OVERCOMING AUGMENTED REALITY TRACKING DIFFICULTIES IN CHANGING LIGHTING CONDITIONS

Matthew Clothier
mclothie@ucsd.edu

Mike Bailey
mjb@ucsd.edu

San Diego Supercomputer Center
University of California San Diego, USA
9500 Gilman Drive
La Jolla, CA 92093-0505

Abstract

Augmented Reality (AR) uses a combination of hardware and software to enhance a person's vision of the world with useful information about his/her surroundings. In recent years, there have been many different implementations of Augmented Reality. The vast majority of these implementations have assumed that lighting conditions are static, usually in an indoor or outdoor environment. Unfortunately, many of these implementations have difficulty adapting to new lighting conditions, especially when the conditions change abruptly. This paper addresses this issue by presenting a robust system that is able to tolerate such changes in lighting conditions. Thus, this paper will show how to construct a vision component of an Augmented Reality system that will help adjust the system automatically to both indoor and outdoor lighting conditions.

Keywords: Augmented Reality, High Dynamic Range, Scientific Visualization, Imaging and Image Processing

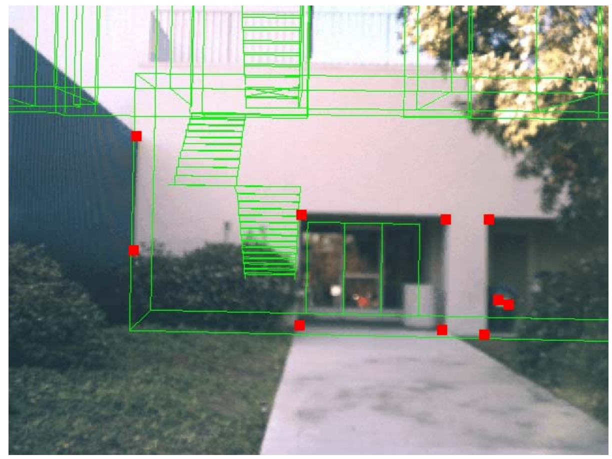

Figure 1: This image shows the feature points that have been dynamically picked for this building. Once the feature points have been found, the structure and its orientation can be found. As a result, useful information such as building schematics can be overlaid on top on the building itself.

1. Introduction

Augmented Reality is rapidly becoming a popular trend in the academic and scientific world. As a result, there has been much research on this topic in recent years. Many of the research papers concerning Augmented Reality have dealt with issues such as implementation, system design, tracking algorithms, user interface design, user interaction, etc. Generally this research has been done in a computer lab, where tracking and communications are easily controlled ([1], [2], [3], [11], [15], [16]). Even when AR research has been in an outdoor environment, it too has been in very controlled environments ([6], [10], [17]). Our research however has focused on the construction of a robust AR system for changing environments. In previous research, we built a system that dynamically picked feature points [14] of articulated structures in order to detect and display information useful about the structure [4]. Thus, the system was able to recognize many different structures regardless of the environment. Results of the system are shown in Figure 1.

Building upon previous research, we discovered that changing lighting conditions had the ability to impair the system from tracking feature points. The system was able to recover after a few seconds by picking new feature points but this could present problems for mission critical applications like fire fighting. In such applications, every second counts and information displayed by the AR system needs to be accurate. Unfortunately, there has not been much research handling rapidly changing lighting conditions. We feel this is an important problem and solving it will greatly aid Augmented Reality systems manage dynamic environments. Thus, this paper demonstrates an AR system that can adjust automatically to changing lighting conditions.

2. Problem: Changing Lighting Conditions

Despite many different variations of Augmented Reality, all implementations have the same basic components. These components include: video/frame input into an AR system, an algorithm/technique that gathers information from selected frames (such as orientation information), and lastly, a means in which to augment the frames with useful data (typically with computer graphics). Considering these three components, the video/frame input is the only one that is dependent on an external source. Usually this external source is from a camera CCD[1] that has taken an image or image sequence from the real world. An image taken from a camera CCD can vary greatly depending on the settings of the camera. This can be difficult to deal with, as there are usually settings to change exposure, brightness, contrast, gamma, etc. Due to the wide variety of images that can be captured, the settings of the camera should be adjusted appropriately depending on the conditions of the environment external to the camera.

Unfortunately, a vast majority of Augmented Reality research has assumed that conditions external to the camera will remain constant. Primarily this has been to make the algorithms easier and to aid the tracking of objects. Vision-based tracking is a crucial part of an AR system as it is used to help gather orientation of objects and various visual queues from the environment. If such tracking data were ever lost, the AR system would be disrupted and might even result in an unrecoverable error. Unfortunately, this situation can occur if the Augmented Reality system does not know how to adjust to the image provided by the camera CCD. As a result, many AR systems only perform as well as the camera CCD adjusts to the lighting conditions of the external environment.

To illustrate further, one of the main ways that an AR system could lose a tracked object would be if the camera CCD is not able to handle a rapid change in light or darkness. Unfortunately, this can be a common occurrence when working with Augmented Reality in the real world. Some examples include: the sun peering in and out of clouds, walking from inside to outside a building (or vice-versa), or even a sudden bright flash from a photographer's camera. Figure 2 shows an example of tracked points that have been lost between two frames when a bright light was introduced into the scene. Thus, it is apparent that this is a major problem that needs to be addressed before Augmented Reality can handle dynamically changing environments.

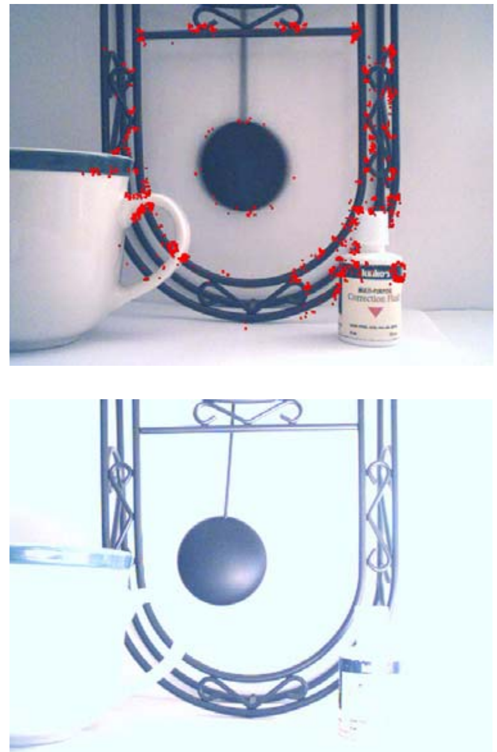

Figure 2: These images are from an image sequence. The top image was taken one frame before the bottom image. Notice how features previously being tracked in the top image are no longer tracked in the bottom image.

3. Idea: High Dynamic Range (HDR)

The primary goal of this paper is to overcome the difficulties that result from a wide range of possible images from a camera CCD. Thus, the first step in overcoming this limitation is to understand that there are many lighting/external conditions that can affect a final image. With this in mind, this paper devises a system that will automatically compensate for these conditions.

In order to help solve the limitations of a camera, an idea is borrowed from the world of high dynamic range (HDR) research. High Dynamic Range imaging is the means of capturing the real world luminosity as it is, without clipping of pixels or bleaching of regions within the image [9]. In recent years, there has been research [5] that has created high dynamic range images using photographs of many different exposures. To illustrate an example, look at Figures 3, 4, and 5 [7]. In Figure 3, the exposure of the image is too high and thus washes out the right side of the image. In Figure 4, the exposure of the image is too low and thus makes the left side of the image black. In order to compensate for these images, a technique was devised that combines the best parts of both images into one, as shown in Figure 5. The features of this image are clear and it is easy to make out all the objects in the scene. Ideally, this is desired result.

[1] CCD is an acronym for Charge-Coupled Device. It is an image sensor that separates the visible spectrum into red, green and blue for digital processing by a camera.

Figure 3: An example image in which the exposure is too high. Notice that the right side of the image is washed out and the features are difficult to see.

Figure 4: An example image in which the exposure is too low. Notice that the left side of the image is all black.

Figure 5: A high dynamic range image using images from Figures 3 and 4. Notice that both of the left and right side of the image are visible and clear.

Unfortunately, the results shown in Figures 3 and 4 are common of all cameras if the exposure is too high or too low. Thus, an Augmented Reality system would have difficulty picking and tracking key feature points in order to do augmentation. However, if a scheme were devised that provided an AR system with a high dynamic range image like the one in Figure 5, this would practically solve this problem entirely. Unfortunately, creating high dynamic range images is processor intensive and it would be difficult to do this in real time. Thus, the following section devises a technique that creates pseudo-high dynamic range images that can be used in real-time, Augmented Reality applications.

4. Setup and Implementation

Before going into the technique that we devised for our Augmented Reality system, let's first look at the hardware and the environment that was used for experimentation. The hardware used were as follows:

1. Dell Inspiron 8200 laptop
2. OrangeMicro iBot2 USB 2.0 web camera

The OrangeMicro iBot2 is able to capture up to 30 frames per second at 640x480. It also has the ability to adjust features such as exposure, brightness, contrast, etc. This camera was used to record the experimental setup. The setup consisted of a scene of a coffee cup, whiteout bottle, and a clock with a moving pendulum. As the camera recorded, various feature points were tracked in the scene (both moving and static). The camera was then exposed to various lighting conditions to test the results. The sample scene was near window blinds so that light adjustment was simple. There were two primary lighting tests. The first test was to start with the blinds closed and then open the blinds quickly thereafter. This would flood the camera with new light coming from outside. The second test was then to quickly drop the blinds so that the sample scene grew dark really fast. Thus, the camera is not able to take in as much light.

One would expect the result of this experiment to be similar to the results of Figure 3 or 4. Indeed, this was the case. When the blinds were opened, the camera was unable to handle all the new light and washed the image out (as was shown in Figure 2). When the blinds were closed, the lack of light produced images that were almost black. The iBot2 camera was able to autocorrect for this, but it took about 5-10 seconds for this to be fully corrected. For Augmented Reality this is not acceptable as important information may not be available at critical times.

It was at this point we realized that high dynamic range research would be useful to solve this problem. The idea is to obtain both a high exposure and a low exposure image back-to-back. If the camera is in auto exposure mode, then we already have either a high or low exposure image. For the next frame, we then set the camera to the opposite exposure to retrieve either the high or low exposure image that is needed. This is considered an exposure pair and is done for the entire frame sequence. The images retrieved from the auto exposure mode will by default act as the "main frame" that is displayed. Since most of the images from auto exposure use a high exposure level, usually the high exposure image becomes the "main frame" while a low exposure image is retrieved to act as the "backup frame." The top images of Figures 6 and 7 show an example high and low exposure pair.

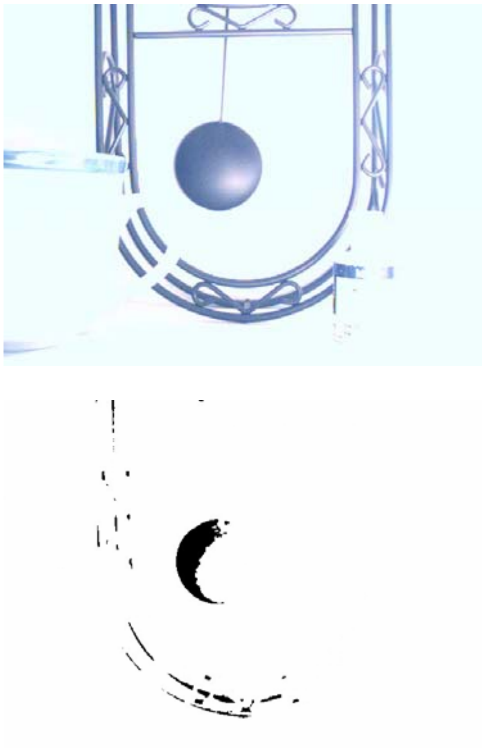

Figure 6: The top image is a high exposure image when new light is introduced into the scene. The bottom image is its resulting threshold at this point. Note that most of the features in the threshold image have disappeared.

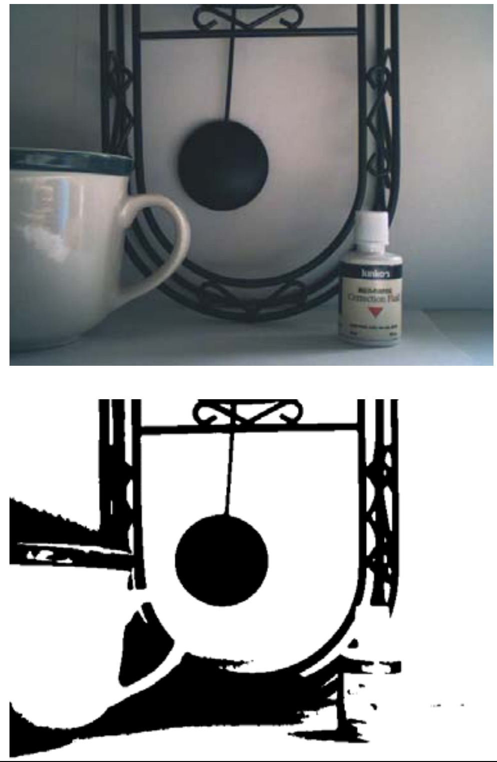

Figure 7: The top image is a low exposure image when new light is introduced into the scene. The bottom image is its resulting threshold at this point. Note that the features in the threshold image are easier to see.

For our experimental setup, we decided to track objects using KLT tracking [14] however any tracking method could have been used. All key features in the scene, such as from the pendulum movement of the clock, were tracked over time. Through experimentation, it was found that if the lighting of the room changed so that the camera was flooded with light, the feature points that were being tracked would be completely lost (as Figure 2 demonstrated). In order to compensate for this, we track the features of the threshold image rather than features of the actual image itself (see the bottom image of Figures 6 and 7 for the threshold and Figure 10 for a tracked threshold). The primary advantage to this is that the threshold of either the low or high exposure image will generate features to track. As it turns out, usually only one of the thresholds is good enough to do tracking rather than both at the same time. Fortunately, this is not a big loss because when there is a rapid change in light, the threshold currently being tracked will switch over to the threshold of the other exposure. In order to switch threshold images, a histogram is taken of the image frame and then analyzed. If the histogram of the image favors one side of the spectrum instead of the whole spectrum, then it is time to switch to the other exposure. This can be seen in Figures 8 and 9 on the following page. Figure 8 is the histogram of Figure 6 and Figure 9 is the histogram of Figure 7. At the moment these images were taken, a new, bright light was introduced into the scene. As a result, Figure 6 looks washed out while Figure 7 provides a good image of the scene. Looking at their histograms, Figure 8 favors the white part of the spectrum because the image is washed out. Figure 9 on the other hand covers a wider range of the spectrum. In this example, it is clear that the algorithm would choose to use the histogram from Figure 9 in order to track feature points. Thus, by analyzing the distribution of the histogram, we can determine which image's threshold is the better one to track.

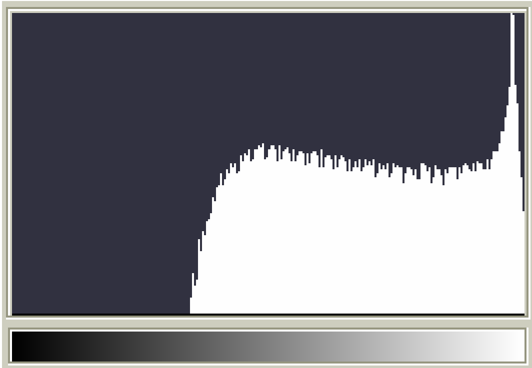

Figure 8: This is the histogram of the high exposure image in Figure 6. Notice that the histogram favors the white part of the spectrum.

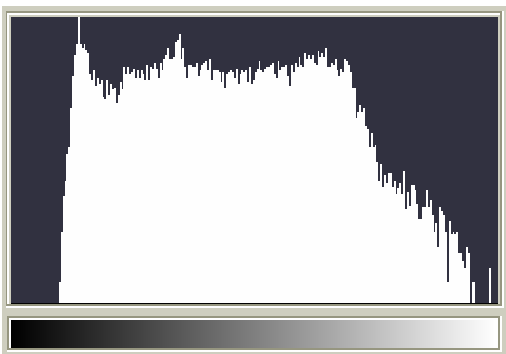

Figure 9: This is the histogram of the low exposure image in Figure 7. Notice that the histogram is well distributed over the spectrum.

Through the use of selecting which threshold to track, we essentially are building a dynamic threshold over time. In essence, this produces a pseudo-high dynamic range in which both high and low exposure images are taken into account. This threshold seamlessly transitions from one exposure to another allowing features to be tracked continuously despite changing lighting conditions. Figure 10 shows an example of a frame from a tracked threshold sequence. At this point, it is easy to take the feature points from the threshold and place them onto the "main frame," as shown in Figure 11. Notice that the features are still visible despite the change in lighting. It is worth noting that this frame is the same frame from Figure 2 that failed to track the feature points before.

The results from this technique are very promising. As a result, the system we have implemented is now able to adjust itself automatically to changing lighting conditions without user interaction. Thankfully, this will help to prevent an Augmented Reality system from failing at critical times and allow objects to be continuously tracked. We feel that this is a step in the right direction for dealing with tough lighting situations and improving the overall robustness of Augmented Reality.

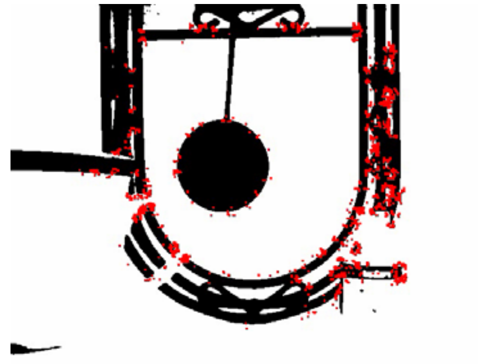

Figure 10: This image shows features being tracked from an image's threshold. Note that this threshold was chosen to be the better threshold to track.

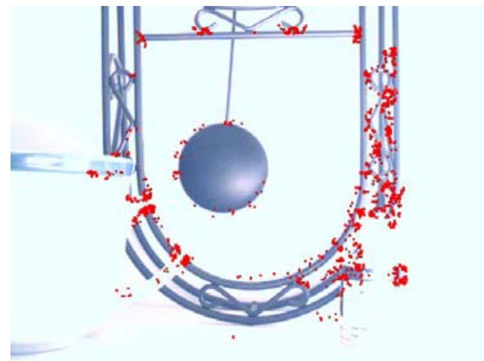

Figure 11: This image is the final result of the proposed method. Note that the scene has been flooded with light but the features are still detected.

5. Future Work

Future work of this research includes determining how to automatically calibrate the exposure of a novel camera and exploring speed improvements for this algorithm. In addition, research will be conducted to see how this technique can be used in a multi-camera environment. Lastly, research investigating real-time high dynamic range images would also be of great benefit.

6. Conclusions

This paper presented a technique that can be used to help an Augmented Reality system become more robust and adjust to changing environments. In order to compensate in dynamic lighting conditions, high and low exposure images are used to create a dynamic threshold that is used for tracking. This allows objects to be continually tracked regardless of the exposure of the image. As a result, this technique will aid Augmented Reality systems to handle many different environments without fear of error.

Acknowledgments

Thanks goes out to the National Science Foundation for supporting this project under grant EIA-9809224. Also thanks to the California Institute for Telecommunications and Information Technology for summer support. Thanks also to Harry Ammons at the San Diego Supercomputer Center for his feedback and experience dealing with camera parameters. Finally, thanks to Dr. Serge Belongie for his encouragement, direction, and guidance on this project.

References

[1] Ronald Azuma, Jong Weon Lee, Bolan Jiang, Jun Park, Suya You, and Ulrich Neumann, “Tracking in unprepared environments for augmented reality systems”, Computers & Graphics, Vol 23, No 6 December 1999, pp. 787-793.

[2] Oliver Bimber, L. Miguel Encarnaçáo, and Dieter Schmalstieg, “Augmented Reality with Back-Projection Systems using Transflective Surfaces”, Computer Graphics Forum, Vol 19, No 3, August 2000.

[3] David Breen, Ross T. Whitaker, Eric Rose, and Mihran Tuceryan, “Interactive Occlusion, Automatic Object Placement for Augmented Reality”, Computer Graphics Forum, Vol 15, No 3, August 1996, pp. 11-22.

[4] Matthew Clothier and Mike Bailey, “Augmented Reality for Viewing Articulated Structures”, IASTED Computer Graphics and Imaging 2003, Hawaii, USA, August 2003.

[5] Raanan Fattal, Dani Lischinski, and Michael Werman.. (2002). Gradient Domain High Dynamic Range Compression. Proc. ACM SIGGRAPH 2001.

[6] Eric Foxlin, Michael Harrington, George Pfeifer, “Constellation: A Wide-Range Wireless Motion-Tracking System for Augmented Reality”, Proceedings of SIGGRAPH 98, pp. 371-378.

[7] http://www.cs.huji.ac.il/~danix/hdr/panorama.html

[8] http://www.hitl.washington.edu/research/shared_space

[9] http://www.spheron.com/products/SpheroCamHDR/spherocam_hdr.html

[10] Tobias Höllerer, Steven Feiner, Tachio Terauchi, Gus Rashid, and Drexel Hallaway, “Exploring MARS: developing indoor, outdoor user interfaces to a mobile augmented reality system”, Computers & Graphics, Vol 23, No 6, December 1999, pp. 779-785.

[11] D. Koller, G. Klinker, E. Rose, D. Breen, R. Whitaker, and M. Tuceryan, “Real-time Vision-Based Camera Tracking for Augmented Reality Applications”, ACM Symposium on Virtual Reality Software, September 1997.

[12] K. C. Ng, H. Ishiguro, and M. M. Trivedi, "Multiple Omni-Directional Vision Sensors (ODVS) based Visual Modeling Approach," IEEE Visualization '99 , San Francisco, California, Oct 1999.

[13] Ulrich Neumann and Youngkwan Cho. (1996). A Self-Tracking Augmented Reality System. Proc. of ACM VRST 96.

[14] Carlo Tomasi and Takeo Kanade, “Detection and tracking of feature points”, Carnegie Mellon University Technical Report CMU-CS-91-132, Pittsburgh, PA, 1991.

[15] http://www.sae.org/ohmag/techinnovations_02-00/12.htm

[16] Andrei State, Mark A. Livingston, Gentaro Hirota, William F. Garrett, Mary C. Whitton, and Henry Fuchs, “Technologies for Augmented-Reality Systems: Realizing Ultrasound-Guided Needle Biopsies”, Proceedings of SIGGRAPH 96, pp. 439-446.

[17] Valerie A. Summers, Kellogg S. Booth, Tom Calvert, and Christine L. MacKenzie, “Calibration for Augmented Reality Experimental Testbeds”, 1999 ACM Symposium on Interactive 3D Graphics, April

[18] Suya You, Ulrich Neumann, and Ronald Azuma, “Orientation Tracking for Outdoor Augmented Reality Registration”, IEEE Computer Graphics & Applications, Vol 19, No 6, Nov-Dec 1999, pp. 36-42.Experimental Testbeds”, *1999 ACM Symposium on Interactive 3D Graphics*, April