Chapter 7

Results

7.1 Test Object

To determine the accuracy of our method we performed several experiments using a special test object. A three dimensional object with seven fiducials was made. The relative positions of the fiducials were accurately measured. Figure 7-1 shows the basic test object. The large pillar near the center is used to measure the accuracy of the enhanced reality visualization. A wire frame corresponding to the edges of the pillar is displayed in the enhanced reality image. The difference between the actual edges and wire frame is a measure of the accuracy of the visualization. For comparison purposes a wire frame is also superimposed on a shorter pillar. Experiments were performed with four slightly different fiducial configurations. The first configuration consists of 6 coplanar fiducials plus a single fiducial 1cm above the plane. The enhanced reality visualizations produced from this configuration are not consistently accurate. Figures 7-2 through 7-5 are typical of the results for this configuration. Notice that the wire frame on the smaller pillar matches in all of the images. Errors occur only significantly away from the volume enclosed by the fiducials. The second configuration moves one of the 6 coplanar fiducials so that it is also 1cm above the plane. Figures 7-6 through 7-8 show typical results for this configuration. The accuracy is greatly improved with the addition of a second noncoplanar fiducial, however occasional slight errors do occur. The next configuration adds a third noncoplanar fiducial. Typical results for this configuration are shown in Figures 7-9 and 7-10. With this configuration, errors rarely occur and when they do they are small. The final configuration is similar to the first except that the noncoplanar fiducial is 4cm out of the plane. Results for this configuration are shown in Figures 7-11 and 7-12. The accuracy of this configuration is comparable to that of the last. Some depth in the model is required to accurately recover the third dimension. Once sufficient depth is present in the model, it appears that the solution is accurate over a large volume.

As shown in Figures 7-2 through 7-12, under reasonable conditions our method produces good results. These figures provide only a subjective measure of accuracy. Next we will attempt to provide a more quantitative measure.

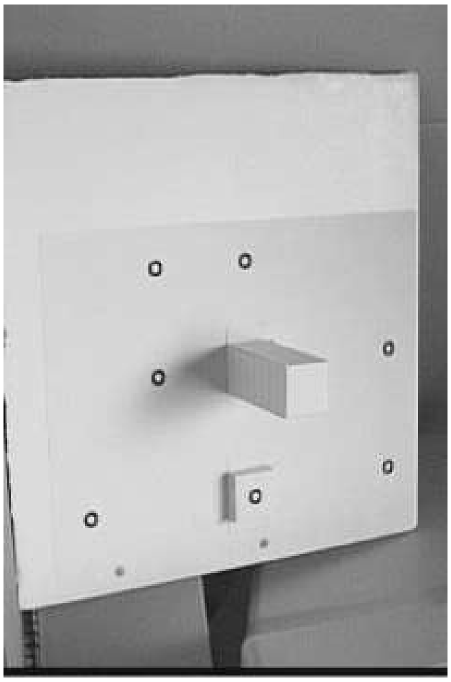

Figure 7-1: Test object.

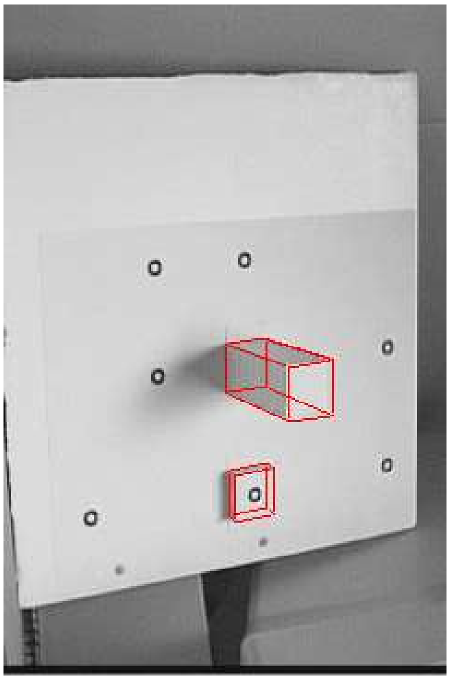

Figure 7-2: Test object view 1.

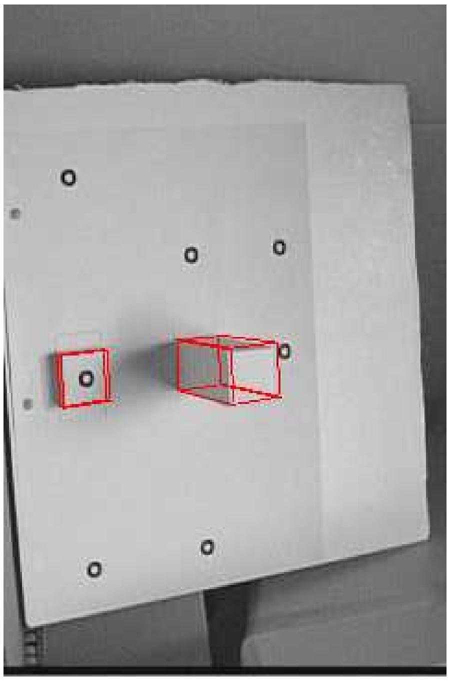

Figure 7-3: Test object view 2.

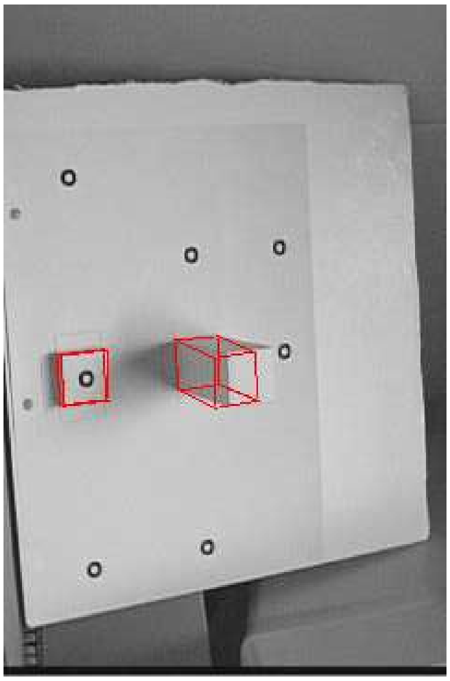

Figure 7-4: Test object view 3.

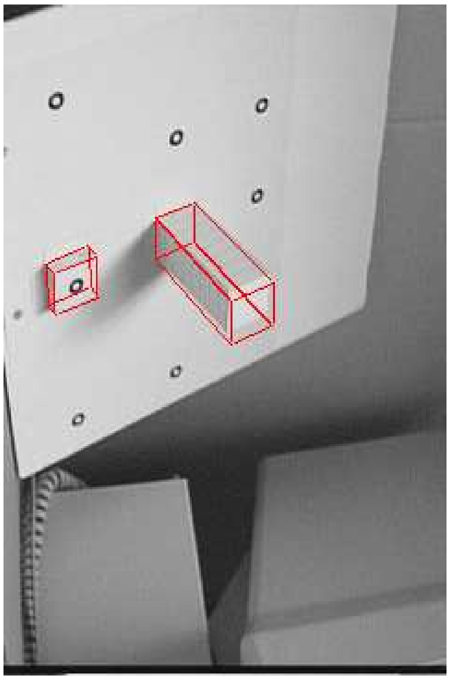

Figure 7-5: Test object view 4.

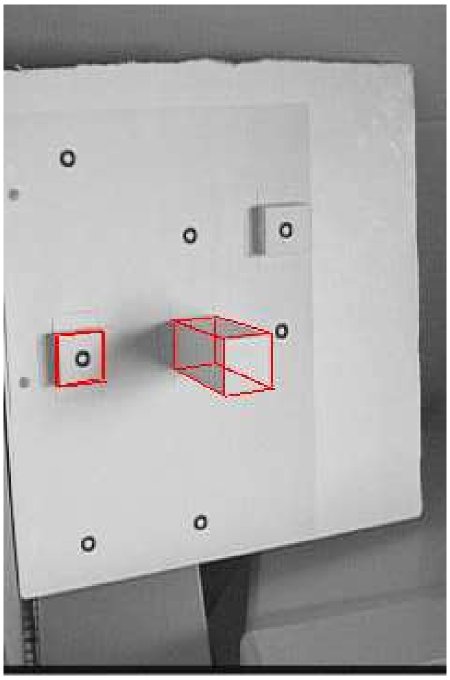

Figure 7-6: Test object view 5.

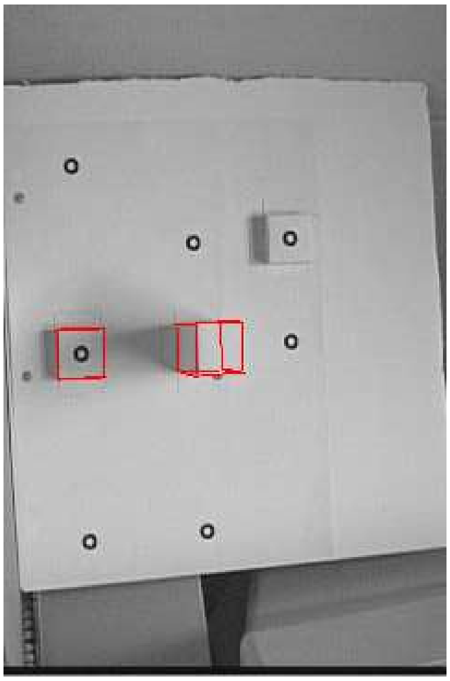

Figure 7-7: Test object view 6.

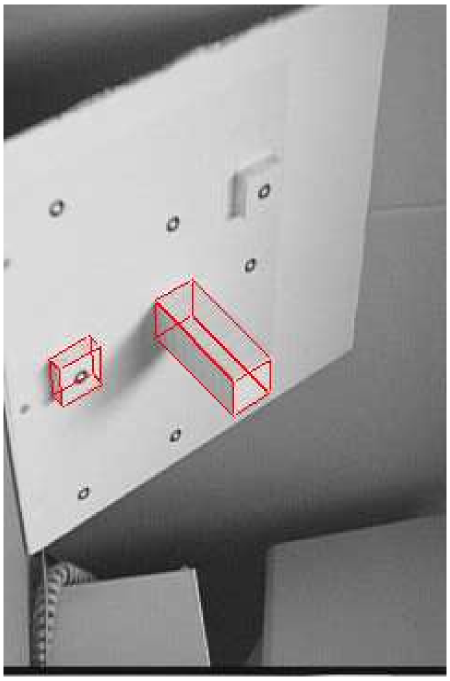

Figure 7-8: Test object view 7.

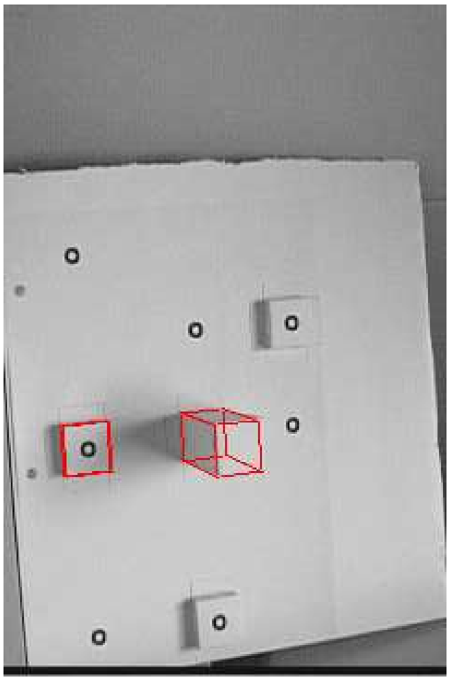

Figure 7-9: Test object view 8.

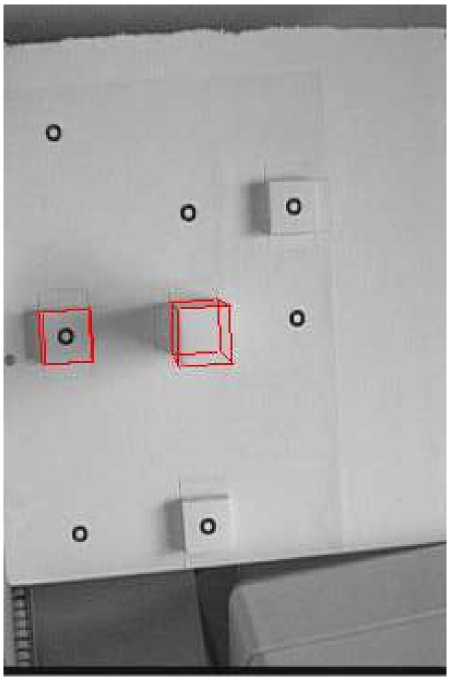

Figure 7-10: Test object view 9.

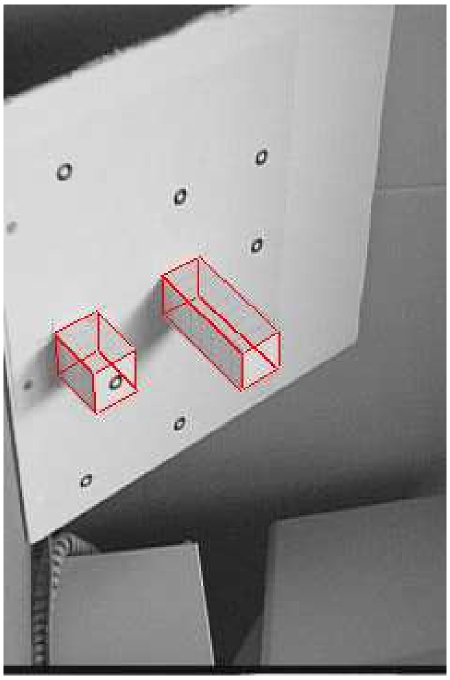

Figure 7-11: Test object view 10.

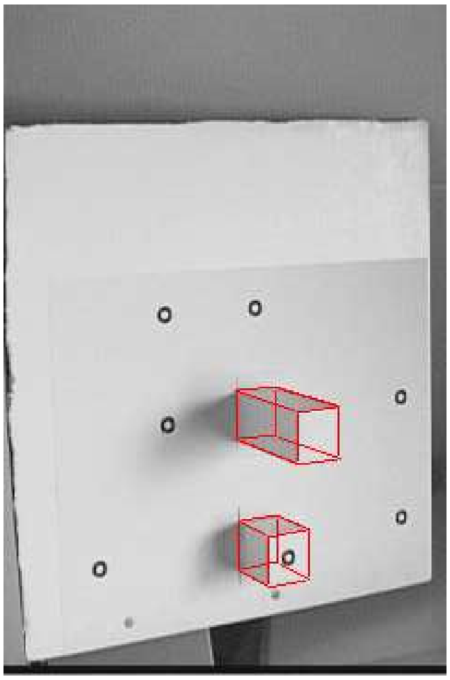

Figure 7-12: Test object view 11.

Average Distance	RMS Distance	Maximum Distance	Median Distance	Number of Points
1.35	1.46	2.94	1.38	848

Table 7.1: Distance between edge-based and enhanced reality vertices.

Average Misalignment	RMS Misalignment	Maximum Misalignment	Median Misalignment	Number of Points
0.88	0.91	1.77	0.87	212

Table 7.2: Misalignment between edge-based and enhanced reality polygons.

Given that the goal of enhanced reality visualization is to produce an *enhanced* image, the proper way to assess accuracy is to consider the deviation of the enhanced image from the *ideal* image. For various reasons we do not have access to the ideal image. However, we can compare the deviation between the image of a physical object and the enhancement produced from a model of the same physical object. For example, we could compare the wire frame to the edges of the test object's central pillar. This is essentially what we will do. Figure 7-13 shows the same test object used above with one change, the top of the central pillar had been darkened to provide high contrast edges. Edge pixels are extracted and chained together using a Canny edge detector. A line is fitted to the chains by minimizing the distance between the edge pixels and the line. These lines are used to compute two measures of accuracy. The first measure is the distance between the vertices of the darkened region as determined by our method (enhanced reality visualization) and those determined by intersecting the lines recovered above. The second measure is the area of misalignment divided by the perimeter or the average distance between the two polygons, see Figure 7-14.

A sequence of 212 images of the test object rotating through 360° was used. The first measure was computed for each of the four vertices in each image. The distance between the vertices in each image are shown in Figures 7-15 through 7-18. The second measure was computed for the darkened polygon in each image and the average distance between the polygons for each image is shown in Figure 7-19. Tables 7.1 and 7.2 summarize these results. The edge-based positions agree well with those produced by our method. It should be noted that edges and vertices recovered using the Canny edge detector are not without error. The most significant source of error is the fact that the implementation used only localizes the edge to the nearest pixel. As a result of

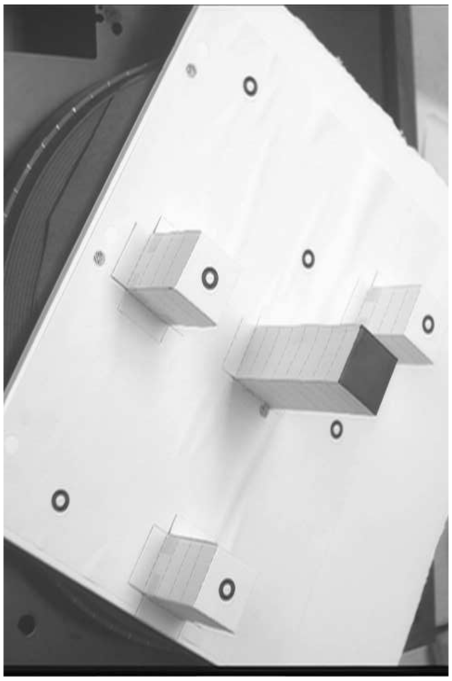

Figure 7-13: Test object used to quantify accuracy.

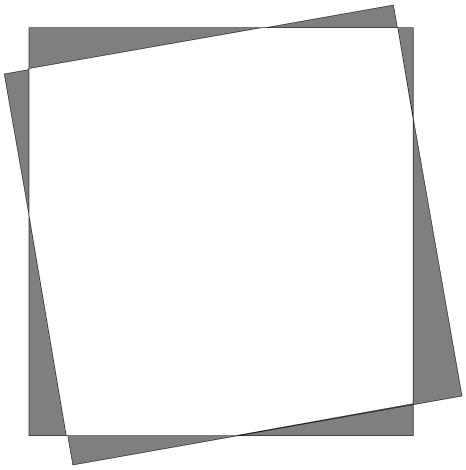

Figure 7-14: Misalignment of edge-based rectangle and enhanced reality rectangle.

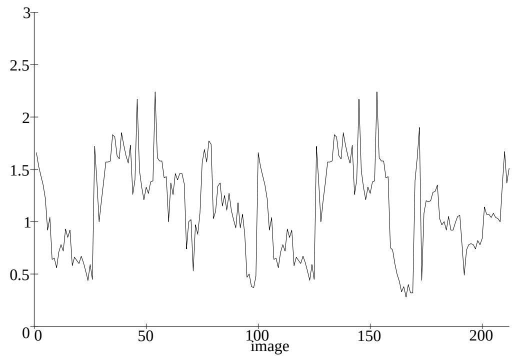

Figure 7-15: Distance in pixels between edge-based and enhanced reality positions for vertex 1.

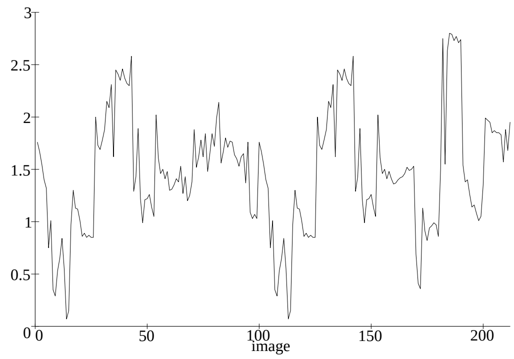

Figure 7-16: Distance in pixels between edge-based and enhanced reality positions for vertex 2.