

Robust acquisition of 3D informations from short image sequences

Sylvain Paris* and François Sillion

ARTIS—GRAVIRIMAG—INRIA¹

Received 7 February 2003; accepted 25 February 2003

Abstract

This paper addresses the problem of 3D reconstruction from a set of viewpoints on a short baseline. Its main contribution is the development of a robust algorithm which can extract 3D informations from a set of images taken with a very small baseline, even in the presence of significant occlusion. To achieve this goal, we use a multi-pass process that works layer by layer. In each pass we begin with a “space carving” step which is made robust through some morphological operations. Then we introduce an original technique to solve the ambiguity inherent to “space carving” in the case of short baseline. Once we have a voxel representation of the objects, we propose a method based on differential geometry to build a smooth mesh. Results are presented, demonstrating the efficiency and usefulness of the method.

© 2003 Elsevier Science (USA). All rights reserved.

Keywords: Structure from motion; 3D reconstruction; Voxel space; Dynamic programming; Mathematical morphology; Partial differential equation filter; Occlusion

1. Introduction

We aim at acquiring 3D information from image sequences such as a short video segment. Our goal is to use this information to add realistic details in virtual environments. For instance, we want to populate a street with passer-bys, post-boxes, and other urban elements. These are not the focus of interest of the scene, but rather

* Corresponding author.

E-mail addresses: Sylvain.Paris@imag.fr (S. Paris), Francois.Sillion@imag.fr (F. Sillion).

¹ ARTIS is a joint research project of CNRS, INPG, INRIA, UJF.

the small add-on that makes the whole street “alive.” Therefore it is not important if the objects are not “complete” for we do not need to look at them very closely nor to see what is behind a post-box. The important point is the realistic look of the result. Moreover we do not want to be restricted to a very specific kind of image sequences—for instance, there may be an interesting object in a movie and then we have no other choice than to deal with that given video segment. Therefore we must cope:

- with viewpoints that are spatially close (we are in the case of *short baseline*) and almost aligned because we cannot assume the image sequence is long nor that it goes all around the object,
- with ordinary images for they can be noisy and we cannot assume there is a “blue screen” in the background,
- and with occlusions because there can be another object that hides the one we are interested in.

We assume that we know all the camera parameters because acquiring them from the images is a full separate problem (see [14] for details on this topic), and that there is no movement in the sequence (acquiring 3D information and movement is a much more complex problem). So we can work from two different types of sequence: either the scene is static and the camera is moving (e.g., for acquiring a statue or post-box), or we have a set of synchronized cameras which shoot simultaneous images [33]. Finally we consider that the only user input is a bounding volume around the objects of interest.

Having defined our goals and the input, we can have an overview of the algorithm we propose. The main structuring point is a multi-pass strategy: we reconstruct the object information in a *front to back order* [27,35]. During a pass, only the closest object still not reconstructed is taken into account. This guarantees all the occluders are known before we try to reconstruct an hidden object.

We need also to choose the 3D structure we use to represent the objects during the process. One way to do this is to use their contours as in [7,21,32,36], but these methods are very dependent on the contour extraction algorithm that does not give satisfying results with noisy images and real background. We can also optimize a surface so it fits to the image contents using shading information as Fua and Leclerc [13]. This approach deals well with untextured objects but cannot cope with occlusions in the general case. Faugeras and Keriven [12] also use a surface-based approach driven by the level set technique, it works well even with occlusions and topologically complex objects but it requires textured surfaces. Recently, a promising approach based on graph cut has been introduced by Roy and Cox [25] and extended by many others [4,16,18,37] but it still lacks of precision and results in objects which are almost flat. Since the working volume is bounded, we prefer a *voxel-based* approach as described in [3,7,19,26–29,31]: we discretize the volume into small cubes (the *voxels*) and the goal then is to select the voxels that belong to the object volume. Thus we are not bound to a contour detection algorithm and we show in Section 4 that we can work without texture. We explain quickly how this voxel space is built in Section 2.

Using this voxel space, we can split each pass of the algorithm into three steps. We begin in Section 3 with the selection of the *consistent* voxels as defined by Seitz and




Fig. 1. Overview of the whole algorithm.

Dyer [27]: we keep the only voxels that have the “same” color in all the images. This first step is quite similar to the *space carving* algorithm [19] during which we take special care to be robust to noise. Kutulakos and Seitz [19] explains that there may be too many voxels left. In our specific configuration of short sequences, we show that this remark is crucial and that there is a very high number of useless consistent voxels and then, as a second step, we introduce in Section 4 an original approach to solve this ambiguity. The third and last step of a pass consists in blocking the lines of sight intersecting the object we have just reconstructed. Section 5 explains this step and how to deal with occlusions.

When all passes are done, we have a voxel representation of the objects that we call *draft surface*. This discretized form is not suitable for further use because real objects are not composed of small cubes. Therefore we propose in Section 6 a new technique based on morphological treatments and differential geometry to transform this representation into a smooth surface.

Fig. 1 gives an overview of the whole algorithm. We illustrate in Section 7 the applicability of our algorithm with some results and propose some extensions.

2. Voxel space geometry

Since the volume of interest is bounded, we can discretize it into voxels. We choose to let their size vary so that they represent a fixed amount of information in image space. More precisely, we impose that all the voxels project to the same area in the camera image planes. This is called by Slabaugh et al. [30] the *constant footprint property*.

In our specific configuration of aligned viewpoints, we build such a voxel space by introducing in the scene the “distance from the cameras.” Geometrically, this distance $\text{dist}_{\text{scene}}(X)$ of a point X in the scene is defined as $d(\Delta_{\text{cam}}, X)$, where $d(x, y)$ is the euclidian distance between x and y and Δ_{cam} is the line on which the cameras lie. Then, as the perspective projection gives a resulting area inversely proportional to the depth, it is sufficient to make the size of a voxel v proportional to $\text{dist}_{\text{scene}}(v)$.

In practice, we build the first voxel plane $z = z_0$ and then we build the other planes using a scaling whose center lies on the camera line (see Fig. 2). This assures the constant footprint property. Note that the global shape of the voxel space is characterised only by the first plane and the position of the homothetic center. This construction is similar to the use of a *virtual camera* [31] with its center aligned with the input viewpoints. Saito and Kanade [26] proposes another approach based only




Fig. 2. Construction of the voxel space using scaling.

on the fundamental matrix of two views that reaches the same property but we will see in Section 4 that our technique helps to define a very useful additional geometric property.

3. Identification of the consistent voxels

In the first step of the algorithm, we aim at reducing the space in which we are working. Actually the whole bounded space given as input is too large, typically involving tens of millions of voxels. The problem has to be alleviated, fewer voxels will allow more complex and more powerful treatments in the next steps of the algorithm.

The criterion we use here is the *voxel consistency* described in [19,27]: a voxel is said *consistent* if its color is the same for all cameras that can see it. It is important to understand that it is a necessary but not sufficient condition. A voxel in the final scene must be consistent but a consistent voxel may not be in the final scene. For instance, if we place all the cameras in front of a red wall, each voxel is seen red by all the cameras, then all voxels are consistent although only the few that represent the wall should be in the final scene.

For now, as suggested in [19,25,27] we use the standard deviation of the colors seen by the cameras to evaluate the consistency. It restricts the algorithm to almost lambertian objects because it implies that the components of the voxel color are the same in all images and hence do not depend of the viewpoint. But we will show that we can still achieve satisfying results with real images even with this hypotheses. Nevertheless it is clear that there is room for improvement using better statistics.

In practice, for each voxel, we proceed as follows:

- we suppose it is the only voxel visible in the scene,
- we compute a color for each camera corresponding to the color of the pixel through which the camera sees the voxel,
- we compute the standard deviation s of the color set we have computed. In practice, we use the *hue, saturation, value* space but this choice is not critical (we only

need an estimation of the distance between two colors to measure the color similarity).

Then we assume a voxel is consistent if $s < s_t$ for some threshold s_t .

At this point, we have built the subset of consistent voxels. All other voxels are inconsistent, for different cameras see them of different colors. In the following steps of our pass, we no longer need to take these voxels into account anymore.

4. Construction of a draft surface

As said before, the voxel consistency is not a sufficient condition to be part of the final reconstruction, it is only necessary. Hence, inside the subset we have just built, there are still too many voxels. The key situation in which unwanted but consistent voxels appear is around a surface of uniform color. If we look at a voxel just in front of such a surface, the lines of sight of all the cameras hit the surface, then all the cameras see the same color and therefore the voxel is consistent, although it is not part of the surface (see Fig. 4).

To study this configuration, we observe that we can consider *epipolar planes*. In classical stereo-vision, epipolar planes are the planes which contain the two optical centers. In our case, as the cameras are aligned, we can straightforwardly extend the notion of epipolar planes to the planes which contain all the optical centers (Fig. 3). These planes are of great help to describe the geometric configuration for they correspond to voxel planes in the voxel space if we choose the x -axis equal to the axis on which the cameras lie (Figs. 2 and 3). This is a consequence of the scaling-based construction.

Then, in these planes, the consistent voxels generated by a surface of uniform color form a quadrilateral. To show this, we only need to consider the two extreme cameras as shown in Fig. 5. We name the four vertices of this quadrilateral: U for the most distant one, D for the closest one, L and R for the left and right remaining vertices, respectively.




Fig. 3. Epipolar planes.




Fig. 4. A voxel in front of an uniform color surface.




Fig. 5. Quadrilateral formed in an epipolar plane by a surface of uniform color. Edges are only due to the two extreme cameras.

This quadrilateral has the following *vertex coordinate property*: in the plane homothetic coordinate system, each vertex has one and only one extremal coordinate. This property is satisfied if and only if the homothetic center lies between the two extremal cameras. To prove this, only L and R have to be considered because U and D are always the most distant and the closest vertices. We give details for the L vertex. L is the intersection of two lines of the sight, one from each extreme camera (Fig. 5). As L is extremal if the whole quadrilateral lies on the right of L , it implies that one of the two lines has a positive slope and the other one a negative slope. Then, we remark that the slope of the lines is directly linked to the angle between the boundary of the voxel space and the line: positive angle means negative slope and conversely. And this angle is linked to the camera position relatively to the homothetic center: if the camera lies on the left (resp. on the right) of the homothetic center, all lines of sight coming from this camera have a negative angle (resp. positive angle). Thus, to summarize, we have the following result: L is extremal if and only if there is one camera on each side of the homothetic center. Fig. 6a illustrates this vertex coordinate property: the homothetic center lies between both extremal camera.

Fig. 6b shows a configuration without the property: neither L nor R are extremal in the homothetic coordinate system.

Thanks to this property, we can now recognize the quadrilaterals formed by surfaces of uniform color. We begin to search the most left consistent voxel in a plane. It is a L vertex. From that voxel, we find the U and D vertices by following the edges of the quadrilateral and then we find R the same way. We mark all the voxels in the quadrilateral as *processed* and we iterate to find another L .

However, because of noise, the above process cannot be applied directly on the subset of consistent voxels. In fact, there are many unwanted voxels in the subset: some voxels can be isolated or there may be some small holes inside the quadrilaterals. To make more robust the quadrilateral recognition, we must perform a pre-treatment on the subset. We chose a morphological filter: we first apply a morphological closure on the subset with a cube as structuring element, followed by an opening. The closure ensures that there is not any small hole and that we do not lose small quadrilaterals between large ones. The opening removes isolated voxels. And to finish, we eliminate small groups of voxels in the subset (using a simple threshold on the number of voxels in a connected group).

Once we have achieved this denoising process, we can run properly our quadrilateral recognition algorithm.

4.1. Computation of the intersection of the surface and the epipolar plane

Now we have characterized the quadrilaterals inside which we have to find the intersection of the surface of uniform color with the epipolar plane. Let us only consider a single quadrilateral with vertices U , R , D , and L .

If we observe Fig. 5, vertices L and R are the extreme points of the intersection we are looking for. This is the starting point of our method: we search for a line whose

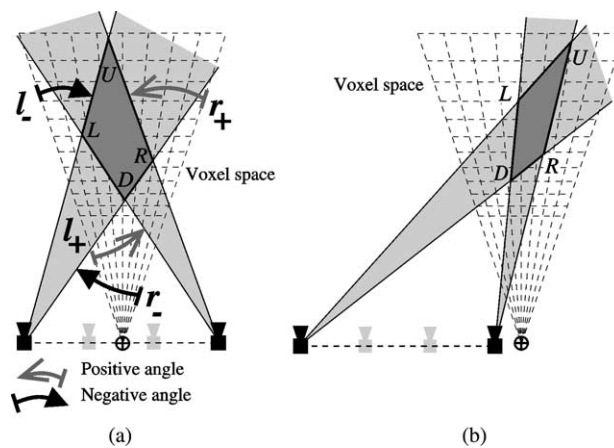


Fig. 6. (a) The homothetic center lies between the extremal cameras: the vertex coordinate property is true; (b) if not, the property is not verified.

extreme points are L and R . Another constraint is quite natural: the voxels of the line should maximize the consistency or equivalently minimize the standard deviation s . The last constraint comes from the fact that the color of the surface is uniform. Since we are able to see the scene, there is some light source somewhere in the scene and therefore there can be some shading variations. The relatively uniform color of the surface therefore implies that the surface is quite smooth.

The classical way to deal with this set of constraints is a snake-like optimization [5]. We represent the line by a function $z = f(x)$ and we want to minimize the following energy:

$$\int_L^R \left[\alpha \cdot s^2(x, f(x)) + \beta \cdot \left(\frac{df}{dx}(x) \right)^2 \right] dx, \quad (1)$$

where α and β control the smoothness of f .

We can solve this optimization problem with the classical technique of dynamic programming [2]. Similar approaches have been often used in computer vision. Our contribution is to work directly in the 3D space and not in the disparity space as in [15,17,22] so we have a more robust localization against ambiguity as shown by Okutomi and Kanade [23]. We also use dynamic programming only in the regions known to represent a smooth surface, unlike other techniques [8,9,15,17,22,23] that work with the whole scan-line and then need some special treatments to cope with occlusions or discontinuities. In addition we introduce the vertex coordinate property to characterize rigorously the search area.

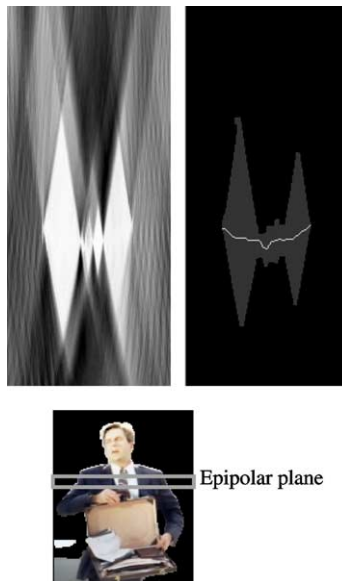


Fig. 7. Views of an epipolar plane.

Fig. 7 shows an epipolar plane: on the left, the standard deviation is represented with grey levels (white for low values and black for high ones) and on the right, the result of the thresholding after the morphological treatment (in grey) and the line (in white) obtained by the snake optimization.

It is important to notice that this step where the surface of the objects is found inside the consistent voxels is crucial. If we only use the space carving algorithm [19] (or a variant [27,28]) in the configuration of a short image sequence, satisfying results cannot be reached because of those quadrilaterals. Since consistency is the only criterion used to carve the voxels, none of the voxels inside the quadrilaterals are eliminated, resulting in large peaks in the uniform color areas.

We now have a collection of lines included in epipolar planes. The lines are the intersection of the model surface with the successive epipolar planes. The voxels in these lines are what we call the *draft surface* of the model.

5. Dealing with occlusions

The main point to deal with occlusions is to know at least approximately the objects that hide, the *occluders*, before reconstructing the objects that are hidden, the *occludees*. If we know this, we can evaluate the consistency of a given voxel by using only the cameras for which the given voxel is not occluded.

To achieve this goal, each pass reconstructs only the surface which is the closest to the cameras and still not reconstructed. This corresponds to the *front to back* order proposed in [27,35]; it guarantees that in the first pass, we are reconstructing only surfaces for which there is no occlusion. And in the following pass, if there is an occluder, we are assured to have already its draft surface for it is closer to the cameras. Note this implies that the user has to include all the occluders in the volume of interest given as input.

In practice, we proceed as follows:

- We evaluate the voxel consistency with only the cameras of which line of sight is not intersected.
- In each epipolar plane we only reconstruct the nearest line corresponding to a surface–plane intersection.
- At the end of each pass, we mark for each camera which area in the image plane corresponds to the lines of sight that intersect the newly built voxel.

The last point to consider carefully is the position of the homothetic center. Since we do not use all the cameras to compute the voxel consistency, each voxel is “seen” by a subset of the cameras, and we want the homothetic center to always fall between the left most and right most cameras of this subset. However that may not be possible simultaneously for all voxels. The homothetic center may not lie between the two extreme cameras for all the voxels in a quadrilateral. This can result in one of the two following situations:

- The ideal case in which there exists a position for which the homothetic center lies between both extreme cameras for all voxels. Then we simply choose such a position as in Fig. 8.

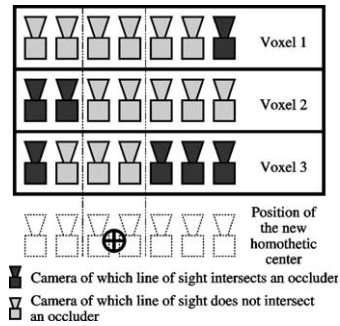


Fig. 8. Choice of the new position of the homothetic center in the ideal case.




Fig. 9. A sample reconstruction in three steps. We use a sequence of images in which a synthetic cube has been added to create a large occlusion. (a–b–c) The draft surfaces resulting from the three steps. (d) The union of all these surfaces.

- The degenerated case in which there is no such position. Then we compute a satisfying position for each voxel and we choose the mean of all these positions. Although we are not in the configuration that assures the vertex coordinate property to be true as demonstrated in Section 4, the property is still valid in practice. In both cases, we have moved the homothetic center to achieve the vertex coordinate property and therefore the quadrilaterals are still recognized properly.

We stop the iterations when a pass does not find anything to reconstruct. Then each pass has produced a piece of the draft surface (Figs. 9a–c) and we merge all these pieces (Fig. 9d).

6. Construction of the final surface

The draft surface is not a true surface: the voxels are only grouped in lines lying in epipolar planes, and do not form a continuous surface. There is no link between two successive lines. Therefore there can be holes between lines, strong discontinuities, etc. We need to include in our reconstruction algorithm some vertical connection between voxels. But we have to be careful: the final surface must be based on the information included in the draft one, we must preserve this information by using only soft filters.

6.1. Morphological operations

First, we apply a morphological dilation along the z -axis. The dilation is purely “in depth”: we do not add any voxel in front of the voxels of the draft surface since we consider the draft surface to be visible from the cameras. We fix a threshold d_{hole} beyond which we consider that there is a hole between two successive lines. So the depth of the dilation corresponds exactly to d_{hole} (in Fig. 10, $d_{\text{hole}} = 3$).

Unfortunately, the process may have failed in some epipolar planes because of noise, resulting in holes, poorly located lines and so forth. We need to have an appropriate treatment to reinforce the process. Again we use an morphological closure and an opening. If we use a cubic structuring element as before, the action will be isotropic and we risk modifying the information inside the epipolar planes, although a lot of work has already been done to acquire this information. This action inside the plane would be an unwanted side effect. What we actually need is an anisotropic action only along the y -axis. Hence we use a purely vertical structuring element. Thus only the action between planes remains and the action internal to a plane disappears: the closure fills the holes due to missing lines and the opening deletes isolated lines. Fig. 10 summarizes the morphological treatment used.

6.2. Creation of a mesh

We have now deep blocks of voxels which front represents the surface of the model. To create a mesh of the surface, we isolate blocks of voxels: once we have found a first voxel of in a block, we straightforwardly find the other voxels of the block by connectivity. For a given block, we adapt a mesh using the positions of the front voxels as vertices.

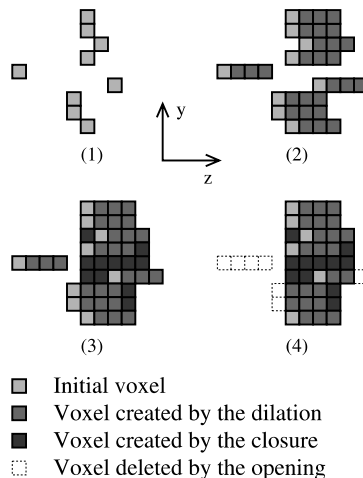


Fig. 10. Morphological treatments on the draft surface.

Now the reconstruction is almost finished: the different objects are isolated, their surface is known. This allows to make a last and object-specific treatment.

6.3. Smoothing filter

The last treatment we perform on our model is a global filter on the whole surface against noise. The snake optimization coupled with the aliasing resulting from the alignment of the vertices on the voxel grid may have left some granularity on the surface. As we have already explained, what we are reconstructing should be quite smooth. Nevertheless, we do not want our final model to look like a gaussian filtered surface: that is to say perfectly smooth with no edge. The goal is then to smooth the surface without damaging the edges.

We choose an approach similar to that of Perona and Malik [24]: we apply an anisotropic filter whose strength is controlled by some criterion we want to preserve. Perona and Malik proposed a denoising filter for images and therefore preserved discontinuities in the image. They used the gradient norm to control their filter. We want to preserve edges on our surface, so we use the curvature to control our filter.

As Perona and Malik, we use a partial derivative equation (PDE) to define the filter. We need to model the surface by a function that can evolve in time t . A natural way to do this is to represent the surface by a height field $z = h(x, y, t)$. Then we exploit the ideas developed by Deriche and Faugeras [10]:

- We define the filter by two orthogonal directions: fortunately, the principal curvatures are orthogonal so we can easily use them.
- The control function expressed in these orthogonal axes has to stay between 0 and 1 as in [1] so as it tends toward 0 for high values of the parameter and then makes the filter have no effects. In our case this aims at preserving sharp edges.
- We compute the controlling parameter on a gaussian filtered surface such that only large features remain to control the filter as in [6].

Hence the equation we propose is:

$$\frac{\partial h}{\partial t} = f(\kappa_1^\sigma)\kappa_1 + f(\kappa_2^\sigma)\kappa_2, \quad (2)$$

where κ_1 and κ_2 are the principal curvatures of the surface at (x, y) (see [20] for details), κ_1^σ and κ_2^σ are the principal curvatures at (x, y) of the surface convolved by a gaussian function of standard deviation σ , f is a decreasing controlling function from $f(0) = 1$ to $f(+\infty) = 0$. For instance, we can use $x \rightarrow 1/(1 + \alpha x^2)$.

We iterate this PDE filter so as to get a smooth surface but without affecting the sharp edges. Fig. 11 shows a sample surface: we have a perfect noiseless surface (on the left) to which we have added some white noise up to 10% of the maximal height (in the middle) and on the right we have the filtered surface. The sharp features of the original surface have been pretty well preserved and especially, the action of the filter is independent of the grid orientation: all edges are sharp and not only the ones parallel to the grid principal directions.

It is important to remark that our goal is to affect the geometric properties of the surface seen as a 3D geometric entity. Our approach is different from [11,34,38] in




Fig. 11. Illustration of the PDE filter.




Fig. 12. Comparison between the filter proposed by Alvarez et al. [1] (left) and our surface filter (right).

which the surface is only an mathematical modelization of an image. In our case, it is in the same time easier and harder to cope with because on one hand, the amount of noise in the surface we get is far less important than the one that can be encountered in noisy images but on the other hand, noise is far more sensitive on a surface, a soft granularity on a surface is immediately caught whereas the equivalent noise in an image is almost not visible. In Fig. 12, we compare our filter to the one proposed by Alvarez et al. [1] which is very efficient for images. We can see that in the case of a surface with a relatively limited noise, our filter achieves better results especially on edges, on the peak and the slopes.

After this operation, we have the surface of the object in the scene. It is smooth almost everywhere but it still conserves sharp edges.

7. Results and conclusions

To solve the 3D reconstruction problem, the algorithm described in this paper makes use of a number of classical techniques: the discretization of the space into voxels, the notion of voxel consistency, the use of epipolar planes or the dynamic programming approach. It also makes several contributions such as:

- The rigorous characterization of the situation in which the voxel consistency is not sufficient to find the surface of the model.
- The dynamic programming approach driven by this characterization and directly used in the 3D space to deal with untextured surface.
- The use of morphological treatments on the voxel space to cope with noisy data and obtain satisfying results even with real images.
- A specific PDE surface filter to achieve robustness and global consistency.
- The adaptation of the “front to back” ordering to this algorithm to solve the occlusion problem.

These techniques make our algorithm able to acquire 3D informations from a set of real photographs even with some occlusions. Nevertheless the method is still quite slow, for instance the man with briefcase (Figs. 15b and c, requires two passes, $140 \times 150 \times 330 \approx 7 \times 10^6$ voxels) needs about 280 min to complete on a MIPS R12000 400 MHz processor but the code is not fully optimized, could be parallelized and above all, during this time there is no need of the user, the only user input is a bounding volume. Then the process is robust enough and uses efficiently the redundancy in the images to achieve similar results with or without occlusion (see Figs. 15 and 13, 40 images with a 692×461 resolution on 1.5 m wide line). And a very satisfying point is that it is able to achieve precise results even if the cameras have a very small view angle on the scene (see the top view in Fig. 13e). This illustrates that we can use redundancy to compensate for small angle and occlusion and then achieve a robust and precise depth estimation. Note that the left shoulder of the running woman (Figs. 15d and e) is not seen by any camera because the flying papers always hide it, so it is not reconstructed since we can deal with occlusions as long as they do not exist for some viewpoints. If something is never seen, the algorithm has no chance to “guess it.” But for instance the head of the woman is hidden in half of the images and the algorithm is able to reach a satisfying results although the head resolution is only about 20×30 pixels.

In Fig. 14 (30 images with 400×300 resolution), we work from synthetic images so as to exactly know the objects in the scene: there are a sphere, a cone and a rectangular parallelepiped. We can see that the reconstruction gives the good geometric properties: we actually have a sphere, a cone and a parallelepiped as results although in the original sequence, the sphere is partially hidden by the cone and the parallelepiped. And if we add noise, the algorithm is robust enough to still achieve a valid reconstruction even if the noise intensity is high (Figs. 14c and d). We have estimated the error on the sphere: we have measured the distance between the reconstructed points and the surface. The mean of these distances divided by the radius of the sphere is: 2.4% with no noise (Fig. 14a), 2.9% with a noise with intensity 0.1 (Fig. 14b), and 4.4% with a noise with intensity 0.2 (Figs. 14c and d).

It is important to note that this process is fairly driven from a very limited a priori knowledge of the observed scene. Moreover it is able to cope with textured and untextured objects.




Fig. 13. Illustration of the robustness to occlusion: (a,b) obtained from the original sequence; (c,d) obtained from the same sequence in which a synthetic cube has been added to create a large occlusion; (e) top view of the scene, from the bottom to the top: the cameras, the cube and the man.




Fig. 14. Reconstruction from a set of synthetic images where white noise has been added to the input images. The top row shows a sample of the input images: noise has been added on each RGB component on a scale between 0 and 1, its maximum intensity is: (a) 0, (b) ± 0.1 , (c-d) ± 0.2 .




Fig. 15. Results from images taken simultaneously [33]: (a) an image with the acquired areas; (b,c) the man on the left (≈ 5 m from the cameras); (d,e) the woman ≈ 6.5 m); (f,g) the man on the right.

7.1. Future work

This algorithm gives satisfying results but we believe it can be still improved. We will continue to work on the voxel consistency estimation to get some more precise and more powerful expression on which we can base the process. Then we want to find a more efficient way to find the surface inside the set of the consistent voxels.

From a more global point of view, we aim at characterizing the arbitrary thresholds that appear during the process. Obviously almost all thresholds depend more or less directly on the resolution and on the scale of the reconstruction. And the most challenging work is to extend the algorithm to deal with specular surfaces and any set of cameras.

References

- [1] L. Alvarez, P.-L. Lions, J.-M. Morel, Image selective smoothing and edge detection by nonlinear diffusion. II, *J. Numer. Anal.* 29 (3) (1992) 845–866.
- [2] R. Bellman, *Dynamic Programming*, Princeton University Press, Princeton, NJ, 1957.
- [3] A.C. Brik, J. Williams, J.J. Connor, Multiresolution, incremental generation of 3D computer models from video data, in: *Solid Modeling and Applications*, ACM, New York, 1993, pp. 95–102.
- [4] C. Buehler, S. Gortler, M. Cohen, L. McMillan, Minimal surfaces for stereo, in: *Eur. Conf. on Computer Vision (ECCV 02)*, LNCS, vol. 2352, Springer, Berlin, pp. 885–899.
- [5] V. Caselles, R. Kimmel, G. Sapiro, Geodesic active contours, *Int. J. Comput. Vision (IJCV)* (1997) 694–699.
- [6] F. Catté, P.-L. Lions, J.-M. Morel, T. Coll, Image selective smoothing and edge detection by nonlinear diffusion, *J. Numer. Anal.* 29 (1) (1992) 182–193.
- [7] R. Collins, A space-sweep approach to true multi-image matching, in: *Computer Vision and Pattern Recognition (CVPR 96)*, 1996, pp. 358–363.
- [8] I.J. Cox, A maximum likelihood n -camera stereo algorithm, in: *Computer Vision and Pattern Recognition (CVPR 94)*, IEEE, New York, 1994, pp. 733–739.
- [9] I.J. Cox, S.L. Hingorani, S.B. Rao, B.M. Maggs, A maximum likelihood stereo algorithm, *Comput. Vision Image Understanding* 63 (3) (1996) 542–567.
- [10] R. Deriche, O. Faugeras, Les EDP en traitement des images et vision par ordinateur, Technical report RR-2697, Institut National de Recherche en Informatique et en Automatique (INRIA), 1995.
- [11] M. Desbrun, M. Meyer, P. Schröder, A.H. Barr, Anisotropic feature-preserving denoising of bivariate data, in: *Graphics Interface (GI 00)*, 2000.
- [12] O. Faugeras, R. Keriven, Variational principles, surface evolution, PDE's, level set methods and the stereo problem, *Trans. Image Process.* (1998).
- [13] P. Fua, Y.G. Leclerc, Object-centered surface reconstruction: combining multi-image stereo and shading, *Int. J. Comput. Vision (IJCV)* (1995).
- [14] R. Hartley, A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge University Press, Cambridge, 2000.
- [15] S.S. Intille, A.F. Bobick, Disparity-space images and large occlusion stereo, in: J.-O. Eklundh (Ed.), *Eur. Conf. on Computer Vision (ECCV 94)*, vol. B, Springer, Berlin, 1994, pp. 179–186.
- [16] H. Ishikawa, Global optimization using embedded graphs, Ph.D. thesis, New York University, 2000.
- [17] R. Koch, M. Pollefeys, L. Van Gool, Multi viewpoint stereo from uncalibrated video sequences, in: *Eur. Conf. on Computer Vision (ECCV 98)*.
- [18] V. Kolmogorov, R. Zabih, Multi-camera scene reconstruction via graph cuts, in: *Eur. Conf. on Computer Vision (ECCV 02)*, LNCS, vol. 2352, Springer, Berlin, 2002, pp. 82–96.
- [19] K. Kutulakos, S. Seitz, A theory of shape by space carving, in: *Int. Conf. on Computer Vision (ICCV 99)*, vol. I, IEEE, New York, 1999, pp. 307–314.
- [20] M. Lipshutz, *Differential Geometry*, McGraw-Hill, New York, USA, 1969.
- [21] P.R.S. Mendonça, K.-Y.K. Wong, R. Cipolla, Camera pose estimation and reconstruction from image profiles under circular motion, in: *Eur. Conf. on Computer Vision (ECCV 00)*, vol. 2, Springer, Berlin, 2000, pp. 864–878.
- [22] Y. Ohta, T. Kanade, Stereo by intra- and inter-scanline search using dynamic programming, *Trans. Pattern Anal. Mach. Intell. (PAMI)* (1985).

- [23] M. Okutomi, T. Kanade, A multiple-baseline stereo, *Trans. Pattern Anal. Mach. Intell. (PAMI)* 15 (4) (1993) 353–363.
- [24] P. Perona, J. Malik, Scale-space and edge detection using anisotropic diffusion, *Trans. Pattern Anal. Mach. Intell. (PAMI)* 12 (7) (1990) 629–639.
- [25] S. Roy, I. Cox, A maximum-flow formulation of the n -camera stereo correspondence problem, in: *Int. Conf. on Computer Vision (ICCV 98)*, IEEE, New York, 1998, pp. 492–499.
- [26] H. Saito, T. Kanade, Shape reconstruction in projective grid space from large number of images, in: *Computer Vision and Pattern Recognition (CVPR 99)*, IEEE, New York, 1999, pp. 49–54.
- [27] S.M. Seitz, C.R. Dyer, Photorealistic scene reconstruction by voxel coloring, *Int. J. Comput. Vision (IJCV)* 35 (2) (1999) 151–173.
- [28] G. Slabaugh, B. Culbertson, T. Malzbender, R. Schafer, Improved voxel coloring via volumetric optimization, Technical report, Center for Signal and Image Processing, Georgia Institute of Technology, 2000.
- [29] G. Slabaugh, B. Culbertson, T. Malzbender, R. Schafer, A survey of methods for volumetric scene reconstruction from photographs, in: K. Mueller, A. Kaufmann (Eds.), *VolumeGraphics*, Springer, Berlin, 2001, pp. 81–100.
- [30] G. Slabaugh, T. Malzbender, W.B. Culbertson, Volumetric warping for voxel coloring on an infinite domain, in: M. Pollefeys, L. Van Gool, A. Zisserman, A. Fitzgibbon (Eds.), *3D Structure from Images (SMILE 00)*, LNCS, vol. 2018, Springer, Berlin, 2000, pp. 109–123.
- [31] R. Szeliski, P. Golland, Stereo matching with transparency and matting, in: S. Chandran, U. Desai (Eds.), *Int. Conf. on Computer Vision (ICCV 98)*, New Delhi, Narosa Publishing House, 1998, pp. 517–526.
- [32] R. Szeliski, R. Weiss, Robust shape recovery from occluding contours using a linear smoother, Technical report DEC-CRL-93-7, Digital Equipment Corporation, Cambridge Research Laboratory, 1993.
- [33] D. Taylor, <http://www.virtualcamera.com>.
- [34] J. Tumblin, G. Turk, LCIS: a boundary hierarchy for detail-preserving contrast reduction, in: A. Rockwood (Ed.), *Siggraph, Annual Conference Series*, ACM Siggraph, Addison-Wesley, Longman, 1999, pp. 83–90.
- [35] M. Ulvklo, H. Knutsson, G.H. Granlund, Depth segmentation and occluded scene reconstruction using ego-motion, in: S.K. Park, R.D. Juday (Eds.), *Visual Information Processing*, SPIE, vol. 3387, 1998, pp. 112–123.
- [36] R. Vaillant, O.D. Faugeras, Using extremal boundaries for 3-D object modeling, *Trans. Pattern Anal. Mach. Intell.* 14 (2) (1992) 157–173.
- [37] O. Veksler, Efficient graph-based energy minimization methods in computer vision, Ph.D. thesis, Cornell University, 1999.
- [38] Y.-L. You, M. Kaveh, Fourth order partial differential equations for noise removal, *Trans. Image Process.* 9 (10) (2000).