

A Semi-Automatic Method for Resolving Occlusion in Augmented Reality

Vincent Lepetit and Marie-Odile Berger
LORIA/INRIA Lorraine
BP 101, 54602 Villers les Nancy, France
{lepetit,berger}@loria.fr

Abstract

Realistic merging of virtual and real objects requires that the augmented patterns be correctly occluded by foreground objects. In this paper, we propose a semi-automatic method for resolving occlusion in augmented reality which makes use of key-views. Once the user has outlined the occluding objects in the key-views, our system detects automatically these occluding objects in the intermediate views. A region of interest that contains the occluding objects is first computed from the outlined silhouettes. One of the main contribution of this paper is that this region takes into account the uncertainty on the computed interframe motion. Then a deformable region-based approach is used to recover the actual occluding boundary within the region of interest from this prediction.

Results and videos are available at the URL:
<http://www.loria.fr/~lepetit/Occlusions>.

1. Introduction

The objective of augmented reality (AR) is to add virtual objects to real video sequences, allowing computer-generated objects to be overlaid on the video in such a manner as to appear part of the viewed 3D scene. Applications include computer-aided surgery, tele-operations, and special effects for the film and the broadcast industries. This paper concentrates on the particular application of video post-production.

Realistic image composition requires that the augmented patterns be correctly occluded by foreground objects. However, solving the occlusion problem for AR is challenging when little is known about the real world we wish to augment. Theoretically, resolving occlusion amounts to compare the depth of the virtual objects to that of the real scene. However, computing dense and accurate depth maps from images is difficult. This explains why the accuracy of the obtained occluding boundary is generally poor. Moreover, in most AR applications, the interframe motion is not a *pri-*

ori known but must be computed. Inaccurate motion estimation thus results in possibly large reconstruction errors.

In order to overcome problems stemming from possibly large reconstruction errors, Ong [6] proposed a semi interactive approach to solve occlusion: the occluding objects are segmented by hand in selected views called key-frames. These silhouettes are used to build the 3D model of the occluding object. The 2D occluding boundary is then obtained by projecting the 3D shape in the intermediate frames. However, due to the uncertainty on the computed interframe motion, the recovered 3D shape do not project exactly onto the occluding objects in the key-frames nor in the intermediate frames.

In this paper, we also use the concept of key-views but we do not attempt to build the 3D model of the occluding objects from all the key-frames. The novelty in this paper is twofold: (i) we do not attempt to recover the 3D model of the occluding objects from all the key-views. We only compute the 3D occluding boundary from two consecutive key views. The projection of this 3D curve is a good prediction of the actual 2D occluding boundary in the intermediate frames. (ii) we recover the actual occluding boundary with a good accuracy using deformable region-based tracking followed by an adjustment stage based on *snakes*. This allows us to compensate easily for the interframe motion error. We then obtain an accurate estimation of the occluding boundary over the sequence.

2. Overview of the system

Theoretically, the 3D shape of the occluding object can be computed from its silhouettes detected in an image sequence. For AR applications however, the interframe camera motion is computed from image/model correspondences or with 2D/2D correspondences over time [4, 7]. The errors resulting from this inaccurate registration makes the 3D reconstruction untractable. That is the reason why we only attempt to recover the 3D occluding boundary from two consecutive key-frames instead of recovering the 3D shape of the occluding object from the whole sequence. Fig. 1

explains the way we compute a first estimation of the 2D occluding boundary in each frame of the sequence. First, the user points out key-frames which correspond to views where aspect changes of the occluding object occur. These key-frames are framed in black in Fig. 1. The user also outlines the occluding object on these key-frames (in white). It is well known that the 3D occluding boundary depends on the camera viewpoint. However, the starting point for our method is to build a good approximation of the 3D occluding boundary which will be used for all the frames between two key-views. This 3D curve is built using stereo-triangulation from the two silhouettes outlined by the user provided that the translation between the two frames is not null (Fig. 1.a and b). The projection of this approximated occluding boundary on the intermediate frames thus provides a fair estimation of the 2D occluding boundary (Fig. 1.c and 1.d).

Due to the uncertainty on the computed interframe motion, this prediction can be relatively far from the actual occluding boundary for at least two reasons (see for instance Fig. 5.a): (i) the computed 3D occluding boundary is only an approximation of the real one because stereo-triangulation is performed from two occluding contours. (ii) more importantly, errors on the camera parameters induce reconstruction errors on the 3D curve and consequently errors on its projection in the considered frame.

One of the main contributions of this paper is to show that the error on the computed camera parameters can be estimated. The uncertainty on the 3D occluding boundary can then be deduced. This allows us to define a region of interest around the predicted contour which is likely to contain the actual occluding boundary (section 3). The refinement stage (section 4) is then carried out within this region: region-based tracking is first used to recover the region whose size and texture only differ from the predicted shape with an affine transformation. Finally, active contour models are used to adjust the occluding boundary.

3. Reconstructing the 3D occluding boundary

3.1. Computing the camera parameters

In this section we first briefly recall how we compute the camera motion over the sequence. Our approach to motion computation takes advantage of 3D knowledge on the scene as well as 2D/2D correspondences over time [7]. Given the viewpoint $[\mathbf{R}_k, \mathbf{t}_k]$ computed in a given frame k , we compute the viewpoint p in the next frame $k + 1$ using the 3D model points M_i whose projections are detected in frame $k + 1$. In addition, we use interest points [5] $(q_k^i, q_{k+1}^i)_{1 \leq i \leq m}$ that are automatically extracted and matched between frames k and $k + 1$. The quality of the viewpoint can be assessed by the distance between q_{k+1}^i and

the epipolar line $ep_{k+1}(q_k^i)$. The viewpoint is therefore recovered by minimizing:

$$\Phi(p) = \frac{1}{n} \sum_{i=1}^n dist^2(m_i, proj(M_i)) + \frac{\lambda}{2m} \sum_{i=1}^m \frac{dist^2(q_{k+1}^i, ep_{k+1}(q_k^i)) + dist^2(q_k^i, ep_k(q_{k+1}^i))}{dist^2(q_{k+1}^i, ep_{k+1}(q_k^i)) + dist^2(q_k^i, ep_k(q_{k+1}^i))} \quad (1)$$

Let C_1 and C_2 be the occluding object outlined in the two key-views. We build the 3D occluding boundary using stereo triangulation: Let m_1 be a point on C_1 . Its corresponding point in the other key-view is the intersection of the epipolar line with C_2 . As usual, the order constraint is used to solve the ambiguity of the correspondence problem.

3.2. Taking into account the error on the estimated motion

The critical role of motion error in scene reconstruction has been pointed out in [8]. Recently, Csurka and Faugeras [3] attempted to compute the covariance on the fundamental matrix recovered from point correspondences between two frames. The uncertainty is computed under the assumption that the matched points are independent. However, this statistical assumption leads to decrease arbitrarily the uncertainty on the fundamental matrix as the number of matched points increases.

That is the reason why we prefer to use the ϵ indifference region [1] to investigate the reliability of the estimated camera parameters. The fact that we have elected to minimize a function $\Phi(p)$ means that we set some store by obtaining a low value of this function. It is reasonable to suppose that values of Φ almost as low as Φ^* would satisfy us almost as much as Φ^* . This gives rise to an ϵ indifference region in p space described by the equation:

$$\epsilon_{region} = \{p \text{ such that } |\Phi(p) - \Phi(p^*)| \leq \epsilon\}$$

In a sufficiently small neighborhood of p^* we may approximate Φ by means of its Taylor equation:

$$\Phi(p) \approx \Phi(p^*) + \nabla \Phi(p^*)^t \delta p + \frac{1}{2} \delta p^t H(p^*) \delta p \quad (2)$$

where H^* is the hessian of Φ computed at $p = p^*$. More details on the computation of H^* are given in Annex A.

As p^* is the minimum of Φ , the gradient is null at the optimum $\nabla \Phi(p^*) = 0$ and equation (2) becomes

$$\Phi(p) \approx \Phi(p^*) + \frac{1}{2} \delta p^t H(p^*) \delta p$$

The ϵ indifference region is then defined by:

$$|\delta p^t H(p^*) \delta p| \leq 2\epsilon$$

which is the equation of a 6-dimensional ellipsoid.

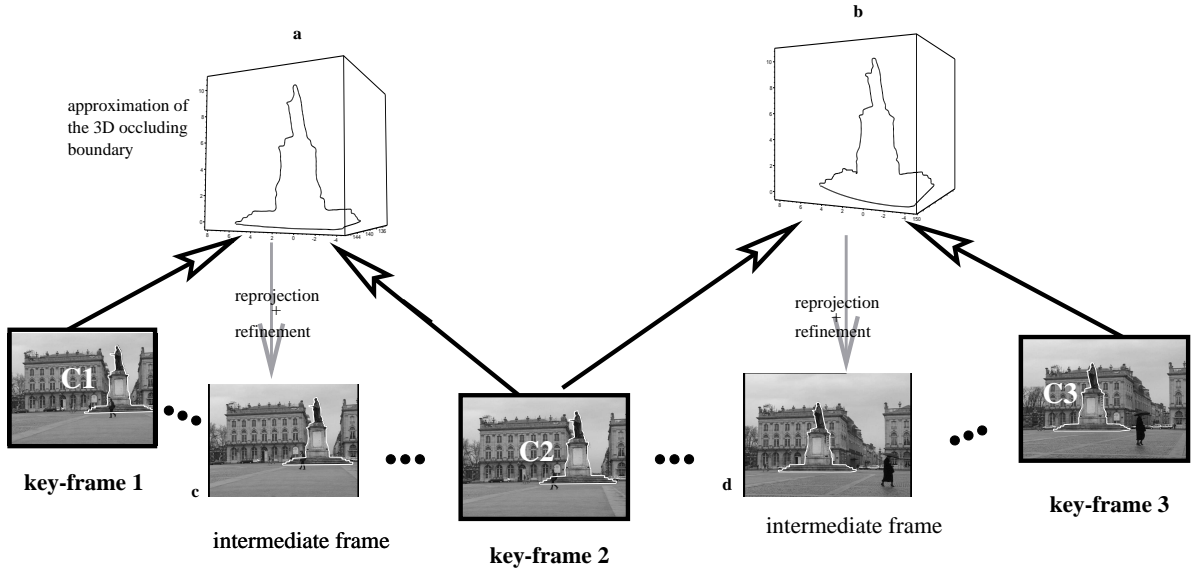


Figure 1. Overview of the system.




Figure 2. The indifference regions for the translation parameters over the Stanislas sequence.

Fig.2 shows these indifference regions computed along the Stanislas sequence (we use $\epsilon = 1$). The building in the background is the 3D model used for registration. For each frame of the sequence, we drew the ϵ indifference region for the translation parameters.

We can now compute the reconstruction error on the occluding boundary from these indifference regions. If point correspondences were available, the reconstruction error could be recovered in an analytical way from viewpoint uncertainties [8]. Unfortunately, as we only have curve correspondences, the matched points depends on the viewpoint and are computed as the intersection of the epipolar line of the point with C_2 . We therefore resort to an exhaustive approach. We consider the *extremal viewpoints*, that are the vertices of the 6-dimensional indifference ellipsoid. Let

$\{p_1^1, \dots, p_1^{12}\}$ (resp $\{p_2^1, \dots, p_2^{12}\}$) the extremal viewpoints in the two key-views. Let m_1 be a point on C_1 . Given an extremal viewpoint p_1 , we can compute the 12 possible reconstructions of m_1 with the 12 extremal views in key-frame 2. Using the 12 extremal viewpoints in key-frame 1, we then obtain 12^2 extremal reconstructions of m_1 according to the uncertainty computed on the two key-views. The convex hull of these 144 points is a good approximation of the 3D reconstruction error on m_1 .

We can now predict the position of the 2D occluding boundary in the in-between frames by simply reprojecting the 3D occluding boundary. To estimate the 2D uncertainty on the projected boundary C , we have to take into account the 3D reconstruction error and the uncertainty on the considered viewpoint. We again resort to an exhaustive method: for each point m_i on C , the 12^2 possible extremal reconstructions are projected onto the current frame using the 12 extremal viewpoints of this frame. We define the spatial uncertainty on the predicted occluding boundary associated to m_i as the convex hull of these 12^3 image points. This area is denoted Λ_i in the following.

The main stages for computing the 2D uncertainty on the predicted occluding boundary are illustrated in Fig. 3: Fig. 3.a exhibits a point on the predicted boundary and Fig. 3.b shows the projection of the corresponding 3D extremal points using the extremal viewpoints and the convex hull Λ_i . Finally, Fig. 5.a shows the 2D uncertainty computed for each point of the predicted boundary (dotted line). The points are drawn with black circles or crosses and the uncertainty is drawn in white. The reader can notice that some points on the steps have no associated spatial uncertainty. Indeed, because the key silhouettes do not match exactly,




Figure 3. Computation of the spatial uncertainty on the predicted occluding boundary.

the epipolar line computed with some extremal viewpoints does not always intersect C_2 . If more than 50% of the epipolar lines computed with the 12^2 extremal viewpoints do not intersect C_2 , the spatial uncertainty is not defined at this point.

4. Refining the occluding boundary

As a result of the prediction stage we get an estimate of the occluding boundary along with its 2D uncertainty in the considered frame. In addition we compute not only the boundary but also the texture of the occluding object so as to get a *predicted template* of the occluding object. The texture $I_{template}$ is computed from the nearest key-view by using 2d local image transformation.

We still have to determine the occluding object from the *predicted template*. Due to the error on the computed motion and also because reconstruction is achieved from occluding contours, the template boundary can be relatively far from the actual occluding object and their shapes can also differ (see for instance Fig. 5.a). However, it is important to note that the actual boundary belongs to the computed uncertainty region. Following previous works on deformable structures [2] we use a hierarchical algorithm; we first compute a global estimation of the shape deformation between the key-frame and the current frame. Then we use a fine tuning deformation to adjust the details. As affine transformations seem to be appropriate to describe shape variations due to motion uncertainties, the affine motion that best matches the occluding template on the considered image is searched for:

$$transf_a(m) = \begin{cases} a_1 m_x + a_2 m_y + a_3 \\ a_4 m_x + a_5 m_y + a_6 \end{cases}$$

The optimal parameter a is defined as the one that yields the best fit between the predicted template $I_{template}$ and the current image I . The best match is defined as the minimum of the correlation measure:

$$\Psi(a) = \sum_i \psi_a(i) \quad (3)$$

$$\psi_a(i) = \sum_{\substack{d^x, d^y = W \\ d^x, d^y = -W \\ (m_i + d) \in R_C}} (I_{template}(m_i + d) - I(transf_a(m_i + d)))^2$$

where the predicted curve C is defined by the set of vertices $\{m_i\}_{1 \leq i \leq n}$, $d = (d^x, d^y)$, W is the size of the correlation window and R_C is the region inside C . Note that only the points which are inside the occluding objects are considered in the estimation. This way, points belonging to the changing background do not affect the matching process.

In addition, we have slightly modified the correlation measure in order to take into account the 2D uncertainty on the predicted curve. A penalty term is used to ensure that the matched point belongs to Λ_i . The penalty has the form αW^2 where α is a constant value. The function to be minimized is therefore defined as:

$$\psi_a(i) = \begin{cases} \sum_d (I_{template}(m_i + d) - I(transf_a(m_i + d)))^2 \\ \text{if } transf_a(m_i) \in \Lambda_i, \\ \alpha W^2 \text{ otherwise.} \end{cases}$$

Note that if Λ_i is not defined, or equivalently if $\Lambda_i = \infty$, the first item of ψ_a is used because the assumption $transf_a(m_i) \in \Lambda_i$ is fulfilled. These points are therefore considered in the correlation function without further constraints. Finally, fine tuning adjustment of the occluding boundary is performed with snakes from $transf_a(C)$.

5. Results and discussion

The effectiveness of our approach is demonstrated on two sequences. The Stanislas sequence was shot from a car which turned around the square. Our aim is to incrust a virtual plane passing behind the statue. Here, the 3D model of the opera is used for registration (the building in the back of the scene) while the 3D model of the statue is unknown. The three key-frames chosen by the user are shown in Fig. 4 (frames 66, 118, 150). Fig. 4 exhibits the recovered occluding boundary in the frames 15, 66 and 130. The overall visual impression is very good though the predicted boundary is sometimes relatively far from the actual one.

Fig. 5 clearly proves the efficiency of incorporating motion error into our process. The uncertainty on the predicted curve is drawn in white. The points m_i that are outside the uncertainty region Λ_i after the region based tracking are shown black crosses, whereas the points inside the region are drawn with black circles. For both images, the predicted 2D curve is shown in dotted lines. If the 2D uncertainty is not considered (Fig. 5.a), the recovered boundary is erroneous, especially near the steps. On the contrary, if points are constrained to be in the uncertainty region, the occluding boundary is successfully recovered (Fig. 5.b).

Finally, our algorithm is demonstrated on the *cow sequence*. Zooms on the cow and the user-defined silhouettes



Figure 4. (first row) : The key-views along with the user-defined silhouette: frame 60, 118 and 150. (second row): The recovered occluding boundary in the frames 15, 66,130 and the augmented scene.

are presented in Fig. 6.b, c and d. The three key-views are relatively far. Moreover, the aspects of the cow are very different in the considered key-frames. Even in this case, the recovered occluding boundary before the snake process is quite convincing (see Fig. 6.e). Finally, adjustment with snakes allows us to recover accurately the occluding object (Fig. 6.f) and to augment the scene (Fig 6.g and 6.h).

6. Conclusion

We have presented a new approach for resolving occlusion for AR tasks. The key concept is that fine detection of occluding boundary can be achieved with moderate user interaction. One of the main strengths of our algorithm concerns its ability to handle uncertainties on the computed motion between two frames. Through judicious choice of key-frames, our approach seems to be more convenient and more accurate than most existing approaches.

Annex A: Computing the Hessian H^*

H^* is the value of the Hessian $H = \frac{\partial}{\partial z} \left(\frac{\partial \Phi}{\partial z} \right)^t$ computed at the minimum p^* of Φ . Φ is defined as $\Phi(p) = \frac{1}{n} \sum_{i=1}^n r_i^2 + \frac{\lambda}{2m} \sum_{i=1}^m v_i^2$ where

$$\begin{aligned} r_i^2 &= \text{dist}^2(m_i, \text{proj}(M_i)) \\ v_i^2 &= \text{dist}^2(q_{k+1}^i, \text{ep}_{k+1}(q_k^i)) + \text{dist}^2(q_k^i, \text{ep}_k(q_{k+1}^i)) \end{aligned}$$

r_i^2 and v_i^2 can be expressed as an analytical function of the 6-dimensional vector $p = (\alpha, \beta, \gamma, t_x, t_y, t_z)$ using the fundamental matrix. Because the analytic expression of the second derivatives of v_i^2 with respect to p are really

untractable, we use an approximation to the first order: $H \approx 2 \sum \frac{1}{n} \left(\frac{\partial r_i}{\partial p} \right)^t \left(\frac{\partial r_i}{\partial p} \right) + \frac{\lambda}{2m} \sum \left(\frac{\partial v_i}{\partial p} \right)^t \left(\frac{\partial v_i}{\partial p} \right)$.

References

- [1] Y. Bard. *Nonlinear Parametric Estimation*. Academic Press, 1974.
- [2] B. Basclé and R. Deriche. Stereo Matching Reconstruction and Refinement of 3D curves Using Deformable Contours. In *Proceedings of 4th International Conference on Computer Vision, Berlin (Germany)*, pages 421–430, 1993.
- [3] Csurka, C. Zeller, Z. Zhang, and O. Faugeras. Characterizing the Uncertainty of the Fundamental matrix. *Computer Vision and Image Understanding*, 68(1):18–36, May 1997.
- [4] A. Fitzgibbon and A. Zisserman. Automatic Camera Recovery for Closed or Open Images Sequences. In *Proceedings of 5th European Conference on Computer Vision, University of Freiburg (Germany)*, pages 311–326, June 1998.
- [5] C. Harris and M. Stephens. A Combined Corner and Edge Detector. In *Proceedings of 4th Alvey Conference*, Cambridge, Aug. 1988.
- [6] K. C. Ong, H. C. Teh, and T. S. Tan. Resolving Occlusion in Image Sequence Made Easy. *The Visual Computer*, 14:153–165, 1998.
- [7] G. Simon, V. Lepetit, and M.-O. Berger. Computer Vision Methods for Registration: Mixing 3D Knowledge and 2D Correspondences for Accurate Image Composition. In *First International Workshop on Augmented Reality, San Francisco*, Nov. 1998.
- [8] J. Thomas, A. Hanson, and J. Oliensis. Understanding Noise: The Critical Role of Motion Error in Scene Reconstruction. In *Proceedings of 4th International Conference on Computer Vision, Berlin (Germany)*, 1993.




Figure 5. The recovered occluding boundary without (a) and with (b) the use of the 2D uncertainty.

The predicted curve is shown with dotted lines. The points that belong to the uncertainty region Δ_i are shown with black circles, whereas the points outside Δ_i are drawn with black crosses.




Figure 6. The cow sequence.

(a): an image of the cow sequence; (b,c,d) zoom on the key-views and the user defined silhouettes (frames 0,30,60); (e): the recovered occluding boundary before snake adjustment in frame 46; (f): the recovered final boundary; (g,h): the augmented scene.