

Image guided reconstruction of un-sampled data: a filling technique for Cultural Heritage models

Matteo Dellepiane · Andrea Venturi · Roberto Scopigno

Received: date / Accepted: date

Abstract Cultural Heritage (CH) is one of the major fields of application of 3D scanning technologies. In this context, one of the main limitations perceived by the practitioners is the incompleteness of the sampling. Whenever we scan a complex artifact, the produced sampling usually presents a large number of unsampled regions. Many algorithmic solutions exist to close those gaps (from specific hole-filling algorithms to the drastic solution of using water-tight reconstruction methods). Unfortunately, adding patches over un-sampled regions is an issue in CH applications: if the 3D model should be used as a master document over the shape (and status) of the artwork, informed CH curators usually do not accept that an algorithm is used to *guess* portions of a surface.

In this paper, we present a low-cost setup and related algorithms to reconstruct un-sampled portions of the 3D models by inferring information about the real shape of the missing region from photographs. Data needed to drive the surface completion process are obtained by coupling a calibrated pattern of laser diodes to a digital camera. Thus, we are proposing a simple active acquisition device (based on consumer components and more flexible than standard 3D scanning devices)

to improve selectively the sampling produced by a standard 3D scanning device.

After acquiring one or more images with the laser-enhanced camera, an almost completely automatic process analyzes the image/s in order to extract the pattern, to estimate the laser projector intersections over the surface and determining coordinates of those points (using the consolidated triangulation approach). Then, the gathered geometric data are used to steer the hole filling in order to obtain a patch which is coherent with the real shape of the object. A series of tests on real objects proves that our method is able to recover geometrical features that cannot be reconstructed using state-of-the-art methods. Consequently, it can be used to obtain complete 3D models without creating plausible but false data.

Keywords Hole filling · Laser triangulation · Image registration

1 Introduction

The last few years of technological development made 3D Scanning a mature technology. This is a due not only to the improvements of the acquisition devices, but also to the development of new algorithms which automatize most of the required raw data processing, and fully exploit the resources of graphics hardware.

One of the most important fields of application of 3D scanning is Cultural Heritage. There are several reasons behind this: the sampling accuracy and speed granted by this technology, the variety of different objects that can be acquired, the possibility to acquire data without touching the objects, the wide number of possible uses of the digital 3D models (from archival and restoration support to rapid prototyping).

M. Dellepiane
Via G. Moruzzi 1, 56124 Pisa (PI) Italy
Tel.: +39-0503152925
Fax: +39-0503152604
E-mail: dellepiane@isti.cnr.it

R. Scopigno
Via G. Moruzzi 1, 56124 Pisa (PI) Italy
Tel.: +39-0503152929
Fax: +39-0503152604
E-mail: scopigno@isti.cnr.it



Fig. 1 A rendering of an incomplete model, where part of the geometry could not be acquired with laser scanner (missing regions are easily detectable as blue holes) and a photograph of the real artifact.

While 3D scanning is able to provide an extremely accurate representation of an object, the final digital 3D model can suffer from a number of geometric artifacts (singular vertices, overlaps/folds, self-intersections, complex edges, holes). Some of them can be generated during the acquisition of the single range maps, while others are introduced by the several post-processing actions performed over the raw data.

Artifacts removal is a quite important subject in the Computer Graphics and Computer Vision research: while some types of artifacts can be easily detected and removed, others can present difficulties in both detection and removal. Holes are a common and unavoidable artifact in scanned meshes. It is very rare to see a *complete* 3D model of a complex object, obtained using 3D scanning, since usually not all the parts of a complex object can be reached by the scanner. Moreover, it is also possible that some parts are not covered due to errors or carelessness during the acquisition campaign.

Presence of holes in a 3D model implies that we have regions with missing data; if we do not want to fill up those regions with interpolated data, their presence can be distracting and annoying during visualization. Hence, several methods to automatically fill the holes were presented in literature, including both methods to fill up selectively one gaps at the time or reconstruction methods that produce water-tight models. Most of these approaches rely on the analysis of the existing geometry of the model: even in the case of holes of medium and large size, the new geometry can be almost indistinguishable from the rest of the model.

Nevertheless, hole filling is a very critical operation in Cultural Heritage applications, since the geometry created is a plausible but essentially *guessed* information. On the contrary, the approach of Cultural Heritage ex-

perts usually tries to avoid the creation of *not reliable* data. Hence, the final geometry is usually kept incomplete in the master representation, unless the hole filling is key for further applications (e.g. rapid prototyping).

We present an approach that aims to fill the holes of an incomplete mesh by extracting information about the missing geometry from one or more images, obtained with a simple and low-cost active lighting setup. The data are obtained by analyzing images of the real object, on which a pre-defined colored pattern (which could be generated by a set of laser diodes or by a portable projector) is projected on the zone which was not covered by the scanner. Using a digital camera, it is much easier to frame partially occluded regions, since the camera and the grid of laser diodes can be arranged with a flexible and much smaller inter-distance. An example of an incomplete model is presented in Figure 1, where a number of small regions of the capital were not be sampled properly (in most cases, since they were hidden either to the emitter or to the sensor of the triangulation-based scanning device). The same regions can be easily framed using a digital camera. Nevertheless, without further information it can be hard to infer data about the geometry of the object. Conversely, the triangulation of the projected laser probes gives additional information that can be used to find a shape-congruent patch to close the hole.

In our approach, each single geometric patch is obtained with a semi-automatic process. The first step is to align the image on the 3D model (this is the only manual step in our current implementation, but automatic solutions are possible also for this phase [7]). Then, the reflected laser pattern is extracted from the image and, given the knowledge of the calibrated projector position wrt.

the camera parameters, we are able to estimate the geometry of several points lying on the object surface and corresponding to the reflected pattern. These geometric data are then used, together with the hole open border and the nearby surface to reconstruct a patching sub-mesh which should more faithfully represent the real surface.

The paper is organized as follows. An overview of the hole filling solutions and issues is presented in Section 2. The requirements and solutions for the image acquisition setup are discussed in Section 3, while Section 4 describes the hole filling process. Some results and a comparison with state-of-the-art methods are shown in Section 5. Finally, conclusions and future work are presented in Section 6.

2 Related Work

Hole filling has been an important field of research in recent years. Closing holes of small size is a trivial operation. Unfortunately, 3D scanning produces many holes (usually from hundreds to thousands in the case of complex models) and many of them have a large extension, a high number of border edges and a complex shape. Examples of the complexity of the possible holes are presented in [8]. In this case, the filling operation can be very long and prone to errors.

Most of the proposed approach assume no information about the missing geometry. Hence, this information is mainly extracted from the shape of the model around the hole. These hole filling techniques can be roughly divided in two groups: volumetric and surface oriented. In the context of the volumetric techniques, the model or the sampled data are usually immersed in a volumetric grid and all the holes are filled while reconstructing the entire shape. These methods essentially use regular [19,12,8] or adaptive [5] grids. They are usually quite robust, able to deal with isles and high resolution models, since they essentially derive from the *Marching Cubes* approach [25]. Nevertheless, there is a substantial remeshing of the geometry, and all the holes are filled (they cannot be managed selectively). These modifications could not be wanted by the user.

The surface oriented algorithms work on a single hole at a time, and the new information is usually inferred by analyzing the shape of the surface near the border of the hole. A simple but effective method was proposed by Varnuška et al. [21]. Other approaches use Radial Basis Functions [6] or Moving Least Squares [2] to obtain a smooth filling surface. Another example is the method proposed by Liepa [16], that is able to provide very good results. It is divided in three main phases: in the first one, a rough triangulation that fills the hole is

found, then the remeshing and smoothing phases produce an almost undistinguishable new surface.

Another subgroup of surface oriented approaches extract information from the analysis of all the geometry of the model. Bendels [4] and Sharf [23] extend the idea of image inpainting and texture synthesis by analyzing the model and finding surfaces whose shape is similar to the ones near the border of the hole. Then, the hole is filled using patches of already existing geometry (a recent work on the topic was presented by Becker [3]). Park et al [20] use a similar idea to reconstruct both shape and appearance from scanned data.

Surface oriented approaches can obtain good results and act selectively on single holes. Unfortunately, the complexity of the algorithm can be very high when the number of border edges increases.

As already stated, the methods cited until now don't assume knowledge of the real shape the missing surface. An alternative idea is to infer missing information from an external source: Akbarzadeh et al [1] combine a large set of different data to complete a model of a city landscape. Analogously, it is possible to use data obtained from photogrammetry [22] to fill the holes in architectural models. Nevertheless, the acquisition of these integration data can be quite cumbersome.

Xu [26] proposed a much simpler source of data: a single image of the object. It is used to estimate a set of normal vectors, which generates a smooth surface which should recover the real shape of the object. Unfortunately, the results shown in the paper are obtained using renderings of an uncomplete model. No example using real images is shown.

3 The acquisition setup

The aim of the proposed work was to provide a method which should be able to fill the holes of a 3D model using a few easily gathered data about the real shape of the missing part. This is a key issue for the use of such an algorithm in the context of Cultural Heritage (but also in other important application fields, like industrial prototyping).

Moreover, to allow easy deployment for CH applications, we decided to design a cheap and flexible setup for the acquisition of additional data. Digital cameras are cheap and flexible devices. Acquiring an image is a very fast operation, and a large amount of data is obtained. Moreover, the easier handling of digital cameras and the possibility to use different optics allows to frame portions of an object which cannot be easily reached by a scanner.

Unfortunately, as the results of the work of Xu [26] showed, gathering 3D information from a plain image




Fig. 2 Three examples of possible projected patterns.

is still an hard task, especially when it is necessary to discriminate between geometry and appearance in uncontrolled environments. For this reason, we decided to add a laser-based pattern projector to the camera, transforming it in a simple active acquisition device: by processing the reflected pattern, we will be able to infer data on the missing surface geometry. Accurate patterns can be generated even by low-end laser diodes. Therefore, it is possible to build a compact and cheap acquisition device formed by a digital camera and a laser diode mounted on a position displaced by the camera optical axis (to allow to compute geometric triangulation). With this device, the user could take photos of the parts of the object which were not covered by the scanner, and then the images can be used to complete the 3D model. We used a laser projector to generate the patterns, but obviously all the processing pipeline is clearly independent on the type of pattern emitter adopted.

3.1 Issues and advantages of the proposed acquisition setup

One of the reasons for the presence of holes in models acquired via 3D Scanning is the fact that part of the geometry cannot be seen by both the emitter and the sensor of the scanning device. Using a system with a smaller distance (or a more flexible assembly) can improve our capabilities to reduce those unsampled regions. Using different 3D scanners with different working volumes is usually not an option, due to the need to reduce the hardware cost. The enhanced active photographic acquisition setup presented above has several advantages with respect to the adoption of a classic scanning device:

- **Smaller cone of blindness:** since the pattern emitter is fixed to a digital camera, their distance could

be small. This generates a smaller cone of blindness (only a few centimeters) with respect to the one of a 3D scanning device.

- **Improved maneuverability:** the simplicity and lightness of the acquisition device makes it much more handy than a traditional 3D scanner. Some not sampled regions are due to the impossibility to move a scanner to some peculiar scanning positions (most scanners require a stable and heavy tripod and cable connection to a PC). When using a digital camera, it is extremely easy to frame the object from different points of view, the device is very light and it can be rotated arbitrarily on the view axe. The example in Figure 1 shows that digital cameras can easily frame portions that are unreachable for a 3D scanner.
- **Redundancy of data:** the use of a generic pattern, the absence of filters on the camera objective and the necessity to estimate the camera and emitter position (see next sections for detail) can influence the accuracy of the device. Nevertheless, since we already dispose of accurate geometric data (result of the previous 3D scanning acquisition) around the unsampled region, this redundancy can be used to correct errors and improve the quality of the data reconstructed from the photographic images (see Section 4.4).

An issue in the design of the acquisition device is the type of pattern to project on the model. Three examples are shown in Figure 2. It was decided to consider only patterns which are inscribed in a pre-defined shape (like the square in the central and righthand examples of Figure 2). This was extremely important for the definition and robustness of the image processing and projector position estimation steps (see next Section).

The output of the acquisition setup is one or more images of the object on which a (pre-defined) pattern is projected. The necessary steps to extract information

which will be used to fill the hole are presented in next Section.

3.2 Alternative techniques limitations

In this subsection we will discuss the limitations of two possible alternative solutions to the proposed approach. In particular, we will focus on their current limitations, which prevents their use for the aim presented in this paper.

Dense stereomatching. Recent advances in reconstruction from images techniques have brought to the creation of extremely reliable and robust system. An example is the web-based reconstruction service presented by Vergawuen [24]. These systems start from a set of uncalibrated images: applying a wide a set of image processing and stereo matching techniques, they are able to reconstruct an accurate geometry. Other recent approaches proved to be robust and reliable [11, 10]. Nevertheless, their possible use for geometry completion is limited by two main factors:

- Dense matching: in order to obtain an accurate reconstruction, every portion of the geometry needs to present detectable features in the highest possible number of images. The intrinsic nature of the un-sampled parts is that there are very few points of view (if not only one) from which they can be framed. This can limit the number of cases in which dense stereo matching can be used.
- Scaling factor: unless a calibration object is present in the scene, or an accurate measure of a known element is available, reconstructed scenes are not scaled to the real measures. Hence it should be necessary to find a reliable way to scale the reconstructed geometry to the real measures, using the data acquired with 3D Scanning.

Z depth cameras. Recent advances brought to the creation of depth cameras which, using essentially a time-of-flight technique, are able to assign a depth value to the pixels of an image. These devices can be extremely useful in the context of tracking, interaction or simplified modeling [15]. Unfortunately, their precision (the mean error is in the order of one or more centimeters) and current cost make them a not usable, while promising, solution.

In conclusion, current technology can't provide a cheap and reliable solution to the proposed problem.

4 Data processing and hole filling

The data processing pipeline needed to fill the holes starting from the image can be divided in the following steps:

- For each image framing the hole region:
 - image registration on the 3D model (image-to-geometry transformation);
 - extraction of the pattern pixels from the image;
 - estimation of the position and orientation of the pattern projector;
- Hole filling, using the information gained with the projected pattern points.

All of the steps are either automatic or semi-automatic. The whole filling operation is completed within minutes, taking into account both user intervention and processing time.

The next subsection will present each processing step and the intermediate data extracted.

4.1 Image-to-geometry alignment

In order to be able to reconstruct the missing geometry, it is necessary to estimate the position of the elements of the acquisition setup. The first step is to find the position of the digital camera wrt. the sampled artifact. This is obtained by registering the image to the digital 3D model.

There are several automatic and semi-automatic methods to do this [14,17]. In our paper, we used a semi-automatic approach [9] which allows to obtain the estimation of the camera parameters (intrinsics and extrinsics) in a couple of minutes. The knowledge of the camera position is critical in order to be able to estimate the projector position and to transform the extracted 2D information (pattern pixels location) into points in 3D space.

4.2 Pattern extraction from image

In addition to the camera position, it is necessary to analyze the image in order to extract further information about the projected pattern. This is done by applying some common image processing filters to find and extract the pattern.

The first operation is to find the pattern pixels in the context of the image. Although this can be done automatically, in the current implementation the user is asked to pick the four external corners of the pattern




Fig. 3 An example of pattern extraction from an image. Top: the original image. Left: the portion of the original image depicting the pattern. Center: the extracted red channel. Right: the extracted pattern after the application of the erosion filter.

on the image. These user-defined data are also used to enhance the projector position estimation. Once that the user has selected the four corners, the bounding square of the pattern is extracted (see Figure 3-left).

Then, the further processing is operated only on the red channel (Figure 3-center). Since we want to associate each line of a pattern to a single line of pixels in the image, an erosion filter is applied in order to extract the central line of pixels for each line of the pattern (Figure 3-right).

4.3 Pattern projector position estimation

One possible design for the acquisition device, as described in Section 3, is to couple a digital camera with a small projector (e.g. a laser diode) firmly mounted on it. In this case, the calibration of the device should be done only once in its lifetime (for example, by taking pictures of a calibration object and calculating the relative distance between the camera and the projector). Known the camera-projector calibration and the image-to-geometry alignment to the 3D model (see previous subsection), the position of the projector can be automatically computed.

Nevertheless, we decided to take into account a more general case, where the projector position is not known in advance. The estimation of the projector position is a key element for the accuracy of the overall geometry reconstruction process. The data extracted in the previ-

ous steps are used in order to perform it. Moreover, it's necessary to have a model of the pattern (essentially the divergence angle of the projected grid). In the case of a laser diode, it can be inferred from the specifications, otherwise it can be easily calculated by measuring the size of the projected grid at different known distances. The projector position is estimated using an implementation [18] of the *Levenberg Marquardt* method, which is a robust MSE estimator for nonlinear problems. The input data to the solver are the model of the projector and a number (16 in our tests) of points chosen from the external lines of the pattern extracted from the image (see previous subsection). The pattern pixels on the image are projected on the 3D model (using the camera data obtained from the image-to-geometry registration phase). Then, the algorithm iteratively fits the projector model to the set of 3D points until convergence. This method proved to be quite robust and fast, since convergence is usually obtained in seconds; its accuracy was assessed with the analysis of the results presented in the next section. As a matter of fact, small errors in position estimation would generate visible inaccuracies in geometry reconstruction.

4.4 Geometry reconstruction and hole filling

The previous steps extracted from the image all the data which are necessary to drive the hole filling process. Knowing the camera and the projector positions and orientations, and the coordinates of the pixels of the image associated to the pattern, it is possible to apply triangulation (in the same fashion as the laser and structured light triangulation 3D scanners) in order to associate a 3D position to each pattern pixel of the image.

Hence, a new 3D point can be created for each pattern pixel which maps on the hole zone. Before starting geometry reconstruction, we take advantage of the redundancy of data in order to improve the accuracy. In a first stage, only the pixels which map on the already existing geometry of the model are taken into account. For each pixel, the 3D coordinates obtained via triangulation and the projection of the pixel on the 3D model are compared. The value of the differences in depth of the two positions (averaged over all the considered pixels) is used as a correction factor that takes into account the possible inaccuracy introduced in the projector position estimation phase.

Finally, the reconstruction of the surface can start. For each pattern pixel not projecting on the 3D model, a 3D position is calculated (taking into account also the correction factor). Starting from the new 3D points,

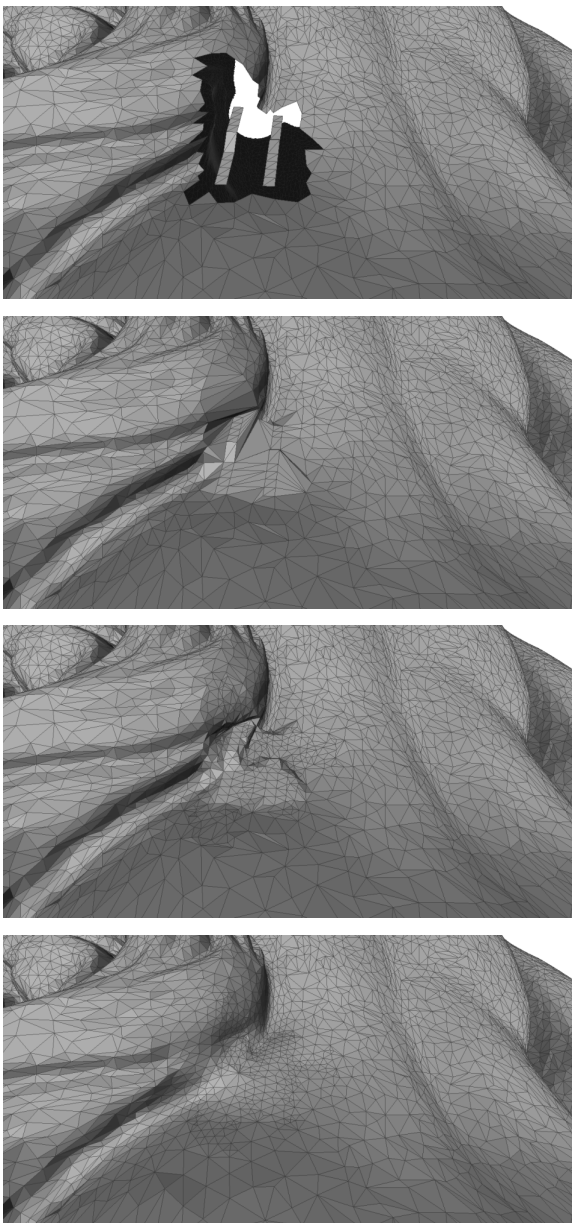


Fig. 4 The hole filling procedure. From top to bottom: (a) the hole with the new triangle strips, (b) rough triangulation, (c) re-meshing, (d) smoothing

several strips of new triangles are created (see Figure 4a) and then connected to the hole borders. This is the key step in the whole procedure: the hole filling problem is split in several simpler hole filling problems. But the splitting is obtained by inserting information about the real shape of the object. In this way, eventual features which were completely absent can be (at least partially) retrieved.

After this very important stage, the hole filling algorithm is very similar to the one proposed by Liepa [16]: in a first stage, a rough filling is obtained by finding minimum-area spanning surface which also takes into

account the dihedral angles between adjacent triangles (4b). Then, a remeshing phase splits the newly created triangles until their average size is similar to the one of the rest of the model (4c). Finally, a smoothing filter makes the new surface almost undistinguishable from the rest of the model (4d).

In conclusion, once that the input data have been analyzed, the hole filling operation is completely automatic, and the quality of the results is extremely good.

4.5 Extension of the system to multiple images

The system presented in the previous sections can be easily extended to the use with multiple images. This can be extremely useful in order to retrieve more information about the missing geometry, especially if a peculiar pattern, like the one used for the comparison (next Section), is available. If it's possible to use multiple images, the user can take several shots of the same portions, rotating the camera and slightly changing the position in order to cover the whole missing surface. Then, the processing pipeline for multiple images is the same, for each image, until the creation of the points in the missing region (first part of Section 4.4). The reconstruction of the geometry from the reconstructed points can be obtained using a surface reconstruction method like the one presented by Kazhdan [13]. In this way, all the data can contribute in the creation of a coherent filling.

5 Results and comparison

The previous Section showed that the quality of the reconstructed surface is comparable to the results obtained from the state-of-the-art hole filling methods. But the main goal of the method was to be able to recover geometric features that couldn't be reconstructed without some knowledge of the shape of the object.

Hence, a set of test cases was created to compare the results of our method with other techniques, not only in terms of visual pleasantness, but especially in terms of geometric accuracy. In order to perform this, a "ground truth" reference was needed. Hence, we created artificial holes in several complete models (acquired with 3D scanning technology) and, then, we compared the results obtained with our method and with two alternative methods: Earcut [21] and Liepa [16]. The first one is a very simple but effective method, where the final triangulation is usually quite rough. The second method, while more complex, provides visually pleasant fillings.




Fig. 5 Four of the tests performed to validate our method. (a) the original mesh, (b) the hole created in the mesh, (c) Liepa [16] result, (d) EarCut [21] result, (e) our method result

Algorithm	Average	Maximum	RMSE
Test 1			
Liepa	0,002083	3,284013	0,052673
EarCut	0,000593	0,687241	0,01462
Our method	0,001201	1,708783	0,0295663
Test 2			
Liepa	0,515288	33,177967	3,379488
EarCut	0,000553	0,79688	0,015343
Our method	0,000222	0,642542	0,007969
Test 3			
Liepa	0,497988	33,177967	3,32634
EarCut	0,00122	1,67301	0,027455
Our method	0,000405	1,048663	0,013196
Test 4			
Liepa	0,002006	3,284013	0,051815
EarCut	0,000294	0,485467	0,008804
Our method	0,000493	1,919416	0,018686

Table 1 Analysis of Hausdorff distance between the original and the reconstructed models

The tests were created in order to cover different possible cases, including the one where no particular geometric feature is missing and the one where important data are partially or totally missing. The results of two tests are shown in Figure 5. In the second one, a visible geometric feature is removed from the model: the result obtained with Liepa has a smooth shape, but information about the geometric features is totally lost. The same can be observed on the EarCut result, with the addition of a lower quality triangulation. Our method is instead able to partially recover the detail. In the last case, all methods obtain similar results.

Visual comparison is not enough to validate the accuracy of our method. For this reason, we calculated the Hausdorff distance between the original and the reconstructed models. This distance allows to measure the actual accuracy of the newly created geometries. The results of four test sets are shown in Table 1. The first observation is that the Liepa results are the less accurate ones: the error is introduced by the remeshing and smoothing phase, which improves the appearance of the surface, but introduces a quite high error. The triangulation obtained by EarCut is quite similar to the original mesh, but it is often visually unpleasant. Our method offers a very good compromise between accuracy and appearance: while producing smooth surfaces, the error is always comparable (if not lower) to the EarCut.

Considering also that our method shares the triangulation, remeshing and smoothing phase with Liepa, this further shows that, since almost no error is added, this is due to the fact that the information extracted from images is really a representation of the shape of the object. Otherwise, the Hausdorff distance value would be

comparable. Given this numerical data, it could be possible to think about different approaches in the remeshing and smoothing phase, in order to preserve the sharp details.

6 Discussion, future work and conclusions

The presented system is a compact and easy to use way to close portions of surfaces which cannot be reached by laser scanner. It is mainly based on triangulation measurement; hence, in principle it suffers from the same weak-points of triangulation devices. But most of them can be overcome by three main advantages: the framed area, the handiness and the redundancy of data. The fact that a much smaller area is framed and reconstructed greatly reduces the cone of occlusion. Moreover, the digital camera is much more handy than a scanner, hence most of the occlusion issues can be solved by moving or rotating it. Its handiness permits also to reach areas which cannot be reached by scanners.

The previous Section showed that the use of our simple acquisition setup can lead to a coherent completion of un-sampled geometry. In particular, it was proved that it is possible to extract accurate geometric information using a simple image. While for the testing part only one image at a time was used, it is possible to integrate the data coming from multiple images in order to improve accuracy in reconstruction.

Other possible future improvements can be found. First of all, considering the decrease in costs of laser diodes with custom patterns, the first goal is to build a low-cost device on which our system can be used in practice. This could lead to a more in-depth test on a set of more complex real cases.

Regarding the improvement of the processing part, the image processing can be made completely automatic, and more robust filters can be experimented in order to extract the pattern from the images. An alternative solution could be the use of a filter on the camera objective (like the ones used by laser triangulation scanner) in order to further facilitate the pattern identification and extraction.

Finally, other applications can be found for this system. For example, the pattern can be used to improve the accuracy of model acquired with lower-accuracy devices (i.e. time-of-flight scanners). Once that an important detailed part is found, several images can be taken in order to improve the quality of the acquired mesh by deforming the geometry in order to fit the data extracted from the images.

In conclusion, a simple, semi-automatic and robust method

to complete un-sampled data (especially in the field of Cultural Heritage) was presented. This system can guarantee the reliable completion of incomplete 3D model, so that they can be used for a wide range of applications, from mere visualization to rapid prototyping.

Acknowledgements The research leading to these results has received funding from the European Community's Seventh Framework Programme (/FP7/2007-2013)/, in the context of 3D-COFORM project, under /grant agreement /n 231809 and from Tuscany Regional Project STArT. We would like to thank Paolo Cignoni for his supervision, and Federico Ponchio and Massimiliano Corsini for their help.

References

1. Akbarzadeh, A., Frahm, J.M., Mordohai, P., Clipp, B., Engels, C., Gallup, D., Merrell, P., Phelps, M., Sinha, S., Talton, B., Wang, L., Yang, Q., Stewenius, H., Yang, R., Welch, G., Towles, H., Nister, D., Pollefeys, M.: Towards urban 3d reconstruction from video. In: 3D Data Processing, Visualization, and Transmission, Third International Symposium on, pp. 1–8 (2006). DOI 10.1109/3DPVT.2006.141
2. Amenta, N., Kil, Y.J.: Defining point-set surfaces. *ACM Trans. Graph.* **23**(3), 264–270 (2004)
3. Becker, J., Stewart, C., Radke, R.: Lidar inpainting from a single image. In: International Workshop on 3-D Digital Imaging and Modeling (3DIM) (2009)
4. Bendels, G.H., Schnabel, R., Klein, R.: Detail-preserving surface inpainting. In: M. Mudge, N. Ryan, R. Scopigno (eds.) The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage, pp. 41–48. Eurographics Association, Pisa, Italy (2005)
5. Bischoff, S., Pavic, D., Kobbelt, L.: Automatic restoration of polygon models. *ACM Trans. Graph.* **24**(4), 1332–1352 (2005)
6. Carr, J.C., Beatson, R.K., Cherrie, J.B., Mitchell, T.J., Fright, W.R., McCallum, B.C., Evans, T.R.: Reconstruction and representation of 3d objects with radial basis functions. In: SIGGRAPH, pp. 67–76 (2001)
7. Corsini, M., Dellepiane, M., Ponchio, F., Scopigno, R.: Image-to-geometry registration: a mutual information method exploiting illumination-related geometric properties. *Computer Graphics Forum* **28**(7), 1755–1764 (2009). URL <http://vcg.isti.cnr.it/Publications/2009/CDPS09>
8. Davis, J., Marshner, S., Garr, M., Levoy, M.: Filling holes in complex surfaces using volumetric diffusion. In: First Int. Symp. on 3D Data Processing, Visualization and Transmission (3DPVT'02), pp. 428–438. IEEE Comp. Soc. (2002)
9. Franken, T., Dellepiane, M., Ganovelli, F., Cignoni, P., Montani, C., Scopigno, R.: Minimizing user intervention in registering 2D images to 3D models. *The Visual Computer* **21**(8-10), 619–628 (2005). URL <http://vcg.isti.cnr.it/Publications/2005/CDFGMS05>
10. Furukawa, Y., Ponce, J.: Accurate, dense, and robust multi-view stereopsis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **32**(8), 1362–1376 (2010). DOI 10.1109/TPAMI.2009.161
11. Goesele, M., Snavely, N., Curless, B., Hoppe, H., Seitz, S.: Multi-view stereo for community photo collections. pp. 1–8 (2007)
12. Ju, T.: Robust repair of polygonal models. *ACM Trans. Graph.* **23**(3), 888–895 (2004). DOI <http://doi.acm.org/10.1145/1015706.1015815>
13. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing, pp. 61–70. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2006)
14. Lensch, H.P.A., Heidrich, W., Seidel, H.P.: Automated texture registration and stitching for real world models. In: PG '00: Proceedings of the 8th Pacific Conference on Computer Graphics and Applications, p. 317. IEEE Computer Society, Washington, DC, USA (2000)
15. Liao, M., Zhang, Q., Wang, H., Yang, R., Gong, M.: Modeling deformable objects from a single depth camera. In: ICCV: Proceedings of the 12th IEEE International Conference on Computer Vision. Kyoto, Japan (2009)
16. Liepa, P.: Filling holes in meshes. In: SGP '03: Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing, pp. 200–205. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2003)
17. Liu, L., Stamos, I., Yu, G., Wolberg, G., Zokai, S.: Multiview geometry for texture mapping 2d images onto 3d range data. *Computer Vision and Pattern Recognition* **02**, 2293–2300 (2006). DOI <http://doi.ieeecomputersociety.org/10.1109/CVPR.2006.204>
18. Lourakis, M.: levmar: Levenberg-marquardt nonlinear least squares algorithms in C/C++. [web page] <http://www.ics.forth.gr/~lourakis/levmar/> (Jul. 2004)
19. Nooruddin, F., Turk, G.: Simplification and repair of polygonal models using volumetric techniques. *IEEE Transactions on Visualization and Computer Graphics* **9**(2), 191–205 (2003)
20. Park, S., Guo, X., Shin, H., Qin, H.: Shape and appearance repair for incomplete point surfaces. In: *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, vol. 2, pp. 1260–1267 Vol. 2 (2005). DOI 10.1109/ICCV.2005.218
21. Parus, M.V.J., Kolingerová, I.: Simple holes triangulation in surface reconstruction. In: A. Handlovicová (ed.) *Algorithmy 2005*. Slovak University of Technology (2005)
22. Przybilla, H.: Fusion of terrestrial laserscanning and digital photogrammetry. In: IEVM06 (2006)
23. Sharf, A., Alexa, M., Cohen-Or, D.: Context-based surface completion. *ACM Trans. Graph.* **23**(3), 878–887 (2004). DOI <http://doi.acm.org/10.1145/1015706.1015814>
24. Vergauwen, M., Gool, L.V.: Web-based 3d reconstruction service. *Mach. Vision Appl.* **17**(6), 411–426 (2006)
25. William E. Lorensen, H.E.C.: Marching cubes: A high resolution 3d surface construction algorithm. *Computer Graphics* **21**(4) (1987)
26. Xu, S., Georghiadis, A., Rushmeier, H., Dorsey, J., McMillan, L.: Image guided geometry inference. In: 3DPVT '06: Proceedings of the Third International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT'06), pp. 310–317. IEEE Computer Society, Washington, DC, USA (2006). DOI <http://dx.doi.org/10.1109/3DPVT.2006.81>