and consists of approximately 60-70 range scans/images.

4.1 Basic Hole-Filler Results

Our experiment for testing the topological effects of the reconstruction stage in the model creation pipeline was designed in the following manner. For each model being tested, we obtained the set of range images corresponding to the model and reconstructed two different versions of each; once using VRIPs internal hole-filler and a second time without it and then filling holes as described in Section 3.2. The results of this process are displayed in Table 1.

Table 1. Hole filling test results

	Voxel Dimension	num of Faces	Genus
Dragon			
VRIP	0.00025	3,538,580	**45**
Our method	0.00025	3,405,232	**17**
Buddha			
VRIP	0.00023	3,130,148	**67**
Our method	0.00023	2,872,186	**28**
Armadillo			
VRIP	0.00023	2,955,588	**1**
Our method	0.00023	2,954,372	**1**

As the table shows, our basic hole-filling algorithm was able to effectively decrease the genus of the final reconstruction by over 60 percent for both the dragon and Buddha models. Interestingly, the armadillo model had negligible topological noise in both cases; we attribute this phenomenon to the fact that the armadillo figurine itself has fairly convex geometry in comparison to the other two models; that is, there are very few regions of the model obscured from the scanner. This helps validate the logical hypothesis that topological noise is most common in the areas of a model where range data is missing; that is, in areas that are difficult to sample using a laser scanner. Our results also validate our initial hypothesis that the reconstruction phase of model creation is responsible for at least a substantial portion of the topological noise present in many complex models created from range data.

Figure 2 and 3 compares the appearance of the two hole-filled models. The visual results of our hole-filler can be observed in various places; the most obvious areas are enclosed by red boxes. As the figure shows, our hole-filler creates noticeably incorrect star-burst shaped patterns in the regions where larger holes were filled. This is a direct result of our simplistic approach in hole-filling. Note that these visual side-effects would not be acceptable for most computer graphics applications. However, our boundary filling method was designed with the intention of assisting our investigation into the source of extraneous topology and therefore a final surface reconstruction algorithm would need to be modified to correct these side-effects.
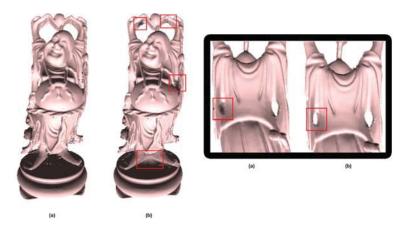
Figure 2. Hole-filled Buddha models. On the left we see a long view of the Buddha statue. Part (a) was hole-filled by VRIP, and part (b) by our own basic hole-filling algorithm. The red boxes point out areas where the visual results of our hole-filler can be most easily observed. On the right we see a close up from a different view of the Buddha's armpit. It is clear that our naive hole filler has incorrectly closed the handle of the Buddha's arm.
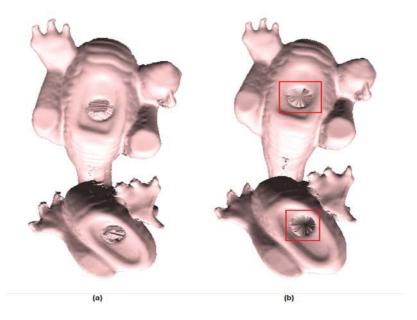

Figure 3. Hole-filled dragon models. Part (a) was hole-filled by VRIP, and parts (b) by our own basic hole-filling algorithm. The star-burst shape of our surface hole filler is clearly visible.

Figure 4. Hybrid hole filled Buddha arm. The Buddha statue with hole filling performed using our hybrid method. On the left we see the entire genus 28 statue. In the middle is a close up of the armpit region hole filled using VRIP (total genus 47), compared with the same region using our method. Note that this surface is an improvement over the pure surface based method (see Figure 2), however, selecting the geometric center for creating a triangle fan to fill this hole, still creates poor geometry for this hole.

4.2 Hole Filler Discussion

As can bee seen in Figure 2 there are potential topological side-effects of a pure surface hole filler. Notably, for large boundaries such as those found in the Buddha's armpit, where there is very little scanner range data, hole filling can actually cause large geometric inaccuracies. This further highlights the importance of this stage in the surface reconstruction pipeline. Given the inaccuracies we experienced with just a surface method, we also explored creating our own hybrid volumetric-surface hole filler. Our motivation was to create a hole filling method that was more faithful to the geometry of the surface by filtering the volume following the geometry of the surface as much as possible. This method first works on filling in surface information in the volume. The method identifies boundary regions in the volume and then applies dilation and expansion similar to Noorudin and Turk,[24] but only to regions of the volume that are mostly perpendicular to the normal of the surface at the boundary. By combining this volumetric filtering with our surface hole filler, we were able to generate models which more faithfully matched the desired surface (i.e. the Buddha's arm handle is left open), however, the hole filling was still imperfect. See Figure 4. The imperfections in Figure 4 are due to an inherent problem with volumetric hole fillers, including volumetric diffusion,[22] which is that it is difficult to know how much diffusion or filtering should be performed in order to fill differing sized holes. In other words, a voxel based approach which is applied uniformly to all holes will either over filter small holes or under fill large holes. The genus for the Buddha model using our hybrid volumetric-surface based hole filler was 27 and for the VRIP hole filler when ran on our filtered volume was 47. Note that although using our filtered volume provided only limited genus improvement over just a surface based hole filler, using the filtered volume improved VRIP's hole filling results by 30 percent.

4.3 Raw data filtering

For our analysis into the topological effects of possibly erroneous or noisy raw range data (as generated by the laser scanning hardware), we compared the output of filtered and unfiltered data sets through the rest of the pipeline in an identical manner, producing two reconstructed models. The results of this investigation are displayed in Table 2.

Table 2. Raw data filtering test results (Note that voxel dimensions are the same as in Table 1)

	Range Data	num of Faces	Genus
Dragon			
VRIP hole fill	Original	3,538,580	**45**
	Filtered (0.0017)	3,552,704	**73**
Our hole fill	Original	3,405,232	**17**
	Filtered (0.0017)	3,401,630	**24**
Buddha			
VRIP hole fill	Original	3,130,148	**67**
	Filtered (0.001)	3,175,546	**100**
Our hole fill	Original	2,872,186	**28**
	Filtered (0.001)	2,831,724	**22**

The modified models described here were filtered using a data point distance threshold of approximately two times the average nearest neighbor distance for the range set; specifically, a threshold of 0.0017 was used for the dragon scans, and a threshold of 0.001 was used for the Buddha. The genus results were somewhat surprising in that the effect of range data filtering was noticeably different depending on the hole-filling mechanism used. For both models, when the VRIP hole-filler was used on both the filtered and non-filtered data, the model resulting from the filtered data had a substantially higher (i.e. worse) genus than that from the original data. Specifically, we saw the models genus increase by 28 in the case of the dragon (a 62 percent increase) and by 37 in the case of the Buddha (a 55 percent increase). When we used our own hole-filling technique, the filtered models genus changed by 7 (or 41 percent) in the case of the dragon, and of 6 (or 21 percent) for the Buddha, thus, removing outlier data points from the range images using our own hole-filler did not effect the final genus to the same extent as it did in the VRIP case.

These results reveal a great deal regarding the source of topological error in these models. First of all, they show that outlier points in range images are not a significant contributor to topological noise in reconstructed models, as we initially hypothesized. Although this in itself does not show us where the problem is, it does at least show us where it is not. Additionally, our results actually support our supplementary hypothesis that hole-filling contributes to topological error. Since the process of filtering range images effectively either creates or enlarges holes or gaps in the original range images, this directly leads to an increased amount of hole-filling during reconstruction.

Although the difference in genus between the filtered and non-filtered models was much smaller when our own hole-filler was used, these cases are noteworthy for a different reason; namely, the fact that the genus increased slightly in the case of the dragon, but decreased slightly in that of the Buddha. One possible explanation for the Buddhas decreased genus is that the filtering process actually removed data points that would have otherwise caused topological handles apart from the hole-filling process (the fact that the models created using our hole-filler still have topological error shows that there is at least one other source, apart from hole-filling). Thus, by re-creating these parts of the model using our hole-filler, we were able to eliminate the error in these areas. That said, the most likely explanation for the dragons increased genus is simply that whatever else is also responsible for creating topological error in these models was exacerbated by the process of filtering the range images. For example, if we assume for the moment that slight misalignments in the range image alignment stage do in fact lead to topological error in some cases, it makes sense that removing parts of the data (especially data around the boundary regions) would worsen the existing misalignment and therefore increase the amount of error generated as a result.

4.4 Conclusion

The ultimate goals of this project were to identify the source of erroneous topological handles in reconstructed 3D models, and to develop a method for alleviating some or all of the erroneous topology. Although we were not able to identify the source of all the error in our experimental data, our analysis of the reconstruction phase demonstrates that hole filling is the cause of over 60 percent of the topological handles in the experimental data. In addition to showing that volumetric hole filling is a major source of topological error, we also demonstrated the equally important result that outlier points in the range image data are not a significant source of topological noise in the corresponding reconstructed models. Finally, in regard to future research efforts, our work also shows that while volumetric hole-filling plays a role in creating topological noise, it is not the only culprit; since the models reconstructed with our simplified mesh hole filling algorithm still have higher-than-expected genus, we may deduce that there is at least one additional unidentified source of error.

5. FUTURE WORK

The work done for this project can be extended in two main ways. The first option for future work is to explore the range image alignment stage in the 3D model generation pipeline with respect to its possible role in creating topological error. In this project, we have focused investigations into the data acquisition and model reconstruction stages, however, we acknowledge that the process of range image alignment is nonetheless a likely source of topological error in reconstructed models. One reason for this is simply because it is one of the most difficult parts of model creation, and still leaves a great deal of room for improvement. Most common alignment techniques used today are based on some variation of the Iterative Closest Points algorithm (ICP),[27] which uses rigid-body transforms to align adjacent range images over a series of incremental steps. However, though this technique is common, recent work has shown that it might not be the best or most logical way to align range image data. According to the work of Brown and Rusinkiewicz,[28] warps in laser scanned range data are often non-rigid in nature; they point out that even small calibration errors in the laser scanner device can result in a low-frequency, non-rigid warp in the acquired data, which cannot accurately be fitted by a rigid-body alignment algorithm like ICP. These authors present a new alignment method using thin-plate splines to perform non-rigid alignment of range images. Therefore, one avenue of research is to align a set of range images using two different methods (such as traditional ICP and the non-rigid technique of Brown and Rusinkiewicz) and then compare the genus results of the models reconstructed from each of these data sets.

A second avenue for future work related to this project is to improve the hole-filling tools that we developed. A hole-filling algorithm which more accurately follows the geometry of the input model, yet does not introduce excess topology is desirable. Our hybrid method which works on the volume and the surface indicates some possible directions of future work.

ACKNOWLEDGMENTS

REFERENCES

1. A. Dolenc, "Software tools for rapid prototyping technologies in manufacturing," in *Acta Polytechnica Scandinavica: Mathematics and Computer Science Series*, **62**, 1993.
2. J. Maintz and M. Viergever, "A survey of medical image registration," in *Medical Image Analysis*, p. 136, 1998.
3. B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH '96*, *ACM Press*, pp. 303–312, 1996.
4. H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface reconstruction from unorganized points," in *SIGGRAPH 92*, **26**, p. 7178, 1992.
5. K. Pulli, T. Duchamp, H. Hoppe, J. McDonald, L. Shapiro, and W. Stuetzle, "Robust meshes from multiple range maps," 1997.

6. W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," in *Proceedings of the 14th Annual Conference on Computer Graphics and interactive Techniques*, pp. 163–169, 1987.

7. I. Guskov and Z. Wood, "Topological noise removal," in *Graphics Interface*, p. 1926, 2001.

8. Z. Wood, H. Hoppe, M. Desbrun, and P. Schrder, "An out-of-core algorithm for isosurface topology simplification," in *ACM Transactions on Graphics*, 2002.

9. I. Guskov, A. Khodakovsky, P. Schroder, and W. Sweldens, "Hybrid meshes: multiresolution using regular and irregular refinement," in *Proceedings of the Eighteenth Annual Symposium on Computational Geometry*, p. 264272, 2002.

10. A. Khodakovsky, P. Schroder, and W. Sweldens, "Progressive geometry compression," in *SIGGRAPH 00*, p. 271278, 2000.

11. P. V. Sander, J. Snyder, S. J. Gortler, and H. Hoppe, "Texture mapping progressive meshes," in *SIGGRAPH 01*, p. 409416, 2001.

12. S. Jaume, B. M. Macq, and S. K. Warfield, "Labeling the brain surface using a deformable multiresolution mesh," in *MICCAI*, p. 451458, 2002.

13. D. W. Shattuck and R. M. Leahy, "Automated graph-based analysis and correction of cortical volume topology," in *IEEE Trans. Med. Imaging*, p. 11671177, 2001.

14. G. Turk and M. Levoy, "Zippered polygon meshes from range images," in *SIGGRAPH '94*, pp. 311–318, 1994.

15. M. Botsch, M. Pauly, L. Kobbelt, P. Alliez, B. Levy, S. Bischoff, and C. Rossl, "Geometric modeling based on polygonal meshes," *SIGGRAPH 2007 Course Notes* , 2007.

16. S. Bischoff and L. P. Kobbelt, "Isosurface reconstruction with topology control," in *Proceedings of the 10th Pacific Conference on Computer Graphics and Applications*, 2002.

17. A. Dale, B. Fischl, and M. Sereno, "Cortical surface-based analysis i: Segmentation and surface reconstruction," *NeuroImage* **9**, pp. 179–194, 1999.

18. B. Fischl, A. Liu, and A. Dale, "Automated manifold surgery: Constructing geometrically accurate and topologically correct models of the human cerebral cortex," in *IEEE Trans. Med. Imaging*, **20**, p. 7080, 2001.

19. X. Han and C. Xu, "Topology correction in brain cortex segmentation using a multiscale, graph-based algorithm," in *IEEE Trans. Med. Imaging*, **21**, p. 109121, 2002.

20. A. Szymczak and J. Vanderhyde, "Extraction of topologically simple isosurfaces from volume datasets," in *Proceedings of the 14th IEEE Visualization 2003*, 2003.

21. T. Ju, "Robust repair of polygonal models," in *ACM Trans. Graphics*, **23**, pp. 888–895, 2004.

22. J. Davis, S. Marschner, M. Garr, and M. Levoy, "Filling holes in complex surfaces using volumetric diffusion," in *Proceedings of the First International Symposium on 3D Data Processing, Visualization, and Transmission*, 2002.

23. J. Wang and M. Oliveira, "A hole filling strategy for reconstruction of smooth surfaces in range images," in *Proceedings of the XVI Brazilian Symposium on Computer Graphics and Image Processing.*, 2003.

24. F. Nooruddin and G. Turk, "Simplification and repair of polygonal models using volumetric techniques," *IEEE Transactions on Visualization and Computer Graphics* , pp. 191–205, 2003.

25. S. Bischoff, D. Pavic, and L. Kobbelt, "Automatic restoration of polygon models," *Transactions on Graphics* , 2005.

26. P. Liepa, "Filling holes in meshes," in *Symposium on Geometry Processing*, pp. 200–205, 2003.

27. S. Rusinkiewicz and M. Levoy, "Efficient variants of the icp algorithm.," in *Proceedings of the International Conference on 3D Digital Imaging and Modeling*, p. 145152, 2001.

28. B. J. Brown and S. Rusinkiewicz, "Non-rigid range-scan alignment using thin-plate splines," in *Proceedings of the 3D Data Processing, Visualization, and Transmission, IEEE Computer Society* **00**, pp. 759–765, 2004.