

Efficient Moving Point Handling for Incremental 3D Manifold Reconstruction

Andrea Romanoni^(✉) and Matteo Matteucci

DEIB, Politecnico di Milano, Via Ponzio 34/5, 20133 Milano, Italy
{andrea.romanoni,matteo.matteucci}@polimi.it

Abstract. As incremental Structure from Motion algorithms become effective, a good sparse point cloud representing the map of the scene becomes available frame-by-frame. From the 3D Delaunay triangulation of these points, state-of-the-art algorithms build a manifold rough model of the scene. These algorithms integrate incrementally new points to the 3D reconstruction only if their position estimate does not change. Indeed, whenever a point moves in a 3D Delaunay triangulation, for instance because its estimation gets refined, a set of tetrahedra have to be removed and replaced with new ones to maintain the Delaunay property; the management of the manifold reconstruction becomes thus complex and it entails a potentially big overhead. In this paper we investigate different approaches and we propose an efficient policy to deal with moving points in the manifold estimation process. We tested our approach with four sequences of the KITTI dataset and we show the effectiveness of our proposal in comparison with state-of-the-art approaches.

1 Introduction

Incremental 3D reconstruction from a sparse point cloud is gaining interest in the computer vision community as incremental Structure from Motion algorithms are consolidating [18]. This is clearly true for those applications where a rough, but dense, surface represents a sufficient and effective representation of the scene, e.g. for traversability analysis in unmanned vehicle navigation. Furthermore, in real-time applications, the map of the environment needs to be updated online, and the surface has to be estimated incrementally.

Most of the existing algorithms [10,13,8,9] bootstrap the reconstruction of a mesh surface from the 3D Delaunay triangulation of a sparse point cloud. Indeed, the 3D Delaunay triangulation is very powerful: the Delaunay property, i.e., no point of the triangulation is inside the sphere circumscribing any tetrahedron, avoids as much as possible the resulting tetrahedra to have a degenerate shape [11]; it is self-adaptive, i.e., the more the points are dense the more the tetrahedra are small; it is very fast to compute, and to update against point removal or addition; off-the-shelf libraries, such as CGAL [16], enable a very simple and efficient management of it.

As soon as a Delaunay triangulation is available, several approaches exist to extract a surface taking into account the visibility of each point. The simplest

algorithm is the Space Carving [6]: it initializes all the tetrahedra as *matter*, then it marks as *free space* the tetrahedra intersected by the camera-to-point *viewing rays*, i.e., the lines from the camera center to the observed 3D points in the triangulation. The boundary between free space and matter represents the final surface of the scene. Pan et al. [13] improve upon this simple procedure by proposing an online probabilistic Space Carving, but this is not an incremental approach: they start from scratch every time new points are added. Lovi et al. [10] present the first incremental Space Carving algorithm which runs real-time, but, as for the previous methods, the estimated surface is not guaranteed to be manifold.

Several reasons lead to enforce the manifold property as explained in [7]. Most Computer Graphics algorithms need the manifold property, for instance smoothing with Laplace-Beltrami operator [12], or the linear mesh parametrization [15]. Moreover the manifold property enables surface evolution in mesh-based Multi-View Stereo, as in [17,1], the manifold property enables a photometric refinement by surface evolution such as with the high accurate Multi-View Stereo mesh-based algorithm as in [17,1]. With these approaches is hard to estimate the surface evolving flow in the presence of non manifold vertices: indeed they compute for each vertex the gradient minimizing the reprojection error, by summing-up the contribution of the incident facets; if the vertex is not manifold, this gradient does not converge. As a further proof of this, [17] needs to manually fix the surface estimated via s-t cut. As in [17], it is possible to fix the mesh as a post-processing step, but reconstructing directly a manifold as in the proposed paper, enables the design of a fully automatic pipeline which do not need human intervention.

In literature, the only algorithm reconstructing a manifold incrementally was proposed by Litvinov and Lhuiller [8,9]. In their work, the authors bootstrap from the Space Carving procedure and, by taking into account the number of intersections of each tetrahedron with the viewing rays, they reconstruct a surface keeping the manifold property valid. The main limitation is that Litvinov and Lhuiller insert a point into the Delaunay triangulation only when its position is definitive, then they cannot move the point position anymore even in the case they could refine their estimate. The main reason of Litvinov and Lhuiller design choice has to be ascribed to the computational cost of updating the visibility information along the viewing rays incident to each moved point, and the computational cost of updating part of the Delaunay triangulation, which in turn induces a new manifold reconstruction iteration step.

Indeed, the very common approach to deal with a point moving in the triangulation, is to remove it and add it back in the new position [16] (Fig. 1). When we remove a point (the point A in Fig. 1(a)) and we want to keep the Delaunay property, we have to remove all the tetrahedra incident to that point (light red triangles in Fig. 1(b)); then, we add a new set of tetrahedra to triangulate the resulting hole (dark green triangles in 1(c)). When we add a new point into the triangulation (the point B in Fig. 1(d)), a set of tetrahedra would conflict with it, i.e., the Delaunay property is broken (light red triangles in Fig. 1(d)); so,

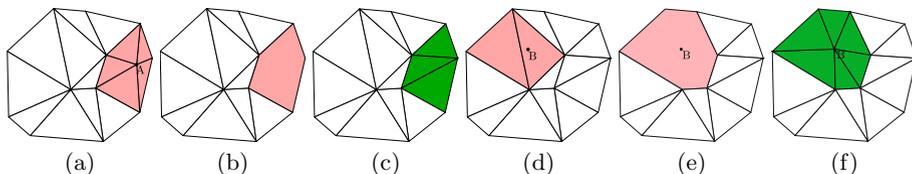


Fig. 1. Example of point removal in 2D case. Light red triangles depict are removed and replaced with the new dark green ones.

we remove this set of tetrahedra again (red triangles in Fig. 1(e)) and we add a new connected set that re-triangulate the hole (dark green triangles in Fig. 1(f)). Whenever a set of tetrahedra is replaced, we have to transfer conveniently the information about the visibility (matter or free space) of the removed tetrahedra to the new one. In addition to this, we have to update the visibility of the tetrahedra crossed by a visibility ray from one camera to the moved point. For these reasons the update of the point position is computational demanding.

To complete the overview of the incremental reconstruction methods from sparse data, we mention here another very different approach was proposed by Hoppe et al. [5] who label the tetrahedra with a random field, and extract the surface via graph-cuts by minimizing a visibility-consistent energy function. This incremental algorithm is effective and handles the moving points, but the manifold property of the reconstructed surface is not yet guaranteed.

In this paper we propose, to the best of our knowledge, the first manifold 3D reconstruction algorithm from sparse data which deals with dynamic point changes. In particular, we show that in this setting the algorithm by Lovi et al. [10] provides a feasible solution, but it is very inefficient and we propose a novel efficient policy to handle the visibility update of Delaunay tetrahedra with moving points.

In Section 2 we summarize a slightly modified version of the approach of [8] we use to reconstruct a manifold surface. In Section 3 we describe Lovi’s approach [10] and our proposal to deal with moving points, together with a complexity analysis that explains why our approach is more efficient. In Section 4 we show the experimental results on the publicly available dataset KITTI [3], while, in Section 5 we point out some future works and in the conclusion of the paper.

2 Manifold Reconstruction

In this paper we reconstruct a manifold surface that represents the observed scene. A surface is manifold if and only if the neighborhood of each point is homeomorphic to a disk. In the discrete case, the points are the vertexes of a mesh, and the neighborhood is represented by the incident triangles (or polygons); a surface is manifold if each vertex v is *regular*, i.e., if and only if the edges opposite to v form a closed path without loops (see [8] for more details).

2.1 Incremental Manifold Extraction with Tetrahedra Weighting

In this section we briefly summarize our variation on the method originally proposed in [8] enhanced by a weighting scheme that avoids the creation of most visual artifact in the final mesh (more discussion about visual artifacts in [9]). In our Space Carving algorithm, a weight roughly represents how many rays intersect a tetrahedron, and in the following, a tetrahedron belongs to *free space* if its weight w is higher than a threshold T_w (in our case $T_w = 1.0$).

Sparse Point Cloud. The input of our algorithm is a sparse 3D point cloud, estimated incrementally by assuming the camera poses to be known. For each keyframe, i.e., every $K = 5$ frames, we extract Edge-point features, i.e. 2D points laying on the image edges [14]; these points represent measurements of 3D points. Frame-by-frame we track these features with the Kanade-Lucas-Tomasi tracker, and, at each keyframe, we estimate the new positions of the 3D points with the new measurements available from the tracking. New estimates are obtained by triangulating the 2D tracked points and by minimizing the reprojection error with a Gauss-Newton algorithm. Once a new estimate of a 3D point is available, we add it to the Delaunay triangulation, i.e., to the reconstruction; then we update its position according to the new measurements: this update induces the motion of the points inside the triangulation.

3D Reconstruction. The reconstruction of the surface bootstraps from the manifold partitioning the 3D triangulation between the set O of *outside* tetrahedra, i.e., the manifold subset of the free space (not all the free space tetrahedra will be part of the manifold), and the complementary set I of inside tetrahedra, i.e. the remaining tetrahedra that represent the matter together with the free space tetrahedra which would invalidate the manifold property.

Let $\delta(O_{t_{\text{init}}})$ be the initial manifold. This initial manifold is obtained with the following steps. *Point Insertion*: add all the 3D points estimated up to time t_{init} and build their 3D Delaunay triangulation. *Ray tracing and tetrahedra weighting*: for each viewing ray, the algorithm traverses the triangulation and adds a weight $w_1 = 1.0$ to the intersected tetrahedra, a weight $w_2 = 0.8$ to the neighbors and a weight $w_3 = 0.2$ to the neighbors of their neighbors. Such weighting scheme acts as a smoother of the visibility and avoids the creation of visual artifacts; it is the main difference between our algorithm and the algorithms proposed in [8,9]. *Growing*: initialize a queue Q starting from the tetrahedron with the higher weight. Then: (a) pick the tetrahedron with highest weight from Q and add it to $O_{t_{\text{init}}}$ only if the resulting surface between $O_{t_{\text{init}}}$ and $I_{t_{\text{init}}}$ remains manifold; (b) if inserted add the neighboring tetrahedra to the queue Q , otherwise discard it; continue iteratively until Q is empty.

Once the system is initialized, a new set of points P_{t_k} is estimated at each $t_k = t_{\text{init}} + k * T_k$ frame, named keyframes, where $k \in \mathbb{N}^+$ and T_k is the inverse of the keyframe rate. The insertion of a point $p \in P_{t_k}$ causes the removal of the set D_{t_k} of tetrahedra breaking the Delaunay property, and, the surface $\delta(O_{t_k}) =$

$\delta(O_{t_{k-1}} \setminus D_{t_k})$ is not guaranteed to be manifold anymore. To avoid this, the authors in [8] define a list of tetrahedra $E_{t_k} \supset D_{t_k}$ and apply the *Shrinking* procedure, i.e., the inverse of Growing: they subtract iteratively from $O_{t_{k-1}}$ the tetrahedra $\Delta \in E_{t_k}$ keeping the manifoldness valid. After this process, it is likely that $D_{t_k} \cap O_{t_k} = \emptyset$. Whenever $D_{t_k} \cap O_{t_k} \neq \emptyset$ the point p is not added to the triangulation, i.e., is dropped. Once all points in P_{t_k} have been added (or dropped), the growing process runs similarly to the initialization procedure, but the queue Q is initialized with the tetrahedra $\Delta \in T \setminus O$ such that $\Delta \cap \delta O \neq \emptyset$.

3 Reconstructing a Manifold with Moving Points

As previously described, Litvinov and Lhuiller [8] algorithm adds points to the triangulation only when their 3D position is completely defined; by doing this, there are no changes in the Delaunay triangulation, induced by moving points. This results in a restriction if we would like to refine the estimation of the position of a point 3D position after its insertion.

Only Lovi et al. [10] presents an incremental Space Carving algorithm which deals with moving points, but their method does not enforce the manifold property. In this paper we verify the approach of Lovi et al. [10] to be very inefficient for manifold reconstruction, and we present a different approach to deal with moving points that leads to a significant faster computation.

3.1 The Straightforward Approach

The simplest way to deal with moving points while reconstructing a manifold surface, is to apply a straightforward modification to the so called *Refinement Event Handler* by Lovi et al. in [10]. The Refinement Event Handler algorithm assumes that, for each tetrahedron in the Delaunay triangulation a list of the intersecting viewing rays is stored. In our voting schema an intersecting ray is each ray that increase the weight of the tetrahedron.

Let p_{old} be a point that moves to position p_{new} , the algorithm in [10] moves the point by removing point p_{old} and adding p_{new} as a new point, according to the classical approach of [2], then for each point they apply the following steps. *Rays collection*: collect in a set U all the rays stored into the tetrahedra incident to p_{old} , i.e., those affected by the p_{old} removal (e.g., the light red triangle in Fig. 1(a)). *Vertex removal*: remove the vertex p_{old} and its neighboring tetrahedra from the triangulation (Fig. 1(b)); then re-triangulate the hole left by the deleted tetrahedra (Fig. 1(c)). *New point insertion*: insert the new point p_{new} into the triangulation and add to the set U all the rays stored in the conflicting tetrahedra (Fig. 1(d-f)). *Rays removal*: for each tetrahedron of the entire triangulation remove the rays ending in p_{old} . *Ray casting*: cast one ray for each ray in U .

In our case, whenever the 3D estimate of a point moves, we apply the Refinement Event Handler, before point addition and region growing, if and only if the point is inside the shrunk volume D_{t_k} (see Section 2), otherwise we do not move the point (this second case happens very rarely [8]).

Table 1. Complexity analysis; “-” means not existing step.

Step	straightforward algorithm	K heuristic	proposed algorithm	window heuristic
Rays collection	$O(F \cdot N^2)$	$O(N)$	-	-
Weight collection	-	-	$O(N)$	$O(N)$
Vertex Removal	$O(N)$	$O(N)$	$O(N)$	$O(N)$
New points insertion	$O(F \cdot N^2 \cdot N)$	$O(N)$	$O(F \cdot N^2 \cdot N)$	$O(N)$
Rays removal	$O(N^2 \cdot F \cdot N^2)$	$O(N^2)$	-	-
Weight Update	-	-	$O(N)$	$O(N)$
Backward ray casting	-	-	$O(N^2 \cdot F)$	$O(N)$
Ray casting	$O(N^2 \cdot F)$	$O(1)$	$O(N^2 \cdot F)$	$O(1)$
Overall complexity	$O(N^4 \cdot F)$	$O(N^2)$	$O(N^3 \cdot F)$	$O(N)$

Complexity. The number of rays involved in space carving algorithms is $O(F \cdot N^2)$ where F and N represent respectively the number of frames and the number of points in the triangulation [10], and the number of tetrahedra in a 3D triangulation grows quadratically with the number of points ($O(N^2)$). In Table 1 we reported the complexities for each of the previous stage; since our implementation exploits the CGAL [16] 3D triangulation data structure, the complexity of a single Ray casting, i.e., a cast of a single ray, is $O(N)$ in the general case, but we bound the size of the viewing ray, to avoid to include too far uncertain 3D points estimates, so the final complexity becomes $O(1)$ (see [19, p.94]).

From the analysis in the table is quite clear that this straightforward solution is not scalable, especially for the dependence between the number of rays and the number of processed frames.

Forgetting Heuristic. Lovi et al. [10] proposed a *forgetting* heuristic to limit the number of rays stored in each tetrahedron to a fixed number K , thus making the complexity independent from the number of the processed frames. However, we show in Section 4 that, when the points are moving, the reconstruction is very inefficient even with this heuristic.

3.2 The Efficient Approach

Our contribution in this paper is an approach to deal with moving points, different from the straightforward variation of [10]. Indeed in our proposal, we avoid storing the list of rays inside each tetrahedron, and we just store the weight associated with it. This allows the incremental reconstruction algorithm of Section 2.1, and, at the same time, we are able to bound the temporal complexity.

The main difficulty in the proposed approach is updating coherently the weights whenever a point moves, i.e., when the point is removed from the triangulation and added as a new point. As soon as the point is removed from the triangulation, we perform a backward ray casting with negative weights for each viewing camera such that the influence of the point is neglected. Then we remove

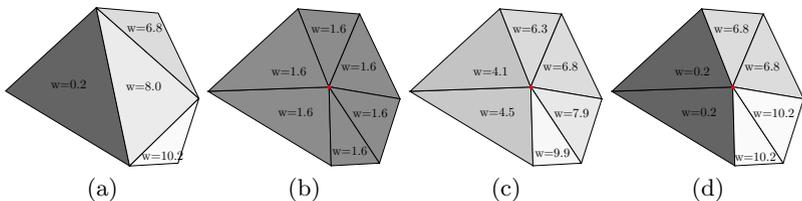


Fig. 2. 2D example of moving point addition in the new position after point removal (a brighter region corresponds to a higher weight, i.e., higher probability to be carved).

the point, and we add a new vertex in the new position. Finally, we perform the ray casting from each viewing camera to the new point.

During both point removal and addition, we have to remove a set of connected tetrahedra from the triangulation and add a new one. Let $R = \{\Delta_1^R, \Delta_2^R, \dots, \Delta_{n_R}^R\}$ be the set of removed off tetrahedra and $A = \{\Delta_1^A, \Delta_2^A, \dots, \Delta_{n_A}^A\}$ the set of the new ones; their associated weights are respectively $W_R = \{w_1^R, w_2^R, \dots, w_{n_R}^R\}$ and $W_A = \{w_1^A, w_2^A, \dots, w_{n_A}^A\}$. The weights W_R are known, while W_A are those to be computed for the new tetrahedra, without recasting the visibility rays related to those tetrahedra.

Different approaches are possible: *Mean value*: $w_i^A = \frac{1}{n_A} \sum_{k=1}^{n_R} w_k$; *Weighted mean*: let $d_{i,j}$ be the Euclidean distances between the centroids of the i -th tetrahedron of A and the j -th of R ; then $w_i^A = \frac{\sum_{k=1}^{n_R} d_{i,k}^{-1}}{\sum_{k=1}^{n_R} d_{i,k}^{-1}}$; *Minimum distance*: $w_i^A = w_{\bar{j}}^R$ such that $\bar{j} = \arg \min_{j \in 1 \dots n_R} (d_{ij})$.

Among these, the third solution gives a non-smooth outcome and, even if this seems counter-intuitive, it results to be more suitable for our purposes. The main reason is that it preserves the discontinuity between matter and free space. For instance in Fig. 2(a) we depict a 2D triangulation where we want to add a new point position; in Fig. 2 (b), (c) and (d) we show the results of weights update after point addition with, respectively, Mean value, Weighted Mean value and Minimum distance approaches. It is clear that only Fig. 2(d) preserved the discontinuity, while in other cases becomes hard to distinguish between matter (lower weights) and free space (higher weights).

In case of very sparse data, the centroids of big tetrahedra, together with the associated visibility information, can be far from the newly added or moved points, and our update policy might lead to results far from the ideal solution, i.e., the straightforward approach discussed in Section 3.1. Our algorithm overcomes this issue thanks to the use of the (so called) Steiner points added to the triangulation before the actual reconstruction is performed; this idea was already introduced in [8]. We add Steiner points to the Delaunay triangulation every 5m along each axis so that they cover all the space that can be represented. The use of Steiner points limits the creation of very big tetrahedra, the visibility information becomes always local, and the update policy avoids drifts. Indeed,

experimental results show good accuracy on varied scenes, even when lack of textures induces very sparse data.

Complexity. The complexity of the steps of our algorithm are reported in Table 1. The main difference with respect to the straightforward algorithm is the replacement of the Rays removal to the weight update and backward casting which are the key of the gaining in computational complexity. The proposed algorithm is thus $O(F \cdot N^2)$, so, in principle, the dependence with F still remains and results in a non scalable solution.

Window Heuristic. We are able to bound the complexity of our algorithm to $O(N^2)$ thanks to the following heuristic: instead of backward casting all the rays connecting the moving point to all the viewing cameras, we consider only the most recent cameras. In this case the complexity of the ray casting becomes $O(W \cdot N^2)$, where W is the (constant) size of the window (in our case $W = 15$), so the final complexity is $O(N^2)$.

4 Experimental Validation

To evaluate our approach, we tested the system on four different sequences of the KITTI dataset [3] on a 4 Core i7-2630QM CPU at 2.2Ghz (6M Cache), with 6GB of DDR3 SDRAM. The video stream was captured by a Point Grey Flea 2, which records 1392x512 gray scale images at 10 fps. The vehicle pose are estimated through a RTK-GPS and they are the initial input of our system together with the video stream.

Among all the sequences we choose the 0095 (268 frames) and 0104 (313 frames) since they depict two different urban scenarios: the former shows a narrow environment where the building façades are close to the camera, the latter captures a wide road. We also tested our approach on sequences 03 (801 frames) and 04 (271 frames) from the odometry dataset: these videos provide a varied landscape mixing natural (trees and bushes) and man-made (houses, cars) features.

To provide a quantitative evaluation we compared the reconstructed meshes with the very accurate point clouds measured by the Velodyne HDL-64E sensor in the KITTI dataset through the CloudCompare tool [4]. This tool computes the reconstruction error as the average of the distances between each Velodyne point and the nearest triangle in the reconstructed mesh.

We evaluated the performance and the accuracy of the Lovi’s approach against our three different updating policy. As explained previously, no manifold incremental reconstruction approach deals with moving points, so a fair comparison results to be between the straightforward approach of Lovi, applied to manifold reconstruction (Section 3.1) and our updating policies. In Fig. 3 we show an example of the reconstruction results before and after the red points has been moved in the Delaunay triangulation (see also the video at http://youtu.be/_-q9sKjcOC0).

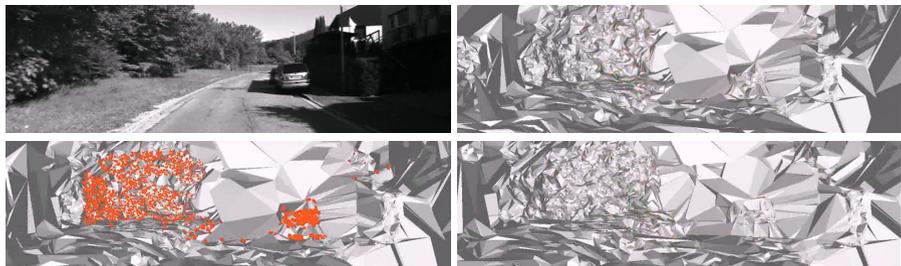


Fig. 3. Incremental reconstruction example. From up left to bottom right: original frame, before point positions update, points moved in the scene (red dots) and manifold updated.

Fig. 4 shows the results of the comparison where we applied the window heuristic (Section 3.2) to all the algorithms. In the case of Lovi’s algorithm we applied the forgetting heuristic with $K = 5$ and $K = 1$, where K is the number of viewing rays stored for each tetrahedron. Fig. 4(a) shows that the accuracy of the proposed approach, i.e, moving point management through minimum distance weight updates, is comparable with respect to Lovi’s proposal outcomes, where the algorithm with $K = 5$ stores more information, so it performs better. We compared our approach with respect to Lovi’s method instead of the other incremental reconstruction algorithm presented in [8]; the reasons are twofold. In [8] Litvinov and Lhuiller does not deal with moving points, which is the main point addressed in this paper. Moreover, Litvinov and Lhuiller point out in [9] that the ideal solution for a manifold reconstruction algorithm is represented by the manifold including as much as free space tetrahedra as possible. Since the solution provided by Lovi et al. coincides with the (non-manifold) mesh containing all the free space tetrahedra, a reconstruction accuracy similar to Lovi’s suggests that the reconstruction is near to the ideal solution. In some cases our algorithm reaches even better accuracy, this is due to the smoothing effect induced by our heuristic.

Fig. 4(a) shows that the Minimum Distance always outperforms the other two updating schema as expected (see the Section 3.2).

In Fig. 4(b) we report the time performance of the algorithms. Let T_{mov} and $T_{\text{non-mov}}$ be the overall processing time with and without moving points, and N_{mov} be the number of the total points moves, e.g., if one point moves three times, $N_{\text{mov}} = 3$. The overhead introduced in the whole reconstruction process for each move of each point has been computed as $\frac{T_{\text{mov}} - T_{\text{non-mov}}}{N_{\text{mov}}}$. The performance of the different update schema we presented in Section 3.2 is very similar since the steps involved are basically the same: for each update on the Delaunay data structure, we iterate over the old tetrahedra to collect the weights, then we iterate over the new tetrahedra to set the new weights. As expected by Section 3.2, our algorithm clearly outperforms Lovi’s approach. Our updating schema is very efficient for two reasons. First, we only need to update locally the visibility, while Lovi’s approach casts a ray for each visibility ray stored

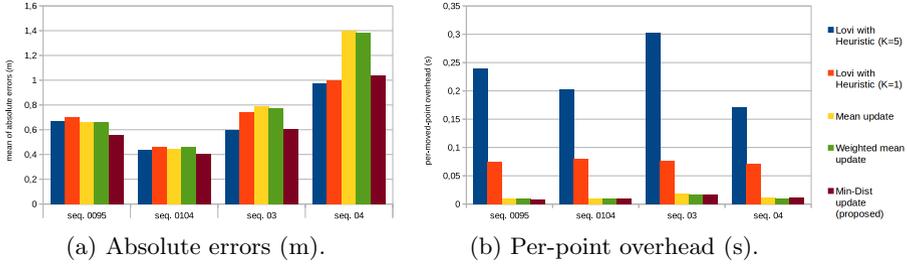


Fig. 4. Experimental evaluation of the proposed approach with respect to Lovi’s [10].

inside the tetrahedra. Second, when we remove a point (first step of moving point management), we perform a ray casting backward to update only the convenient tetrahedra, instead of iterating over the whole triangulation to remove the visibility rays involving the point moved as in [10].

5 Conclusion and Future Work

In this paper we have shown that manifold reconstruction from sparse data with moving points is not a trivial task. To keep the Delaunay property valid when a point moves inside the Delaunay triangulation, we have to remove it and add a new point in the new position. This induces the removal of a set of tetrahedra, with the associated visibility information; then, we have to add a new set of tetrahedra with coherent visibility information; finally we have to update the visibility information in all the tetrahedra affected by the point move.

Existing solutions successfully applied for classic Space Carving, result to be inefficient and slow when applied in the manifold reconstruction setting. In this setting, we investigated different approaches to handle visibility information propagation, by updating the weight for each tetrahedron, which roughly represents the number of ray intersections, and we proposed an efficient algorithm to conveniently update it. We tested our system with the KITTI dataset and it clearly outperforms the existing approach of Lovi et al. [10] for incremental manifold reconstruction.

Future works would include a photometric refinement of the manifold extracted incrementally, and an evaluation of the manifold quality on-the-fly, relying on the uncertainty information carried by the estimation of 3D points. A natural extension could also deal with the reconstruction of non-rigid shapes whose 3D points are moving.

Acknowledgments. This work has been partially funded by the SINOPIAE project, from the Italian Ministry of University and Research and Regione Lombardia, and the MEP: Maps for Easy Paths project funded by Politecnico di Milano under the POLISOCIAL program.

References

1. Delaunoy, A., Prados, E., Piracés, P.G.I., Pons, J.P., Sturm, P.: Minimizing the multi-view stereo reprojection error for triangular surface meshes. In: *BMVC 2008-British Machine Vision Conference*, pp. 1–10. BMVA (2008)
2. Devillers, O., Teillaud, M.: Perturbations and vertex removal in a 3d delaunay triangulation. In: *14th ACM-Siam Symp. on Algorithms*, pp. 313–319 (2003)
3. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3354–3361. IEEE (2012)
4. Girardeau-Montaut, D.: Cloud compare (access March 22, 2015). <http://www.cloudcompare.org/>
5. Hoppe, C., Klopschitz, M., Donoser, M., Bischof, H.: Incremental surface extraction from sparse structure-from-motion point clouds. In: *Proc. BMVC (2013)*
6. Kutulakos, K.N., Seitz, S.M.: A theory of shape by space carving. *International Journal of Computer Vision* **38**(3), 199–218 (2000)
7. Lhuillier, M.: 2-manifold tests for 3d delaunay triangulation-based surface reconstruction. *Journal of Mathematical Imaging and Vision*, 1–8 (2014)
8. Litvinov, V., Lhuillier, M.: Incremental solid modeling from sparse and omnidirectional structure-from-motion data. In: *BMVC (2013)*
9. Litvinov, V., Lhuillier, M.: Incremental solid modeling from sparse structure-from-motion data with improved visual artifacts removal. In: *International Conference on Pattern Recognition (ICPR) (2014)*
10. Lovi, D.I., Birkbeck, N., Cobzas, D., Jagersand, M.: Incremental free-space carving for real-time 3d reconstruction. In: *Fifth International Symposium on 3D Data Processing Visualization and Transmission (3DPVT) (2010)*
11. Maur, P.: Delaunay triangulation in 3d. Tech. rep., Dep. of Computer Science and Engineering. University of West Bohemia, Pilsen, Czech Republic (2002)
12. Meyer, M., Desbrun, M., Schröder, P., Barr, A.H.: Discrete differential-geometry operators for triangulated 2-manifolds. In: *Visualization and Mathematics III*, pp. 35–57. Springer (2003)
13. Pan, Q., Reitmayr, G., Drummond, T.: Proforma: probabilistic feature-based on-line rapid model acquisition. In: *BMVC*, pp. 1–11 (2009)
14. Rhein, S., Lu, G., Sorensen, S., Mahoney, A.R., Eicken, H., Ray, G.C., Kambhamettu, C.: Iterative reconstruction of large scenes using heterogeneous feature tracking. In: *2013 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 407–412. IEEE (2013)
15. Saboret, L., Alliez, P., Lévy, B.: Planar parameterization of triangulated surface meshes. In: *CGAL User and Reference Manual*. CGAL Editorial Board, 4.4 edn. (2000)
16. The CGAL Project: CGAL User and Reference Manual. CGAL Editorial Board, 4.5 edn. (2014)
17. Vu, H.H., Labatut, P., Pons, J.P., Keriven, R.: High accuracy and visibility-consistent dense multiview stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34**(5), 889–901 (2012)
18. Wu, C.: Towards linear-time incremental structure from motion. In: *2013 International Conference on 3D Vision-3DV 2013*, pp. 127–134. IEEE (2013)
19. Yu, S.: Automatic 3d modeling of environments: a sparse approach from images taken by a catadioptric camera. Ph.D. thesis, Université Blaise Pascal-Clermont-Ferrand II (2013)