

Indirect Tracking to Reduce Occlusion Problems

Peter Keitler, Michael Schlegel, and Gudrun Klinker

Technische Universität München, Fakultät für Informatik, Boltzmannstraße 3, 85748 Garching
b. München, Germany
{keitler, schlegel, klinker}@in.tum.de

Abstract. There exist many infrared inside-out 6-DOF pose tracking configurations with cameras mounted rigidly to the environment. In such a setup, tracking is inherently impossible for IR targets inside/below/behind other opaque objects (occlusion problem).

We present a solution for the integration of an additional, mobile IR tracking system to overcome this problem. The solution consists of an indirect tracking setup where the stationary cameras track the mobile cameras which in turn track the target. Accuracy problems that are inherent to such an indirect tracking setup, are tackled by an error correction mechanism based on reference points in the scene that are known to both tracking systems. An evaluation demonstrates that, in naive indirect tracking without error correction, the major source of error consists in a wrong detection of orientation of the mobile system and that this source of error can be practically eliminated by our error correction mechanisms.

Keywords: Augmented Reality, indirect tracking, sensor fusion, absolute orientation problem.

1 Motivation

One of the currently most common tracking setups for AR and VR applications consists of an outside-in configuration with a number of infrared cameras mounted rigidly to the environment, observing a fixed volume within their midst. The camera arrangement imposes restrictions toward tracking moveable objects inside/below/behind other opaque objects in the scene. We call this the occlusion problem. It is not generally solvable by adding additional cameras to the classical outside-in setup since first, occlusions generated by scene objects cannot always be known in advance and second, the scene may offer only small and varying viewing angles to the outside, which cannot simply be covered by adding some more cameras. This is especially true for trackable objects surrounded by other objects, e.g. a tool inside a car body.

Our *indirect tracking* approach adds an additional, mobile IR tracking system which can be placed in the scene on-the-fly such that it can see trackable objects that are hidden to the stationary cameras. The mobile setup itself is equipped with a marker so that its pose can be tracked by the stationary setup (see Figure 1(a)). Figure 1(b) shows how the proposed solution would look like in AR-stud-welding, one of our industrial scenarios where the stud-welding gun has to be tracked inside the car body so that navigational information about the next welding position can be shown on a display attached to the welding gun. This application is already in productive use but until now, it suffers from the restrictions of classical outside-in tracking described above.




Fig. 1. Proposed indirect tracking approach

The spatial relationships of the proposed tracking approach are depicted in Figure 2. The nodes of the graph represent coordinate systems and the arrows denote transformations between them. The desired pose of the object to be tracked (T) with respect to the coordinate system used by the stationary cameras (S) is then given by a concatenation of several transformations via the mobile target (MT) and the mobile cameras (M). Of course, the direct transformation between S and T can still be used, when available. It provides a useful reference value for evaluating the quality of indirect tracking in Section 4. In theory, indirect tracking works properly. In practice, however, the indirection via the mobile system can result in serious errors estimating the pose of the object with respect to the stationary system which often represents the world coordinate system. Our goal is to reduce those errors and thereby ideally close the gap in tracking accuracy between indirect and direct tracking.

Positional errors in the concatenated transformations just accumulate when transforming the measurement of the mobile cameras back into our reference frame defined by the stationary cameras. Any orientational error, however, results in positional errors at the region of interest which increase linearly with the distance [1]. This particularly holds for the orientation of the mobile target, the *orientation* estimation of which we assume to be rather imprecise. To clarify the importance of accuracy for the proposed setup, one may imagine that a deviation of 0.1° in the estimated pose of the mobile tracker results in a displacement of 1.7 (3.4) mm between estimated and effective position of the tracked object in a distance of 1 (2) m from the mobile setup. This follows directly from the sine function. The general idea for reduction of this error is to use *common reference points* somewhere in the scene which are visible to both, the stationary and the mobile tracker, in order to estimate the orientation of the mobile tracker as good as possible.

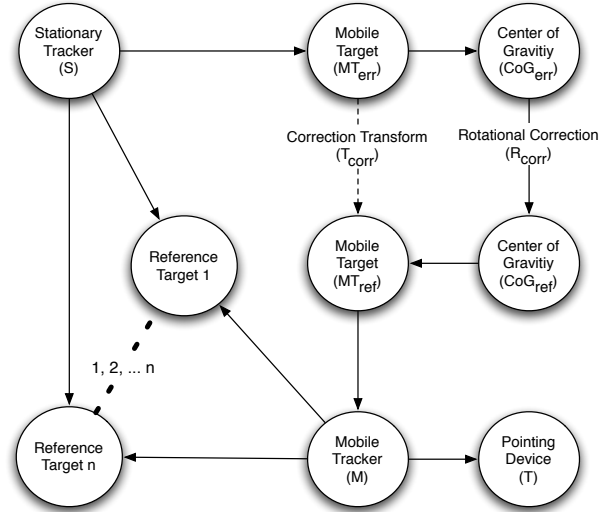


Fig. 2. Proposed indirect tracking approach with additional reference points in the scene

2 Related Work

The relationship between the mobile setup and the tracked object could be described as *inside-in tracking* [12] since both entities are moved more or less dynamically in the scene. This notion has been used first by Mulder [11], synonymously with the notion *relative tracking*. It is not clear, however, whether the relationship between the mobile setup and common reference points as described above should be denoted as inside-in or rather inside-out tracking. This certainly depends on how mobile the reference target actually is, whether it is considered part of the scene or rather of the background. We will stick to the term inside-out tracking for this kind of relationship.

Many researchers have worked on the combination of optical inside-out and outside-in tracking systems and some of them investigated accuracy issues. Bichlmeier et al. [2] used a reference target to combine outside-in instrument tracking with inside-out estimation of HMD pose. In this setup, a monocular tracking system, rigidly mounted above the HMD, tracks the reference target. The latter is simultaneously tracked by the outside-in tracking system, so that the pose of the HMD can be calculated with respect to the outside-in system. Fischer et al. [4] presented a hand-eye calibration algorithm which utilizes the combination of a monocular inside-out and an n-ocular outside-in optical tracking system. Both works, however, did not address the error resulting from their combined setup.

Hoff et al. [7] have also combined inside-out and outside-in optical tracking to improve the accuracy of tracking an HMD. They used two rigidly combined reference targets which are simultaneously tracked by an inside-out and outside-in optical tracking system. Based on accuracy specifications delivered by the manufacturer, they estimated the 6 DoF tracking accuracy for each involved target and for each tracking system. The

estimated error covariances are used for sensor fusion of the inside-out and outside-in tracking system using Kalman filtering [10]. This approach decreases the pose uncertainty by using two independent tracking informations.

3 Correcting Rotational Errors

As already stated above, to improve the accuracy of indirect tracking setups, it is crucial to minimize the error in orientation estimation of the mobile tracking setup. In this section, methods are presented which allow for on-the-fly compensation of this error without extensive recalibration. The first method assumes that the mobile setup is being tracked and constantly corrects the orientation by one or two additional reference points in the scene. The second method gets by without tracking the mobile setup. It needs at least three reference points for pose estimation. Even though the latter approach is not new, we incorporate it in our evaluation for the sake of completeness and also as a reference for our correction methods.

3.1 Tracking of Mobile Setup

One (implicit) reference point correspondence is already given by the center of gravity of the mobile target if we assume that *position* tracking of the mobile target with respect to the stationary tracking system is not error-prone and that the mobile target is registered well with the mobile cameras. Additional reference points may come from marker balls or complete 6 DoF targets that are suitably placed in the scene and that are seen by both tracking systems.

Especially when using a big mobile target (such as shown in Figure 1(a)), it is important to apply the rotational correction in a coordinate frame with the center of gravity of the target marker balls as its origin because it probably is the center of rotation where the rotational error gets added to the true orientation of the target in the pose estimation algorithm used by the vendor [3]. Usually, however, the target frame defined by the vendor differs. The situation is depicted in Figure 2. Via the static transformation $T_{MT_{ref}2M}$ between the mobile target and the mobile system and the translation from the target frame to the *center of gravity* frame T_{MT2COG} , we are able to transform the measurements of the reference point(s) made by the mobile system to the mobile target's center of gravity coordinate frame. The latter transformation can be obtained from the body calibration of the mobile target. Estimation of the former transformation is a bit more complicated (cf. section 4.1). Concatenation of both transformations results in a 3D position of the reference point(s) (as seen by inside-out tracking) in the center of gravity coordinate system of the mobile target. We call these the *reference points*

$$p_{ref} = T_{MT2COG}^{-1} \circ T_{MT_{ref}2M} \circ p_{mobile}. \quad (1)$$

The corresponding measurement made by the stationary tracking system $p_{stationary}$ can also be transformed to the center of gravity frame of the mobile target. However, since the transformation between the stationary tracker and the mobile target $T_{S2MT_{err}}$ is perturbed by rotational error by assumption, this yields a wrong 3D position of the reference point(s) (as seen by outside-in tracking) in the center of gravity coordinate system

of the mobile target. We call these the *erroneous points*

$$p_{err} = T_{MT2COG}^{-1} \circ T_{S2MT_{err}}^{-1} \circ p_{stationary}. \quad (2)$$

Based on the reference point correspondence(s), a rotational correction is computed using one of the methods described in the following subsections. The indirect tracking equation then becomes

$$p_{stationary} = T_{S2MT_{err}} \circ T_{corr} \circ T_{MT_{ref}2M} \circ p_{mobile} \quad (3)$$

with the additional correction transform

$$T_{corr} = T_{MT2COG} \circ R_{corr} \circ T_{MT2COG}^{-1} = \begin{bmatrix} R & -Rt + t \\ 0 & 1 \end{bmatrix}. \quad (4)$$

One Reference Point Correspondence Let us assume for now that rotational error occurs only in directions that are orthogonal to the line joining the center of gravity of the mobile target and a known reference point in the scene. We call this line the *direction of the reference point*. Let us further assume that the position of the mobile target is tracked perfectly and that the 3D position of the reference point is known perfectly in both tracker frames.

The geometry of this problem is shown in Figure 3(a). The center of gravity of the mobile target lies in the center of a unit sphere. By normalization, a general point in the scene can be converted to a point on the surface of the sphere. It is sufficient to consider this directional vector. Every rotation about some axis through the rotation center lets this direction vector describe a circular orbit on the surface of the sphere.

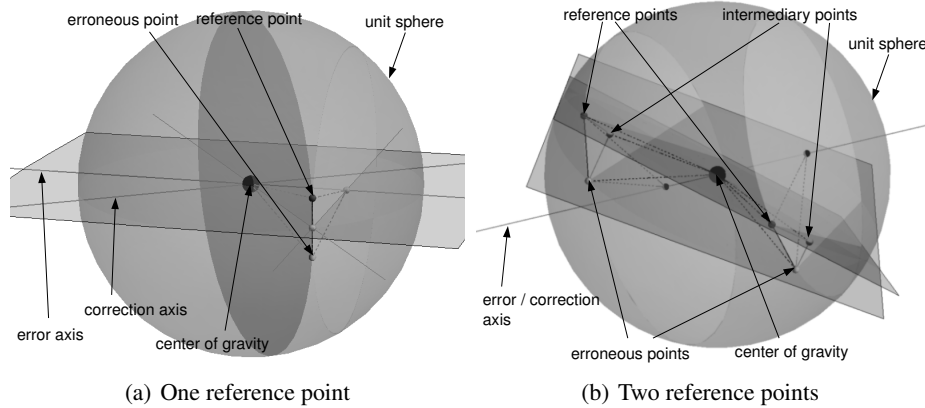


Fig. 3. Correction of rotation using reference point correspondences in the scene. At least two additional correspondences are needed for reconstructing the *true* axis of rotational error.

First, the erroneous point as well as the reference point are projected onto this sphere. Thereby, they define the normal vector to a plane through the center of gravity. It contains all the axes that would be suitable for mapping the erroneous point back

to its reference position. Furthermore, all axes pass through the rotation center. The rotation axis about which the original error rotation took place, is shown in the Figure 3(a) but its orientation in the plane is unknown. However, we obtain another (normalized) rotation axis from the vector cross product of the two measurements of the reference point.

$$a = \frac{p_{ref} \times p_{err}}{\|p_{ref} \times p_{err}\|} \quad (5)$$

With a rotation about this axis, by the angle between the two directions of the erroneous point and the reference point, it is always possible to map the erroneous point back to its reference position. The rotation angle ϕ is obtained by

$$\phi = \cos^{-1}(p_{ref} \cdot p_{err}). \quad (6)$$

The rotational correction can be expressed as a rotation matrix using Rodrigues' rotation formula [6].

For points different from the erroneous point, this correction yields wrong results as soon as the error rotation axis contains the direction of the reference point in its linear combination. In most cases, the error can still be reduced. Due to its measures, e.g. our mobile target (see Figure 1(a)), probably exhibits the greatest rotational uncertainty in its pitch axis (parallel to the long edge). If the reference target is placed in front of the cameras, the orthogonality assumption might be save. On the other hand, an error with its axis being aligned with the direction of the reference point can not even be detected using a single reference point since the point is mapped onto itself by the error.

Two Reference Point Correspondences As alternative, we use two additional point correspondences A and B , thus an overall amount of three if the implicit center of gravity is also counted. The spatial relationships are again depicted in Figure 2. The availability of three non-collinear point correspondences principally enables us to solve the absolute orientation problem (7 DoF similarity transformation between two point sets) between the stationary and mobile tracking system using any absolute orientation algorithm (cf. section 3.2, see also [3]). However, the problem at hand allows us to make various simplifications which lead to a neat solution of the problem.

- It is not necessary to find a scaling factor because both tracking systems already provide metric measurements. This reduces the problem by one degree of freedom.
- If we again transform the known 3D measurements of the two reference targets from their respective coordinate frames of the stationary and the mobile tracking system into the frame of the mobile target's center of gravity, then both coordinate systems already share the same origin. Thus, it is not necessary to solve for the translational part. This reduces the problem by another three degrees of freedom.

The remaining problem is a 3 DoF rotation about the center of gravity of the mobile target. A nice solution has been presented by Horn for estimating the orientation in the special case of coplanar point sets [8]. It is particularly suitable for three point correspondences which are always coplanar. A sketch is given in Figure 3(b). Since one

correspondence is given by the common origin of the center of gravity reference frame, the remaining two correspondences result in two independent position vectors p_{err_A} and p_{err_B} for the erroneous points and p_{ref_A} and p_{ref_B} for the reference points. p_{err_i} and p_{ref_i} define two planes, one containing the erroneous points and one containing the reference points. An informal description of the algorithm shall be given in the following.

1. First, normals to both planes are computed using the vector cross product. From these normals, a rotation axis is deduced, again by the cross product. The angle between the plane normals is the searched for rotation angle. By using Rodrigues' formula for rotating one plane by the computed angle about the computed axis, both sets of coplanar points are mapped into a single plane. This step of course only yields correct results if either all the points are lying in an exact plane or only three points are used. In Figure 3(b), this corresponds to a mapping of the plane spanned by the erroneous points and the origin to the plane spanned by the reference points and the origin. Thereby, the erroneous points are mapped to the intermediary points.
2. The remaining problem is a rotation in the reference plane about the normal to that plane, which maps the intermediate points from the last step to the reference position. Estimating the angle which yields the least-squares solution is now only a one DoF problem. In contradiction to Horn's solution, however, we do not consider the distance between the original points, but the distance of the normalized points on the unit sphere. Thus, the optimal angle is given by the mean of the angles between points A and B. In Figure 3(b), this corresponds to a mapping of the intermediary points to the reference points.

3.2 General Absolute Orientation Approach

Another approach to solve the indirect tracking problem is to abandon the idea of a tracked mobile setup in combination with one or two additional reference points in the scene and use a large amount of point correspondences instead. This solution would require a brief calibration process after movement of the mobile setup and before the on-going tracking process in which the tracked object is being panned within the close-up range of the mobile system which is also visible by the stationary system. The resulting point correspondences could be used to solve the general absolute orientation problem each time the pose of the mobile setup has changed ¹.

Several closed-form least-squares solutions exist for the absolute orientation problem. A comparison conducted by Eggert et al. [3] did not yield a particular preference for any of them. For our evaluation (cf. section 4), we used Horn's quaternion-based solution in a slightly stripped-down version which does not optimize for the scaling factor [8].

4 Evaluation and Results

The methods for correction of rotational error described in the previous section have been evaluated on real tracking data. All tests were carried out using a stationary *A.R.T.*

¹ The mobile target might still be useful in order to detect these position changes.

system consisting of three *ARTrack 1* cameras and a mobile *A.R.T. SmARTtrack* system mounted on a tripod, consisting of two *ARTrack 2* cameras with a baseline of 50 cm.

The acquisition of measurement data was carried out with our Ubitrack framework [9] with the aid of the trackman configuration tool [13]. Data was evaluated offline with *Mathematica*².

4.1 Calibration

For calibration of the camera setups and the involved tracking targets, the mechanisms prescribed by the vendor were used. Each target was calibrated in the stationary tracking setup. It is supposed to yield better results than a two-camera setup with a small baseline of 50 cm. The mobile target was mounted to the mobile setup before calibration of extrinsic camera parameters in order to avoid negative effects of this mechanical intervention with regard to both, camera and target geometry.

It is crucial to calibrate the transformation between the mobile target and the mobile coordinate system very accurately. Assuming that orientation tracking of the mobile target suffers from pose-dependent systematic errors, we decided to integrate measurements from several poses of the mobile setup into the calibration process. There might be better calibration approaches, this is still an open issue.

For each pose of the mobile setup, points were measured with a pointing device in both tracker coordinate systems. Pointing at a rigid surface ensures that the tool tip remains constant during the acquisition of its position by the two measurement systems even under the influence of misaligned sensor timestamps. In a proper setup allowing for continuous measurements, one would have to synchronize the cameras using a common sync source which should not be a problem with quality tracking equipment.

Next, the absolute orientation problem was solved independently for each pose of the mobile setup (cf. section 3.2). The solutions corresponding to the different poses were then combined to a single transformation in order to obtain a single calibration of higher precision. This was done by averaging the translational offsets in each axis as well as the Euler rotation angles.

4.2 Evaluation of the Indirect Tracking Setup

The pointing device was also used to make measurements. Unlike in the intended use-case, it was always visible in both tracking systems. By this means, a relative comparison was possible of 3D points measured indirectly via the mobile system against those measured directly by the stationary system.

For the acquisition of evaluation data, the mobile system was sequentially set up in two distinct poses *one* and *two* within the tracking volume of the stationary cameras. For each pose, 35 point correspondences were measured on a rigid surface in the common tracking volume of both systems. No averaging was performed for the point measurements. Figure 4 shows the measurements of the stationary tracking system as well as the center of gravity of the mobile target for pose *one* as spheres.

² Wolfram Research, Inc




Fig. 4. Evaluation setup for pose *one*. The vectors show for each point the deviation of indirect tracking from direct tracking whereas larger deviations result from indirect tracking without rotational correction. Smaller deviations in this example result from the simple correction using reference point 18.

Instead of placing reference targets in the scene, different subsets of the measured points were used as reference positions in the offline evaluation of the data.

For both poses of the mobile system, the three approaches described in section 3 were tested and compared to the naive indirect tracking approach which works without any correction at all. In Figure 4 the distance of the indirectly tracked points obtained from the naive approach from the corresponding directly tracked points are depicted by light-colored lines, lengthened by a factor of 50 for better visibility. Results for other poses of the mobile setup feature a similar systematic behavior, and the orientation of deviation vectors varies depending on the chosen pose. This suggests a significant amount of systematic error which we cannot fully explain at the moment. The shorter dark lines depict - exemplarily for the simple correction using the point with ID 18 as reference point - the distance of the indirectly tracked points to their directly tracked correspondences, again lengthened by a factor of 50. All conducted evaluations will be described in the following.

1. *Simple correction (1 additional point)*: Based on one reference point correspondence, a correction of the orientation of the mobile target was carried out following the approach described in section 3.1. As stated there, this method only corrects those parts of the original rotational error that are orthogonal to the line joining the center of gravity of the mobile target and the reference point. The correction results are presented in table 1 for six reference points with varying positions relative to the mobile cameras.
2. *Full correction (2 additional points)*: The information of two of the measured points was used to compute an orientation correction according to the method described in section 3.1. The correction result is presented for six pairs of reference points with different positions relative to the mobile cameras. As a test for robustness, the first three pairs were chosen to contain points very near together. The last three pairs each contain points on opposite sides of the volume.
3. *General absolute orientation (5 points)*: A subset containing five of the measured 35 points is used to compute a rigid transformation with an absolute orientation algorithm (cf. section 3.2). The correction result is presented for three sets of five

points each, one near and one far to the mobile cameras as well as one random sample of all points.

4. *General absolute orientation (all points)*: All point correspondences are used to compute the absolute orientation. Since the resulting rigid transformation is optimal in a least-squares sense (cf. section 3.2), the measurement residuals give information about how well the actual mapping can be described by a rigid transformation.

For each indirect tracking approach, the RMS was computed in mm for the deviation of indirectly tracked points from their directly measured counterparts.

Table 1. Performance of rotational correction approaches. For all evaluated constellations of correction approach, mobile setup pose, and reference points, the RMS euclidean distance between indirectly and directly tracked points is given in mm.

No corr.		Simple corr.			Full corr.			Abs. orient.			Gen. abs. orient.	
Pose 1	Pose 2	Pt.	Pose 1	Pose 2	Pts.	Pose 1	Pose 2	Pts.	Pose 1	Pose 2	Pose 1	Pose 2
3.1	2.7	1	1.2	1.8	4,11	1.8	1.6	1,4,7, 9,13	1.2	1.6	1.0	1.2
		15	1.1	1.4	18,25	1.2	1.5					
		29	1.1	1.4	25,32	1.2	3.4	22,25,28,	1.0	1.3		
		4	1.1	1.7	1,14	1.4	1.8	31,34				
		18	1.1	1.3	15,21	1.1	1.4	1,10,22,	1.1	1.4		
		32	1.1	1.6	22,35	1.1	1.4	32,35				

4.3 Discussion

The naive indirect tracking approach yields very bad results. This is not unexpected according to the propagation rules for orientation error (cf. 3).

The residuals from the absolute orientation with all points represent the greatest lower bound for all error corrections based on optimizing a rigid transform. Lower residuals could only be achieved by fitting a more general, maybe even non-linear mapping. However, for a metric measurement system, this is not desirable. There must be a considerable amount of probably systematic error left that can only be explained by discrepancies between the two tracking systems. Possible reasons are the following:

- Deficiencies in the room calibration algorithms of the manufacturer which determine the intrinsic and extrinsic camera parameters
- Errors introduced during body calibration
- Errors introduced during tip calibration of the measurement tool
- Insufficient number of point correspondences during those two types of calibration
- Inaccurate 3D point measurements due to merging markers which is when two marker balls come close to the same line of sight and converge to a single point measurement on the image plane
- Tolerances in the manufacturing process, each marker ball is laminated with retro-reflective foil by hand.

All correction mechanisms perform quite well in converging to the theoretical lower bound. In most constellations, at least two third of the residuals can be explained and corrected by the assumption of erroneous orientation estimation of the mobile target. The full correction seems to have problems with point correspondences that are chosen too close together, especially, if they are far away (points 25 and 32) and therefore result in very similar direction vectors. Thus, reference points should be chosen carefully in order to avoid such degenerate configurations.

Astonishingly, the simple correction is not worse than the full correction. This could be explained by the layout of our mobile target which suggests that pitch movements of the mobile setup can be detected worst by the stationary cameras. Since this pitch axis is roughly orthogonal to the line joining the center of gravity of the mobile target and the reference point, these errors can be corrected quite well. In general, the full correction with a reasonable pair of reference points should be preferred, if it is feasible. A good alternative is the general absolute orientation, if enough point correspondences are available or can be obtained without much effort.

5 Conclusion and Future Work

We have presented an indirect tracking approach as extension to classical outside-in tracking. It is beneficial whenever the scene becomes more complex and occlusion cannot be avoided in a robust manner using a classic stationary outside-in setup. We further showed that it is crucial to minimize error in estimating the orientation of the mobile tracking setup in order to avoid huge degradation of accuracy. The presented methods for correcting this error can be categorized by the amount of point correspondences needed for their application. We showed that even a single correspondence can be enough under some conditions with respect to target layout, to reduce the resulting positional error of the tracked object a by factor of three.

There are some issues which deserve further investigation. The evaluation results for indirect tracking without error correction suggest that the orientational error systematically depends on the current pose of the mobile target. It would be interesting to know more about this phenomenon. Irrespective of the concrete reasons, this systematic error affects the calibration of the static transformation between the mobile cameras and the mobile target (see 4.1). A method is needed to perform this calibration accurately despite this systematic error.

Another interesting issue is the remaining error after fitting an optimal rigid transformation. The evaluation results again suggest a systematic error. This could offer the opportunity for further improvements, either in the algorithms of the tracking system vendor or afterwards by additional calibration.

In addition to the above-mentioned stud-welding scenario, our approach might for example be useful in a surgery where many people are involved and unforeseeable circumstances often do not allow for optimization of the lines of sight in advance. A mobile tracking system could be mounted rigidly to a mechanical arm or put on a tripod so that the physician could move it into arbitrary poses depending on the current situation.

Using reference points which are a *natural* part of the scene, would make indirect tracking even more feasible because additional reference targets become superfluous. This could be achieved on the one hand by mounting a camera on the mobile setup and performing inside-out tracking of features in the scene. The pose of the mobile target obtained via the outside-in IR tracking could serve as a good initialization for the markerless inside-out tracker [5]. On the other hand, there might already be retro-reflective markers in the scene. This is the case for example in our stud-welding scenario where the car body is tracked by the stationary cameras and is equipped with approximately 15 marker balls for this purpose.

References

1. M. Bauer, M. Schlegel, D. Pustka, N. Navab, and G. Klinker. Predicting and estimating the accuracy of vision-based optical tracking systems. In *Proc. IEEE International Symposium on Mixed and Augmented Reality (ISMAR'06)*, Santa Barbara (CA), USA, October 2006.
2. C. Bichlmeier, T. Sielhorst, and N. Navab. The tangible virtual mirror: New visualization paradigm for navigated surgery. In *AMIARCS - The Tangible Virtual Mirror: New Visualization Paradigm for Navigated Surgery*, Copenhagen, Denmark, October 2006. MICCAI Society.
3. D. Eggert, A. Lorusso, and R. Fisher. Estimating 3-d rigid body transformations: a comparison of four major algorithms. *Machine Vision and Applications*, 9(5):272–290, 1997.
4. J. Fischer, M. Neff, D. Freudenstein, and D. Bartz. Medical Augmented Reality based on Commercial Image Guided Surgery. *Proceedings of Eurographics Symposium on Virtual Environments (EGVE)*, pages 83–86, 2004.
5. P. Fua and V. Lepetit. *Emerging Technologies of Augmented Reality, Interfaces and Design*, chapter Vision Based 3D Tracking and Pose Estimation for Mixed Reality, pages 1–22. Idea Group Publishing, 2007.
6. R. Hartley and A. Zisserman. *Multiple View Geometry in computer vision*. Cambridge University Press, 2000.
7. W. Hoff and T. Vincent. Analysis of head pose accuracy in augmented reality. In *IEEE Transactions on Visualization and Computer Graphics*, volume 6, pages 319–334. IEEE Computer Society, 2000.
8. B. K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America*, 4(4), Apr. 1987.
9. M. Huber, D. Pustka, P. Keitler, F. Echlter, and G. Klinker. A system architecture for ubiquitous tracking environments. In *Proceedings of the 6th International Symposium on Mixed and Augmented Reality (ISMAR)*, Nov. 2007.
10. R. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.
11. A. Mulder. Human movement tracking technology. Technical Report 94-1, School of Kinesiology, Simon Fraser University, Burnaby, B.C., Canada, 1994.
12. S. Nölle. *Augmented Reality als Vergleichswerkzeug am Beispiel der Automobilindustrie*. PhD thesis, Technische Universität München, 2006.
13. D. Pustka, P. Keitler, M. Huber, F. Echlter, and G. Klinker. Ubiquitous tracking for quickly solving multi-sensor calibration and tracking problems. Demonstration at IEEE International Symposium on Mixed and Augmented Reality (ISMAR'07), November 2007.