

Magnetic Tracker Calibration for Improved Augmented Reality Registration

Mark A. Livingston

Andrei State

Department of Computer Science
University of North Carolina at Chapel Hill*

August 30, 2001

Abstract

We apply a look-up table technique to calibrate both position and orientation readings from a magnetic tracker for use in virtual environments within a defined working volume. In a test volume of 2.4 cubic meters, the method reduced the tracker's average position error by 79% and its average orientation error by 40%. We test the correction table against the tracker's performance outdoors (a metal-poor environment) and show that readings taken in our lab and corrected by our method exhibit less error than uncorrected readings taken outdoors. We demonstrate that such reduction in position error visibly improves registration in an augmented reality system, whereas the (lesser) reduction in orientation error does not visibly improve registration. We show that the model we used for the orientation error function was incorrect, preventing our method from achieving better correction of orientation error. We discuss future directions for correction of orientation error.

1 Introduction

1.1 Purpose and Motivation


In virtual reality (VR) and augmented reality (AR) systems, magnetic tracking devices are often used to determine the position and orientation of the user's head within a virtual world. This data is passed to an image generator which renders a picture matching the reported values. Magnetic trackers are subject to inaccuracy in their reports, however. In AR environments, inaccurate reports lead to misregistered synthetic elements. To avoid this, data from the magnetic tracker should be corrected before being passed to the image generator. One such method for correcting reports is to store a table of correction factors corresponding to locations and look up a correction factor to apply for a given position. We will refer to this as the look-up table (LUT) technique.

1.2 Background

1.2.1 Calibration in Augmented Reality

In AR applications, virtual objects are superimposed on the real world, either with lenses or half-silvered mirrors (*optical see-through* [Furness86]) or electronics (*video see-through* [Bajura92]). For the merged environment to be convincing, the virtual objects must be aligned with the real objects. This requires an extremely high degree of precision in calibrating the components of the AR system: the head tracker, the head-mounted display (HMD), the real-world objects, the world coordinate system, and, in video see-through, the cameras and lenses.

* University of North Carolina at Chapel Hill, Department of Computer Science, CB#3175 Sitterson Hall, Chapel Hill, NC 27599-3175. {livingst, state}@cs.unc.edu



1.2.2 Magnetic Tracking

Magnetic tracking systems generate a magnetic field from a transmitter and measure the field with a receiver attached to the object being tracked. These measurements determine the position and orientation of the receiver in relation to the transmitter. Metal and electromagnetic devices (e.g. computers, CRTs, metal desks, metal floors, electrical conduits, and ventilation ducts) in the environment distort the transmitter's field. Such interference results in incorrect data being fed to the position and orientation computations. This leads to poor registration of real and synthetic objects, as is the case for the grid shown in Figure 1a. Any metal that is closer to the transmitter than the receiver will distort the readings from the receiver, whether the metal is between the transmitter and receiver or on the opposite side of the transmitter from the receiver.

A widely-used magnetic tracking device is the Ascension Technology Flock of Birds (FOB) with Extended Range Transmitter (ERT) [Ascension95]. The ERT emits a pulsed DC magnetic field and tracks within a hemisphere¹ with a radius of eight feet. The numbers in the FOB reports have 14 bits of precision and a position resolution of 0.45 mm. A factory calibration technique is available for the FOB with ERT.


Twelve measurements are taken by the FOB-ERT system in order to compute the position and orientation of the receiver in the magnetic field. For each of four magnetic field environments (Earth's field alone, Earth's field plus a field in the transmitter's x direction, Earth plus y , and Earth plus z), the receiver measures the field's effect in its local x , y , and z directions. The FOB then computes the distance to the receiver from the ERT, the direction (specified by two angles) to the receiver from the ERT, and the relative orientation angles (Euler angles) that give the orientation difference between the ERT and the receiver. These angles are converted to the requested output format. (We use a vector and a quaternion.) The algorithms used for determining position and orientation of the receiver relative to the ERT follow the SPASYN system [Kuipers80].

1.2.3 Mechanical Trackers

Mechanical tracking devices use high-precision analog or digital measuring devices attached to jointed arms to give position and orientation reports. Since the devices do not use electromagnetic, sonic, or light waves, they avoid interference from the environment. However, there are practical limits on their range.

One such mechanical tracker is the Faro Metrecom IND-1 (Faro arm). The Faro arm is a six-degree-of-freedom tracker with an approximate six-foot spherical measuring range [Faro93]. The Faro arm is constructed of anodized aircraft aluminum. Each joint has a hybrid analog-digital rotary transducer. The Faro arm returns a vector for

¹The user can select one of six hemispheres within a sphere; each hemisphere is centered around either the positive or negative half of one of the transmitter's axes. As noted by Ascension Technology, "If the application notes the crossing between hemispheres, the tracking is full sphere." [Work96]



position and Euler angles for orientation of the tip relative to a base coordinate system located at the center of a reference sphere. The Faro arm reports have eight decimal digits of precision. The arm has a position accuracy of 0.46 mm (2 Sigma). The orientation accuracy of the Faro arm is not specified by Faro Technologies. The high degree of accuracy of the Faro arm is demonstrated in Figure 1b. This figure indicates that our AR system is a valid test and that our calibration techniques work with success proportional to the accuracy of the device that tracks the camera. Figure 1b has good registration because the mechanical tracker is very accurate, while Figure 1a has poor registration because the magnetic tracker suffers from distortion of its magnetic field. Figure 1b also indicates that the orientation reports from the Faro arm are indeed quite accurate.

1.3 Contribution

We apply a look-up table calibration method to a magnetic tracker and shows that it greatly reduces the average magnitude of the position error (as shown by [Bryson92] and [Ghazisaedy95], see Section 2) and moderately reduces the orientation error (a new finding). We also demonstrate that the orientation error depends on the measured orientation (a new finding). This contradicted our initial assumptions and prevented the method from achieving better orientation correction. It also makes a LUT technique appear impractical for orientation correction, due to the exponential increase in both the number of samples required and in the storage cost of the table.

The method performs well compared to the tracker's performance outdoors, in a virtually metal-free environment. We show with an AR system that this correction method improves the registration of the computer-generated imagery with the real world noticeably, but not sufficiently for registration in an AR environment (Figure 2).

Section 2 briefly reviews similar experiments. We present in Section 3 a brief analysis of calibration methods and error metrics for position and orientation readings, and a detailed description of our method. In Section 4, we present numerical and visual results. We draw conclusions in Section 5 and discuss future work in Section 6.

2 Previous Work

Bryson [Bryson92] calibrated a Polhemus Isotrak with three different calibration methods, one by mapping the distortion with a least-squares fit to a polynomial function, and two by using a LUT (one with linear interpolation and one with a Gaussian-weighted interpolation). Bryson measured errors on a rectilinear grid defined by a pegboard. He found that the Gaussian-weighted LUT and function-fit methods performed best. He also found that extreme noise in his measurements beyond 1200 mm from the source made a calibration method unreliable beyond that distance. Finally, he found that the measurements were not repeatable from day to day, which limited the potential

for calibration to within the repeatability of the measurements (25-50 mm)².

The CAVE VR system [Cruz-Neira93] at the Electronic Visualization Laboratory of the University of Illinois at Chicago uses a FOB with ERT to determine position and orientation. Ghazisaedy et al [Ghazisaedy95] calibrated their system with a LUT method. They used an ultrasonic measuring device to determine position correction vectors at sample locations. They took data on an approximate rectilinear grid, at a resolution of one foot, to calibrate a ten-foot cubic volume, and trilinearly interpolated intermediate points. They reduced long-range (three-meter tracking range) maximum errors from 1200 mm to 82 mm. They reduced short-range (one-meter tracking range) maximum errors from 180 mm to 40 mm. The latter is virtually the same relative correction of position error as our method gives, though our data has smaller pre-correction errors. The maximum pre-correction error for our data (comparable to the short-range experiment in the CAVE) was 108.0 mm, and the maximum post-correction error was 23.1 mm.

3 Calibration Method

Calibrating a magnetic tracking device involves mapping the distortion of the magnetic field, characterizing the distortion function, and providing a correction algorithm (which follows directly from the characterization). This requires a reliable method of collecting data and the definition of an error metric to analyze the calibration. It also requires some familiarity with resampling and averaging transformations, particularly rotations. We show how we set up a transformation loop to determine the tracker error. Since our method relies on auxiliary devices, we want to make sure that we use them in a way that will not interfere with the calibration.

3.1 Data Collection

To map the distortion of the magnetic field, one must collect data tuples that consist of the measured data and the “truth values” – what an accurate tracker should report. The accuracy of these truth values is crucial to the accuracy of the calibration method as a whole, since the goal of the calibration procedure is to map measured data into (or, at least, as close as we can get to) truth values. The measured data is read directly from the magnetic tracker being calibrated. Averaging multiple readings for each sample can reduce the noise and dynamic error in measured readings. The truth values must be generated by either a precise auxiliary tracker or by a sampling lattice in which the receiver can be mounted in a rigid and repeatable manner for both position [Bryson92] and orientation.

Sample points can be collected either on a rectilinear grid or as irregularly distributed set. The choice of method does not affect the final representation of the distortion function, and both methods must allow samples to be taken at arbitrary rotations if orientation error is to be corrected. (In fact, we shall see in Section 4.7 that too many orientations are needed per sample point for orientation correction via the LUT method to be practical.) Sampling on a rectilinear grid does not require a resampling algorithm to form a LUT, but is very difficult to perform in a precise, automatic way without metallic structures. Collecting an irregularly distributed set of samples is somewhat easier.

3.2 Error Metrics

After collecting data, we need a way to measure its quality. We define error metrics in terms of the local coordinate systems (CS) of the devices used. For the transmitter and receiver of the FOB, these CS are defined in the user manual [Ascension95]. For the tip and base of the Faro arm, these CS are defined in its user manual [Faro93].

For position error, we compute the magnitude of the distance between the two local CS. This metric can be computed with vector algebra. For orientation error, we measure the angle through which the measured or corrected local CS must rotate to match the true local CS. This metric can be computed with quaternion algebra. We are less concerned with the axis of rotation. A small angle around any axis implies a small error, whereas a large angle around any axis implies a large error.

We created a noise metric for the position and orientation samples. For each sample, we take multiple readings to reduce the effects of noise. Some points exhibit excessive noise and must be excluded from the computations. This noise metric also detects samples for which the receiver was not still. For the position portion of samples, the noise metric is defined as the length of the diagonal of the bounding box enclosing the multiple readings of a sample. For the orientation portion of the samples, the noise metric is computed by rotating a unit vector by each reading’s quaternion in turn. The length of the diagonal of the bounding box of the respective endpoints of the rotated vectors is the noise metric. (This is a pessimistic measure, but easy to compute.) We derived a maximum allowable length

²All measurements cited in this section were reported in feet and/or inches. We have converted them for consistency. The accuracy of the reported numbers varies. We have used approximate values where few significant digits were provided.

for that diagonal for each tracker from the tracker's resolution specifications and practical experience with the tracker in our lab. The maxima used in the experiments described in Section 4 were 20 mm for FOB position, 7.5 mm for Faro position, 0.05 mm for FOB orientation (2.9 degrees), and 0.018 mm for Faro orientation (1.0 degrees).

When we collect samples of position and orientation error, we are sampling a distortion function of the magnetic field. It is well-understood in sampling theory that, in order to accurately characterize a function, one must sample at twice the frequency of the highest frequency component of the function. In this application, undersampling will result in widely varying correction factors for nearby samples. The noise metrics above can also be used to detect this condition. We can also visually inspect the samples and resulting LUT to see that the error changes smoothly between nearby samples. If the metric for either position or orientation is above an error threshold (derived not from the precision, as described above, but from the alignment accuracy required and/or from the accuracy of the tracker), the samples are not dense enough to produce a smooth approximation to the error at intermediate locations.

3.3 Function Fitting vs. Look-up Table

Distortion of the magnetic field is the major source of position and orientation errors from the tracker. This distortion can be approximated by a function or sampled into a LUT. To obtain a function, tuples of measured and true points are fed into an optimization procedure such as least squares fit, with a list of basis functions (constants, linear terms, transcendentals, etc.). One problem with such an optimization is that it is unknown whether a good approximation function, or set of piecewise functions, can be found. One might have to sacrifice local accuracy in a part of the working volume in order to obtain better global accuracy. In addition, the functional approximation cannot be guaranteed to capture all of the fine detail in the distortion, or it may introduce variations of frequency higher than the frequencies which exist in the data.

In contrast, a LUT uses local information to compute the table entry for a given point. It still cannot guarantee the capture of the fine details of a high degree distortion function, but the solution to this is to increase the sampling density and the resolution of the table. Increasing the sampling density increases the amount of data used to generate each table value. To determine whether the table accurately captures the data in the sample set, one can measure the error in approximating a table entry by interpolation from neighboring table entries.

For either function fitting or LUT, there must be a defined set of dependencies for the corrected values. An auxiliary program available from Ascension Technology [AscensionFTP] assumes that both position and orientation error are dependent solely on the measured x , y , and z positions. The consequence of this assumption is that the error correction method requires only three parameters instead of six. This reduces the number of terms and degree for polynomial correction functions, or the dimension of (and hence the memory required for) a LUT. To test whether this assumption is correct, one must take readings at a fixed position with varying orientations. After failing to adequately correct orientation error, we determined that this assumption does not hold (Section 4.7).

3.4 Representing and Averaging Rotations


Quaternions have been used to represent rotations of space in computer graphics for about a decade [Shoemake85]. Tracking systems now commonly report orientation with quaternions. They are favored over Euler angles because they do not rely on an assumed order of application of three rotations, as Euler angles do. They are favored over matrices because of their compactness (four numbers for a quaternion versus nine for a rotation matrix).

In building a correction table for orientation reports, we must average or interpolate between the computed orientation error at several points. Averaging rotations is done by summing the matrices that correspond to the rotations and computing the singular-value decomposition (SVD) of the summed matrix (a method described in more detail in [Curtis93]).

3.5 Coordinate Systems

As noted above, our calibration method relies on having two readings for each measured sample. One is the reading from the tracker, and the other is the truth value. We generate the truth value by calibrating the two trackers to each other and forming a closed loop of transformations (that would compose, if all were accurate, to the identity). The difference from this composite transformation to the identity is then the (composite) error. We attribute all of this error to the magnetic tracker – a reasonable approximation given the accuracy of the mechanical tracker and of our calibration techniques. Figure 3 depicts the transformation loop used for our sample collection; a different set of tools should produce a similar loop.

At the top of Figure 3 is the ERT, supported by a wooden frame as it hangs from the ceiling in our lab. The underside of this wooden frame defines the world CS (Frame). We calibrate the Faro arm to the world system by



touching points on the frame with the Faro arm tip to form a square in the Frame CS; the readings taken from the Faro arm during this procedure were used to compute the Frame-to-Faro transformation. We calibrate the FOB to the Frame CS by taking a series of FOB readings while rotating the receiver (mounted on a special calibration tool) around the edges of the Frame; we then compute the Frame-to-FOB transformation [Bajura95]. We built a custom mount for the Faro arm and attached a 0.3-meter plastic rod to it. This rod provides separation between the stainless steel mount of the Faro arm and the receiver of the FOB. Without this separation, the mount would interfere with the readings from the FOB. We calibrate the plastic rod to the Faro arm using a procedure provided by the arm's controller [Faro93]. We then affixed the FOB receiver to the plastic rod.

In a separate experiment, we determined the center of the receiver with a series of readings in FOB space, in a small volume where we assumed the distortion to be constant (thus relative measurements are valid). Since we only need relative positions to measure the center, this assumption is sufficient. We then computed the rod-to-receiver transformation.

3.6 Interference

We cleaned the environment of unnecessary metal and electromagnetic devices, but we had to keep the (aluminum) Faro arm in the environment to determine the truth value for the readings. We therefore wanted to test whether the arm would interfere with the FOB readings.

To do this, we fixed the FOB receiver's position and orientation in the world and took a series of samples. For the first set of samples, nothing in the environment moved. For the second set, the researcher moved the Faro arm to varied positions that occur during sample collection. For each set, we measured the variation in position and orientation readings using the metrics described in Section 3.2. As shown in Figure 4, neither the position or orientation noise or error changes significantly when the aluminum Faro arm moves in the environment.




Figure 4: Readings from the FOB with no motion of Faro arm (light) and motion of Faro arm (dark). Left: Graph of relative position error (mm) versus time (sec); Right: Graph of relative orientation error (deg) versus time (sec). This amount of noise and error may seem high, but the important note is that the light and dark lines are not significantly different, showing that the motion of the Faro arm has minimal, if any, effect on the readings from the FOB.





Figure 5: Views of the collected samples. Left: 3D view of the position error at the collected samples. Center: Noise metric (in meters) for position readings from the FOB, plotted as a function of distance from the FOB transmitter (in meters). Right: Noise metric (in meters—see text) for orientation samples from the FOB, plotted as a function of distance from the FOB transmitter (in meters). The black horizontal line represents the cutoff limit used in construction of the LUT.

4 Experiment and Results

4.1 Sample Collection

We collected an irregularly distributed set of samples (Figure 5) and resampled the data into a rectilinear grid (Figure 6). This made data collection easier and allowed multiple samples to contribute to a single grid point. Samples were collected in an approximate 3D grid, but for ease of sampling, not of necessity. After clearing the environment of those sources of distortion that would not be present during use of the tracker, we collected a total of 12801 samples. Remaining sources of interference were a metal drop floor (2.25 meters below the transmitter), various CRTs (all at least 3.0 meters from the transmitter), an electrical conduit (1.0 meters above the transmitter), and the Faro arm (on average, one meter below the transmitter, but not stationary).

For each sample, we took eleven readings from the FOB and three readings from the Faro arm, reflecting the relative speeds at which we can read the two trackers. Application of the noise metrics discussed in Section 3.2 eliminated 1787 of the 12801 samples. The remaining samples were contained within a sphere of radius 2.0 meters. We then averaged the eleven FOB readings into a single FOB report and the three Faro readings into a single Faro report. The position error is the (vector) difference between the composed transformation from the Receiver to the FOB (via the Faro arm) and the Receiver-to-FOB transformation as reported by the FOB. The orientation error is the (quaternion) difference between the composed transformation from the Receiver to itself (around the loop, via the FOB and Faro arm) and the identity.



4.2 Building the Table

We built a 3D table containing position and orientation *corrections* at each grid point (Figure 6). We selected the extent of the LUT for each dimension and a resolution of 5 cm for each dimension. The resampling process used a Gaussian kernel (σ of 5 cm, cutoff at 15 cm) centered at each grid point in turn to perform a weighted average of nearby samples. We required that every grid point have at least two samples within the cutoff radius; most had several and some had dozens. We also set a minimum total weight (0.01) for the samples, so that at least one sample was close to the grid point, in case there were few samples near the point.

We computed the error metric for table resolution described in Section 3.3. The average position error between the interpolated grid points and the computed grid points was 4.7 millimeters. The average orientation error was 0.56 degrees. This is a small percentage of the average error at the samples, indicating that we captured most of the information contained in the sample set. That is, there is little to be gained by increasing the resolution of the table without also collecting additional samples.

4.3 Correcting with the Table

Position correction is performed by trilinear interpolation of the error vectors for the corners of the cell into which the sample falls. Orientation correction is performed by an analogous interpolation of the error quaternions for the corners of the cell. We reduced the trilinear problem to a sequence of simple linear interpolations, each between two quaternions; we used spherical linear interpolations [Shoemake85]. (Use of the SVD-based averaging method for rotations does not significantly change the results.) Using a Gaussian kernel to weight correction vectors and quaternions corresponding to the corners of the cell also does not significantly change the results.


To check that this correction procedure is valid, we extracted from the original samples that we used to build the table a set of test points. This yielded 5334 valid samples inside the table. When these *dependent* points are corrected using the table, the error should be zero. Any remaining error must be due to resampling error, interpolation error, calibration error, noise in the data, or bad data that went unnoticed. The results of this “sanity check” are graphed in Figure 7.

The table can also be used for extrapolation – that is, correction of samples not in the volume of the LUT. While the accuracy of extrapolation is less than extending the volume and falls off quickly as the sample get further from the calibrated volume, it does ensure C^0 continuity in the corrected readings. This prevents a jump in the head position or orientation when the user moves out of the calibrated volume. This is helpful for an AR or VR system.

4.4 Testing the Table

To measure the numerical results, we took three independent sets of test points (each with eleven FOB and three Faro readings per test sample). We then corrected the readings from the FOB using the table and compared the result with the truth value, obtained as before using the Faro arm and the transformation loop. The data collection yielded 720 valid points inside the table. The results are graphed in Figure 8.

The orientation error correction does not have the success of the position error correction. This lead us to conduct an experiment challenging our assumption that orientation error depends solely on the measured position (see Section 4.7). We have also noted that of the 720 correctable samples, 23 have greater position error after applying



the correction method, 105 have greater orientation error after “correction,” and 1 (additional) point has both greater position error and greater orientation error after “correction.”

4.5 Tracker performance outdoors


As an additional verification of our methods, we collected samples from the tracker outside to test its accuracy in a metal-poor environment. A wooden structure required to keep the pieces in the closed transformation loop rigid is pictured in Figure 9. The data set we collected consists of 87 points that were less than one meter from the transmitter. The average position error was 9.1 mm, the maximum position error was 17.8 mm, the average orientation error was 2.3 deg, and the maximum orientation error was 4.4 deg. It is interesting to note that the corrected tracker readings in our lab are more accurate than the uncorrected readings taken outside.

4.6 Applications

While numerical results are important for quantifying results and giving insights into the reasons why the method does or does not perform, our goal is to improve registration in an AR system. The position correction we obtained appears to be sufficient for noticeable improvement in an AR system, while orientation correction does not. Recent work in our lab has shown that tracker error in position adds to the registration error by the magnitude of the position error, but that orientation error adds 9 mm of registration error per degree of orientation error at a working distance of 500 mm [Holloway95]. It is therefore difficult to demonstrate improved registration due only to position correction in the presence of significant remaining orientation error. Thus we will use an alternate method of orientation correction to display visual results, so that we can demonstrate the usefulness of the position correction achieved by our method.

Our alternate method is vision-assisted; it detects and tracks two circular landmarks in the video image. Using the 2D video image coordinates as well as the known real-world 3D coordinates of the two landmarks, we correct the orientation report from the tracker in roll, pitch, and yaw of the camera [Bajura95][State96a]. Using this alternate correction method for orientation does not diminish the significance of the position correction; rather it makes the effect of the position correction noticeable. The visual effect is to correct the perspective of the computer-generated elements to more closely match the perspective of the real world.

The visual results are depicted by viewing images captured from the HMD view in our AR application. The HMD is fitted with a camera and a FOB receiver. The transformation between the camera and the receiver is determined using the method described in [Bajura95] (where the transformation is referred to as Head-to-Camera). This procedure uses an iterative refinement of the transformation. The initial estimate is obtained from estimates of the transformation from the tracking sensor to the transmitter, and from the transmitter to a calibration fixture.



Then a series of corrections are applied according to the following heuristics used when trying to align synthetic imagery to the calibration fixture. If the camera is relatively far from the fixture, misregistration is mostly due to rotation error. If the camera is relatively near the fixture, misregistration is mostly due to translation error. Using these observations as the basis for iterative manual adjustments, the synthetic imagery can be visually aligned to the calibration fixture, producing a correct transformation.

For Figure 2a the FOB reported the position and the landmark tracking system corrected the orientation error. Note that the grid is drawn with an error in perspective. This error is due almost entirely to FOB position error. For Figure 2b our LUT method corrected the position readings from the FOB and the landmark tracking again corrected the orientation error. Note the improved registration due to correct, matching perspective, which results from the more accurate position readings.

Such accurate position also improves the performance of hybrid tracking techniques. Accurate position reports from a calibration of the magnetic tracker help the vision-based correction methods to perform useful heuristic corrections in underdetermined cases [State96a].

4.7 Orientation error dependence

To experimentally determine whether orientation error was independent of the measured orientation, we collected samples while holding the receiver in a nearly constant position (within a sphere of radius 6 mm) but at various receiver orientations. To do this required moving the Faro arm (which is attached to the receiver via the plastic rod) into various attitudes surrounding the (fixed) position of the receiver. To ensure that we maintained the receiver in a (nearly) fixed position, we used the apparatus shown in Figure 10.

We ran this test five times. The first two trials were in our lab early in our work with the ERT. The third trial was performed with a standard-range transmitter from Ascension and a different receiver. The fourth trial was performed in our lab with the same hardware as in the first two trials, but only after a factory test and a new factory calibration technique were performed³. The fifth and final trial was performed outdoors with the calibrated ERT. Each trial used a different receiver position with respect to the transmitter.

A plot of the orientation error in degrees versus the azimuth and elevation angles of the receiver reveals a coherent function (Figure 11). The data from the first trial is black; the data from the second is white. This graph does not show the axis around which to rotate in order to correct the orientation. The apparatus that enables us to maintain a fixed position unfortunately prevents us from considering roll. Still, this graph shows that orientation error is a function of measured orientation. Similar graphs result from the other trials.

Thus we have shown that a 6D LUT or a six-parameter function would be needed to perform orientation error

³The reported accuracy with our ERT and receiver was 1.47 mm root-mean-square in position and 0.15 deg root-mean-square in orientation in an environment free of distorting metal or interfering electromagnetic signals.




Figure 9: Picture of support structure used to conduct the experiment outside.

correction. Collecting enough samples to build such a table or compute such a function is a daunting task. Consider taking samples at every 45° at each sample point – a rather coarse sampling frequency. This would multiply the number of samples by 26. (We took 12801 samples to build our current table, a task requiring approximately 14 hours for an experienced user.) Also, the 6D LUT will require even more memory and the correction operations for either method will take longer to perform.

5 Conclusions

The LUT calibration method works well for position error, correcting 78.6% of the position error on average. The position correction achieved by our method thus offers a visible improvement in registration in an AR system. The correct position improves the performance of hybrid tracking methods. However, the method does not work as well for orientation error, correcting only 40.1% of the orientation error on average. This is due to the fact that we assumed a model for the orientation error that did not depend on the measured orientation, but the orientation error does indeed depend on the measured orientation. Were we to use all of the parameters upon which orientation error depends in our correction method, we would expect to see improvement in the performance of the method in correcting our test data.

We know that severe orientation error will result in extreme misregistration; we clearly need a better solution to the orientation error problem in order to significantly improve registration. We can see that extending this method to orientation error will require a 6D LUT, which is not a practical solution.

6 Future Work

A general 6D LUT applicable to all applications is impractical, since we cannot hope to collect enough data at all orientations and the resulting table would be very expensive for many applications to store. However, certain applications require a user to observe an object on a tabletop, implying that only certain orientations are of interest (“outside-looking-in application”). Thus a straightforward extension to this work is to build a table from samples that are taken in these restricted orientations. An option to this idea is to use the alternate orientation correction




Figure 10: Apparatus allowing (limited) freedom for the receiver's orientation while maintaining constant position.




Figure 11: Graph of orientation error as a function of the azimuth and elevation angles of the receiver. The longer axis of the reference grid is the azimuth angle, ranging from -180 degrees to 180 degrees with lines every 20 degrees. The shorter axis is elevation, ranging from -90 degrees to 90 degrees, with lines every 20 degrees. The hole in the graph corresponds to the position of a wooden support structure used.

algorithm of vision-assisted tracking [Bajura95][State96a] to determine the orientation error.

The dependence of orientation error on the true orientation leads us to wonder if there is any dependence of the position error on the true orientation. To verify whether this is true would require an apparatus that is extremely precise in its position while allowing free movement of orientation. We have not performed such a test, but we suspect the dependence is minimal if it exists at all, since we have had good results in correcting position error. We therefore believe the apparatus shown in Figure 10 may not be accurate enough for such a test.

Our method requires that we calibrate the magnetic tracker to the system. This calibration must be done with an uncorrected tracker on the first pass. Clearly, we would like to have precise readings to perform this auxiliary calibration. We built a table using these uncorrected readings. Since this table does offer some improvement, we can use this table to correct these readings used for auxiliary calibration. This should increase the accuracy of our auxiliary calibration, and thus improve the accuracy of the calibration table built. This process has not yielded any improvement, however, and we are investigating the reasons.

We have introduced a number of parameters to the resampling process: maximum noise in position and orientation samples, the radius and σ parameter of the Gaussian resampling kernel, and the minimum amount of data (samples or weight of samples) needed to build a grid point. We have also experimented with adapting the shape of the kernel to ensure that the center of mass of the samples used to create an error factor stored at a grid point is as close to the grid point as possible. (This implies that a kernel can become non-symmetric.) A full exploration of this parameter space might yield some "best" set of parameters.

In this paper, we have dealt only with static registration. In a working system, however, latency is an important issue and must be managed. Data collection would be less taxing and much faster if the collection method accounts for the relative latency between the two trackers. We have yet to see any evidence of "drift" in the calibration of the

tracker that would render the LUT incorrect, as previous authors have noted [Bryson92]. However, we will need to be wary of this as we make use of the LUT we constructed. Finally, we would like to understand why certain samples have greater error after application of our correction method. This could indicate bad data or insufficient data in certain portions of the calibrated volume.

The data we collected during our experiments is available at <http://www.cs.unc.edu/~us/magtrack.html> for those who may wish to experiment with similar methods.

7 Acknowledgements

We extend our gratitude to Ernest Blood, Steve Work, Jack Scully, and Per Hansen from Ascension Technology Corporation for their helpful insights and pointers. We would also like to thank Gary Bishop, Henry Fuchs, Mary Whitton, Vern Chi, David Harrison, Fred Jordan, Kurtis Keller, John Thomas, Peggy Wetzel, Richard Holloway, Michael Bajura, Arthur Gregory, Stephen Wendel, Bill Garrett, Gentaro Hirota, Shankar Krishnan, Anantha Kancherla, Marq Singer, Merlin Hughes, Bill Yakowenko, and Andrew Thall. Funding was provided by ARPA (ISTO DABT 63-93-C-0048). Approved by ARPA for Public Release - Distribution Unlimited.

We thank the anonymous reviewers for their comments and suggestions.

A Additional Note

Just before press time, we received private correspondence from Ascension Technology [Blood96] in which they confirm that the orientation error does indeed depend on the orientation of the sensor. We are continuing to work with Ascension to improve upon the performance of our correction method and have identified potential areas of investigation.

References

- [AscensionFTP] Ascension Technology Corporation. Correcting distorted position and orientation data using the curve fit method, `curvefit.exe`. Ascension Technology anonymous FTP site, <ftp://ftp.std.com/ftp/vendors/Ascension>.
- [Ascension95] Ascension Technology Corporation, Burlington, VT. *The Flock of Birds Installation and Operation Guide*, January 1995.
- [Bajura92] Michael Bajura, Henry Fuchs, and Ryutarou Ohbuchi. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. In *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 203–210, July 1992.
- [Bajura95] Michael Bajura and Ulrich Neumann. Dynamic registration correction in video-based augmented reality systems. *IEEE Computer Graphics and Applications*, 15(5):52–60, Sep 1995.
- [Blood89] Ernest B. Blood. Device for quantitatively measuring the relative position and orientation of two bodies in the presence of metals utilizing direct current magnetic fields. United States Patent 4,945,305, April 1989.
- [Blood96] Ernest B. Blood. Private communication, October 1996.
- [Bryson92] Steve Bryson. Measurement and calibration of static distortion of position data from 3D trackers. *SPIE Stereoscopic Displays and Applications III*, pages 244–255, 1992.
- [Curtis93] W. Dan Curtis, Adam L. Janin, and Karel Zikan. A note on averaging rotations. In *Virtual Reality Annual International Symposium '93*, pages 377–385, 1993.
- [Cruz-Neira93] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 135–142, August 1993.
- [Faro93] Faro Technologies, Incorporated, Lake Mary, FL. *Industrial Metrecom Manual, v2.0*, 1993.
- [Furness86] T. Furness. The super cockpit and its human factors challenges. In *Proceedings of the Human Factors Society*, pages 48–52, 1986.
- [Ghazisaedy95] Morteza Ghazisaedy, David Adamczyk, Daniel J. Sandin, Robert V. Kenyon, and Thomas A. DeFanti. Ultrasonic calibration of a magnetic tracker in a virtual reality space. In *Virtual Reality Annual International Symposium '95*, pages 179–188, March 1995.

- [Holloway95] Richard L. Holloway. *Registration Errors in Augmented Reality Systems*. Ph.D. dissertation, University of North Carolina at Chapel Hill, 1995.
- [Kuipers80] Jack B. Kuipers. SPASYN—an electromagnetic relative position and orientation tracking system. *IEEE Transactions on Instrumentation and Measurement*, IM-29(4):462–466, Dec 1980.
- [State96a] Andrei State, Gentaro Hirota, David T. Chen, William F. Garrett, and Mark A. Livingston. Superior augmented reality registration by integrating landmark tracking and magnetic tracking. *Proceedings of SIGGRAPH 96* (New Orleans, Louisiana, August 4–9, 1996). In *Computer Graphics Proceedings, Annual Conference Series*, 1996, ACM SIGGRAPH, pages 429–438.
- [State96b] Andrei State, Mark A. Livingston, William F. Garrett, Gentaro Hirota, Mary C. Whitton, Etta D. Pisano M.D., and Henry Fuchs. Technologies for augmented reality systems: Realizing ultrasound-guided needle biopsies. *Proceedings of SIGGRAPH 96* (New Orleans, Louisiana, August 4–9, 1996). In *Computer Graphics Proceedings, Annual Conference Series*, 1996, ACM SIGGRAPH, pages 439–446.
- [Shoemake85] Ken Shoemake. Animating rotation with quaternion curves. In *Computer Graphics (SIGGRAPH '85 Proceedings)*, volume 19, pages 245–254, July 1985.
- [Work96] Steve Work. Private communications, July 1995–April 1996.