

Quantitative Comparison of Hole Filling Methods for 3D Object Search

M. A. Rojas^{1,2} F. M. Sukno^{1,2} J. L. Waddington² and P. F. Whelan¹

¹Center for Image Processing and Analysis, Dublin City University, Dublin, Ireland

²Molecular and Cellular Therapeutics, Royal College of Surgeons in Ireland, Dublin, Ireland

Abstract

Retrieval of 3D models has become an important issue due to the increase in the number of digitized objects that are available in many different fields. When stored data presents defects such as holes, accurate and reliable repairing tools are needed to solve these issues. In this work we present a comparative evaluation of hole filling algorithms from the local and global perspective, measuring quantitatively the quality of the repaired meshes and describing the impact these tools have on the models. We do this by mapping holes from one mesh onto another in order to create a synthetic dataset with realistic holes and ground truth and use the Hausdorff and RMS distance, as well as the mean angular deviation, to quantify the errors. The results show that the performance at a local level is similar for all compared methods, but large differences (up to 20%) appear when viewed at a global level, where algorithms that use volumetric representations introduce significant changes in the original models.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representation

1. Introduction

Digital 3D modelling has become an important component in a variety of fields ranging from purely scientific to industrial sectors. This in turn has led to the generation of search tools that are capable of organizing, managing and retrieving samples of models in these areas. Retrieving a sample from a database can be done using global or partial matching or a combination of both. Global matching methods are robust against small changes in the model and to topological degeneracies, but may not be able to capture subtle shape differences that might be needed in certain applications [TV08]. Local matching methods are capable of encoding shape details and can perform local matching, but local shape descriptors are sensitive to the model geometric and topological conditions. Furthermore, the reliability and accuracy of the results of the queries performed on search tools depends on specific domains. For instance, applications of morphometrics that study the relation between craniofacial structure and biological disorders, require model descriptors with sufficient discriminative power to separate models that exhibit differences in the order of millimetres [HBB*10, CBN*12].

Our interest lies in the analysis of facial structure and the subtle differences that appear therein as a consequence of

mental disorders associated with early developmental disruption [HBB*10]. To this end we have acquired a set of 3D meshes of the face with a hand-held laser scanner. While this technology allows very accurate reconstruction of most surfaces [STD09, BF05], it also suffers from some specific limitations. Holes on the object model can be generated in cases where there is poor reflection of the laser (e.g. the eyebrows) or in regions where the surface bends in such a way that the reflections can not be captured by the camera (e.g. the ears).

In this case, given the subtle differences between models (3D surfaces of patients and controls), global matching approaches do not have the accuracy required to discriminate samples from one another. On the other hand, local approaches require local shape descriptors, which may be sensitive to geometric and topological differences. For instance, descriptors such as Spin Images [JH99] are affected by holes because the missing points cannot contribute to the computed histogram. Figure 1 shows an example of the variation between the histogram in the presence and absence of a hole where the difference between descriptors can reach up to 30% of the maximum value of the descriptor. For this reason the availability of complete, clean models is an impor-

tant requirement for applications that require highly accurate data.

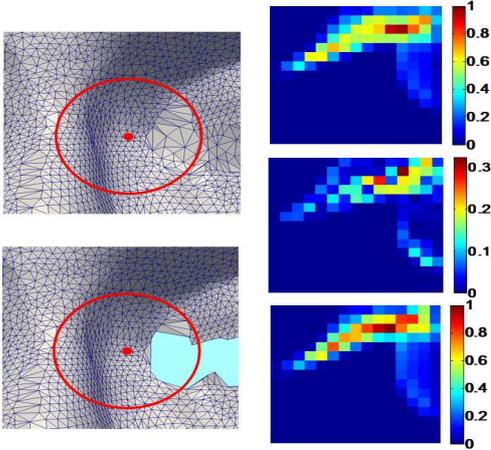


Figure 1: Example of the Spin Image descriptor for the inner eye corner. Right column: Descriptor in absence of hole (top), difference (middle) and descriptor in the presence of hole (bottom). Regions on which descriptors were extracted are illustrated in the left column.

In this work we are interested in comparing and evaluating algorithms for hole filling on triangulated meshes acquired by a laser scanner. Our goal is to quantitatively measure the difference between the solutions generated by different algorithms and the original meshes. This is made feasible by a strategy of hole mapping, where real holes from one captured surface are mapped onto another surface that has been captured in the same way, but is free of holes in the region of interest.

Despite the large amount of hole filling algorithms in the literature, the evaluation of these methods is usually restricted to aesthetic aspects (such as subjective evaluations of visual quality), to the level of correspondence of the generated patch with the sample density of the surrounding mesh (enumerating the number of points, triangles and their size) or processing time. Many of these algorithms are typically evaluated on a reduced number of models and, although these models would enable direct comparison with ground truth data, this type of evaluation is rarely done. Furthermore, in synthetically generated models it might be difficult to reproduce holes in a realistic manner.

With this work we intend to contribute to the research community by providing an insight into the features and limitations of state-of-the-art hole filling algorithms with respect to their impact on the final mesh, and provide a clear quantitative measure of performance of the compared algorithms in the context of real, accuracy-sensitive data.

The work is presented as follows. Firstly, we establish a

context of the state-of-the-art in hole filling and its evaluation. Secondly, we describe the data set employed and we proceed to generate the synthetically enhanced dataset by mapping holes from one mesh to another. Thirdly, we evaluate the performance of the algorithms on a set of surfaces that contain real holes. Finally we present the conclusions.

2. Related Work

Different hole filling algorithms have been presented in the literature which can be classified in two main groups: global and local approaches.

Global approaches use an intermediate representation of the model to perform hole filling, usually consisting of a volumetric grid. Other possible representations include Octrees [Ju04] and their combinations with graphs [PR05]. In general, the idea is to associate a signed value distance to each voxel depending on whether it is inside or outside of the model’s volume. The surface is reconstructed by separating the voxels with different signs [CL96, DMGL02, NT03, FIMK07].

Subsequently the hole filling operation can be performed by iterative diffusion principles (Davis et al. [DMGL02]), applying Marching Cubes on the mesh (Guo et al. [GLWZ06]) or fitting quadratics to the signed distance function (Masuda et al. [Mas04]). It can also be done by shrinking the area of the hole using anisotropic partial differential equations (Verdera et al. [VCBS03]) or iteratively changing the signs of the voxels of an interpolated surface (Sakawa and Ikeuchi [SI08]).

Global approaches have few requirements with respect to the input meshes; a condition these methods may have is that of a single connected object to guarantee a global solution. Furthermore, by using an intermediate representation these methods can solve non-manifoldness and topological issues with no ambiguities. However, the same intermediate representation may lead to aliasing degrading the appearance of the model and causing internal structures to be replaced by a coarse outer hull.

In local approaches the first step is usually to detect boundary edges and define the holes. Once the boundaries have been defined, the contours of the holes may be projected on 2D to perform the hole filling [TC04, WO07, BWS*09]. These approaches work in cases where holes do not have a complex geometry at the boundary. When the projections present intersections, holes can be divided using criteria based on fitted polynomial blending [LYZ08, LMW10].

The patching triangulation can then be generated using a variety of mechanisms. Liepa [Lie03] minimizes the maximal dihedral angle, Hu et al. [HWLL12] apply the principle of minimum angle between adjacent edges of the hole boundary, Zhao et al. [ZGL07] apply the Advancing Front Mesh technique (AFM), Branch et al. [BPB06]

use Radial Basis Functions and Pfeife and Seidel [PS96] use B-Splines. Finally, the generated patches are smoothed using fairing procedures like thin plate energy functional minimization (Pfeife and Seidel [PS96]), umbrella operator (Liepa [Lie03]), energy minimization based on mechanical models (Pernot et al. [PMV06]), Poisson PDE (Zhao et al. [ZGL07]) or bilateral filtering (Hu et al. [HWLL12]).

In the case of local approaches, a common assumption is that of input manifold meshes, with no singularities on the hole boundaries (i.e. no two adjacent holes share a vertex) and no islands. An important requirement of these methods is that of orientable surfaces, at least in the vicinity of the hole. Some of the approaches are capable of dealing with complex-geometry holes, but holes with irregular boundaries represent a common limitation. Local methods may generate triangulations with intersections, degenerate elements and may even be unable to fill holes completely (i.e. when the generated triangulation generates non-manifolds). By operating on local regions of the surface, they modify the models far less than global methods. For a more exhaustive and thorough survey on general 3D mesh repairing algorithms the reader is referred to [Ju09, ACK13].

2.1. Hole Filling Evaluation

The evaluation of repairing algorithms, and in particular of hole filling, is usually done in terms of the number of elements and/or points introduced [PR05, LMW10, WO07, HWLL12, WLG03], the areas or quality of the triangles created [FBG98, ZGL07, HWLL12], visual quality of the resulting mesh [GLWZ06, DMGL02, Lie03, PMV06, LMW10] or processing time [SOS04, NT03, CL96, LMW10, WLG03]. To the best of our knowledge, to date there has not been a quantitative evaluation of hole filling algorithms in terms of the numeric quality of the patches produced (e.g. metric differences between surfaces). This point is important for a number of applications in medicine such as forensics and morphometrics or reverse engineering, where a reliable estimation of missing information can be important for further processing of the data.

3. Methodology

The pipeline of the work presented here for comparison and evaluation consists of four blocks: Data acquisition, hole mapping, hole filling and evaluation.

3.1. Data Acquisition

The data used in this work consists of 144 3D laser scans of the face acquired with a hand held Polhemus scanner at the Royal College of Surgeons in Ireland [HS02]. This scanner can achieve resolutions of 0.1 mm (0.5 when scanning at a distance of 200 mm) with an accuracy of 0.13 mm [STD09, BF05].

3.2. Hole Mapping

We wish to compare the patches generated by hole filling algorithms to some reference or ground truth. Evidently, in the presence of a hole, the actual surface data is missing. As a solution, we generate our ground truth by mapping holes between pairs of facial surfaces. Let $S_1 \in \mathbb{R}^3$ and $S_2 \in \mathbb{R}^3$ be two surfaces that can be related by a mapping $f_{12} : S_1 \rightarrow S_2$. When captured, these surfaces generate the triangulated meshes \mathcal{M}_1 and \mathcal{M}_2 , respectively. Let $h_1 \subset S_1$ be a portion of the surface S_1 that is not captured, becoming a hole of \mathcal{M}_1 , which we indicate by:

$$\mathcal{M}_1 \cap h_1 = \partial h_1, \quad S_1 \cap h_1 = h_1 \quad (1)$$

where ∂h_1 is the boundary of the hole, which is the only information available about it from \mathcal{M}_1 . If the hole h_1 maps into a hole-free region of \mathcal{M}_2 , i.e. if

$$\mathcal{M}_2 \cap f_{12}(h_1) = f_{12}(h_1) \quad (2)$$

then we get a synthetically generated hole $h_1^{(2)} = f_{12}(h_1)$ with ground truth $\mathcal{M}_2 \cap h_1^{(2)}$. If the mapping is accurate, the generate hole $h_1^{(2)}$ has both realistic shape and location on \mathcal{M}_2 as much as h_1 does on \mathcal{M}_1 .

To estimate the mapping f_{12} we combine the conformal mappings of S_1 and S_2 . Given that the face is (approximately) a genus-0 surface, it can be mapped conformally into the 2D domain. The conformality condition ensures that the angles are locally preserved, hence minimizing mapping distortion. At least two corresponding points are needed to make the mapping unique, but additional points can be added to increase robustness and produce a Least Squares Conformal Mapping [LPRM02].

In this work, we used a set of 26 anatomical landmarks there were manually annotated on each of the input surface by experts in human anthropometry [HKW02]. By setting the 2D coordinates of these landmarks to a fixed position for all meshes, we obtain a common 2D domain $\Omega \in \mathbb{R}^2$ and piecewise linear mappings from it to each surface, e.g. $g_1 : \Omega \rightarrow S_1$ and $g_2 : \Omega \rightarrow S_2$.

Thus, we can relate S_1 and S_2 by $f_{12} = g_2 \circ g_1^{-1}$. This is possible because if a mapping is conformal, its inverse is also conformal. In case a hole from \mathcal{M}_1 is partially mapped to a hole on \mathcal{M}_2 , this mapping is discarded.

From our initial set of 144 facial scans, we randomly selected pairs of surfaces and produced the mapping of its holes as explained above. We obtained 200 scans which contain over 800 holes. Figure 2 shows an example of the process described above.

3.3. Hole Filling

We have chosen examples of tools and algorithms that use local and global approaches. We compared tools that are publicly available and are able to process all images with

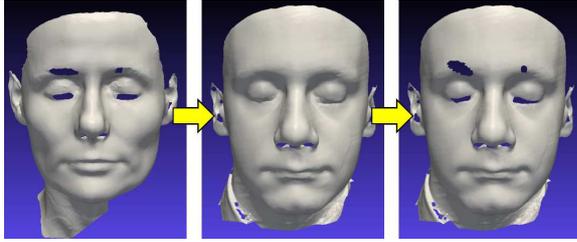


Figure 2: Mapping the holes from one surface to another. Left: Source mesh from which to map the holes; Middle: Ground truth (target) mesh; Right: Result of the hole mapping process. Hole deformation in the target mesh results from the difference in the geometry (the target face is larger and more elongated), positions of the sampled points and triangulations between source and target meshes.

no user interaction. We have also implemented two state-of-the-art algorithms [ZGL07] and [HWLL12] based on local approaches.

- **Volumetric Diffusion - VF:** The method by Davis et al. [DMGL02] first converts the input model into a volumetric grid defined in the immediate vicinity of the surface by a distance function d and an associated weighting function which provides a measure of confidence on the values of d . The hole filling is done by an iterative diffusion process that extends d inwards across the holes until they are closed. The final surface is obtained applying the marching cubes [LC87] algorithm. This method redefines the triangulation and the sampling of the original mesh. The authors have made the **VolFill** tool available at <http://graphics.stanford.edu/software/volfill/>.
- **Reconstruction by Contouring - PM:** Tao [Ju04] presented a method that generates a closed volume following three steps: scan conversion, sign generation and surface reconstruction. In the first step, the model is converted to an Octree that is built incrementally and contains information on intersection edges. The following step proceeds to generate cells with signs that are consistent with intersection edges. In this way, cells that cross the model show a change in sign. The final surface reconstruction is achieved separating grid points with opposite signs, generating a closed object model with redefined sampling density and triangulation. The Author has made the **PolyMender** tool available at <http://www.cse.wustl.edu/~taoju/code/polymender.htm>.
- **MeshFix - MF:** Attene and Falcidieno [AF06] presented a tool to repair triangulated meshes that generates closed models representing a single connected manifold with sampling density varying as little as possible from the input. The algorithm for hole filling is based on that of Liepa [Lie03], where triangulation is obtained from the weighting function that minimizes the maximal dihedral

angle. Subsequently, a refining step is applied by relaxing the edges to comply with Delaunay criteria. The generated triangulation is smoothed by modifying the positions of the vertices to minimize the normal-field variation. The Authors have made their **MeshFix** tool available at <http://sourceforge.net/projects/meshfix/>.

- **Advancing Front Mesh - Z12:** Zhao et al. [ZGL07] presented an algorithm that fills holes locally in three steps: Hole identification, AFM and fairing based on the Poisson equation. Once the hole has been defined, the advancing front mesh technique adds iteratively a given number of triangles, based on the angle between adjacent edges at each vertex. The generated patching mesh is smoothed by manipulating the gradient field of the mesh instead of directly manipulating the coordinates. The diffusion equation is used to deform the gradient field and the Poisson linear system is solved to derive a surface that matches the modified gradient field. This method does not modify the original triangulation or the sampling density.
- **Minimum Angle - H12:** Hu et al. [HWLL12] proceed to locally fill holes in a three step scheme: minimum angle triangulation, refinement and smoothing. After identifying the holes, triangles are iteratively added based on the angle between adjacent edges for each vertex. The difference from the method of Zhao et al. is that triangles are added considering the consistency of the normals of neighbouring vertices. This triangulation is subsequently refined by iteratively splitting the triangles at the centroid and applying edge relaxation based on Delaunay criteria. Finally, bilateral filtering is applied on the direction of normals to smooth the patching mesh. The values for the angles and filter parameters were kept as those proposed in the paper. The density control parameter of the refining step was set to 1.75 instead of 1.41 given the resolution of the meshes used. This method does not modify the original triangulation or the sampling density.
- **Mesh Processing Tool - MP:** Grimm and Phan [GP] presented a tool for mesh processing (i.e. simplification, denoising, curvature computation) and repair that can fill holes locally with a simple patching scheme without manipulating the rest of the surface. It can take as parameter the maximum size of hole to fill and patches the hole by fitting a polygon with the same number of vertices as the boundary of the hole, but with as few new vertices as possible; in fact, unless the hole is too complex, the algorithm generates patches that triangulate the hole from the centroid of the boundary to each of its vertices. The Authors have made their **Manifold Mesh Processing** tool available at <http://sourceforge.net/projects/meshprocessing/>.

4. Evaluation

This work aims to compare the quality of hole filling algorithms at a quantitative level. To compare the patched surface with the original, we adopt the concepts of discrete differ-



Figure 3: Example of captured surface of a mannequin and results of the described methods. Top Row from left to right: Original Object, **VF**, **MF**, **Z12**. Bottom row from left to right: Original Surface, **PM**, **H12**, **MP**.

ential geometry presented by Hildebrandt et al. [HPW06], where the notions of total normal convergence are shown to be equivalent to the convergence of surface area, intrinsic metrics and Laplace-Beltrami operators.

In particular, we are interested in defining error measures and, according to the theoretical framework presented in [HPW06], point wise convergence of a polyhedral surface can be measured by the Hausdorff distance.

Let S be the original surface and S' be the patched surface. Given a point p in S , the distance between p and S' can be defined as:

$$d(p, S') = \min_{p' \in S'} \|p - p'\|_2 \quad (3)$$

where $\|*\|_2$ denotes the Euclidean distance. The forward Hausdorff distance between S and S' can be defined as:

$$d(S, S') = \max_{p \in S} d(p, S') \quad (4)$$

This is not a symmetric distance, i. e., $d(S, S') \neq d(S', S)$; thus the symmetric Hausdorff distance can be defined as:

$$d_s(S, S') = \max\{d(S, S'), d(S', S)\} \quad (5)$$

In our experiments we used the **Metro** tool [CRS98] to compute the Hausdorff distance between meshes and a MATLAB function to measure the distance between patches. In practice the Hausdorff distance is sensitive to outliers and may provide a disrupted measure of disparity between the surfaces. In contrast, the Root Mean Square distance is less affected by this type of situations. For every point p in S and a corresponding matching point p' in S' it can be defined as:

$$RMSE(S, S') = \sqrt{\frac{\sum_{i=1}^N (p_i - p'_i)^2}{N}} \quad (6)$$

A super-sampling scheme was implemented for comparison of the methods at patch level, in order to better quantify the differences between surfaces. This is particularly important given that some of the methods change the triangulation of the repaired mesh. An area-based scheme was implemented in which a given triangle is divided if its area is K times larger than a predefined threshold α . Two cases are considered: if the triangle can be divided in three triangles of area approximately equal to α , only one new point at the centroid is added; if it can be divided in $4, 9, \dots, (k+1)^2$ triangles of an area approximately equal to α , then $3k + \sum_{i=1}^k i - 1$ new points are added in such a way that a regular grid is formed by triangles of equal area. Figure 4 shows examples of the generated grid with 3, 4 and 25 triangles. Taking into account that the average maximum area of the triangles in the data set used is around $15mm^2$, we have set $\alpha = 1.5mm^2$. Furthermore, as the average face



Figure 4: Sampling scheme performed on triangles of $4.5, 6$ and $37.5 mm^2$

has a surface area of $560cm^2$, we have focused the analysis on holes that are larger than 2 average size triangles (an area of $30mm^2$ or 0.05% of the facial surface), which in turn represent 45% of the total amount of holes generated. The sampling scheme was applied to both the original and the reconstructed meshes. Moreover, given the definitions of the distance measures applied, we have used a *KD-tree* and performed a nearest-neighbor search on the super-sampled point clouds to find correspondences between points in the ground truth and the reconstructed meshes.

5. Results

All the algorithms analyzed fill the synthetically generated holes by the procedure described in section 3.2 as well as those present in the original meshes, which are product of the capture process. We present results from two perspectives. The first considers the performance at patch level comparing only the set of generated triangles that overlap with those removed from the ground truth. The second considers the performance at the mesh level comparing the entire generated mesh against the original one. Figures 5, 6 and 9 show the box-plots for all six methods for both measures.

In general, the performance at patch level is similar between the methods, with the exception of **PM**. This is consistent for both the HD and RMS measures. In the case of **PM** the use of an Octree representation may cause self-intersections and generate large local distortions. Furthermore, the comparatively large Inter-Quantile-Range (IQR)

in both measures, as well as in the mean angle deviation (shown in figure 9), indicates low stability and consistency in the solutions generated by the method (Figure 5). **Z12** also underperforms when compared to other methods as well. This can be explained by the sensitivity of the algorithm with respect to the parameters that regulate the size of the triangles and, in turn, the density of the generated patching mesh. If the value is too low with respect to the triangulation of the boundary of the hole, small triangles will be created and the AFM algorithm may generate rippled surfaces. Moreover, if the length of the boundary edges of the hole is not uniform, the algorithm may not converge. In the case of

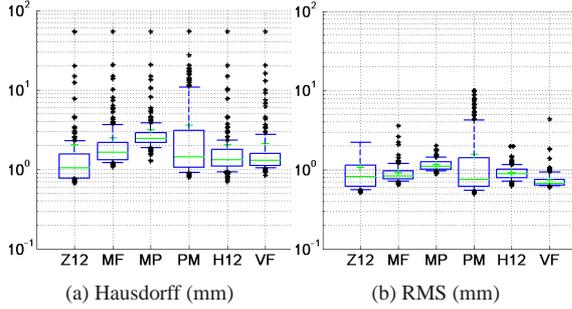


Figure 5: Measures per Patch

performance at mesh level, two groups can be clearly distinguished representing the local (**Z12**, **MP** and **H12**) and the global (**MF**, **PM** and **VF**) approaches - see figure 6. Two of the local methods show a similar performance in terms of the consistency of their solutions (**MP** and **H12**). This can be explained by the initial rough estimation they make of the patching solution. However, as **MP** stops there, **H12** proceeds to the refinement of the triangulation, which is still bounded to the initial estimation. This further processing explains the small differences in performance between these two methods. On the other hand, the difference in perfor-

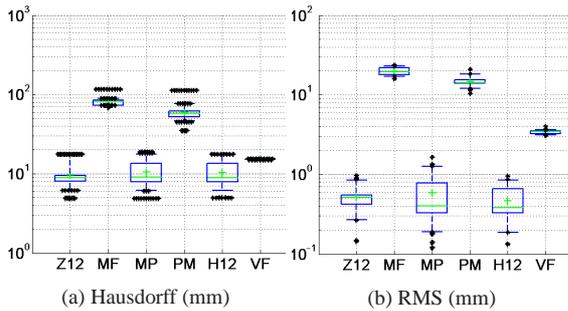


Figure 6: Measures per Mesh

mance between the local and global approaches is explained by the form of the final solution, which in two cases (**MF** and **PM**) is a closed object. Between these, the difference lies in the added portion of the the mesh. **PM** closes the surface

with a piece-wise planar approximation, reducing the area of the added portion, whereas **MF** generates an additional larger surface that renders a smooth transition between the boundaries of the original object, as shown in figure 7. In the case of **VF** the intermediate performance can be explained by the fact that, although the algorithm uses a volumetric intermediate representation, the diffusion process is only applied to boundary regions. As this covers holes within the surface, it also includes other boundary regions where the diffusion may not converge. Thus, as the algorithm modifies these regions, they show significant differences with the ground truth, as shown in figure 8. The quality of the patches

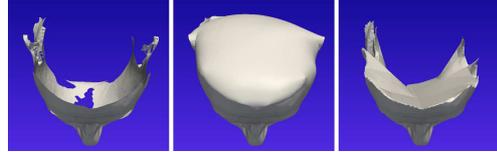


Figure 7: Original reverse side of a mesh and additional surface added by **MF** (center) and **PM** (right)

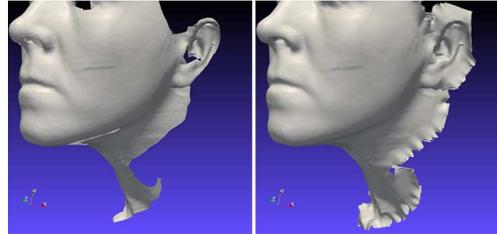


Figure 8: Original and output of **VF**. Notice the excess surface produced

can also be evaluated in terms of the angular deviation between the normals of the ground truth and the solution generated by each method, as shown in figure 9. Given the localized nature of the patches, this can be observed better at a region level. Again, the performance of **MP** is lower than for the other methods. This is due to the basic approach used to generate the patching mesh. For instance, in those cases where the hole is not in a planar region, the method estimates the centroid of the hole and generates a poor quality triangulation that has little resemblance to the original surface. This can be seen as the highest median of the angular deviation among methods. In the case of **PM** the large IQR in the angular deviation reflects the instability and lack of consistency mentioned before. Regarding the performance of the methods with respect to the hole size, it can be seen that there is a gradual degradation as the regions requiring a patch grow larger. This is particularly well captured by the Hausdorff distance - see figure 10. Specifically it can be seen that **MP** performs rather poorly from this perspective.

In terms of angular deviation the performance is slightly

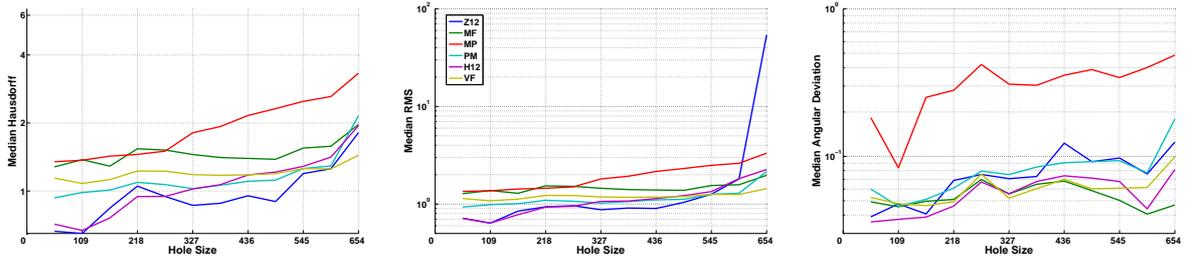


Figure 10: Hausdorff, RMS and angular deviation as a function of the size of the holes (mm^2).

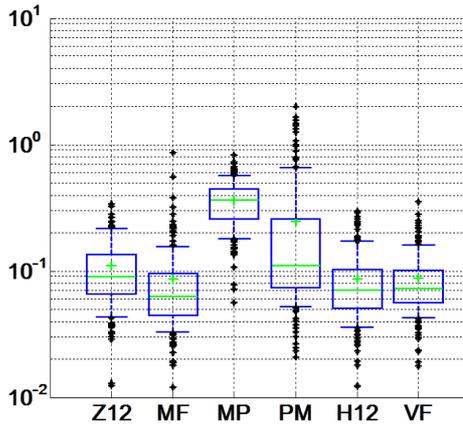


Figure 9: Mean angle deviation per patch (rads)

more stable for **MF** than for the other methods, which is consistent with the results shown in figures 5 and 9. It should be pointed out that these results reflect the quality of the output of the algorithms at a local level; thus, deformations on the rest of the object are not considered.

6. Conclusions

We have presented a comparative evaluation of state-of-the-art methods for hole filling of triangulated meshes, including our implementation of two recent methods. To compare and evaluate the methods, we have generated a set of ground truth surfaces by mapping holes between meshes. We implemented a sampling scheme which is a versatile and realistic tool to measure quantitatively the quality of hole filling algorithms.

The measures used consider the performance of the methods from the local and global point of view. This provides quantitative means of assessing the capabilities of a given method, and their suitability for a specific application. The results obtained show that the performance of all methods is similar at a local level, with small differences due to the simplicity of the approach or the sensitivity to the param-

eters. At the global level, the large differences are the result of the amount of variation to which the original model is subjected. Those algorithms that produce closed object models are seen as more disruptive for the types of models used because, by changing the triangulation and sampling density, they introduce significant changes in the original data.

Among the analyzed methods, the algorithm by Hu et al. [HWLL12] has the best balance between performance and triangulation manipulation. The method is capable of filling complex geometry holes and generates a triangulation similar to that of the neighbourhood of the hole, while introducing a limited amount of distortion. The second best option is that of Attene and Falcidieno [AF06] which generates a triangulation similar to one of the hole neighbourhood. This method generates a closed model that, although does not alters the existing original triangulation, modifies the overall geometry of the object by adding the part of the surface that "closes" the object.

As part of future work, we expect to make publicly available a set of models similar to the one used in this work, where real holes can be mapped on equivalent meshes. In this way researchers will be able to compare the performance of repairing algorithms with realistic ground truth.

References

- [ACK13] ATTENE M., CAMPEN M., KOBELT L.: Polygon mesh repairing: An application perspective. *ACM Computing Surveys* 45, 2 (2013), 15. 3
- [AF06] ATTENE M., FALCIDIENO B.: Remesh: An interactive environment to edit and repair triangle meshes. In *SMI* (2006), IEEE, pp. 41–41. 4, 7
- [BF05] BOEHNEN C., FLYNN P.: Accuracy of 3d scanning technologies in a face scanning scenario. In *3DIM* (2005), IEEE, pp. 310–317. 1, 3
- [BPB06] BRANCH J., PRIETO F., BOULANGER P.: Automatic hole-filling of triangular meshes using local radial basis function. In *3D Data Processing, Visualization, and Transmission, Third International Symposium on* (2006), IEEE, pp. 727–734. 2
- [BWS*09] BRUNTON A., WUHRER S., SHU C., BOSE P., DE-MAINE E. D.: Filling holes in triangular meshes by curve unfolding. In *Proceedings SMI*, (2009). 2

- [CBN*12] CHINTHAPALLI K., BARTOLINI E., NOVY J., SUT-
TIE M., MARINI C., FALCHI M., FOX Z., CLAYTON L. M.,
SANDER J. W., GUERRINI R., ET AL.: Atypical face shape and
genomic structural variants in epilepsy. *Brain* 135, 10 (2012),
3101–3114. 1
- [CL96] CURLESS B., LEVOY M.: A volumetric method for build-
ing complex models from range images. In *Proceedings of the
23rd annual conference on Computer graphics and interactive
techniques* (1996), ACM, pp. 303–312. 2, 3
- [CRS98] CIGNONI P., ROCCHINI C., SCOPIGNO R.: Metro:
Measuring error on simplified surfaces. In *Computer Graphics
Forum* (1998), vol. 17, Wiley Online Library, pp. 167–174. 5
- [DMGL02] DAVIS J., MARSCHNER S. R., GARR M., LEVOY
M.: Filling holes in complex surfaces using volumetric diffusion.
In *3D Data Processing Visualization and Transmission, Proceed-
ings. First International Symposium on* (2002), IEEE, pp. 428–
441. 2, 3, 4
- [FBG98] FREY P. J., BOROUCI H., GEORGE P.-L.: 3d
delaunay mesh generation coupled with an advancing-front ap-
proach. *Computer methods in applied mechanics and engineer-
ing* 157, 1 (1998), 115–131. 3
- [FIMK07] FURUKAWA R., ITANO T., MORISAKA A.,
KAWASAKI H.: Improved space carving method for merg-
ing and interpolating multiple range images using information
of light sources of active stereo. In *ACCV*. Springer, 2007,
pp. 206–216. 2
- [GLWZ06] GUO T.-Q., LI J.-J., WENG J.-G., ZHUANG Y.-T.:
Filling holes in complex surfaces using oriented voxel diffusion.
In *MLC*. (2006), IEEE, pp. 4370–4375. 2, 3
- [GP] GRIMM C., PHAN L.: Man-
ifold mesh processing tool. URL:
<http://sourceforge.net/projects/meshprocessing/>.
4
- [HBB*10] HENNESSY R. J., BALDWIN P. A., BROWNE D. J.,
KINSELLA A., WADDINGTON J. L.: Frontonasal dysmorphol-
ogy in bipolar disorder by 3d laser surface imaging and geometric
morphometrics: comparisons with schizophrenia. *Schizophrenia
research* 122, 1 (2010), 63–71. 1
- [HKW02] HENNESSY R. J., KINSELLA A., WADDINGTON
J. L.: 3d laser surface scanning and geometric morphometric
analysis of craniofacial shape as an index of cerebro-craniofacial
morphogenesis: initial application to sexual dimorphism. *Biolog-
ical psychiatry* 51, 6 (2002), 507–514. 3
- [HPW06] HILDEBRANDT K., POLTHIER K., WARDETZKY M.:
On the convergence of metric and geometric properties of poly-
hedral surfaces. *Geometriae Dedicata* 123, 1 (2006), 89–112.
5
- [HS02] HENNESSY R. J., STRINGER C. B.: Geometric morpho-
metric study of the regional variation of modern human cranio-
facial form. *American journal of physical anthropology* 117, 1
(2002), 37–48. 3
- [HWLL12] HU P., WANG C., LI B., LIU M.: Filling holes in tri-
angular meshes in engineering. *Journal of Software* 7, 1 (2012),
141–148. 2, 3, 4, 7
- [JH99] JOHNSON A. E., HEBERT M.: Using spin images for ef-
ficient object recognition in cluttered 3d scenes. *TPAMI* 21, 5
(1999), 433–449. 1
- [Ju04] JU T.: Robust repair of polygonal models. *ACM TOG* 23,
3 (2004), 888–895. 2, 4
- [Ju09] JU T.: Fixing geometric errors on polygonal models: a sur-
vey. *Journal of Computer Science and Technology* 24, 1 (2009),
19–29. 3
- [LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high
resolution 3d surface construction algorithm. In *ACM Siggraph
Computer Graphics* (1987), vol. 21, ACM, pp. 163–169. 4
- [Lie03] LIEPA P.: Filling holes in meshes. In *Proceedings Eu-
rographics/ACM SIGGRAPH* (2003), Eurographics Association,
pp. 200–205. 2, 3, 4
- [LMW10] LI Z., MEEK D. S., WALTON D. J.: Polynomial
blending in a mesh hole-filling application. *Computer-Aided De-
sign* 42, 4 (2010), 340–349. 2, 3
- [LPRM02] LÉVY B., PETITJEAN S., RAY N., MAILLOT J.:
Least squares conformal maps for automatic texture atlas gen-
eration. *ACM TOG* 21, 3 (2002), 362–371. 3
- [LYZ08] LI G., YE X.-Z., ZHANG S.-Y.: An algorithm for
filling complex holes in reverse engineering. *Engineering with
Computers* 24, 2 (2008), 119–125. 2
- [Mas04] MASUDA T.: Filling the signed distance field by fitting
local quadrics. In *Proceedings 3DPVT*. (2004), IEEE, pp. 1003–
1010. 2
- [NT03] NOORUDDIN F. S., TURK G.: Simplification and repair
of polygonal models using volumetric techniques. *Visualization
and Computer Graphics, IEEE Transactions on* 9, 2 (2003), 191–
205. 2, 3
- [PMV06] PERNOT J.-P., MORARU G., VÉRON P.: Filling holes
in meshes using a mechanical model to simulate the curvature
variation minimization. *Computers & Graphics* 30, 6 (2006),
892–902. 3
- [PR05] PODOLAK J., RUSINKIEWICZ S.: Atomic volumes for
mesh completion. In *Symposium on Geometry Processing*
(2005), Citeseer, pp. 33–41. 2, 3
- [PS96] PFEIFLE R., SEIDEL H.-P.: Triangular b-splines for
blending and filling of polygonal holes. In *Proceedings of the
conference on Graphics interface '96* (1996), Canadian Informa-
tion Processing Society, pp. 186–193. 3
- [SI08] SAGAWA R., IKEUCHI K.: Hole filling of a 3d model by
flipping signs of a signed distance field in adaptive resolution.
TPAMI 30, 4 (2008), 686–699. 2
- [SOS04] SHEN C., O'BRIEN J. F., SHEWCHUK J. R.: Interpo-
lating and approximating implicit surfaces from polygon soup.
In *ACM TOG* (2004), vol. 23, ACM, pp. 896–904. 3
- [STD09] SANSONI G., TREBESCHI M., DOCCHIO F.: State-of-
the-art and applications of 3d imaging sensors in industry, cul-
tural heritage, medicine, and criminal investigation. *Sensors* 9, 1
(2009), 568–601. 1, 3
- [TC04] TEKUMALLA L. S., COHEN E.: A hole-filling algorithm
for triangular meshes. *School of Computing, University of Utah,
UUCS-04-019, UT, USA* (2004). 2
- [TV08] TANGELDER J. W., VELTKAMP R. C.: A survey of con-
tent based 3d shape retrieval methods. *Multimedia tools and ap-
plications* 39, 3 (2008), 441–471. 1
- [VCBS03] VERDERA J., CASELLES V., BERTALMIO M.,
SAPIRO G.: Inpainting surface holes. In *Proceedings. ICIIP 2003*.
(2003), vol. 2, IEEE, pp. II–903. 2
- [WLG03] WAGNER M., LABSIK U., GREINER G.: Repairing
non-manifold triangle meshes using simulated annealing. *Inter-
national Journal of Shape Modeling* 9, 02 (2003), 137–153. 3
- [WO07] WANG J., OLIVEIRA M. M.: Filling holes on locally
smooth surfaces reconstructed from point clouds. *Image and Vi-
sion Computing* 25, 1 (2007), 103–113. 2, 3
- [ZGL07] ZHAO W., GAO S., LIN H.: A robust hole-filling algo-
rithm for triangular mesh. *The Visual Computer* 23, 12 (2007),
987–997. 2, 3, 4