Physically Valid Triangulation of Sparsely Matched Images Using Texture Information: Application to View-Synthesis*

J.S. Perrier, G. Agam, and P. Cohen

Perception and Robotics Laboratory
Department of Electrical and Computer Engineering
Ecole Polytechnique, Montreal H3C 3A7, Canada

Abstract

Image-based view synthesis is an emerging research topic, with numerous applications. The ability to synthesize new views from existing images, enables the generation of enhanced and synthetic environments. Approaches for view synthesis relay upon dense or sparse matching of the views. In both cases, some parts of the images are inevitably unmatched. Triangulation of sparsely matched points present a possible solution for the handling of those unmatched regions. However, such triangulation should respect the underlying geometry of the scene.

In this paper, a novel approach is proposed for physically valid triangulation of sparsely matched points. The proposed approach is based on the maximization of a physical validity criterion which is supported by textured regions in the images. The produced triangulation is such that each triangle corresponds approximately to a planar surface in the scene. Given an arbitrary initial triangulation, the proposed approach refines it by flipping the edges of triangles. Furthermore, since missing matched points may preclude the correct triangulation of the scene, an additional stage handles the addition of matched points inside low score triangles. Inherent to our approach, the support region which is used for the evaluation of the correctness of the added matches is normally much larger than the one used in a local match evaluation. The paper contains examples of image-based view synthesis of real and virtual scenes, produced by the proposed approach.

1 Introduction

In recent years, image-based view synthesis has become the focus of intensive research for applications that range from the efficient photo-realistic rendering of complex virtual or real scenes to the efficient compression of video sequences. The different approaches to image-based view synthesis could be classified into methods that include an explicit 3D reconstruction or some other complete representation, and methods that do so only implicitly. The various methods differ in their calibration needs, the required number and closeness of the input views, the density and accuracy of required matched points between the views, and the way in which the information is re-projected to create new views. When the input views are strongly calibrated and densely matched, a complete 3D model of the scene can be estimated and used to generate arbitrary new views. Arbitrary views may also be produced from a voxel [11] or plenoptic [8] representations, which could be generated based upon a dense set of calibrated and densely matched input views. With three densely matched input views, the trilinear tensor can be retrieved to generate new views by either reconstructing a 3D model or re-projecting the input views directly [2]. With two densely matched input views, the fundamental matrix can be computed and used to generate new views by re-projecting the pixels of the input views onto a new image plane [6]. Finally, without calibration, by using only a set of matched points between two images, view interpolation and view morphing techniques [4, 10, 12] can be used to generate new images. A thorough review and classification of image-based view synthesis techniques may be found in [1].

While many methods for image-based view synthesis are based upon the establishment of dense matchings between corresponding views, the sole reliance on densely matched points is sensitive to local match errors and cannot produce results for uniform (non-textured) regions in which dense matching cannot be reliably obtained. Consequently, such methods do not necessarily guarantee improved results with respect to methods which are based on a sparse, cleverly selected set of matched points. Moreover, sparse matchings may be obtained more efficiently and with greater reliability by using robust techniques. Scenes, particularly man-made

*This work is supported by the IRIS-3 Network of Centers of Excellence (Project SMART NR-COH).

ones, can often be successfully approximated by a set of relatively small planar surfaces in image space. For instance, a relatively flat object seen from far away can safely be approximated by a single planar surface. Hence, a correct triangulation of a carefully selected set of sparsely matched points can be sufficient for view synthesis.

A triangulation approach to image-based view synthesis is capable of handling unmatched regions by representing them with a set of triangles where the vertices of the triangles lie on the boundaries of the unmatched regions. Moreover, a triangulation approach is considerably more efficient in terms of number of operations and memory utilization, and may be accelerated by standard graphics hardware. Existing approaches for joint triangulation of matched views, such as the one described in [7], are based on an initial dense set of matched points which are grouped into triangles by merging neighboring matched points that respect the same planar homography. The disadvantage of such methods lie, again, in their reliance on dense matchings, which are inherently unreliable. The joint triangulation of a sparse set of matched points, constrained by linear segments corresponding to boundaries of planar surfaces in the image, is described in [5, 3]. The correctness of the produced triangulation, in this approach, depends upon the correct and complete identification of linear segments at the boundaries of planar surfaces in the image. However, since an automated algorithm can hardly distinguish between edges inside textured surfaces and edges which bound them, and since some surface boundaries may be missing, the approach resorts to the manual specification of constraints by the user.

Contrary to the approaches described above, our proposed method is based on the direct triangulation of a sparse set of matched points using a texture-based triangulation constraint. The sparse set of matched points is automatically triangulated to produce triangles approximating planar surfaces in the scene where the evaluation of the approximation correctness is based upon the texture inside the triangles. Since the support region used to evaluate the correctness of a triangle is normally much larger than the one used for evaluating a match between individual points, the results are less sensitive to local ambiguities in the data and local match errors are eliminated.

Figure 1 presents the flow chart of the proposed approach. Given a sparse set of matched points in two views, a Delaunay triangulation is used as an initial approximation of the correct triangulation. This initial triangulation is then refined based on an evaluation of the correspondence of each triangle to a planar surface in the scene. This evaluation is obtained by measuring the correlation between the textures of matched triangles after rectifying them to a common viewpoint where one can expect triangles relating to a real planar surface in the scene to have a large correlation value. Since the initial set of matched points could lack some points precluding a correct triangulation in some regions, an additional stage targets the extraction of new matched points in triangles with low texture correlation. A global verification for the added matched points is obtained through the estimation of the ensuing triangulation, thus reducing the susceptibility to match errors.

The following sections provide further details of the proposed approach. The process of rectifying matched triangles in order to assess the validity of the planarity assumption for the corresponding scene surface is described in Section 2. The actual evaluation of the planarity assumption is discussed in Section 3. The algorithms for modifying the triangulation are detailed in Section 4. Results and conclusions are provided in Section 5.

2 Triangles Rectification

Given matched triangles in two views, it is possible to verify that they correspond to a planar surface in the scene by rectifying them to a common viewpoint and measuring the correlation between them. Since the rectification process assumes that the matched triangles represent a planar surface, it will create discrepancies if this assumption is incorrect. The rectification of the triangles may be obtained by an affine or perspective mapping. Different mapping approaches are discussed in Section 2.1, and their evaluation is presented in Section 2.2.

2.1 Warping the triangles

Given matched triangles in two views, their rectification to a common viewpoint is obtained by mapping the texture from one triangle to the other, which may be obtained by well known affine or perspective warping techniques [13].

Representing the three vertices of a triangle in homogeneous coordinates $A = (A_x, A_y, 1)^T$, $B = (B_x, B_y, 1)^T$ and $C = (C_x, C_y, 1)^T$, a point $P = (P_x, P_y, 1)^T$ within it may be expressed by a linear combination of its vertices: $P = \alpha A + \beta B + \gamma C = M(\alpha, \beta, \gamma)^T$, where M is a 3×3 matrix having A, B, and C as its columns. Using the linear combination coefficients $(\alpha, \beta, \gamma)^T$, the corresponding point P' in the second triangle is obtained by: $P' = M'(\alpha, \beta, \gamma)^T = M'M^{-1}P$ where M' is composed of the vertices of the second triangle. Hence, in the affine case, the transformation of a point is given by: $H = M'M^{-1}$.

The perspective transformation relates two different projections of a planar region. It is defined as a one to one mapping such that: $\lambda P' = HP$ where H is a 3×3 non-singular matrix. Since the plane homography H is known only up to a scale factor λ, it has only 8 degrees of freedom. Hence, four matched pairs of points are needed in order to solve the linear system.

Since images are perspective projections of a scene, the perspective transformation is expected to produce more real-

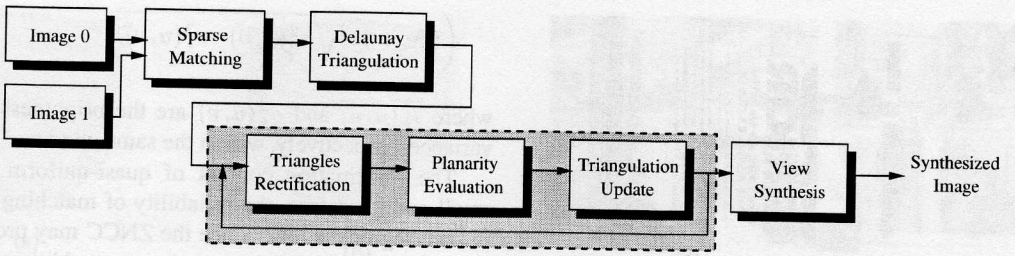

Figure 1: Flow chart of the proposed approach.

istic results with respect to the affine transformation. Since at least four pairs of points are required to uniquely determine the perspective transformation, a possibility consists of pairing neighboring triangles in order to have four points. Another possibility consists of using the epipoles of the image planes as the fourth pair of points. Any planar homography will map the epipole of one image onto the epipole of the second image [9, Chapter 5].

The first approach to selecting four pairs of points consists of selecting a pair of triangles, assuming it corresponds to the same plane in the scene. The three triangles adjacent to each triangle should be considered as pairing candidate. The evaluation of the correctness of the assumption that two adjacent triangles correspond to the same 3D plane, may be obtained by using an affine transformation to map the texture within the pair of triangles onto the second image, and measuring the match between them using equation 3 in section 3.1. In addition to its increased computational complexity, a second drawback of this approach is the imposed need to pair the triangles again whenever their configuration is modified during the re-triangulation process.

The second approach to that problem, assumes that the epipolar geometry between the images has been recovered through the computation of the fundamental matrix. Since matching points between the two views are available, this may be accomplished by using one of the methods described in [14]. Nevertheless, depending on the configuration of the cameras, small inaccuracies in matched points may lead to large errors in the epipole positions. Furthermore, since the epipoles may be far from the triangles, the perspective texture mapping may become highly sensitive to numerical instabilities. Particularly, as the images become co-planar, the epipoles migrate to infinity, thus precluding their use.

Consequently, affine transformations should be considered as an approximation to the rectification of the triangles, despite the fact that they do not describe the real planar transformation. A comparison between rectification results obtained by affine and perspective mappings are described in the following section.

2.2 Evaluating affine and perspective mappings

In order to evaluate the validity of affine warping with respect to perspective warping, a controlled test has been conducted on several sequences of images. Each test sequence contained 6 images of an object rotated by about 10 degrees per images. Hence, there is a rotation of about 50 degrees between the first and last images in each sequence. Five triangles, corresponding to the same planar surface of the object, were defined in the first image and tracked in the consequent images. The triangles of each image in the sequence were warped to their corresponding ones in the first image and a normalized zero-mean correlation was applied to measure their similarity. The compared methods include affine warping and perspective warping which is based on the pairing of triangles.

Figure 2-a presents the 6 images from one of the test sequences. Figure 2-b shows the triangles used. The left and right graphs in Figure 3 show the texture correlation results for affine and perspective warping respectively. The triangle pairing algorithm used in perspective warping tries the pairing of each triangle with each of its neighbors and selects the one that produces the highest texture correlation after affine warping. As may be observed, the results of the two texture mapping approaches are very similar. The effect of the perspective correction is mostly noticeable around a rotation of 30 and 40 degrees. The fact that the perspective warping does not produce much higher correlation with respect to the affine warping when the rotation is large is due to the large difference in resolution of the triangles. For instance, triangle #1 in the first image of the sequence has 3395 pixels, and only 950 pixels in the last image of the sequence. The above experiment shows, as expected, that the distortions introduced by the affine warping are significantly smaller compared to the large differences in areas of corresponding triangles.

3 Planarity Evaluation

Based upon the rectification process detailed in the previous section, the validity of the planar approximation assumption

(a) (b)

Figure 2: **(a)** The 6 frames of the rotated box scene used in the perspective correction test. **(b)** The triangle configuration used for the perspective correction test.

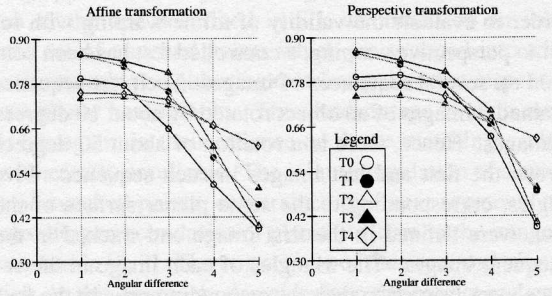

Figure 3: The left and right graphs show texture correlation results for the affine and perspective warpings respectively.

may be verified by measuring the correlation between the rectified triangles in one image and the triangles in the second image. It is expected that the correlation between triangles belonging to an actual planar surface in the scene will be high whereas, otherwise, the rectification process would produce discrepancies which reduce the correlation between the triangles. The approach for measuring triangle matchings is discussed in Section 3.1 and a detailed assessment of the influence of 3D rotation on the match evaluation is outlined in Section 3.2.

3.1 Measuring triangles match

Given two views, I_0 and I_1, the rectification process produces a third image, I'_0, containing all the rectified triangles warped from I_0 to I_1. The evaluation of the texture correlation is performed between I'_0 and I_1. The zero-mean normalized cross-correlation measure (ZNCC) is used. The ZNCC measure, between images I_0 and I_1, over a window of $(2n+1) \times (2n+1)$ centered at (u, v), is defined by:

$$\mathrm{ZNCC}_n(u,v) = \sum_{i=-n}^{n} \sum_{j=-n}^{n} ([I_0(u+i, v+j) - \overline{I_0(u,v)}] \cdot [I_1(u+i, v+j) - \overline{I_1(u,v)}]) \Big/ \left((2n+1)^2 \sqrt{\sigma_0^2(u,v) \cdot \sigma_1^2(u,v)}\right) \quad (1)$$

where $\overline{I_k(u,v)}$ and $\sigma_k^2(u,v)$ are the brightness mean and variance, respectively, within the same window.

The information content of quasi-uniform regions is small and, therefore, the reliability of matching evaluation in such regions is low. Since the ZNCC may produce good correlation results for such regions, an additional measure, called Normalized Joint Variance (NJV), is introduced. The NJV at point (u, v) is based on the local variance over a $(2n+1) \times (2n+1)$ window around that point. It is defined as:

$$\mathrm{NJV}(u,v) = \frac{\max(\sigma_0^2(u,v), \sigma_1^2(u,v))}{\sigma_m^2} \quad (2)$$

where σ_m^2 is the variance maximal value within the images. The normalization by σ_m^2 guarantees that $\mathrm{NJV}(u,v) \in [0,1]$.

The NJV measure is used to weight the ZNCC correlation measure so that low-variance regions are given less weight in the matching evaluation. Let T_{ABC} be the set of pixels contained in the triangle formed by vertices A, B and C. The matching score that measures the similarity between the textures of that triangle in images I'_0 and I_1, is defined as:

$$\mathrm{Match}(T_{ABC}) = \frac{\sum_{(u,v)\in T_{ABC}} \mathrm{ZNCC}_n(u,v) \cdot \mathrm{NJV}(u,v)}{\sum_{(u,v)\in T_{ABC}} \mathrm{NJV}(u,v)} \quad (3)$$

The matching score is normalized by the total amount of activity in the triangle. Since $\sum_{(u,v)\in T_{ABC}} \mathrm{NJV}(u,v) \geq \sum_{(u,v)\in T_{ABC}} \mathrm{NJV}(u,v) \cdot \mathrm{ZNCC}_n(u,v)$, we get that $\mathrm{Match}(T_{ABC}) \in [0,1]$. The matching score, as defined here, gives a good assessment of matching for highly correlated textured regions only.

Finally, an additional confidence measure of the reliability of the information contained in a triangle is defined as the average NJV over that triangle:

$$\mathrm{Conf}(T_{ABC}) = \frac{\sum_{(u,v)\in T_{ABC}} \mathrm{NJV}(u,v)}{\#T_{ABC}} \quad (4)$$

where $\#T_{ABC}$ is the cardinality of the set T_{ABC}. This confidence measure is used in the refinement step, described in section 4.2.

Figure 4 summarizes the sequence of operations performed during the evaluation of the planarity assumption. The original views I_0 and I_1 are presented in Figures 4-a and 4-b, respectively. The warp I'_0, created by rectifying the triangles of I_0, is presented in Figure 4-c. The discrepancies created by the rectification in the texture of two incorrect triangles are clearly visible. Figures 4-d and 4-e show the result of the NJV and weighted ZNCC between I'_0 and I_1 respectively. As can be observed, the weighted ZNCC

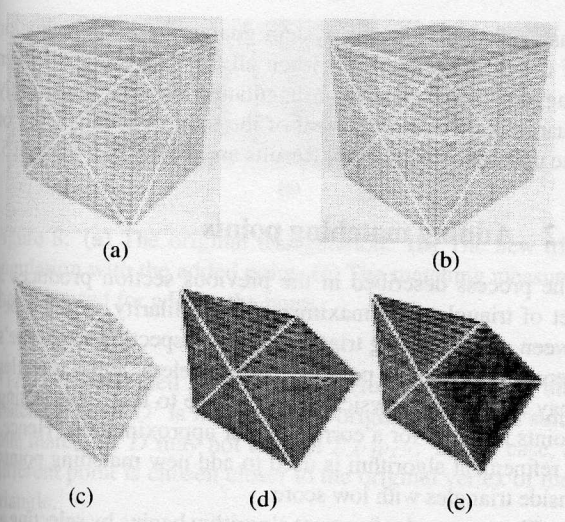

Figure 4: **(a)–(b)** The two views with the initial Delaunay triangulation. **(c)** The result of warping the triangles of the first image. **(d)** The result of the NJV measure. **(e)** The result of the NJV · ZNCC_2 measure.

measure is lower in the regions corresponding to incorrect triangles.

The matching measure defined in equation 3 is sensitive to small misalignments between the triangles. Due to numerical errors in the rectification process, even small correspondence errors between triangle vertices may systematically diminish the correlation results. In order to cope with this problem, it is not assumed that $I'_0(u,v)$ is in exact correspondence with $I_1(u,v)$. Instead, the ZNCC measure is evaluated between $I'_0(u,v)$ and $I_1(u+k, v+l)$, where $|k|, |l| \leq m$. The maximum ZNCC result, obtained in the $(2m+1) \times (2m+1)$ search region around (u,v), replaces $\text{ZNCC}_n(u,v)$ in equation 3.

3.2 Evaluating the influence of 3D rotations

Incorrect results obtained by the texture similarity measure defined in equation 3 could mislead the triangulation process. Therefore, an evaluation of its robustness to scene transformations such as 3D rotation, is essential.

In order to evaluate the performance of the measure under 3D scene rotations, several tests where conducted on real and virtual scenes. In each test, several images of a gradually rotated scene were taken. The goal of the tests was to evaluate the behavior of physically valid triangles with respect to invalid ones. The images used in one of the tests are presented in Figure 5. Figure 5-a, shows 4 out of the 13 frames in the sequence. The triangles used in the test are illustrated in Figure 5-b. The simplicity of the scene facilitates the tracking of the various types of triangles in the sequence. The different types of triangles in this scene

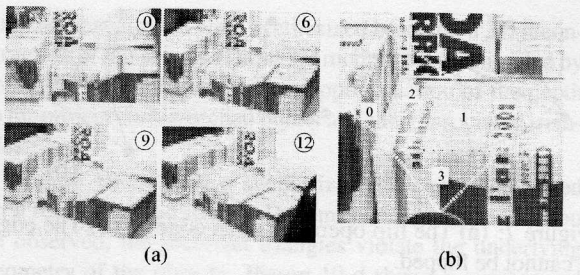

Figure 5: **(a)** Some frames of the test sequence. **(b)** The triangles used in the test.

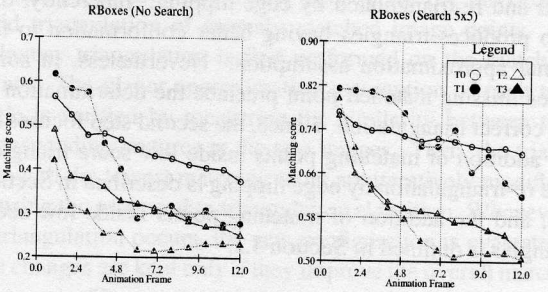

Figure 6: Results of the similarity measure of the four triangles shown in Figure 5-b.

are: T_0 – a vertically rotated physically valid triangle; T_1 – a horizontally rotated physically valid triangle; T_2 – an invalid triangle connecting objects by a plane over the scene; T_3 – an invalid triangle connecting objects by a plane through the scene.

The results of the similarity measure for the four triangles tracked in the 13 images are shown in Figure 6. Figure 6-a shows the similarity measure obtained without a search region, whereas Figure 6-b presents the measure obtained with a search region 5×5 pixels. As could be observed, the separation between the physically valid and invalid triangles is increased with the increase of the search region. Additionally, when using a search region, the matching scores of the triangles are globally higher, and the separation between valid and invalid triangles for very small rotations is lower. Consequently, the search region should be used only for views of large rotations. This evaluation demonstrates the good behavior of the matching measure. Even without a search region, the score of the invalid triangles is usually lower than that of the valid ones.

4 Triangulation Update

After ranking the validity of all the triangles in the images, the next stage of our approach is to modify triangles that do

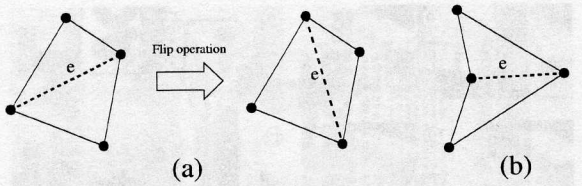

Figure 7: **(a)** The flip operation of an edge e. **(b)** The edge e cannot be flipped.

not conform to the planar approximation assumption. In this stage there are two steps. First, low score triangles are coupled and re-triangulated by edge flipping. Inherently, this step produces triangles having better conformation to the planar approximation assumption. Nevertheless, in some cases missing matched point preclude the determination of the correct triangulation. Hence, the second step focuses on the addition of matching points inside low score triangles. The re-triangulation by edge flipping is described in Section 4.1, and the addition of matching points inside low score triangles is outlined in Section 4.2.

4.1 Updating the triangulation through edge flipping

Given an initial Delaunay triangulation of the scene and a matching score for each of the triangles, an edge flipping re-triangulation algorithm is used to generate a modified triangulation with an improved total matching score. Edge flipping for re-triangulation is based upon an iterative process in which edges offending some property are flipped, provided that the shape formed by the two triangles sharing the edge is convex. An example of such algorithm is the conversion of an arbitrary triangulation into a Delaunay triangulation. The flipping operation is illustrated in Figure 7. In Figure 7-a two neighboring triangles are re-triangulated by flipping their common edge. In Figure 7-b, the neighboring triangles do not form a convex quadrilateral, so the flipping is not possible.

Re-triangulation by edge flipping is used here to transform a Delaunay triangulation into a physically valid triangulation by flipping edges that are not physically valid. The triangles are paired for flipping based on their matching score. Pairs of triangles in which flipping is tested are the lowest score triangle and its lowest score neighbor, provided that the quadrilateral formed by them is convex. After each flip, the matching score of the new triangles is evaluated, and the flip is kept only if the total matching score is improved. A flagging system is used to mark triangles for which a tested flip did not improve the matching score and triangles that do not form a convex quadrilateral. In order to allow the re-triangulation to propagate properly, flagged triangles are unflagged if at least one of their neighbors has changed. This flagging system guarantees the termination of the flipping algorithm when all the triangles have been flagged. The proposed re-triangulation algorithm inherently guarantees the improvement of the total matching score of the updated triangulation. Results are provided in Section 5.

4.2 Adding matching points

The process described in the previous section produces a set of triangles that maximizes the similarity measure between corresponding triangles, thus, respecting the scene's geometry whenever possible. Nevertheless, some triangles may still remain physically invalid due to lack of matching points required for a correct planar approximation. Hence, a refinement algorithm is used to add new matching points inside triangles with low score.

The proposed refinement algorithm begins by selecting a triangle with low matching score and high confidence score (see Section 3). The high confidence score is necessary in order to verify correctly the added match. Let T_{ABC} be the selected triangle in image I_0 and $T_{A'B'C'}$ its corresponding triangle in image I_1. A match candidate point D is selected as a point with high local variance, inside T_{ABC}. A potential matching point D' is computed by using the warping equation between T_{ABC} and $T_{A'B'C'}$. The point matching D is searched for in the vicinity of D'. In order to increase the likelihood of the prediction D', and to reduce the required search in the vicinity of D', the point D is selected in the proximity of A, B, or C.

The evaluation of the matching between two image points, based on their local neighborhood, is ambiguous and unreliable. This is particularly true for D and D', since they were not detected by the initial, robust, feature-based matching algorithm. The resulting matched pair (D, D') is used to split each of the original triangles into three triangles. The correctness of the match is then verified by measuring the combined matching score of the generated triangles given by: $F(D, D') = Match(T_{ABD}) + Match(T_{ACD}) + Match(T_{BCD})$.

The measure $F(D, D')$ is then computed for points in a growing neighborhood of D', until a global maximum of $F(D, D')$ is found inside the triangle $T_{A'B'C'}$. The constraint of looking for the added matched point D' inside the triangle $T_{A'B'C'}$ is imposed in order to maintain consistency with the underlying triangulation. Since the domain of possible positions of D' is bounded, it is possible to find the maximum value by a systematic exploration. Unconstrained optimization methods, such as the steepest descent approach, can be used to speedup the process. Figure 8 illustrates the process of adding matching points. Figure 8-a presents the original triangulation in which a missing matched point precludes the correct triangulation of the sides of the cube. In Figure 8-b, an additional match has been added based upon the evaluation of the measure

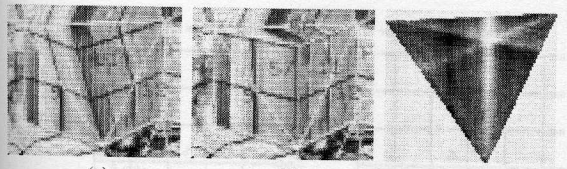

(a) (b) (c)

Figure 8: **(a)** The original triangulation. **(b)** The new triangulation with the added point. **(c)** The matching measure that was used for adding the point.

$F(D, D')$ presented in Figure 8-c. However, because the matching point D' is limited to the original triangle, a suitable match for D may not exist in $T_{A'B'C'}$. In that case, a different point is chosen closer to the original vertex of the triangle.

After adding a new matching point inside an existing triangle, the re-triangulation process described in Section 4.1 is applied locally, in order to get the best possible triangle configuration before proceeding to add new matching points.

5 Results and Conclusion

In this section, some results of the re-triangulation algorithm are presented. The examples provided are composed of simple scenes in order to clearly demonstrate the effect of a physically valid triangulation. The first scene, shown in Figures 10-a and 10-b, consists of a pile of boxes with different textures, where the matching points between the views are located at the visible corners of the boxes. The result of warping one view into the other based on the initial Delaunay triangulation of the matching points is shown in Figure 9-a. The discrepancies generated by the invalid triangles are visible on each of the three large boxes. Figure 9-b presents the warping results obtained by using the triangulation produced by the proposed algorithm. As can be observed, the previously invalid triangles have been corrected. In addition, some of the valid triangles on the small boxes where changed to a better configuration. The changes in valid triangles occur due to the fact that the re-triangulation algorithm managed to increase the matching score. Although less visible, this change in triangulation reduces the affine texture mapping distortion. The quantitative improvement on the matching scores is presented in the graphs in Figure 9. In these graphs each point represents the matching and confidence scores of a triangle. Figures 9-c and 9-d show the distribution of the matching scores before and after the re-triangulation, respectively. As expected, the matching score of the individual triangles is improved.

Further illustration of the effects of correct triangulation is presented through the application of the proposed approach to image-based view synthesis, in Figure 10. The results in these figures were synthesized based on a 3D reconstruction of the scenes. Those 3D models where obtained by using the stereo disparity as an approximation of the depth to un-project each matched points and their associated triangles.

Figure 10-c shows a synthesized view which is based on a reconstruction using a Delaunay triangulation. As can be observed, the incorrect triangles violate the underlying geometry of the objects. Figure 10-d shows a synthesized view which is based on a reconstruction using the physically valid triangulation of the proposed approach. In this case all the triangles respect the geometry the scene.

This paper proposed a new method for the physically valid triangulation of sparsely matched image points. A Delaunay triangulation is first performed on the matched points. The planar approximation assumption of each triangle is verified by measuring the similarity between the corresponding textures in the two images. Then, the triangles with the lowest matching score are iteratively modified by using the proposed re-triangulation algorithm. Whenever re-triangulation occurs, the new configuration is evaluated. The changes are kept only if they improve the overall matching scores. The re-triangulation process stops when a configuration that can no longer be improved is reached. Since some triangles may still violate the planarity assumption, a matching refinement algorithm based on the triangle similarity measure is used. Results for real and synthetic scenes demonstrate the advantages of the proposed approach.

The proposed approach can be adapted to a large number of view synthesis techniques ranging from simple view interpolation to complete 3D reconstruction. It can be used in automated or semi-automated 3D CAD modeling systems and various computer graphics algorithms such as view-dependent mesh optimization. In augmented reality applications, where the correct positioning of graphical models in an image is required, the problem may be considerably simplified by using the proposed approach to generate a view-dependent mesh of the scene.

References

[1] G. Agam, G. Michaud, J. S. Perrier, J. L. Houle, and P. Cohen. A survey of image based view synthesis approaches for interactive 3D sensing. Technical Report GRPR-RT-9901, Perception and Robotics Laboratory, Ecole Polytechnique, Montreal, Canada, April 1999.

[2] S. Avidan and A. Shashua. Novel view synthesis in tensor space. In *Proceedings. 1997 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1034–1040, San Juan, Puerto Rico, June 1997.

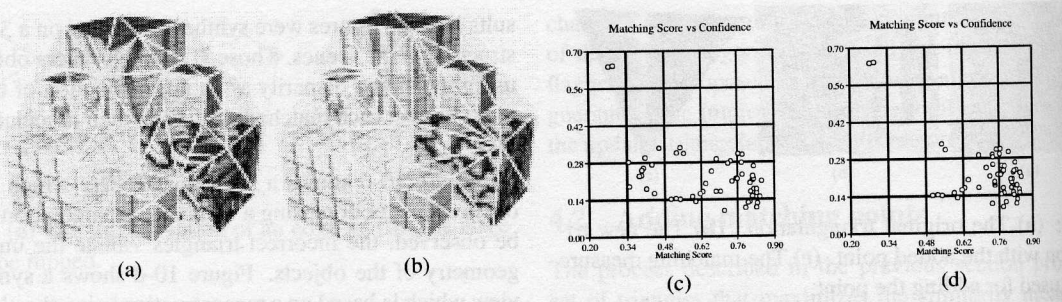

Figure 9: Re-triangulation results for the scene in Figures 10-a and 10-b. **(a)** Initial Delaunay triangulation of the matching points. **(b)** Physically valid triangulation generated by the proposed algorithm. **(c)–(d)** Matching score and confidence measure of each triangle before and after the re-triangulation, respectively.

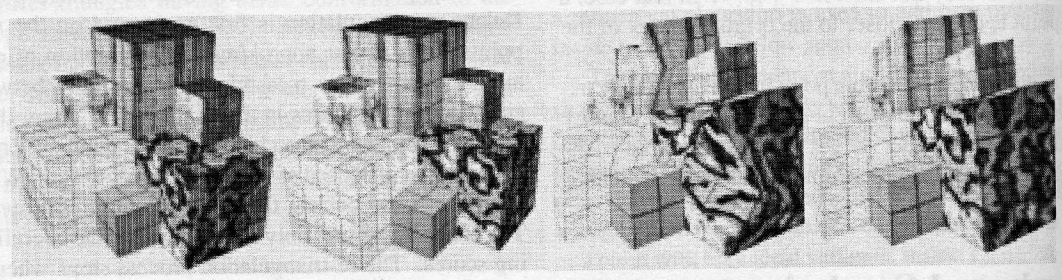

Figure 10: **(a)–(b)** Two views used for the synthesis of new views. **(c)–(d)** 3D models resulting from Delaunay and the proposed triangulation, respectively.

[3] Qian Chen and Gerard Medioni. Image synthesis from a sparse set of views. In *Visualization97*, pages 269–275, 1997.

[4] S.E. Chen and L. Williams. View interpolation for image synthesis. In *Computer Graphics Proceedings*, pages 279–288, Anaheim, CA, USA, Aug. 1993.

[5] P. Havaldar, Mi-Suen Lee, and G. Medioni. View synthesis from unregistered 2-d images. In W.A. Davis and R. Bartels, editors, *Proceedings. Graphics Interface '96*, pages 61–9, Toronto, Ont., Canada, May 1996.

[6] Stephane Laveau and Olivier Faugeras. 3-D scene representation as a collection of images and fundamental matrices. This paper shows how to synthesize new images without a 3-D model, February 1994.

[7] Maxime Lhuillier. Towards automatic interpolation for real and distant image pairs. Technical report, INRIA, Feb 1999.

[8] L. McMillan and G. Bishop. Plenoptic modeling: an image-based rendering system. In R. Cook, editor, *Computer Graphics Proceedings. SIGGRAPH 95*, pages 39–46, Los Angeles, CA, USA, Aug. 1995.

[9] Leonard Jr. McMillan. *An Image-Based Approach to Three-Dimensional Computer Graphics*. PhD thesis, Computer Science Dept, University of North Carolina, Chapel Hill, 1997.

[10] S.M. Seitz and C.R. Dyer. Physically-valid view synthesis by image interpolation. In *Proceedings IEEE Workshop on Representation of Visual Scenes*, pages 18–25, Cambridge, MA, USA, June 1995.

[11] S.M. Seitz and C.R. Dyer. Photorealistic scene reconstruction by voxel coloring. In *Proc. Computer Vision and Pattern Recognition Conference*, pages 1067–1073, 1997.

[12] T. Werner, R.D. Hersch, and V. Hlavac. Rendering real-world objects using view interpolation. In *Proceedings. Fifth International Conference on Computer Vision*, pages 957–62, Cambridge, MA, USA, June 1995.

[13] G. Worlberg. *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, California, 1990.

[14] Zhengyou Zhang. Determining the epipolar geometry and its uncertainty: A review. Technical report, INRIA, July 1996.