

Hand-held acquisition of 3D models with a video camera

M. Pollefeys¹, R. Koch², M. Vergauwen and L. Van Gool
ESAT-PSI, K.U.Leuven
Kardinaal Mercierlaan 94
B-3001 Heverlee, Belgium
Marc.Pollefeys@esat.kuleuven.ac.be

Abstract

Modeling of 3D objects from image sequences is a challenging problem and has been a research topic for many years. Important theoretical and algorithmic results were achieved that allow to extract even complex 3D scene models from images. One recent effort has been to reduce the amount of calibration and to avoid restrictions on the camera motion. In this contribution an approach is described which achieves this goal by combining state-of-the-art algorithms for uncalibrated projective reconstruction, self-calibration and dense correspondence matching.

1. Introduction

Obtaining 3D models from objects is an ongoing research topic in computer vision. A few years ago the main applications were robot guidance and visual inspection. Nowadays however the emphasis is shifting. There is more and more demand for 3D models in computer graphics, virtual reality and communication. This results in a change in emphasis for the requirements. The visual quality becomes one of the main points of attention.

The acquisition conditions and the technical expertise of the users in these new application domains can often not be matched with the requirements of existing systems. These require intricate calibration procedures every time the system is used. There is an important demand for flexibility in acquisition. Calibration procedures should be absent or restricted to a minimum.

Additionally, the existing systems are often built around specialized hardware (e.g. laser range scanners or stereo rigs) resulting in a high cost for these systems. Many new applications however require robust low cost acquisition

¹Postdoctoral Fellow of the Fund for Scientific Research - Flanders (Belgium) (F.W.O.)

²Now at the Multimedia Information Processing, Inst. for Computer Science, Universität Kiel, Germany

systems. This stimulates the use of consumer photo- or video cameras.

Other researchers have presented systems for extracting 3D shape and texture from image sequences acquired with a freely moving camera. The approach of Tomasi and Kanade [28] used an affine factorization method to extract 3D from image sequences. An important restriction of this system is the assumption of orthographic projection.

Debevec, Taylor and Malik [4] proposed a system that starts from an approximate 3D model and camera poses and refines the model based on images. View dependent texturing is used to enhance realism. The advantage is that only a restricted number of images are required. On the other hand a preliminary model must be available and the geometry should not be too complex.

In this paper we present a system which retrieves a 3D surface model from a sequence of images taken with off-the-shelf consumer cameras. The user acquires the images by freely moving the camera around the object. Neither the camera motion nor the camera settings have to be known. The obtained 3D model is a scaled version of the original object (i.e. a *metric* reconstruction), and the surface texture is obtained from the image sequence as well.

Our system uses full perspective cameras and does not require prior models nor calibration. Some prior versions of this system have been presented previously [20, 21, 22].

Some important enhancements have been added to it. The images are not only related to the previous and the following image of the sequence, but are related to all the images which are close. This allows a more stable computation of the preliminary reconstruction and calibration. A second enhancement consist of a generalized rectification scheme. The traditional rectification schemes are not usable when the motion of the camera contains an important forward component (i.e. epipoles in the images). The new approach works in all cases and additionally provides minimal image sizes.

The complete system combines state-of-the-art algorithms of different domains: *projective reconstruction, self-*

calibration and dense depth estimation.

Projective Reconstruction: It has been shown by Faugeras [6] and Hartley [10] that a reconstruction up to an arbitrary projective transformation was possible from an uncalibrated image sequence. Since then a lot of effort has been put in reliably obtaining accurate estimates of the projective calibration of an image sequence. Robust algorithms were proposed to estimate the fundamental matrix from image pairs [29, 31]. Based on this, an algorithm which sequentially retrieves the projective calibration of a complete image sequence has been developed [1].

Self-Calibration: Since a projective calibration is not sufficient for many applications, researchers tried to find ways to automatically upgrade projective calibrations to metric (i.e. Euclidean up to scale). Typically, it is assumed that the same camera is used throughout the sequence and that the intrinsic camera parameters are constant. This proved a difficult problem and many researchers have worked on it (e.g. [7, 11, 18, 30]). One of the main problems is that critical motion sequences exist for which self-calibration does not result in a unique solution [27]. We proposed a more pragmatic approach [19] which assumes that some parameters are (approximately) known but which allows others to vary. Therefore this approach can deal with zooming/focusing cameras.

Dense Depth Estimation: Since the calibration of the image sequence has been estimated we can use stereoscopic triangulation techniques between image correspondences to estimate depth. The difficult part in stereoscopic depth estimation is to find dense correspondence maps between the images. The correspondence problem is facilitated by exploiting constraints derived from the calibration and from some assumptions about the scene. We use an approach that combines local image correlation methods with a dynamic programming approach to constrain the correspondence search [14]. This technique was first proposed by Cox [3] and further developed by Koch [13] and Falkenhagen [5].

This paper is organized as follows: In section 2 a general overview of the system is given. In the subsequent sections the different steps are explained in more detail: projective reconstruction (section 3), self-calibration (section 4), dense matching (section 5) and model generation (section 6). Section 7 concludes the paper.

2. Overview of the method

The presented system gradually retrieves more information about the scene and the camera setup. The first step is to relate the different images. This is done pairwise by retrieving the epipolar geometry. An initial reconstruction is then made for the first two images of the sequence. For the subsequent images the camera pose is estimated in the

projective frame defined by the first two cameras. For every additional image that is processed at this stage, the interest points corresponding to points in previous images are reconstructed, refined or corrected. Therefore it is not necessary that the initial points stay visible throughout the entire sequence. For sequences where points can disappear and reappear, it can be necessary to match the images to other images than the previous one. The algorithm was adapted to efficiently deal with this. The result of this step is a reconstruction of typically a few hundred to a few thousand interest points and the (projective) pose of the camera. The reconstruction is only determined up to a projective transformation.

The next step is to restrict the ambiguity of the reconstruction to a metric one. In a projective reconstruction not only the scene, but also the camera is distorted. Since the algorithm deals with unknown scenes, it has no way of identifying this distortion in the reconstruction. Although the camera is also assumed to be unknown, some constraints on the intrinsic camera parameters (e.g. rectangular or square pixels, constant aspect ratio, principal point in the middle of the image, ...) can often still be assumed. A distortion on the camera mostly results in the violation of one or more of these constraints. A metric reconstruction/calibration is obtained by transforming the projective reconstruction until all the constraints on the camera's intrinsic parameters are satisfied.

At this point the system effectively disposes of a calibrated image sequence. The relative position and orientation of the camera is known for all the viewpoints. This calibration facilitates the search for corresponding points and allows us to use a stereo algorithm that was developed for a calibrated system and which allows to find correspondences for most of the pixels in the images.

From these correspondences the distance from the points to the camera center can be obtained through triangulation. These results are refined and completed by combining the correspondences from multiple images.

A dense metric 3D surface model is obtained by approximating the depth map with a triangular wireframe. The texture is obtained from the images and mapped onto the surface.

In figure 1 an overview of the system is given. It consists of independent modules which pass on the necessary information to the next modules. The first module computes the projective calibration of the sequence together with a sparse reconstruction. In the next module the metric calibration is computed from the projective camera matrices through self-calibration. Then dense correspondence maps are estimated. Finally all results are integrated in a textured 3D surface reconstruction of the scene under consideration.

Throughout the rest of this paper the different steps of the method will be explained in more detail. An image se-

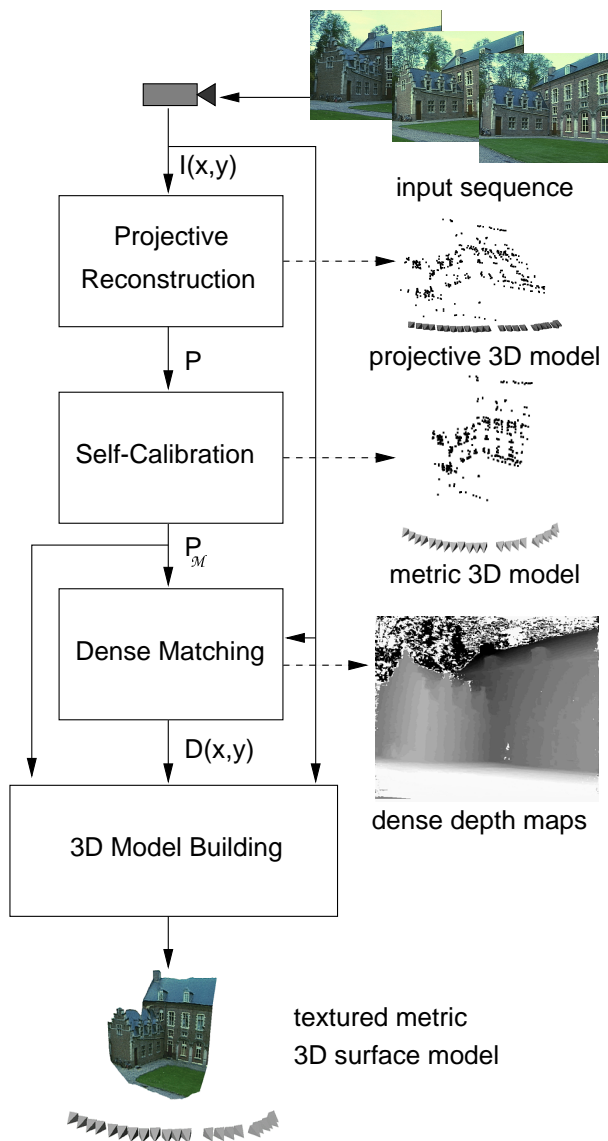


Figure 1. Overview of the system: from the image sequence $I(x, y)$ the projective reconstruction is computed; the projection matrices P are then passed on to the self-calibration module which delivers a metric calibration P_M ; the next module uses these to compute dense depth maps $D(x, y)$; all these results are assembled in the last module to yield a textured 3D surface model. On the right side the results of the different modules are shown: the preliminary reconstructions (both projective and metric) are represented by point clouds, the cameras are represented by little pyramids, the results of the dense matching are accumulated in dense depth maps (light means close and dark means far).




Figure 2. First and last image of the Arenberg castle sequence. This sequence is used throughout this paper to illustrate the different steps of the reconstruction system.

quence of the Arenberg castle in Leuven will be used for illustration. Some images of this sequence can be seen in Figure 2. The full sequence consists of 24 images recorded with a video camera.

3. Projective reconstruction

At first the images are completely unrelated. The only assumption is that the images form a sequence in which consecutive images do not differ too much. Therefore the local neighborhood of image points originating from the same scene point should look similar if images are close in the sequence. This allows for automatic matching algorithms to retrieve correspondences. The approach taken to obtain a projective reconstruction is very similar to the one proposed by Beardsley et al. [1].

3.1. Relating the images

It is not feasible to compare every pixel of one image with every pixel of the next image. It is therefore necessary to reduce the combinatorial complexity. In addition not all points are equally well suited for automatic matching. The local neighborhoods of some points contain a lot of intensity variation and are therefore easy to differentiate from others. The Harris corner detector [9] is used to select a set of such points. Correspondences between these image points need to be established through a matching procedure.

Matches are determined through normalized cross-correlation of the intensity values of the local neighborhood. Since images are supposed not to differ too much, corresponding points can be expected to be found back in the same region of the image. Therefore at first only interest points which have similar positions are considered for matching. When two points are mutual best matches they are considered as potential correspondences.

Since the epipolar geometry describes the complete geometry relating two views, this is what should be retrieved.

- repeat
 - take minimal sample (7 matches)
 - compute \mathbf{F}
 - estimate $\%inliers$
- until $P_{OK}(\%inliers, \#trials) > 95\%$
- refine \mathbf{F} (using all inliers)

Table 1. Robust estimation of the epipolar geometry from a set of matches containing outliers using RANSAC (P_{OK} indicates the probability that the epipolar geometry has been correctly estimated).

Computing it from the set of potential matches through least squares does in general not give satisfying results due to its sensitivity to outliers. Therefore a robust approach should be used. Several techniques have been proposed [29, 31] based on robust statistics [25]. Our system incorporates the RANSAC (RANDOM SAMPLING CONSENSUS) [8] approach used by Torr et al. [29]. Table 1 sketches this technique.

Once the epipolar geometry has been retrieved, one can start looking for more matches to refine this geometry. In this case the search region is restricted to a few pixels around the epipolar lines.

3.2. Initial reconstruction

The two first images of the sequence are used to determine a reference frame. The world frame is aligned with the first camera. The second camera is chosen so that the epipolar geometry corresponds to the retrieved \mathbf{F}_{12} (see [1] for more details).

$$\begin{aligned} \mathbf{P}_1 &= \begin{bmatrix} \mathbf{I}_{3 \times 3} & | & 0 \\ \mathbf{e}_{12} \times \mathbf{F}_{12} + \mathbf{e}_{12} \mathbf{1}^\top & | & \sigma \mathbf{e}_{12} \end{bmatrix} \quad (1) \\ \mathbf{P}_2 &= \end{aligned}$$

where $[\mathbf{e}_{12}]_\times$ indicates the vector product with \mathbf{e}_{12} . Equation 1 is not completely determined by the epipolar geometry (i.e. \mathbf{F}_{12} and \mathbf{e}_{12}), but has 4 more degrees of freedom (i.e. l_X, l_Y, l_Z, σ). $\mathbf{1} = [l_X l_Y l_Z]^\top$ determines the position of the plane at infinity and σ determines the global scale of the reconstruction. To avoid some problems during the reconstruction (due to the violation of the quasi-Euclidean assumption), it is recommended to determine l_X, l_Y, l_Z in such a way that the plane at infinity does not cross the scene. Our implementation uses an approach similar to the quasi-Euclidean approach proposed in [1], but the focal length is chosen so that most of the points are reconstructed in front of the cameras. This approach was inspired by Hartley's




Figure 3. Image matches (m_{k-1}, m_k) are found as described before. Since the image points, m_{k-1} , relate to object points, M_k , the pose for view k can be computed from the inferred matches (M, m_k).

chirality [11] and the oriented projective geometry introduced by Laveau [16]. Since there is no way to determine the global scale from the images, σ can arbitrarily be chosen to $\sigma = 1$.

Once the cameras have been fully determined the matches can be reconstructed through triangulation. The optimal method for this is given in [12]. This gives us a preliminary reconstruction.

3.3. Adding a view

For every additional view the pose towards the pre-existing reconstruction is determined, then the reconstruction is updated. This is illustrated in Figure 3. The first step consists of finding the epipolar geometry as described in Section 3.1. Then the matches which correspond to already reconstructed points are used to compute the projection matrix \mathbf{P}_k . This is done using a robust procedure similar to the one laid out in Table 1. In this case a minimal sample of 6 matches is needed to compute \mathbf{P}_k . Once \mathbf{P}_k has been determined the projection of already reconstructed points can be predicted. This allows to find some additional matches to refine the estimation of \mathbf{P}_k . This means that the search space is gradually reduced from the full image to the epipolar line to the predicted projection of the point. This is illustrated in Figure 4.

Once the camera projection matrix has been determined the reconstruction is updated. This consists of refining, cor-




Figure 4. (a) a priori search range, (b) search range along the epipolar line and (c) search range around the predicted position of the point.

recting or deleting already reconstructed points and initializing new points for new matches.

After this procedure has been repeated for all the images, one disposes of camera poses for all the views and the reconstruction of the interest points. In the further modules mainly the camera calibration is used. The reconstruction itself is used to obtain an estimate of the disparity range for the dense stereo matching.

3.4. Relating to other views

The procedure to add a view described in the previous section only relates the image to the previous image. In fact it is implicitly assumed that once a point gets out of sight, it will not come back. Although this is true for many sequences, this assumptions does not always hold. When the system is used to record the appearance of an object from all around, the camera is often moved back and forth. In this case points will continuously disappear and reappear.

Although this is true for many sequences, it is certainly not always the case. For some hand-held sequences the camera can be used to record the appearance of the object from all around. Often the camera will be moved back and forth to scan the viewing sphere of the object. With the traditional scheme this would imply that new 3D points keep being instantiated. This will not immediately cause problems, but since for the system these two 3D points are unrelated nothing enforces their position to correspond.

This is especially crucial for longer image sequences where the errors accumulate. It results in a degraded calibration or even causes the failure of the algorithm after a certain number of views.

A possible solution consists of relating every new view with all previous views using the procedure of Section 3.1. It is clear that this would require a considerable computational effort. We propose a more pragmatic approach. This approach worked well on some cases we worked on.

Let $\tilde{\mathbf{P}}_i$ be the initial estimate of the camera pose obtained as described in the previous section. A criterion is then used to define which views are close to the actual view. All these close views are matched with the actual view (as described in Section 3.1). For every close view a set of potential 2D-




Figure 5. Sequential approach (left) and extended approach (right). In the traditional scheme view 8 would be matched with view 7 and 9 only. A point M which would be visible in views 2,3,4,7,8,9,12,13 and 14 would therefore results in 3 independently reconstructed points. With the extended approach only one point will be instantiated. It is clear that this results in a higher accuracy for the reconstructed point while it also dramatically reduces the accumulation of calibration errors.

3D correspondences is obtained. These sets are merged and the camera projection matrix \mathbf{P}_i is reestimated using the same robust procedure as described in the previous section. Figure 5 illustrates this approach.

We applied a very simple criterion to decide if views were close or not. It worked well for the applications we had in mind, but it could easily be refined if needed. The position \mathbf{t}_i of the camera is extracted from $\tilde{\mathbf{P}}_i$ and the distance d_{ij} to all the other camera positions is computed. Close views are selected as views for which $d_{ij} < 1.6d_{i(i-1)}$. Note that strictly speaking such measure is meaningless in projective space, but since a quasi-Euclidean initialization was carried out and only local qualitative comparisons are made the obtained results are good. More details on this approach can be found in [15].

4. Upgrading the reconstruction to metric

The reconstruction obtained as described in the previous paragraph is only determined up to an arbitrary projective transformation. This might be sufficient for some robotics or inspection applications, but certainly not for visualization.

The system uses the self-calibration method described in [19, 23] to restrict the ambiguity on the reconstruction to metric (i.e. Euclidean up to scale). This flexible self-calibration technique allows the intrinsic camera parameters to vary during the acquisition. This feature is especially useful when the camera is equipped with a zoom or with auto-focus.

It is outside the scope of this paper to discuss this method in detail. The general concept consist of translating constraints on the intrinsic camera parameters to constraints on



Figure 6. Reconstruction before (top) and after (bottom) self-calibration.

the absolute conic. Once this special conic is identified, it can be used as a calibration pattern to upgrade the reconstruction to metric.

Some reconstructions *before* and *after* the self-calibration stage are shown. The top part of Figure 6 gives the reconstruction before self-calibration. Therefore it is only determined up to an arbitrary projective transformation and metric properties of the scene can not be observed from this representation. The bottom part of Figure 6 shows the result after self-calibration. At this point the reconstruction has been upgraded to metric.

5. Dense depth estimation

Only a few scene points are reconstructed from feature tracking. Obtaining a dense reconstruction could be achieved by interpolation, but in practice this does not yield satisfactory results. Small surface details would never be reconstructed in this way. Additionally, some important features are often missed during the corner matching and would therefore not appear in the reconstruction.

These problems can be avoided by using algorithms which estimate correspondences for almost every point in the images. Because the reconstruction was upgraded to metric, algorithms that were developed for calibrated stereo rigs can be used.

5.1. Rectification

Since we have computed the calibration between successive image pairs we can exploit the epipolar constraint that restricts the correspondence search to a 1-D search range. It is possible to re-map the image pair to standard geometry with the epipolar lines coinciding with the image scan lines [13]. The correspondence search is then reduced to a matching of the image points along each image scan-line. This results in a dramatic increase of the computational efficiency of the algorithms by enabling several optimizations in the computations.

For some motions (i.e. when the epipole is located in the image) standard rectification based on planar homographies is not possible and a more advanced procedure should be used. Recently Roy, Meunier and Cox [26] presented such a method. The idea of the method is to rectify the images on a cylinder in stead of a plane. The method is however relatively complex and has some disadvantages.

The system described in this paper uses a new approach. The method combines simplicity with minimal image size and works for all possible motions. The idea is to use polar coordinates with the epipole as origin. Since the ambiguity for the location of matching points in a pair of images is restricted to half epipolar lines [16] we only have to use positive values for the longitudinal coordinate. Since (half-)epipolar line transfer is fully described by an oriented fundamental matrix, this is all our method needs. The necessary information is also easily extracted from (oriented) camera projection matrices. The angle between two consecutive half-epipolar lines is computed to have the worst case pixels preserve their area. This is done independently for every half-epipolar line. This therefore results in a minimal image size. A more complete description of this method is given in [24].

As an example a rectified image pair from the castle is shown for both the standard technique and our new generalized technique. Figure 7 shows the rectified image pair for both methods.

A second example shows that the methods works properly when the epipole is in the image. Figure 8 shows the two original images while Figure 9 shows the two rectified images. In this case the standard rectification procedure can not deliver rectified images. This rectified image pair was used with success by our standard stereo matching algorithm.

5.2. Dense stereo matching

In addition to the epipolar geometry other constraints like preserving the order of neighboring pixels, bidirectional uniqueness of the match, and detection of occlusions can be exploited. These constraints are used to guide the



Figure 7. Rectified image pair for both methods: standard homography based method (top), new method (bottom).



Figure 9. Rectified pair of images of the desk. It can be verified visually that corresponding points are located on corresponding image rows. The right side of the images corresponds to the epipole.




Figure 8. Image pair of a desk a few days before a deadline. The epipole is indicated by a white dot (top-right of 'Y' in 'VOLLEYBALL').

correspondence towards the most probable scan-line match using a dynamic programming scheme [5].

For dense correspondence matching a disparity estimator based on the dynamic programming scheme of Cox et al. [3], is employed that incorporates the above mentioned constraints. It operates on rectified image pairs (I_k, I_l) where the epipolar lines coincide with image scan lines. The matcher searches at each pixel in image I_k for maximum normalized cross correlation in I_l by shifting a small measurement window (kernel size 5×5 to 7×7 pixel) along the corresponding scan line. The selected search step size ΔD (usually 1 pixel) determines the search resolution. Matching ambiguities are resolved by exploiting the ordering constraint in the dynamic programming approach (see Koch [13]). The algorithm was further adapted to employ extended neighborhood relationships and a pyramidal estimation scheme to reliably deal with very large disparity ranges of over 50% of image size (see Falkenhagen [5]). This algorithm that was at first developed for calibrate stereo rigs (See Koch [13]) could easily be used for our purposes since at this stage the necessary calibration information had already been retrieved from the images.

5.3. Multi view matching

The pairwise disparity estimation allows to compute image to image correspondence between adjacent rectified image pairs, and independent depth estimates for each camera viewpoint. An optimal joint estimate is achieved by fusing all independent estimates into a common 3D model. The fusion can be performed in an economical way through controlled correspondence linking. The approach utilizes a flexible multi viewpoint scheme which combines the advantages of small baseline and wide baseline stereo (see Koch, Pollefeys and Van Gool [14]).

The result of this procedure is a very dense depth map. Most occlusion problems are avoided by linking correspondences from up and down the sequence. An example of such a very dense depth map is given in Figure 10.

6. Building the model

The dense depth maps as computed by the correspondence linking must be approximated by a 3D surface representation suitable for visualization. So far each object point was treated independently. To achieve spatial coherence for a connected surface, the depth map is spatially interpolated using a parametric surface model. The boundaries of the objects to be modeled are computed through depth segmentation. In a first step, an object is defined as a connected region in space. Simple morphological filtering removes spurious and very small regions. Then a bounded thin plate model is employed with a second order spline to smooth the




Figure 10. Dense depth map (light means near and dark means far).

surface and to interpolate small surface gaps in regions that could not be measured.

The spatially smoothed surface is then approximated by a triangular wire-frame mesh to reduce geometric complexity and to tailor the model to the requirements of computer graphics visualization systems. The mesh triangulation currently utilizes the reference view only to build the model. The surface fusion from different view points to completely close the models remains to be implemented. Sometimes it is not possible to obtain a single metric framework for large objects like buildings since one may not be able to record images continuously around it. In that case the different frameworks have to be registered to each other. This could be done using available surface registration schemes [2].

Texture mapping onto the wire-frame model greatly enhances the realism of the models. As texture map one could take the reference image texture alone and map it to the surface model. However, this creates a bias towards the selected image and imaging artifacts like sensor noise, unwanted specular reflections or the shading of the particular image is directly transformed onto the object. A better choice is to fuse the texture from the image sequence in much the same way as depth fusion.

The viewpoint linking builds a controlled chain of correspondences that can be used for texture enhancement as well. The estimation of a robust mean texture will capture the static object only and the artifacts (e.g. specular reflections or pedestrians passing in front of a building) are suppressed. The texture fusion could also be done on a finer grid, yielding a super resolution texture [17].

An example of the resulting model can be seen in Figure 11. Some more views of the reconstruction are given in Figure 12.




Figure 11. 3D surface model obtained automatically from an uncalibrated image sequence, shaded (left), textured (right).



Figure 12. More views of the castle model.

7. Conclusion

An automatic 3D scene modeling technique was discussed that is capable of building models from uncalibrated image sequences. The technique is able to extract metric 3D models without any prior knowledge about the scene or the camera. The calibration is obtained by assuming a rigid scene and some constraints on the intrinsic camera parameters (e.g. square pixels).

Work remains to be done to get more complete models by fusing the partial 3D reconstructions. This will also increase the accuracy of the models and eliminate artifacts at the occluding boundaries. For this we can rely on work already done for calibrated systems.

Acknowledgments

We wish to acknowledge the financial support of the Belgian IUAP 24/02 'ImechS' project and of the EU ESPRIT project LTR 23515 IMPROOFS.

References

- [1] P. Beardsley, A. Zisserman and D. Murray, "Sequential Updating of Projective and Affine Structure from Motion", *International Journal of Computer Vision* (23), No. 3, Jun-Jul 1997, pp. 235-259.
- [2] Y. Chen, and G. Medioni, "Object Modeling by Registration of Multiple Range Images", *Proc. Int. Conf. on Robotics and Automation*, 1991.
- [3] I. Cox, S. Hingorani and S. Rao, "A Maximum Likelihood Stereo Algorithm", *Computer Vision and Image Understanding*, Vol. 63, No. 3, May 1996.
- [4] P. Debevec, C. Taylor and J. Malik, "Modeling and Rendering Architecture from Photographs: A Hybrid Geometry- and Image-Based Approach", *Siggraph*, 1996.
- [5] L. Falkenhagen, "Hierarchical Block-Based Disparity Estimation Considering Neighbourhood Constraints". *Proc. International Workshop on SNHC and 3D Imaging*, Rhodes, Greece, 1997.
- [6] O. Faugeras, "What can be seen in three dimensions with an uncalibrated stereo rig", *Computer Vision - ECCV'92*, LNCS, Vol. 588, Springer-Verlag, pp. 563-578, 1992.
- [7] O. Faugeras, Q.-T. Luong and S. Maybank. "Camera self-calibration: Theory and experiments", *Computer Vision - ECCV'92*, LNCS, Vol. 588, Springer-Verlag, pp. 321-334, 1992.
- [8] M. Fischler and R. Bolles, "RANDOM SAMPLING CONSENSUS: a paradigm for model fitting with application to image analysis and automated cartography", *Commun. Assoc. Comp. Mach.*, 24:381-95, 1981.
- [9] C. Harris and M. Stephens, "A combined corner and edge detector", *Fourth Alvey Vision Conference*, pp.147-151, 1988.
- [10] R. Hartley, "Estimation of relative camera positions for uncalibrated cameras", *Computer Vision - ECCV'92*, LNCS, Vol. 588, Springer-Verlag, pp. 579-587, 1992.
- [11] R. Hartley, "Euclidean reconstruction from uncalibrated views", in : J.L. Mundy, A. Zisserman, and D. Forsyth (eds.), *Applications of Invariance in Computer Vision*, LNCS, Vol. 825, Springer-Verlag, pp. 237-256, 1994.
- [12] R. Hartley and P. Sturm, "Triangulation", *Computer Vision and Image Understanding*, 68(2):146-157, 1997.
- [13] R. Koch, *Automatische Oberflächenmodellierung starrer dreidimensionaler Objekte aus stereoskopischen Rundum-Ansichten*, PhD thesis, University of Hannover, Germany, 1996 also published as Fortschritte-Berichte VDI, Reihe 10, Nr.499, VDI Verlag, 1997.
- [14] R. Koch, M. Pollefeys and L. Van Gool, Multi Viewpoint Stereo from Uncalibrated Video Sequences. *Proc. European Conference on Computer Vision*, pp.55-71. Freiburg, Germany, 1998.
- [15] R. Koch, M. Pollefeys, B. Heigl, L. Van Gool and H. Niemann, "Calibration of Hand-held Camera Sequences for Plenoptic Modeling", *Proc. International Conference on Computer Vision*, IEEE Computer Society Press, 1999.
- [16] S. Laveau and O. Faugeras, "Oriented Projective Geometry for Computer Vision", *Computer Vision - ECCV'96*, LNCS, Vol. 1064, Springer-Verlag, pp. 147-156, 1996.

- [17] E. Ofek, E. Shilat, A. Rappoport and M. Werman, "Highlight and Reflection Independent Multiresolution Textures from Image Sequences", *IEEE Computer Graphics and Applications*, vol.17 (2), March-April 1997.
- [18] M. Pollefeys and L. Van Gool, "A stratified approach to self-calibration", *Proc. 1997 Conference on Computer Vision and Pattern Recognition*, IEEE Computer Soc. Press, pp. 407-412, 1997.
- [19] M. Pollefeys, R. Koch and L. Van Gool, "Self-Calibration and Metric Reconstruction in spite of Varying and Unknown Internal Camera Parameters", *Proc. International Conference on Computer Vision*, Narosa Publishing House, pp.90-95, 1998.
- [20] M. Pollefeys, R. Koch, M. Vergauwen and L. Van Gool, "Flexible acquisition of 3D structure from motion", *Proc. IEEE workshop on Image and Multidimensional Digital Signal Processing*, pp.195-198, Alpbach, 1998.
- [21] M. Pollefeys, R. Koch, M. Vergauwen and L. Van Gool, "Metric 3D Surface Reconstruction from Uncalibrated Image Sequences", *Proc. SMILE Workshop (post-ECCV'98)*, LNCS, Vol. 1506, pp.138-153, Springer-Verlag, 1998.
- [22] M. Pollefeys, R. Koch, M. Vergauwen and L. Van Gool, "Automatic Generation of 3D Models from Photographs", *Proc. Virtual Systems and MultiMedia*, 1998.
- [23] M. Pollefeys, *Self-calibration and metric 3D reconstruction from uncalibrated image sequences*, Ph.D. Thesis, ESAT-PSI, K.U.Leuven, 1999.
- [24] M. Pollefeys, R. Koch and L. Van Gool, "A simple and efficient rectification method for general motion", *Proc. International Conference on Computer Vision*, IEEE Computer Society Press, 1999.
- [25] P. Rousseeuw, *Robust Regression and Outlier Detection*, Wiley, New York, 1987.
- [26] S. Roy, J. Meunier and I. Cox, "Cylindrical Rectification to Minimize Epipolar Distortion", *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp.393-399, 1997.
- [27] P. Sturm, "Critical Motion Sequences for Monocular Self-Calibration and Uncalibrated Euclidean Reconstruction", *Proc. 1997 Conference on Computer Vision and Pattern Recognition*, IEEE Computer Soc. Press, pp. 1100-1105, 1997.
- [28] C. Tomasi and T. Kanade, "Shape and motion from image streams under orthography: A factorization approach", *International Journal of Computer Vision*, 9(2):137-154, 1992.
- [29] P. Torr, *Motion Segmentation and Outlier Detection*, PhD Thesis, Dept. of Engineering Science, University of Oxford, 1995.
- [30] B. Triggs, "The Absolute Quadric", *Proc. 1997 Conference on Computer Vision and Pattern Recognition*, IEEE Computer Soc. Press, pp. 609-614, 1997.
- [31] Z. Zhang, R. Deriche, O. Faugeras and Q.-T. Luong, "A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry", *Artificial Intelligence Journal*, Vol.78, pp.87-119, October 1995.