Online Real-Time Presentation of Virtual Experiences for External Viewers

Kevin Ponto,

University of Wisconsin-Madison, Madison, WI 53715

Hyun Joon Shin^{*},

Ajou University, Suwon, Korea

Joe Kohlmann, and

University of Wisconsin-Madison, Madison, WI 53715

Michael Gleicher

University of Wisconsin-Madison, Madison, WI 53715

Kevin Ponto: kbponto@wisc.edu; Hyun Joon Shin: joony@ajou.ac.kr; Joe Kohlmann: jkohlmann@wisc.edu; Michael Gleicher: gleicher@cs.wisc.edu

Abstract

Externally observing the experience of a participant in a virtual environment is generally accomplished by viewing an egocentric perspective. Monitoring this view can often be difficult for others to watch due to unwanted camera motions that appear unnatural and unmotivated. We present a novel method for reducing the unnaturalness of these camera motions by minimizing camera movement while maintaining the context of the participant's observations. For each time-step, we compare the parts of the scene viewed by the virtual participant to the parts of the scene viewed by the camera. Based on the similarity of these two viewpoints we next determine how the camera should be adjusted. We present two means of adjustment, one which continuously adjusts the camera and a second which attempts to stop camera movement when possible. Empirical evaluation shows that our method can produce paths that have substantially shorter travel distances, are easier to watch and maintain the original observations of the participant's virtual experience.

Keywords

Virtual Reality; Viewpoint Similarity; Camera Motion; Stabilization; Observation

1. INTRODUCTION

Monitoring the experiences of a user in a virtual environment is generally useful for spectators, scientists, architects, and designers. While some virtual reality hardware allow external viewers to see both the individual and the projection, enclosed environments such as a C6 CAVE or a HMD provide little means to ascertain insight into the participant's virtual experience. In these cases, the traditional means in which viewers can gain an

^{*}Work performed as a visiting professor at the University of Wisconsin-Madison.

Publisher's Disclaimer: Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

understanding of the user's experience is through the replication of the user's egocentric perspective. Unfortunately, this egocentric view is often difficult to watch as it is filled with movements that feel unnatural and unmotivated to an outside viewer [16].

In this paper, we introduce methods for creating camera views for external viewers based on the experience of a participant in a virtual environment. The methods operate online¹ in real-time², creating a camera path that meets two goals: conveying what the participant is seeing and providing camera movement that is easy to watch for external viewers. These two goals are often in conflict as the participant's control can lead to movements that they can anticipate (since they control them), but will appear as unnatural, unmotivated, and jerky to an outside viewer. Therefore, methods for creating views for external viewers must balance between being faithful to the participant's path, which fully conveys what they are seeing but may be difficult to watch, and simplifying the camera path to create an easy to watch video, at the expense of not conveying the participant's full view information.

Our approach to external view synthesis balances fidelity (conveying what the participant sees) with watchability by using a contentdependent metric that limits the amount of motion of the camera based on how similar what it sees is with what the participant's camera sees. This effectively removes camera motions that are not necessary to show what the participant is seeing. Unlike a pure filter-based approach, our method will produce sharp camera motions if they are necessary for conveying content. The method operates on a per-frame basis, moving the camera towards the participant's view at each time step. The amount that the camera is allowed to move in each step is controlled by the view content similarity metric.

The contribution of this paper is a method for creating a watchable external view that conveys the experience of a virtual environment participant, online and in real-time. While the problem of effective external view creation has been considered in prior work, our approach is the first that operates online in real-time, and thus can serve important applications where observation occurs during the participant's experience. Our method builds on a prior view-dependent metric to provide an adaptivity that would not be possible with simple filtering. We demonstrate the effectiveness of our approach through quantitative assessment and through a crowd-sourced user study. Our study shows that our method is able to produce paths that are easier to watch while maintaining the content of the original participant's observations.

2. RELATED WORK

Our work is inspired by recent results in video stabilization that show that optimization can explicitly balance between motion smoothness and faithfulness to the original camera [8, 9]. Unlike our work, these approaches use off-line optimizations and focus on challenges unique to 2D video stabilization.

Techniques to stabilize imagery in an online real-time fashion have been used in a variety of fields. Hansen et al. developed methods to stabilize satellite imagery in order to generate a larger photo mosaic [10]. Many researchers have used filtering techniques to stabilize video sequences in real-time [2, 6, 18, 21]. This type of filtering has become a particular focus for researchers dealing with mobile platforms [12, 15]. Others have used real-time stabilization as a means of treatment. For example, Pothier et al. used real-time image stabilization and augmented reality eyewear in the treatment of Oscillopsia for patients with bilateral

¹We define an online algorithm as an algorithm which must process input sequentially without knowledge of future inputs.

²We define real-time as an algorithm that runs at interactive rates.

vestibular loss [17]. Unfortunately, these types of image stabilization techniques generally do not handle the high-frequency movements of egocentric data gracefully [16].

Sensor fusion along with Kalman filtering has proven to be extremely useful to help stabilize the view of the participant in a virtual environment [7, 24]. These techniques are now common for commercial grade tracking equipment [14]. This type of filtering is unfortunately incredibly dependent on the conditions of the environment [6] and may therefore require parameters to be tuned for every different environment in order to be effective.

Researchers have studied changing the virtual perspective as a means to present different views of a virtual environment. Salamin et al. studied whether a participant inside of a Virtual and Augmented Reality situation would prefer to see themselves from a 3rd or 1st person perspective [19]. The authors found that the preferred perspective was very dependent on the task the participant was asked to accomplish. Yang and Olson tested whether 3rd party commanders were able to direct more efficiently with either ego- or exocentric viewpoints [23]. Their results also showed very mixed results, with ego-, exo-, or tailing cameras performing better depending on the situation. As the viewpoints were not being generated from a tracking device, no stabilization was needed.

A variety of methods to automatically position virtual cameras have been developed based on different types of constraints. Christie et al. provide a survey paper outlining these various techniques [3]. The major motivation of our work comes from Ponto et al. who were able to show that egocentric viewpoints could be processed in order to create replay videos that were more effective at communicating participants' experiences in virtual environments for outside observers [16]. The authors' method decoupled the egocentric viewpoint from the replay camera path in order to minimize the camera movement and more effectively place the camera. Salient camera positions were determined using clustering by comparing viewpoints using a content dependent Viewpoint Similarity metric [16]. The authors demonstrated that their methods were able to outperform simple filtering techniques in terms of their ability to convey content. Due to the offline nature of their algorithms, the techniques could not be implemented for an online application where future data are unavailable. We use the authors' similarity metric and presentation style as motivation for the development of our own online real-time algorithms.

3. OUR METHOD

Our basic idea is to produce a novel camera path in an online real-time manner by adjusting the camera only when necessary to maintain the content of the participant's observations. To do this, our method consists of three components as shown in Figure 2. First we determine the similarity between the viewpoints of the participant and viewer. From this similarity, we determine how much viewer's viewpoint should be adjusted. Finally camera is repositioned by interpolating its previous position with the current participant's viewpoint.

3.1 Viewpoint Similarity

We chose to use the Viewpoint Similarity metric, described by Ponto et al. [16], as it provide a means to rapidly compare the content seen by two different viewpoints using a flashlight analogy. The method is based on a variant of shadow mapping [22], thus providing an efficient means for implementation on the GPU.

The method first computes the visibility of one viewpoint (B) in the context of the other viewpoint (A) as shown in Equation (1).

$$V(A, B) = \sum_{p=0}^N \frac{L(A_p, B_p)}{N}, \quad (1)$$

where $L(A_p, B_p)$ is a visibility function for a pixel, p of B in A and N is total number of pixels in the view. This function can be easily implemented on the GPU using occlusion queries [16]. As $V(A, B)$ is not equivalent to $V(B, A)$, the Viewpoint Similarity metric between A and B is defined as a weighted sum of the two viewpoints' relative visibilities:

$$S(A, B) = w_A V(A, B) + w_B V(B, A). \quad (2)$$

We use this computation to determine similarity between the viewpoints of the participant and the viewer. From Equation (2), we set A to the participant's viewpoint and B to the current state of viewer's viewpoint, with the weights according to the scheme shown in [16]. This similarity is used to adjust the camera position as shown below.

3.2 Adjustment Function

After computing the Viewpoint Similarity, we create an adjustment function which indicates how much the camera's viewpoint should be altered. If the two viewpoints are highly similar, we can assume that camera needs very little adjustment. Conversely, if the two viewpoints are highly dissimilar, the camera viewpoint will need to be greatly repositioned. For our purposes, we have created two adjustment profiles, the Continuous Adjustment Profile (CAP) and the Stopping Adjustment Profile (SAP).

The CAP is designed to prioritize maintaining a camera position that is able to convey the content that the participant seeing over easing the camera movement. For the CAP, we define u , the interpolation parameter as

$$u = k(1 - S(A, B)). \quad (3)$$

For this equation, the variable k defines a sensitivity parameter, which is adjustable by the participant. This value determines the maximum step size for each time step, meaning if two viewpoints are entirely dissimilar, this would be the max amount the camera should be adjusted. In practice we found the inverse of the rendering frame-rate to be a useful default value.

While this first function constantly adjusted the camera position to increase the view similarity, others have shown that is useful to stop the camera movement when possible to help external viewers understand the virtual experience more effectively [8, 16]. Therefore, we introduce another profile, SAP, which prioritizes generating easy to watch camera movement over maintaining the precise view that the participant is seeing. For SAP, the interpolation parameter u , is defined as:

$$u = nk(1 - S(A, B))^n. \quad (4)$$

As the similarity value is normalized between zero and one, raising it to a large power will assure that highly dissimilar values will still adjust the camera, while marginally dissimilar values will not. As this function decreases camera movement, we also multiply the function by the exponent to ensure proper sensitivity.

In practice, both profiles have their advantages and disadvantages. In Section 4.2 we evaluate both methods to inform the choice between them and the selection of their parameters.

3.3 Camera Interpolation

After determining the adjustment amount, we move the camera by interpolating between its previous position and the participant's viewpoint.

We adopt exponential maps for performing these interpolations as they rely on directly manipulating the viewing transforms, as opposed to interpolating the various camera components independently. Exponential maps accomplish this by performing linear interpolation operations in the logarithm space of the matrices [1]. We chose to use exponential maps as opposed to parameter interpolation as, in common cases, they do a better job of keeping points of interest centered in view, as shown by Hawkins and Grimm [11]. It is worth noting that as the interpolated amounts are often quite small, the difference between parameter interpolation and exponential maps is likely negligible. After interpolation we remove camera roll to produce views that are level in the horizontal axis, following the practices of Ponto et al. [16].

Equation (5) demonstrates how $A(t)$, the next position of viewer's viewpoint, is calculated given the viewpoint of the viewer ($A(t - \Delta t)$) updated previously, participant's viewpoint ($B(t)$), and the adjustment parameter (u) determined from Section 3.2.

$$A(t) = e^{(1-u)\log A(t-\Delta t) + u\log B(t)}. \quad (5)$$

While this method does not have any explicit easing parameters as were shown in [8] and [16], some smoothing will happen naturally assuming that the viewpoints do not change too rapidly. For instance, as the participant slowly changes their view, the Viewpoint Similarity will also decrease, thereby increasing the interpolated step-size. Conversely, as the viewpoints become more similar, the Viewpoint Similarity metric will also increase, thus decreasing the interpolated step-size.

3.4 Presentation

As shown by Ponto et al., manipulating a camera path may reduce an external viewer's understanding of the participant's experience [16]. To reduce this artifact, the authors drew a rectangular light representing the participant's viewpoint as shown in Figure 3-1 to provide the viewer with a clear representation of where the participant is looking. However, this rectangular light with a sharp cutoff was shown to be distracting when the light and the camera were moving at the same time. Furthermore, the field of view represented by the rectangle did not precisely match that of the participant's perspective. We instead created a virtual light with a gradual fade-off to provide probabilistic information as to where the participant was looking, as shown in Figure 3-2. This proved to also be less distracting for the viewer when the camera was in motion.

4. RESULTS

The goal of our method was to produce new camera paths in an online real-time fashion that reduce the camera movement, thus making them easier to watch without obscuring content from the raw camera path. In order to determine if our methods achieved these objectives, we first generated camera paths by asking participants to walk around in a fully enclosed C6 CAVE and find a particular type of objects within one of several virtual environments. This

task was asked only as a means to generate paths and the performance of the participant to find these objects in the virtual environment was not evaluated.

The eight scenes in which participants were asked to locate objects are shown in Figure 4. The scenes could be grouped into four classes: small scale walk-around, small interior environments, large interior environments, and larger outdoor spaces. In the first group, the participants were tasked with counting apples (Figure 4-1) and birds in a dense bush (Figure 4-2). The small interior environments included an office with coffee mugs (Figure 4-3) and a living room with books (Figure 4-4). In the two larger interior spaces, the participant was tasked with locating toys in a small restaurant (Figure 4-5) and cans in a store (Figure 4-6). For the last group, we created a park scene with light posts (Figure 4-7) and an old western town with barrels (Figure 4-8).

During these tasks, participants' head positions and orientations were recorded for further analysis. To assess the relative performance of our method to other approaches, we considered four different on-line methods.

The first path consisted of the raw egocentric viewpoint, which is used for external viewers in most of practical implementations. For the second path, we adopted a Kalman filter with constant velocity models for camera position and orientation [20]. The filter parameters including measurement errors and process errors, were set to be enough to suppress high-frequency components in raw camera movements while preserving its important features without significant delay. The third path was constructed using the CAP method (Equation (3)) with k set at 0.016. The fourth path was constructed using SAP method (Equation (4)) with k set at 0.016 and n set at 64. Each of the four methods was applied to the traces of from all of the scenes.

4.1 Performance

While the Kalman filter requires very little computation overhead to process, the CAP and SAP methods due incur some computational penalties. To compute the Viewpoint Similarity, the scene must be rendered from both the camera and the participant positions, thus requiring a second render of the geometric models. Secondly, the Viewpoint Similarity metric requires a screen-space computation to be completed for each comparison (Equation (2)). Finally, the flashlight overlay (Figure 3) requires one final screenspace pass. We rendered each of the eight scenes into a 960×540 sized window on a computer equipped with an Intel Xeon 2.67 GHz CPU, 24 GB RAM, and an Nvidia Quadro 5000 GPU. We were able to achieve a rendering rate greater than 60 Hz for both the CAP and SAP method.

While the end goal of our method is to produce camera paths that are easier to watch, the total path length traveled by the camera provides an accessible, quantitative assessment of camera motion. Videographers are generally encouraged to use camera motion sparingly [8] and it has been shown that simpler and shorter camera paths are desirable because they are easier to watch [16].

To test whether the CAP and SAP methods did reduce the camera movement, we monitored the change in camera position at each time step for each of the methods. Figure 5 shows an example velocity profile demonstrating that both the CAP and SAP methods tend to have a lesser velocity compared to Raw or Kalman filtered camera paths.

Table 1 shows the distance the camera traveled for each of the methods, as well as its ratio compared to the original camera path. As shown, the Kalman filter does not affect the amount the camera travels compared to the Raw path on average for the listed experiences. This makes sense as the Kalman filter is targeted to remove the high-frequency components

of the camera movement and thus can overshoot the targeted camera motions. The CAP and SAP methods both reduce the distance the camera traveled in all eight scenes, with the SAP method able to reduce the camera motion by over 50% for some scenes. Further discussion of this result will be given in Section 5.

4.2 Study

We designed a study to test whether the CAP and SAP methods were able to create a better viewing experience without obscuring information. Our hypothesis was that viewers should be able to perform the same counting tasks asked of the participant in the virtual environment based solely the visual images being shown to them. We hypothesized that our method could improve the viewing experience for our study subjects, allowing them to complete the counting tasks more easily and with more confidence, and without reducing their performance compared to viewing the raw camera path.

We ran the experiment through Amazon Mechanical Turk, following the practices described by Kittur et al. [13] and Downs et al. [5]. Specifically, we chose a standard rate of pay and created questions to verify that the participant was human, had the technical ability to participate, and was actively engaged. Subjects were compensated based on an estimate that the experiment would require approximately 15 minutes of their time.

After completing questions about demographic information, the subjects were shown eight videos corresponding to a randomly selected scene and method. The order of the videos and methods was randomly selected so that each subject would see each scene once and each method twice. The study subjects were asked to count the same objects that the participant was asked to count inside of the virtual environment. After watching each video, the participant was asked to input their counted value and input a prominent floating word from the video as an engagement check. The participant was then asked a series of questions about the camera movement, the difficulty of the task, and their confidence in their answers. Each question was answered on a five point Likert scale. The study attracted 24 subjects and all subject's data were used.

We found that subjects were generally able to count the objects accurately, with subjects able to produce the correct result 57% of the time. We found no significant difference in the error of the object count between the methods. We also found no significance between the error count and the scene being shown. The result confirms our hypothesis that subjects can complete the same counting task and that the CAP and SAP methods did not obscure the content of the original observations.

Figure 6 shows graphs for each of the subjective questions asked in the study. We note that while these questions used a five-point Likert scale, research has shown Likert scales to be an approximate of a continuum [4]. As shown, the CAP and SAP methods were able to perform well in all of the categories listed.

We found that subjects rated the calmness of the camera movement to be statistically significant, $F(3, 188) = 18.105$, with Tukey post-hoc analysis showing the CAP and SAP methods to be significantly better than the Kalman and Raw methods. Subjects' ratings of the smoothness of the camera movement also proved to be statistically significant, $F(3, 188) = 19.385$, with Tukey post-hoc analysis showing the CAP and SAP methods to be significantly better than the Raw method and the CAP method to be significantly better than the Kalman method. Subjects' ratings of the motivation of camera movement proved to be statistically significant, $F(3, 187) = 4.4988$, with Tukey post-hoc analysis showing the CAP and SAP methods to be significantly better than the Raw method.

We also found that subjects rated the ease with which they could accomplish the counting task to be statistically significant, $F(3,188) = 5.9002$, with Tukey post-hoc analysis showing the CAP and SAP methods to be significantly better than the Raw method. Finally, we found subjects' ratings of the ease to watch the video were statistically significant, $F(3,188) = 7.1315$, with Tukey post-hoc analysis showing the CAP and SAP methods to be significantly better than the Raw method. We did not find a significant difference for subjects' confidence for answering the questions about the number of objects in a scene. These results confirmed our hypothesis that viewers would find the CAP and SAP methods easier to watch and find the task easier to accomplish. While we did see a higher confidence for answering the questions with the CAP and SAP methods, the result was not statistically significant.

5. DISCUSSION

From the results of our study, we feel that both the CAP and SAP methods are superior to the standard method of presenting views using raw egocentric camera motion. The SAP method slightly out-performed the CAP method in most metrics with the exception of the smoothness test, which is to be expected. From this result it would be interesting to study a variety of settings for the exponential parameter, n , in the SAP method.

The results of our study show that the CAP and SAP methods perform equally well in comparison to the Raw path in terms the counting of objects, thus indicating that our method does not obscure the content of the original path. While it was outside of the scope of this paper, it would be interesting to determine if any method would be able to outperform the original observations. Although we did not see any significant difference between the methods for the correctness of the counting task, we think it is worth mentioning that the variance in the error of the normalized object count is smaller with the SAP method (0.010) compared to the Raw, Kalman and CAP methods (0.032, 0.026, and 0.020 respectively).

As shown in Figure 6, the subjects were not significantly more confident in their answers with the path from either CAP or SAP method than the others. However, we found no correlation between the subject's confidence in their answer and the correctness, suggesting that self-reported confidence is not a good predictor of performance for this task.

Another interesting finding is that geometric smoothness does not always coincide with perceptual smoothness. While the velocity profile of the Kalman filter appears to be smoother than either the CAP or SAP methods, the subjects rated the camera movements of both CAP and SAP methods to be more calm and/or smooth.

The fact that CAP and SAP methods reduced the distance the camera traveled is of interest when coupled with the fact that the CAP and SAP methods did not cause a reduction in observed content. In this regard, many of the camera motions in the Raw path can be viewed as unnecessary in terms of communicating content and can be removed as was done in the CAP and SAP methods. Furthermore, the CAP and SAP paths produced significantly higher ratings of the camera movement when compared to the Raw path, matching the results of Ponto et al. which found that simplified paths were generally easier to watch [16].

While this distance metric only accounts for changes in position, a variety of other distance metrics could take orientational changes in the camera as well. While future work may analyze different methods to quantify camera movement, these new distance metrics would only provide a new source for comparison and would not affect the resulting paths from the CAP or SAP methods as they used screen space constraints.

It should also be noted that the experience for the participant in the virtual environment and the subject viewing the paths on a 2D monitor were quite different. Our method aims to

present the experience of a participant in a highly immersive environment to an external observer with substantially less immersive capabilities. Future work will explore how to convey virtual experiences to an external participant in a second highly immersive environment.

For all results of our paper we used a GPU implementation that utilized an offscreen buffer with the same sized window. However, this may be undesirable in situations with extremely complicated scenes, ultra-high-resolution displays, or lower-end graphical hardware. By reducing the size of the buffers used to store the depth maps for the participant's view and the camera's view, both the scene rendering and pixel-wise comparisons can be done more rapidly. As shown in Ponto et al., reducing the size of these buffers can greatly increase performance without affecting results [16].

6. CONCLUSIONS

We created a novel method for presenting the experience of a participant in a virtual environment to external viewers in an online real-time fashion. Our method uses a content dependent metric to determine how much the camera should be adjusted on a per-frame basis. Our method is able to greatly reduce the camera movement over time and is thusly able to provide camera paths that are easier and more pleasing to watch while maintaining the content of the participant's original observations. Future work will explore using other viewing perspectives, such as an over the shoulder viewpoint, and will investigate methods to change the camera's field of view dynamically.

Acknowledgments

We would like to acknowledge the support of the the Living Environments Laboratory, the Wisconsin Institute for Discovery, and the CIBM training program. This project was supported in part by NLM award 5T15LM007359 and NSF awards IIS-0946598, CMMI-0941013, and IIS-1162037. We would specifically like to thank Patricia Brennan, Kendra Jacobsen, Andrew Morland, and Ross Tredinnick for their support and assistance.

References

1. Alexa M. Linear combination of transformations. *ACM Trans Graph.* Jul.2002 21:380–387.
2. Chang H-C, Lai S-H, Lu K-R. A robust real-time video stabilization algorithm. *Journal of Visual Communication and Image Representation.* 2006; 17(3):659–673.
3. Christie, M.; Machap, R.; marie Norm, J.; Olivier, P.; Pickering, J. *Proceedings Smart Graphics.* Springer; 2005. Virtual camera planning: A survey; p. 40-52.
4. Dawes J. Do data characteristics change according to the number of scale points used? *International Journal of Market Research.* 2008; 50(1)
5. Downs, JS.; Holbrook, MB.; Sheng, S.; Cranor, LF. Are your participants gaming the system?: screening mechanical turk workers. *Proceedings of the 28th international conference on Human factors in computing systems CHI '10*; New York, NY, USA. 2010. p. 2399-2402.
6. Ertürk S. Real-time digital image stabilization using kalman filters. *Real-Time Imaging.* 2002; 8(4): 317–328.
7. Foxlin, E. Inertial head-tracker sensor fusion by a complementary separate-bias kalman filter. *Virtual Reality Annual International Symposium, 1996, Proceedings of the IEEE 1996*; mar-3 apr 1996; p. 185-194.
8. Gleicher ML, Liu F. Re-cinematography: Improving the camerawork of casual video. *ACM Trans Multimedia Comput Commun Appl.* Oct.2008 5:2:1–2:28.
9. Grundmann, M.; Kwatra, V.; Essa, I. Auto-directed video stabilization with robust L1 optimal camera paths. *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on; IEEE;* 2011. p. 225-232.

10. Hansen M, Anandan P, Dana K, van der Wal G, Burt P. Real-time scene stabilization and mosaic construction. *Applications of Computer Vision, 1994.*, Proceedings of the Second IEEE Workshop on. Dec.1994 :54–62.
11. Hawkins A, Grimm CM. Camera keyframing using linear interpolation of matrices. *journal of graphics, gpu, and game tools.* 2007; 12(3):55–69.
12. Kim Y, Jayanthi V, Kweon I. System-on-chip solution of video stabilization for cmos image sensors in hand-held devices. *IEEE transactions on circuits and systems for video technology.* 2011; 21(10):1401–1414.
13. Kittur, A.; Chi, EH.; Suh, B. Crowdsourcing user studies with mechanical turk. *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, CHI '08*; New York, NY, USA. 2008. p. 453-456.
14. Liang, J.; Shaw, C.; Green, M. On temporal-spatial realism in the virtual reality environment. *Proceedings of the 4th annual ACM symposium on User interface software and technology, UIST '91*; New York, NY, USA. 1991. p. 19-25.
15. Lin C-T, Hong C-T, Yang C-T. Real-time digital image stabilization system using modified proportional integrated controller. *Circuits and Systems for Video Technology, IEEE Transactions on.* Mar; 2009 19(3):427–431.
16. Ponto K, Kohlmann J, Gleicher M. Effective replays and summarization of virtual experiences. *Visualization and Computer Graphics, IEEE Transactions on.* Apr; 2012 18(4):607–616.
17. Pothier DD, Hughes C, Dillon W, Ranalli PJ, Rutka JA. The use of real-time image stabilization and augmented reality eyewear in the treatment of oscillopsia. *Otolaryngology – Head and Neck Surgery.* 2012
18. Ratakonda, K. Real-time digital video stabilization for multi-media applications. *Circuits and Systems, 1998. ISCAS '98. Proceedings of the 1998 IEEE International Symposium on;* may-3 jun 1998; p. 69-72.
19. Salamin, P.; Thalmann, D.; Vexo, F. The benefits of third-person perspective in virtual and augmented reality?. *Proceedings of the ACM symposium on Virtual reality software and technology, VRST'06*; New York, NY, USA. 2006. p. 27-30.
20. Shin HJ, Lee J, Shin SY, Gleicher M. Computer puppetry: An importance-based approach. *ACM Trans Graph.* Apr; 2001 20(2):67–94.
21. Wang C, Kim J-H, Byun K-Y, Ni J, Ko S-J. Robust digital image stabilization using the kalman filter. *Consumer Electronics, IEEE Transactions on.* Feb; 55(1):6–14. 2009.
22. Williams L. Casting curved shadows on curved surfaces. *SIGGRAPH Comput Graph.* Aug.1978 12:270–274.
23. Yang, H.; Olson, GM. Exploring collaborative navigation: : the effect of perspectives on group performance. *Proceedings of the 4th international conference on Collaborative virtual environments, CVE '02*; New York, NY, USA. 2002. p. 135-142.
24. You S, Neumann U. Fusion of vision and gyro tracking for robust augmented reality registration. *Virtual Reality, 2001. Proceedings IEEE.* Mar.2001 :71–78.

Figure 1.

External viewers often find it difficult to observe the experience of a virtual participant due to the perceived unnaturalness of the egocentric camera movements. Our method is able to make these external views easier to watch by minimizing camera movement while maintaining the context of the participant's observations.

Figure 2.

The system consists of three components for each point in time. First the similarity between the camera and the participant's viewpoint is determined. From this, an adjustment function is computed to determine the amount the camera should be repositioned. Finally the camera is moved in order to achieve a better view by interpolating it towards the participant's viewpoint.

Figure 3.

(Top) The rectangular light used by Ponto et al. with sharp edges of the light was shown to be distracting for viewers when both the camera and light were moving [16]. (Bottom) Our method using a gradual falloff to provide probabilistic information as to where the participant was looking proved to be less distracting for the same situation.

Figure 4.
Images taken from the eight different scenarios for the participant to locate virtual objects.

Figure 5. Velocity profile for approximately 10 seconds of a participant walking around a virtual tree (Scene 1). As shown, our method reduces the magnitude of the velocity compared to the Raw or Kalman filtered camera path.

Figure 6.

Graphs for each of the subjective questions asked in the study. The results proved to be statistically significant, with the only exception for the subjects' confidence in their answers. The bars show the first and third quartiles.

Table 1

Camera movcf each scene for each method. Note that while the Kalman filter has very mixed results in reducing camera travel distance, our methods are able to consistently reduce the distance the camera traveled. This reduction in the camera travel distance did not result in a loss of content as shown by our user study.

Camera Travel Distance [m] (Distance / Raw)				
Scene	Raw	Kalman	CAP	SAP
1	5.0	4.4 (.88)	3.2 (.64)	2.8 (.56)
2	14.3	12.6 (.88)	11.2 (.78)	9.4 (.66)
3	9.8	9.1 (.93)	6.1 (.62)	4.5 (.46)
4	11.3	11.3 (1.0)	9.3 (.82)	7.3 (.65)
5	31.1	31.0 (1.0)	26.5 (.85)	25.8 (.83)
6	27.9	30.3 (1.1)	25 (.90)	24.1 (.86)
7	33.5	33.5 (1.0)	29.5 (.88)	26.9 (.80)
8	90.1	100.1 (1.1)	85.0 (.94)	85.7 (.95)
Average Ratio		.99	.81	.72