

Using a Panoramic Camera for 3D Head Tracking in an AR Environment

B. Giesler, T. Salb, R. Dillmann
University of Karlsruhe (TH)
IAIM

T. Weyrich
ETH Zürich

June 4, 2002

1 Introduction

Augmented Reality (AR, for short) is the layering of Virtual Reality elements (such as 3D models or markers) over a viewer's image of the real world. This can be achieved in a number of ways, the most popular of which is to have the user wear a pair of semi-transparent 3D goggles. In Karlsruhe, we are constructing an AR system for human-robot interaction, to make it easy for the user to immediately see the interpretation of his or her actions by the computer. To be able to overlay elements of simulation with the view through the 3D glasses, it is necessary for the system to very precisely keep track of the viewer's head position and direction of gaze.

The tracking must meet real-time requirements to prevent lagging of the virtual elements. Since it is very difficult to model human head motions, the benefit of using motion-prediction techniques is very limited. For real-world usability, the tracking should also be able to cover as wide an area as possible, should be very easy to set up, require little or no modification to the environment and, last but not least, should not be prohibitively expensive.

Some existing approaches to this problem make use of external camera systems, such as the commercially available *POLARIS* [NDI02]. Such systems, while often very precise, have a limited working space and are therefore more suitable for usage in stationary applications, such as surgical aid [Sal00]. Others require extensive modification to the environment, such as the University of North Carolina's *HiBall* tracking system [WB+99].

Our novel approach consists of a panoramic camera that is affixed to the AR glasses and tracks artificial environmental features. This approach, using only a single camera, requires relatively little computation and therefore comes close to real-time requirements; it is also inexpensive. We are currently tracking artificial targets that can be mounted very quickly and easily, so very little environmental modification is necessary.

2 Properties of the System Components

We have decided to use a system that uses artificial landmarks instead of natural environmental features, because we consider 6-DOF position reconstruction using natural features to be still too difficult and most of all too costly in terms of processing time to meet with the required real-time constraints. Therefore, we can design both the features that we are recognizing and the sensors that we recognize them with to be perfectly adapted to each other.

2.1 Properties of the Panoramic Camera

A panoramic camera is a system using a CCD camera taking pictures of a convex mirror that reflects a distorted view of the environment. This mirror can be a half-sphere, cone, paraboloid or any other regular convex body.

Using panoramic cameras for position reconstruction is not in itself a novel approach; such cameras have been successfully used, for example, for position estimation of mobile robots (with a conical mirror: [FS+98], [FS+97]; with a spherical mirror: [G+V99]). However, so far work has been limited to two-dimensional reconstruction. For three dimensions, a novel approach is needed, and the camera should fulfill certain conditions.

For two-dimensional position reconstruction, it is sufficient to locate environmental features that lie in the plane of reconstruction, or in a plane parallel to that. For a panoramic camera, this means that the feature images must lie on a cross-section of the camera's mirror. For reconstruction, it is therefore sufficient to know the optical characteristics of that single cross-section in the resulting camera image.

To make three-dimensional position reconstruction possible, it is necessary to find environmental features distributed in all three dimensions around the camera; for a panoramic camera, this means that the features must be found over the entire mirror, i.e. the optical characteristics of the entire camera image must be known.

This last criterion is only given if the camera's mirror has the shape of a paraboloid. A paraboloidal mirror has an important property: All rays of sight that hit the mirror parallel to the mirror's symmetry axis are reflected in such a way that they all seem to originate in the focal point of the paraboloid. Therefore, it is possible to tell the elevation ϵ and azimuth α in real space of a certain point in the image directly from its image coordinates (see figure 1).

A parabolic mirror has the disadvantage, however, that images of objects are distorted in a non-trivial way; while a circle reflected in a conical mirror becomes an ellipse, a paraboloidal mirror distorts it into an egg shape which is not so simple to reconstruct, and whose center of gravity does not correspond with the center of the original circle. But as shown in section 3, reconstruction of the shape is not strictly necessary for high accuracy.

The camera currently in use (figure 2) has a mirror covering an angular area of 360° azimuth and 60° elevation. The mirror's image is taken by a 640×480 pixel CCD and encoded as a PAL signal. The resolution is very low considering the large section of the environment that is depicted; figure 3 shows an image taken with the camera and its cartesian reconstruction. However, simulation and experiments both show that the low resolution mainly poses problems for actual landmark recognition; the camera resolution has only marginal effect on the accuracy of the position reconstruction if sufficiently many landmarks are successfully recognized.




Figure 1: Deflection of rays of sight by a paraboloidal mirror.



Figure 2: The omnidirectional camera.



Figure 3: A 360° image and its cartesian reconstruction

2.2 Properties of the Artificial Landmarks

To meet with the real-time requirements outlined in section 1, landmarks should be designed in such a way that we can find them very easily and quickly. We are currently using circular landmarks in the primary colors red and blue, as shown in figure 4. The landmarks are not uniquely coded in any way, because owing to the limited resolution of the panoramic camera they should be extremely simple in structure to be easily recognizable even if their image is only a few pixels large.




Figure 4: Circular landmarks in red and blue.

The landmark colors make it easy to find the exterior and interior of the landmarks by examining only the red and blue channels of the RGB image stream (thus working on two separate binary images) and finding blobs. This results in two sets of blobs, one for each channel. We then calculate the size of each blob's bounding box and compare the bounding boxes in the red channel with those in the blue channel. If a 'red' bounding box lies within a 'blue' one or vice versa, and the sizes of the bounding boxes fulfil, within certain limits, the known size ratio of the inner and outer circles, a landmark has been found.

Figure 5 shows the basic process of landmark identification.

3 Landmark Recognition and Pose Reconstruction

The outlined process makes it possible to quickly find landmarks in the camera image. However, their *precise* location relative to the viewer is still not known. The camera's limited resolution necessitates the development of some scheme to achieve sub-pixel accuracy. As stated in section 2.1, the camera's parabolic mirror makes simple center-of-gravity calculation difficult; but in the course of this work, a method has been developed that is both more precise and makes excellent use of the camera's properties.

3.1 Conic matching

Since working with the 2D representation of the landmarks is not practical due to the distorting properties of the parabolic mirror, we have decided to use the landmarks' 3D shape. Therefore, the border pixels are detected where a landmark's outer ring adjoins to the inner circle. It is known that these pixels must (within some margin of error) lie on a cone with an elliptical cross-section whose origin is the optical center of the camera. Therefore, rays from the optical center are constructed that go through the border pixels, and a least-squares error minimization algorithm is used to construct a conic through the optical center that approximates the rays most closely.

As can be seen in figure 6, one final ambiguity must be resolved: The conic matching process delivers a very good estimate of the landmark's perimeter, but not of the landmark's center. If the landmark is not seen straight-on, there are two possible surface normals and therefore two possible centers. A simple heuristic is used to resolve this ambiguity, based on the fact that in human-inhabited areas almost all angles are 90° angles.

Up to this point, all calculation takes place in ray space, that is, the three-dimensional rotational coordinate system rooted at the optical center of the camera. Due to lack of stereo information, it is not possible to determine the distance to the landmark at this point; therefore, all rays are represented as infinite lines. This method delivers a high sub-pixel accuracy while operating directly on the camera image, which is given in ray space. No conversion to the cartesian coordinate system is necessary, which makes this approach both fast and accurate.

Unfortunately, the method only works for landmarks that are close enough for a sufficiently large number of border pixels to be detected. Therefore, we use a simple center-of-gravity approach as a fallback if the




Figure 5: Classification by color and bounding box ratios




Figure 6: Border pixel detection and conic matching

landmark's image occupies less than 16x16 pixels. If the image is so small, the distorting effect of the parabolic mirror can be neglected.

3.2 Pose reconstruction

After the landmarks have been detected and identified, the viewer's pose can be reconstructed: As stated above, the angles of the landmark normals are now known, and therefore the angle that the camera takes relative to each landmark are known as well. When the positions of the landmarks in cartesian space are known, the reconstruction becomes trivial. Theoretically three recognized landmarks would suffice to determine the camera's pose, but if more can be found in the image, the resulting over-determined system can be used to enhance accuracy.

The pose is reconstructed via translatic and rotatoric adaption. The adaption steps are also performed in ray space and are repeated until the resulting error lies below a certain threshold¹. The translatic adaption calculates the approximate section point of the viewing rays of the landmark centers, and then matches the viewer's position to that point. The rotatoric adaption uses the method described by B.K.Horn in [Hor87].

4 Results and further work

The system has been implemented initially in a simulation environment, to be able to experiment with several types of virtual cameras and use the results in choosing the optimal real camera. Therefore, we present two sets of experimental results; one for the simulation system and one for the actual camera.

4.1 Results in simulation

In [Wey01], a simulation system has been developed that transforms a 3D scene into a simulated camera image. This system has made it possible to test the effects of changes in camera resolution and interlaced/non-interlaced sensors on the accuracy of position reconstruction. Figure 7 shows an image obtained from the camera simulation.




Figure 7: A simulated camera image

The simulation environment allows the recording and playback of trajectories in the simulated 3D scene. In playback, the position reconstruction works on the simulated camera image. Figure 8 shows a recorded trajectory

¹It can be shown that the employed method does not converge in some rare cases. Therefore, the reconstruction process is aborted after a maximum of 48 repetitions.

from two different angles (*a* and *b*). The graphs on the right display the rotational (*c*) and translational (*d*) errors corresponding to the reconstructed trajectory.

The graphs show that the rotational error on this (fairly typical) run is always well below one degree, and the translational error below 5mm. These results were achieved by simulating a NTSC camera of 640x480 pixels interleaved resolution. Moving toward a simulated camera of 1024x768 pixels and “progressive scan” technology improves the medium errors to 4.25mm (10.9% improvement) and 0.14° (4.4% improvement). This shows that an increase in resolution alone does not significantly increase the accuracy of the method.

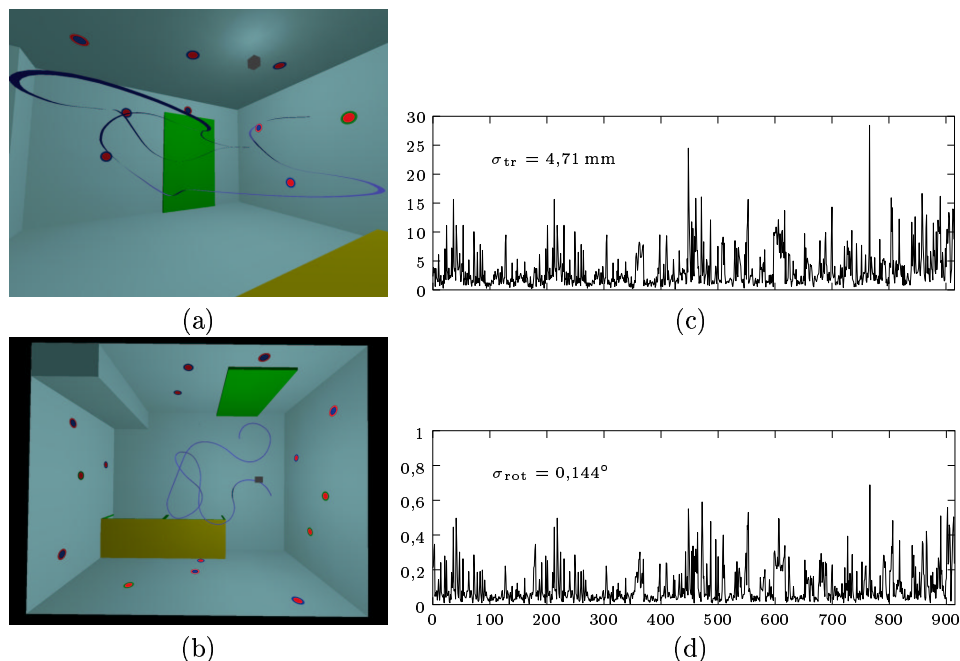


Figure 8: A simulated trajectory and the corresponding translational and rotational errors

4.2 Results with a real camera

The actual panoramic camera is an NTSC camera yielding pictures at 640x480 pixel interleaved resolution, equalling the camera simulation. The results with this real camera mirror those achieved with the simulated one; however, image noise (which was disregarded in simulation) poses a large problem; the color-separated red and blue images are too noisy to effectively recognize the more distant landmarks. Figure 9 shows the results of a test run with this real camera; the dashed lines mark a time segment when the camera was moved out of the range of several landmarks. Naturally, in this time the error increased well below the tolerable measure, which reflects in the medium accuracy over this test run.

4.3 Further work

The camera module used at the moment is far too noisy to be really useful for position reconstruction in larger areas. Therefore, our main focus right now is twofold: First, we are examining image enhancement techniques to investigate whether images from the existing camera can be improved sufficiently; secondly, other more adequate camera modules are being evaluated.




Figure 9: Errors from a trajectory using the actual camera

Once a sufficiently good camera has been found, the system will be enhanced to support automatical calibration, so that the landmark positions do not have to be initially known. Afterwards, we will use the system as the main head tracking component in the AR system used at the University of Karlsruhe for robot programming by demonstration.

References

- [FS+98] M.O. Franz, B. Schölkopf, H.A. Mallot, H.H. Bühlhoff, A. Zell. *Navigation mit Schnappschüssen*. In: Proceedings of the 20th DAGM Symposium, Springer Verlag, Berlin 1998.
- [FS+97] M.O. Franz, B. Schölkopf, H.H. Bühlhoff. *Homing by Parameterized Scene Matching*. In: Proc. of the 4th European Conference on Artificial Life, MIT Press, Cambridge 1997.
- [G+V99] J. Gaspar, J.S. Victor. *Visual Path Following with a Catadioptric Panoramic Camera*. In: Proceedings of the 7th International Symposium on Intelligent Robotic Systems (SIRS'99), University of Coimbra, 1999.
- [Hor87] B.K. Horn. *Closed-form solution of absolute orientation using unit quaternions*. In: Journal of the OSA, Issue 4, April 1987.
- [NDI02] Northern Digital Inc. *Product website for the Polaris Optical Tracking System*. At <http://www.ndigital.com/polaris.html>
- [Sal00] T. Salb et. al. *Intraoperative presentation of surgical planning and simulation results using a stereoscopic see-through head-mounted display*. In: Proceedings of Stereoscopic Displays and Virtual Reality Systems VII, part of SPIE / Photonics West 2000, San Jose, CA, January 2000
- [Sal01] T. Salb, O. Burgert, T. Gockel, B. Giesler, R. Dillmann. *Comparison of tracking techniques for Intraoperative Presentation of medical data using a see-through head-mounted display*. In: Proceedings of Medicine Meets Virtual Reality 9, Newport Beach, CA, January 2001
- [WB+99] G. Welch, G. Bishop et al. *The HiBall Tracker: High-Performance Wide-Area Tracking for Virtual and Augmented Environments*. In: Proceedings of the ACM VRST 99, University College London, December 20-22, 1999.

- [Wey01] T.Weyrich. *Entwicklung eines Kopfverfolgungssystems auf der Basis einer Panoramakamera und künstlicher Landmarken*. Master's Thesis, University of Karlsruhe (TH), 2001.