

Improving 3D Reconstruction for Digital Art Preservation

Jurandir Santos Junior, Olga Bellon, Luciano Silva, and Alexandre Vrubel

Department of Informatics, Universidade Federal do Parana,
IMAGO Research Group, Curitiba, Brazil
{jurandiro,olga,luciano,alexandrev}@inf.ufpr.br
<http://www.imago.ufpr.br>

Abstract. Achieving a high fidelity triangle mesh from 3D digital reconstructions is still a challenge, mainly due to the harmful effects of outliers in the range data. In this work, we discuss these artifacts and suggest improvements for two widely used volumetric integration techniques: VRIP and Consensus Surfaces (CS). A novel contribution is a hybrid approach, named IMAGO Volumetric Integration Algorithm (IVIA), which combines strengths from both VRIP and CS while adds new ideas that greatly improve the detection and elimination of artifacts. We show that IVIA leads to superior results when applied in different scenarios. In addition, IVIA cooperates with the hole filling process, improving the overall quality of the generated 3D models. We also compare IVIA to Poisson Surface Reconstruction, a state-of-the-art method with good reconstruction results and high performance both in terms of memory usage and processing time.

Keywords: range data, 3D reconstruction, cultural heritage.

1 Introduction

The preservation of natural and cultural assets is one of the most challenging applications for digital reconstruction. It requires high fidelity results at the cost of high-resolution range acquisition devices, along with a set of algorithms capable of dealing with noise and other artifacts present in range data.

Several stages compose a complete 3D reconstruction pipeline, [1], [2], [3]. In this work, we focus on mesh integration within the pipeline proposed in [3]. Several approaches perform this task; we categorize them as presented in [1]. *Delaunay-based methods* [10] work on point clouds, but they are usually sensitive to noise and outliers since they interpolate the data points. *Surface-based methods* [11], [12] create surfaces directly, but some of them can catastrophically fail when applied on high curvature regions [6]. Also, topologically incorrect solutions can occur due to outliers in the input views.

Parametric surfaces methods [13], [14], [15] deform an initial approximation of the object through the use of external forces and internal reactions and constraints. Kazhdan *et al.* [9] calculate locally supported basis functions by solving

sparse linear systems that represent a Poisson problem. A parallel and an out-of-core implementations of [9] are proposed in [16], [17]. One limitation of such algorithms is finding a balance between over smoothing effect and non-elimination of noisy surfaces. These methods may also fill holes incorrectly.

Volumetric methods [22], [23] create an implicit volumetric representation of the final model, which is also called a SDF (*Signed Distance Field*). The object surface is defined by the isosurface at distance value 0, that is usually extracted with the MC (*Marching Cubes*) [18]. Curless and Levoy [6] calculate the distance function according to the line-of-sight of the scanner and uses weights to assess the reliability of each measurement. Wheeler *et al.* [7] discard outliers through the calculation of a consensus surface. Volumetric methods use all available information and ensure the generation of manifold topologies [6]. The main limitation is their performance, both in terms of memory usage and processing time.

All integration approaches have their limitations. We have chosen volumetric methods because they impose fewer restrictions to the reconstructed objects; offer an easy way to change the precision of the output; can easily support the space carving technique [6], and can work in the presence of noisy input data. We implemented, tested and modified VRIP [6], [4] and Consensus Surface (CS) [7], [5]. So, in the attempt to solve their limitations, we developed a novel hybrid integration method, named IMAGO Volumetric Integration Algorithm (IVIA), which improves the fidelity of the generated 3D models. The first pass of our method approximate normals and eliminate outliers based on a modified VRIP. Then, a second pass reconstruct a bias-free surface using CS with our suggested improvements, and fully cooperates with the hole filling algorithm of [8].

IVIA was developed to handle with artifacts present on range data, such as noise, occlusions and false data. We show how it overcomes these problems and generates more accurate results. We also compare it to the Poisson Surface Reconstruction (PSR) [9]. Although PSR is not a volumetric approach, it is considered the state-of-the-art method for 3D reconstruction. Also, PSR presents superior results when compared to other methods and good performance both in terms of memory usage and processing time.

The remainder of this paper is organized as follows: first we brief analyze the artifacts present on range data in section 2. Then, in section 3 we suggest improvements for VRIP [6] and CS [7]. In section 4, we present our novel hybrid volumetric method, named IVIA. We discuss our results compared to VRIP, CS and PSR in section 5, followed by conclusions in section 6.

2 Range Data Analysis

To better understand the challenges of building a high quality integration algorithm, we analyzed the types of artifacts that appear in the input range data. In this section, we discussed data captured with a laser triangulation 3D scanner (*e.g.* Vivid 910 from Konica Minolta).

This analysis is important because those artifacts are complex and harder to automatically detect and discard. For instance, noise on the object surface is one

of the most common artifact on range data. Even for very smooth objects, the captured data is usually rough. Another common artifact is data returned by the scanner that do not exist in the object. They usually appear as small groups of triangles in regions that should be empty. It is hard to eliminate this type of artifact without leading to the elimination of valid data. In [6] the authors proposed the space carving technique that was used to fill holes. However, they do not eliminate outliers as we do in our proposed approach.

Deformed surfaces is a more serious and relatively common artifact of difficult detection because they consist of large areas, usually connected to valid data, but completely wrong compared to the original object. Automatic detection and elimination of such artifacts is particularly difficult. Our IVIA approach successfully discards most of these artifacts, as discussed in section 5.

Another problem is due to the triangulation of the point cloud captured by the scanner. That introduces false internal and external silhouettes surfaces. Our solution is to check the angle between the face normal and the line-of-sight from the scanner position to the center of the face. If this angle is above some threshold, the face should be discarded. By using a threshold of 75° we achieved good results in all of our experiments.

Impossible to be captured data is a problem caused due to occlusions. For many applications, such holes are unacceptable and hole filling techniques complement the captured data [8]. The more information we can extract from the captured data, the better the completed surface will be. While other integration techniques usually discard helpful data that could be used in the hole filling process, IVIA was designed to fully cooperate with the hole filling stage.

3 VRIP, CS and PSR Approaches

As our hybrid method was developed to overcome the limitations of the VRIP and the CS, to better understand it, we assess them and suggest some improvements. As we compared our IVIA to PSR we shortly discuss this approach too.

3.1 Volumetric Range Image Processing (VRIP)

Like all volumetric integration methods, the goal of VRIP (Curless and Levoy [6]) is to calculate the signed distance from each voxel of the volume to the integrated surface. VRIP performs fast and allows incremental addition of views. Besides, it generates a smoother integrated surface, reduces noise and does not require all views being loaded in memory simultaneously.

The main limitation of VRIP is not comprise any specific process to discard outliers. It is even difficult to reduce their weight as each view is processed individually. It also integrates metrics from different viewpoints, resulting the SDF to be non-uniform. In addition, a flaw causes artifacts in corners and thin surfaces, as pointed by the authors [6]. As the algorithm combines positive and negative values in order to find the average surface, it creates creases on thin surfaces where one side would interfere with the other, increasing the thickness.

To reduce both the influence of outliers and the flaw near corners and thin surfaces, we propose a new weight attenuation curve according to the value of the signed distance. Instead of using a linear reduction from half of the range, as originally suggested in [6], we adopted a non-symmetric curve, giving more weight to measurements we are sure of (*i.e.* outside voxels). Although better visual results are obtained, this “trick” is not yet a definitive solution for a high fidelity reconstruction; however, it will be useful to our IVIA approach later.

3.2 Consensus Surface (CS)

Wheeler *et al.* [7] proposed to use consensual distances averaged to create the signed distance to the surface from the current voxel. Therefore, in consensual regions from several views, CS eliminates outliers and provides a relatively smooth integrated surface. However, it performs slowly, their magnitude and sign of the distance may be incorrectly calculated around the borders of views and when it automatically fills holes, the results may present inconsistencies.

The majority of the visible artifacts on CS result from incorrect signs being calculated for the distance values, as pointed by Sagawa *et al.* [20], [21]. However, they assumed that the consensus criterion successfully removes all outliers. We noticed in our experiments that even if the SDF signs are repaired, their magnitudes can still be wrong, and we cannot guarantee that the reconstruction will have high fidelity.

Trying to address the main limitations of CS, we propose modifications on it. First, we choose to use equal sized voxels to achieve the highest possible precision. As CS calculates distance measurements incorrectly when the nearest point in the view is on a mesh border, we choose to discard all of them. After the integration, we also perform an additional data validation by discarding voxels whose distances are not compatible with their neighbors. This process removes any eventual incorrect remaining data after the previous stages. Our modified CS still bears some limitations of the original algorithm such as: difficulty in choosing thresholds, the necessity to have all views loaded in memory simultaneously and outliers may appear on the integrated results. Our new IVIA algorithm combines the strengths of both the VRIP and our modified CS to overcome these problems aiming the generation of more accurate results.

3.3 Poisson Surface Reconstruction (PSR)

Kazhdan *et al.* [9] express surface reconstruction as the solution to a Poisson equation. Their idea is to compute a function χ that has value one inside the surface and zero outside (*i.e.* indicator function), and by extracting an appropriated isosurface they reconstruct the object surface. The oriented point set can be viewed as samples of the gradient from that function. They solve the indicator function by finding the scalar function whose gradient best matches the vector field \vec{V} , a variational problem of a standard Poisson equation.

PSR performs good both in terms of memory usage and processing time, and has the benefit to increase the resolution by increasing the octree depth. However,

whenever resolution is increased the time performance becomes lower. It is also difficult finding a balance between over smoothing effect and non-elimination of noisy surfaces (see section 5). Besides, PSR does not ensure the generation of manifold topologies or when it creates manifold meshes it may occurs that outliers as vertices appear on the final model (see section 5). In additional, PSR does not use all available information as well (*e.g.* line-of-sight information) what may cause incorrectly connections of some regions, a problem pointed out by the authors [9].

4 The Proposed IVIA Approach

The greatest challenge to develop a high quality reconstruction is detecting and discarding several types of artifacts that usually appear on range data. We developed the a novel approach, named IMAGO Volumetric Integration Algorithm (IVIA), that automatically detects and eliminates most of them. We perform the integration in two passes. The first one creates a volumetric representation of the integrated model by using our modified version of the VRIP algorithm; the second one uses the representation created to detect and discard outliers during the generation of the final volumetric integration.

The goal of the first pass is to build an approximated volumetric representation through the modified VRIP presented in section 3.1. Besides, a binary volume *empty* is created to represent the space carving operation [6]. For each voxel there is a corresponding bit in *empty*; if the bit is set, the corresponding voxel is considered empty (*i.e.* outside the object). As the distance calculated is over the line-of-sight of the scanner, any negative values are considered “empty” and the corresponding bit of *empty* is set. It is important to notice that we use our modified distance weight curve, to help eliminate surface outliers.

After all views are processed, we apply a 3D mathematical morphology erosion operation on the binary volume *empty*. This is necessary for two reasons: first, to prevent any incorrect measurement “dig holes” in the object. Though unusual, we already noticed this type of error on some range images. The other reason is that *empty* is obtained through the union of empty spaces of all views. Therefore, it tends to represent the lowest measurement for each surface point, and not an average measurement. When reducing the empty space by a distance $+D_{max}$ we set aside a space near the surface to integrate the measurements from several views, and at the same time we keep an empty space representation to eliminate outliers far from the surface. Several artifacts presented in section 2 are eliminated by using *empty*.

Finishing the first pass, we smoothen the values of the volumetric representation, using a $3 \times 3 \times 3$ filter with larger weights to central voxels on the mask (*i.e.* 2D smoothing filter). This smoothing completes individual voxels with plausible values and attenuate noise and surface outliers. This attenuation is important because this volumetric representation will be used to estimate the integrated object normals in the second pass. After that, the volumetric representation created is saved.

The second pass of the algorithm does the definitive integration, discarding outliers. This pass possesses elements from CS, like the Euclidian distance calculation. We used them to prevent the biased surfaces generated when VRIP distances are used. Besides, Euclidian distances improve the hole filling algorithm we used [8]. Only voxels near the view surfaces and not set in *empty* are evaluated. To find these voxels, we loop on each vertex of the view marking the voxels near them. This second pass has linear complexity on the number of views, instead of original CS, which is quadratic.

We calculate an estimated normal n for the current voxel v from the volumetric representation generated in the first pass. This normal will be used to validate the view data, discarding incorrect measurements and outliers. The normal n is calculated from the (x, y, z) gradients of volumetric representation. If the voxel does not have valid neighbors, n cannot be estimated. There are two possibilities in this case: if the boolean parameter *flagDiscardNoNormal* is true, the voxel is skipped; otherwise, this voxel will not be validated by the normal n , which can lead to outliers being accepted on this voxel. The rare cases that require *flagDiscardNoNormal* to be false are when there are very thin surfaces compared to the voxel size. In these cases, if *flagDiscardNoNormal* is true, holes are usually created in the reconstructed surface.

Next, the nearest point p' , its normal n' , weight w and signed distance d are evaluated. These values are calculated as in CS by using a *kd-tree* of view i to find the nearest vertex. The search for the nearest point is done on the faces incident on this vertex by calculating point to triangle distances. This process also tells us if the nearest point belongs to the border of view i . If p' is located on a border, this measurement is discarded to avoid the problem discussed on section 3.2. Measurements larger than $+D_{max}$ are also discarded. Finally, if an estimated normal n was calculated, the angle between n and n' is calculated. If this angle is larger than the threshold *consensusAngle*, the measurement is discarded. In our experiments, a value from 30° to 45° for the threshold *consensusAngle* returned good results. This procedure solves the flaw of VRIP presented in section 3.1. Besides, most of the deformed surfaces presented in section 2 are eliminated.

Following, the basic reliability w of the measurement (which depends on the angle between the scanner line-of-sight and the normal n') is altered by other 3 factors. The first one is the border weight w_{border} , where measurements near the borders of the view have lower weights than interior measurements. This weight is used to discard p' if it is too small. The second factor is w_{angle} , due to the angle between n and n' . This factor, ranging from 0.0 to 1.0, therefore $1.0 \geq (n.n') \geq \cos(\text{consensusAngle})$. The third factor $w_{distance}$ varies according to the signed distance d , like the weight curves from VRIP. Our IVIA approach uses a new weight curve as described in section 3.1. This curve allows a smooth and unbiased integration of the distance values. Finally, we integrate the measurements of all views through a weighted sum using the VRIP formula [6].

After the integration is completed, we have a last error elimination step. Neighboring voxels should have similar distance measurements because of the

use of Euclidian distances. Therefore, two neighbor voxels should have a maximum distance difference equivalent to the distance between the voxel centers. Using 26-neighborhood, the distances between voxel centers can be $voxelSize$, $\sqrt{2} * voxelSize$ or $\sqrt{3} * voxelSize$. However, we cannot be so restrictive, because the integration slightly violates this condition, due to the several weights used. Nevertheless, neighboring voxels with excessively different values should be discarded to avoid the generation of bad surfaces on the final model. Therefore, we use a threshold *compatibleFactor* (usually 1.5), which multiplied by the voxel center distance gives us the maximum allowed difference between neighboring voxels. An important detail is that this elimination cannot be done in a single pass, because a “wrong” voxel ends up spoiling its neighbors, since their difference would be too large. Because of that, we first gather all the “suspect” voxels, and sort them (using *bucketSort*) by the number of wrong neighbors (which can be 26 at maximum). Next, we eliminate the “most suspect” candidate, updating the suspect list to ignore neighborhood to voxels already eliminated. Therefore, the elimination is more selective, discarding only the really incompatible voxels. These voxels are tagged as unknown (distance value of $+D_{max}$ and weight 0).

Our IVIA algorithm deals with noise and outliers in several ways. With space carving, outliers far from the surface are eliminated. Both outliers near the surface and measurements out of consensus with the first pass are also eliminated. Measurements have their reliability assessed through several parameters to help decrease the influence of any outliers that had not been removed previously. Finally, after integration we eliminate the remaining incompatible data.

5 Experimental Results

We performed several experiments to assess the effectiveness of our method for applications that demand high fidelity 3D reconstructions, such as digital preservation of natural and cultural assets. The dataset used in the experiments is composed of art objects (from the Metropolitan Museum of Curitiba), fossils (from the Natural Science Museum of UFPR), insects (from the Biological Collections of UFPR), Baroque masterpieces (Aleijadinho’s sculptures) and personal objects. The results show how our method deals with difficult situations.

The ability of IVIA to detect and discard outliers can be seen in Fig. 1. The captured views have several artifacts because the statue material (marble) is not optically cooperative. Fig. 1b shows false data and deformed surfaces. In Fig. 1c we show the influence of the space carving on the view data, and the successful detection and discarding of artifacts. Fig. 1a shows result without hole filling.

In Fig. 2 we show a comparison of results from the three volumetric algorithms when applied in a difficult case. IVIA generated the most reliable result. VRIP and our modified CS returned outliers on the final result, even after a postprocessing step that discarded disconnected geometries. We disabled this post-processing on IVIA to show its robustness in removing artifacts. The result from our modified CS (see section 3.2) is similar to the ones from IVIA; however, the result from the original CS has lots of outliers on the final result. VRIP result

is good, except from the connected outliers and the biased surfaces, as well as the artifacts near corners and thin surfaces. In our IVIA result, the surface is as smooth as the VRIP one, but without the problems previously mentioned. The price we pay is slightly larger holes on regions of unreliable data.




Fig. 1. Outlier removal of the IVIA: (a) statue integrated with IVIA; (b) one of the views, with several types of bad data; (c) space carving technique where red/green as influenced regions by the w_{angle} , and purple as discarded regions




Fig. 2. Comparison of 3 integrations: (a) VRIP; (b) our modified CS; (c) IVIA. Several artifacts can be seen in (a) and (b) (arrows). In (c) IVIA eliminates almost all incorrect data and kept the surface smooth.

As mentioned IVIA was developed to cooperate with the holefilling algorithm of [8]. For that, IVIA performs a more effective outlier removal, uses Euclidian distances and space carving technique, producing more accurate results. In Fig. 3 we activate holefilling aiming to compare our results to PSR (original authors'). The IVIA's reconstructed model of Aleijadinho's Prophet Joel is showed in Fig 3a. Fig 3b shows the legs of Joel in details, where we can see that PSR reconstructed model (Fig 3b down) is over smoothed compared to our IVIA (Fig 3b up), we also noticed in the highlighted rectangle area that PSR creates outliers as vertices (see section 3.3). In Fig 3c we can see two detailed regions, Joel's hat (Fig 3c up) and Joel's vestment (3c down). We also tested IVIA on objects with smaller data sizes and two of them are shown in Fig 3d.

The computer used in all experiments was a 2.20GHz Core2 Duo PC, with 2 GB of RAM. Duck model had 67,772 faces, Wolf model had 206,627 faces, the marble statue (Fig 2) had approximately 1,300,300 faces and Joel model had 3,487,652 faces. To reconstruct Joel PSR took 3002s, VRIP took 2718s

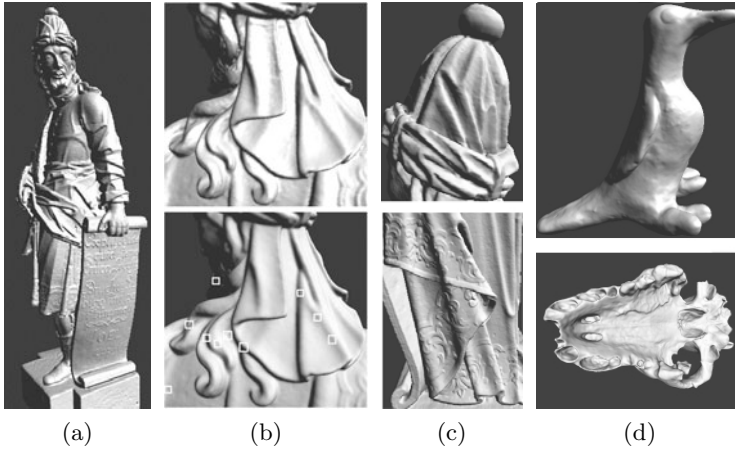


Fig. 3. Reconstructed models: (a) Prophet Joel; (b) comparison between IVIA and Poisson reconstructions of Joel's legs; (c) details of IVIA reconstruction of Joel's hat (up) and vestment (down); (d) Duck and Wolf models

and IVIA 11265s. However, the time performance of IVIA, can still be further optimized using 3D scan conversion on the space carving stage, the stage when the algorithm spends most of its execution time.

6 Final Remarks

Despite the good experimental results, IVIA may still be improved. Its use of VRIP on the first pass can lead to outliers being accepted later, since outliers can survive the first pass. What we did was to reduce the outlier influence by mixing them with good measurements. However, in very noisy regions, this methodology may fail. One possibility is to use feedback and more passes to improve outlier detection and removal between passes. An alternative is to combine parametric surfaces [14] with a volumetric representation and our outlier detection and removal techniques. Our work shows that more research is needed to guarantee high quality digital reconstruction results. We showed that real input range data has complex types of artifacts, and that widely used volumetric integration methods may still have major limitations. We proposed enhancements to both VRIP and CS, besides presenting our novel IVIA algorithm, which achieved precise reconstruction results even when compared to PSR.

Acknowledgments. The authors would like to thanks to CNPq, CAPES, UNESCO and IPHAN for supporting this research.

References

1. Bernardini, F., Rushmeier, H.: The 3D model acquisition pipeline. *Computer Graphics Forum* 21(2), 149–172 (2002)
2. Ikeuchi, K., Oishi, T., et al.: The Great Buddha project: Digitally archiving, restoring, and analyzing cultural heritage objects. *IJCV* 75, 189–208 (2007)

3. Vrabel, A., Bellon, O., Silva, L.: A 3D reconstruction pipeline for digital preservation of natural and cultural assets. In: Proc. of CVPR, pp. 2687–2694 (2009)
4. Levoy, M., Pulli, K., et al.: The Digital Michelangelo project: 3D scanning of large statues. In: SIGGRAPH, pp. 131–144 (2000)
5. Miyazaki, D., Oishi, T., Nishikawa, T., Sagawa, R., Nishino, K., Tomomatsu, T., Takase, Y., Ikeuchi, K.: The Great Buddha project: Modelling cultural heritage through observation. In: Proc. of VSMM, pp. 138–145 (2002).
6. Curlless, B., Levoy, M.: A volumetric method for building complex models from range images. In: Proc. SIGGRAPH, pp. 303–312 (1996)
7. Wheeler, M.D., Sato, Y., Ikeuchi, K.: Consensus surfaces for modeling 3D objects from multiple range image. In: Proc. of ICCV, pp. 917–924 (1998)
8. Davis, J., Marschner, S.R., Garr, M., Levoy, M.: Filling holes in complex surfaces using volumetric diffusion. In: Proc. of 3DPVT, pp. 42–438 (2002)
9. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: Proc. of SIGGRAPH, pp. 61–70 (2006)
10. Edelsbrunner, H.: Shape reconstruction with Delaunay complex. In: Proc. of Latin Amer. Symp. Theoretical Informatics, pp. 119–132 (1998)
11. Turk, G., Levoy, M.: Zippered polygon meshes from range images. In: Proc. of SIGGRAPH, pp. 311–318 (1994)
12. Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C., Taubin, G.: The ball-pivoting algorithm for surface reconstruction. In: IEEE TVCG, pp. 349–359 (1999)
13. Sharf, A., Lewiner, T., Shamir, A., Kobbelt, L., Cohen-Or, D.: Competing fronts for coarse-to-fine surface reconstruction. In: Computer Graphics, pp. 389–398 (2006)
14. Ohtake, Y., Belyaev, A., Alexa, M., Turk, G., Seidel, H.P.: Multi-level partition of unity implicits. *ACM Transactions on Graphics* 22, 463–470 (2006)
15. Fleishman, S., Cohen-Or, D., Silva, C.T.: Robust moving least-squares fitting with sharp features. *ACM Transactions on Graphics* 24, 544–552 (2005)
16. Bolitho, M., Kazhdan, M., Burns, R., Hoppe, H.: Parallel poisson surface reconstruction. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Kuno, Y., Wang, J., Wang, J.-X., Wang, J., Pajarola, R., Lindstrom, P., Hinkenjann, A., Encarnaç o, M.L., Silva, C.T., Coming, D. (eds.) ISVC 2009, Part I. LNCS, vol. 5875, pp. 678–689. Springer, Heidelberg (2009)
17. Bolitho, M., Kazhdan, M., Burns, R., Hoppe, H.: Multilevel Streaming for Out-of-Core Surface Reconstruction. In: Proc. Eurographics (2007)
18. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. In: Proc. of SIGGRAPH, pp. 163–169 (1997)
19. Hilton, A., Stoddart, A.J., Illingworth, J., Windeatt, T.: Reliable surface reconstruction from multiple range images. In: Buxton, B.F., Cipolla, R. (eds.) ECCV 1996. LNCS, vol. 1064, pp. 117–126. Springer, Heidelberg (1996)
20. Sagawa, R., Nishino, K., Ikeuchi, K.: Robust and adaptive integration of multiple range images with photometric attributes. In: Proc. of CVPR, pp. II:172–II:179 (2001)
21. Sagawa, R., Ikeuchi, K.: Taking consensus of signed distance field for complementing unobservable surface. In: Proc. of 3DIM, pp. 410–417 (2003)
22. Masuda, T.: Object shape modelling from multiple range images by matching signed distance fields. In: Proc. of 3DPVT, pp. 439–448 (2002)
23. Rocchini, C., Cignoni, P., Ganovelli, F., Montani, C., Pingi, P., Scopigno, R.: The marching intersections algorithm for merging range images. *Visual Computer* 20(2-3), 149–164 (2004)