

Using Saccadic Suppression to Hide Graphic Updates

J. Schumacher^{1,2} and R. S. Allison¹ and R. Herpers^{1,2}

¹Centre for Vision Research, Department of Computer Science and Engineering, York University, Toronto, Canada

² Department of Computer Science, Bonn-Rhein-Sieg University of Applied Sciences, Sankt Augustin, Germany

Abstract

In interactive graphics it is often necessary to introduce large changes in the image in response to updated information about the state of the system. Updating the local state immediately would lead to a sudden transient change in the image, which could be perceptually disruptive. However, introducing the correction gradually using smoothing operations increases latency and degrades precision. It would be beneficial to be able to introduce graphic updates immediately if they were not perceptible. In the paper the use of saccade-contingent updates is exploited to hide graphic updates during the period of visual suppression that accompanies a rapid, or saccadic, eye movement.

Sensitivity to many visual stimuli is known to be reduced during a change in fixation compared to when the eye is still. For example, motion of a small object is harder to detect during a rapid eye movement (saccade) than during a fixation. To evaluate if these findings generalize to large scene changes in a virtual environment, gaze behavior in a 180 degree hemispherical display was recorded and analyzed. This data was used to develop a saccade detection algorithm adapted to virtual environments. The detectability of trans-saccadic scene changes was evaluated using images of high resolution real world scenes. The images were translated by 0.4, 0.8 or 1.2 degrees of visual angle during horizontal saccades. The scene updates were rarely noticeable for saccades with a duration greater than 58 ms. The detection rate for the smallest translation was just 6.25%. Qualitatively, even when trans-saccadic scene changes were detectible, they were much less disturbing than equivalent changes in the absence of a saccade.

Categories and Subject Descriptors (according to ACM CCS): I.3.3, I.3.7, I.3.8 [Computer Graphics]: Graphic Updates

1. Introduction

In interactive graphics it is often necessary to introduce large changes in the image in response to updated information about the state of the system. For example, distributed applications such as networked simulations often rely on dead reckoning techniques to minimize bandwidth requirements and to alleviate the detrimental effects of network latency [Mat97, CH97]. Local estimates of the state of remote entities can drift from the true state requiring a correction to be transmitted. Updating the local state immediately would lead to a sudden transient change in the image which can be perceptually disruptive. Thus, smoothing operations are applied to correct the estimated state but imply a tradeoff with accuracy and latency.

Similarly many motion tracking devices used in interac-

tive systems rely on hybrid sensor systems to estimate the user's pose [Azu93, HRJA03, WBV*99]. In many systems, these hybrids consist of a rapidly responding, high bandwidth sensor, such as accelerometer, that can track relative motion with low latency but that cannot provide an absolute pose estimate due to drift. Another, often slower, system is used to estimate absolute pose and correct the estimates obtained from the faster system. As in the dead reckoning example, sudden correction of the estimates could be perceptible and disruptive. However, introducing the correction gradually using smoothing operations increases latency and degrades accuracy.

It would be beneficial to be able to introduce graphic updates (or updates in the models of the simulated worlds) immediately if they were not perceptible. In this paper we con-



Figure 1: Eye-tracking system used in IVY, a six sided projective display located at York University [RLZ*02].

sider the use of saccade-contingent updates to hide graphics updates during the period of visual suppression that accompanies a rapid, or saccadic, eye movement.

Two experiments have been performed in order to evaluate the possibilities of trans-saccadic scene updates in virtual environments. In the first one, a spherical display was used to present the virtual environment and the eye-movements of subjects were recorded for later analysis. In the second experiment, we studied the sensitivity of observers to large trans-saccadic image translations using a CRT projector.

2. Saccadic Suppression

A distinguishing feature of the human eye is that highest acuity is restricted to small portion of the central retina called the fovea. To appreciate a natural or computer-generated scene, fast eye movements called saccades must be used to direct this foveal region to areas of interest in the scene. Thus, while scanning a scene, the high acuity fovea is successively sampling different regions of the scene and rapid saccadic eye movements connect these samples. During these saccades, images of objects will be streaming across the retina at hundreds of degrees per second. Despite this disjoint motion of the retina, the world does not appear disjoint or unstable and motion blur during saccades is not appar-

ent. Part of the reason is that, during saccades, sensitivity to visual stimuli is reduced - an effect known as saccadic suppression [LZ99]. Thus, the visual world appears stable even as its image moves at speeds of up to 800 deg/s across the retina. In contrast, if a subject views an object or scene that moves with a velocity of about 800 deg/s across the retina when she is not making a saccade, she will still be able to see lower spatial frequencies in the image [LZ99].

This suppression is not obvious since the brain combines information from successive eye fixations to create an apparently continuous view of the visual field. It is assumed that this mechanism masks the actual motion blurred picture and maintains space constancy over the retinal image shift in order to provide a stable view of the world [LZ99]. However, it appears that the visual system does not maintain a complete, rich internal map of the scene. Numerous studies have demonstrated that, during visual interruptions such as saccades, human beings can be insensitive to even large changes in the scene. This phenomenon is known as change blindness. Change blindness can occur during transient blanking of a screen, 'mud-splash' patterns briefly overlaid on the image, or eye blinks [SL97, SL98, ORC99] as well as during saccadic eye movements [MC96, MJA03].

2.1. Using Saccadic Suppression to Hide Graphic Updates

Our approach is conceptually simple. If a graphics update is pending and a saccadic eye movement occurs then introduce the update during the saccade when saccadic suppression should prevent it from being apparent. If no saccade occurs then blend in the correction gradually using standard techniques. The approach is related to multi-resolution gaze contingent displays [RTW89, TSHB02, Duc03] that change the location of a high resolution window into the larger display during the eye movement. However, in our proposal, changes in the graphics model (in particular the viewer's pose or a modeled object's location) are introduced instead of simply the rendering detail.

The remainder of this paper investigates the feasibility of this approach during interaction with rich virtual environments. Three topics are addressed. First, we ask how often one could expect to see saccades that are suitable for hiding graphics updates when viewing wide field-of-view computer displays of complex scenes. Second, real-time algorithms to detect such saccades using video eye-tracking hardware are developed and evaluated. Video-based eye trackers are the most common and flexible systems but typically suffer from low measurement rates. Finally, the results of an experiment studying saccadic suppression while viewing a large projection display are described. The goal is to determine the degree to which saccadic suppression reduces sensitivity to changes that are typical of simulations in an immersive virtual environment.

3. Saccade Frequency and Detectability

Unfortunately most previous studies of saccade-contingent updates have been performed using normal monitors [MC96, MJA03] and simple objects [DBS98] as saccade targets. It is not clear how well these results generalize to complex real world scenes or a virtual environment. It is also unclear whether a video-based eye tracking system with a temporal resolution of 120 Hz is able to detect saccades fast enough to perform scene changes in a realistic environment. To investigate these issues we studied gaze behaviour while viewing complex scenes.

This initial experiment was concerned with three questions: How many saccades do subjects perform while analyzing a complex scene in a virtual environment? What is the distribution of saccade amplitudes and durations? How well does our saccade detection algorithm work? The experiment was designed to collect representative data to answer these questions and to provide a standard data set for comparison of saccade detection algorithms.

3.1. Subjects

Nine normal subjects participated in the first experiment, aged from 21 to 34 years. They were all naive with respect to the task and inexperienced with the equipment. Three of them performed the experiment twice in order to evaluate whether the behavior of subjects was repeatable across two different sessions for the same subject.

3.2. System Setup

The system was based on the Vision 2000 video eye tracking system by EL-MAR Inc. (Toronto, Canada), which uses an adaptive real-time image processing technique to obtain accurate measurements of the eye movement. The eye position is measured at a frequency of 120Hz and sent to the workstation over a RS232 serial port at a baud-rate of 38400. A standard Linux workstation (dual AMD Athlon MP 1900+ processors, 1024 MB System RAM, GeForce4 TI4600, Red-Hat Linux with kernel version 2.4.18) is used to display the stimuli, process eye movement data and perform the changes in the image for the experiments.

To get the most natural display we used a VisionStation hemispherical display by Elumens to present the pictures. With a screen height of 1.45m, a width of 1.63m and a depth of 0.53m it offered a screen surface area of 2.4 square meters. The projector used to project the image on the surface of the VisionStation was a Hitachi CP-SX5600W, with native resolution of 1200x1024 pixels at 75 Hz. The 16.7 million colors could be displayed with a brightness of 1,800 ANSI Lumens and a contrast ratio of 600:1.

3D-rendered pictures were used for this experiment. We used the "True Frame" Software by Elumens to display the correct spherical projection.

3.3. System End-to-End Latency

The time between the saccade onset and the actual display change was measured for two different output devices. One was the projector-based display used in the main experiment (Section 4) and the second was a common View Sonic CRT Monitor with a refresh rate of 100Hz. A photodiode detector attached to the graphical output device indicated changes of the display luminance. The horizontal eye position, provided by the analog output of the Vision 2000 eye-tracker, indicated the progression of a saccade. A large luminance change in the display was triggered by the saccade crossing a known threshold. The end-to-end latency between the saccade crossing this threshold and the display luminance change could be monitored with a Kikusui 5020A oscilloscope with sufficient precision. The time-base used was 20 ms/div (2 ms under magnification) with vertical sensitivity of 5 volt/div on the eye-tracker position signal and 1 volt/div on the photoreceptor. Requiring the subject to switch fixation between two targets separated by approximately 20 deg controlled the amplitude of the saccades triggered during this test. The average duration of such a movement was 67 ms. The average delay for the CRT monitor was 42 ms and for the Barco 808 CRT projector was 35 ms (due to the higher refresh frequency of 120Hz). These measurements do not include the internal delay of the eye-tracker which is uniformly distributed between 8 and 16 ms depending at what point in the video frame the saccade started. Hence we can estimate the total end-to-end latency as between 50 and 58 ms for the desktop monitor and between 43 to 51 ms for the Barco projector.

3.4. Calibration

Each session started with a calibration procedure where the subject sequentially fixated ten points, 5 horizontal and 5 vertical. The horizontal points were evenly distributed between -15 and +15 degrees and the vertical points between -11 and +11 degrees of viewing angle. If the internal routine of the eye tracker indicated that the calibration was successful, we started the trial. The quality of the calibration was also manually checked on a plot of the fit of the calibration points to the estimated data.

3.5. Method

In each session, five arbitrarily chosen scenes were sequentially shown to the subjects for 30 seconds each while their eye movements were recorded. Subjects began the session watching a black screen with a white point centered in the middle. This screen was displayed at the beginning and after each trial for 2 seconds. The subjects were asked to fixate the point in the middle of the screen after each trial to ensure that the system was still correctly calibrated.

The distance from the eyes to the screen was about 80

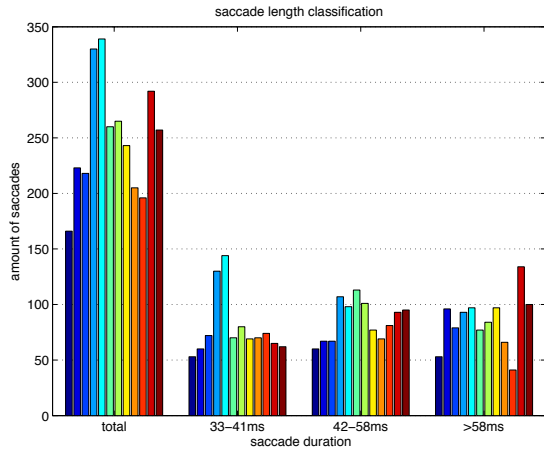


Figure 2: Classification of saccade duration for saccades elicited during free viewing of images. Each bar shows the number of saccades made by a subject (color coded) in a session of 150 s.

cm and viewing was binocular. Subjects were told to analyze and memorize the scenes for a later memory task. Even though no scene changes were imposed, we asked them to watch out for possible scene changes and press a button as soon as they detected one. We did this to avoid possible different eye movement behavior when this additional task was required in the second experiment. While the subjects moved their gaze over the picture, the eye position data was recorded.

3.6. Data Analysis

The recorded eye movement data consisted of the eye position and a flag indicating the occurrence of an eye blink for each sample instant. A offline program calculated the current velocity using both a first-order backward difference and a smoothed five point differentiator (for more details see section 3.7).

The collected data were evaluated to determine the frequency of saccades made during free viewing of scenes in a virtual environment and to develop and optimize our detection algorithm. Based on their duration, the measured saccades were classified into three categories (short, medium, long). Originally, two methods of detecting saccades were tested.

While viewing the five scenes (total time 150 s), on average subjects made 249.5 saccades with a standard deviation of 52.4 (see figure 2). The distribution of saccades in the different classes was consistent across subjects with 79.08 small ($\sigma = 28.1$), 85.67 ($\sigma = 17.7$) medium and 84.75 ($\sigma = 24.3$) long duration saccades. Since we are primarily interested in the class of long-duration saccades, this would

indicate that scenes could be changed about 0.56 times/s or every 1.7 seconds assuming that all saccades are detectable.

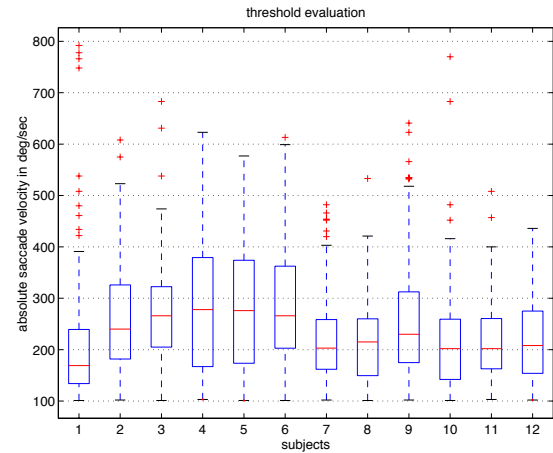


Figure 3: Average horizontal saccade velocity of long duration saccades at the second frame after the threshold of 10 deg/s was exceeded. The lines of the box are at the lower quartile, median and upper quartile values. Lines extending from the box show the extent of the rest of the data. Points beyond the lines are outliers.

Only the category of long saccades, with duration longer than 58 ms, were sufficiently long to be used for saccade-contingent changes in our experimental setup. Due to system latency, a saccade-contingent change needed to be initiated by the third eye tracker sample (frame) from the start of such a large saccade. In order to determine a reasonable threshold value that has to be exceeded to trigger a change, we evaluated the velocity at this period in time for long-duration saccades. The results shown in figure 3 indicate that a value of 200 deg/s would be an appropriate threshold. These values are consistent with values determined from earlier studies of saccadic profiles [LZ99]. Our results are generally consistent with scan path behavior in other studies that used complex images [Yar67].

3.7. Detection Algorithm Evaluation

The eye position was sampled with a frequency of 120Hz; successive samples were 8.33 ms apart. Two algorithms were tested offline with the collected data.

The first algorithm tested was a simple 2-point differentiator with 2 different thresholds in two successive frames of the data. If and only if both thresholds were met, a saccade event was flagged. The thresholds of 10 deg/s for the first and 200 deg/s for the second frame were based on the velocity profile of saccades described in standard texts such as Leigh and Zee [LZ99].

The second method tested used was based upon a five

point differentiator that worked on a window of estimates from five successive frames and calculates the filtered velocity for the value in the middle. This filter is symmetric about its center and thus adds a delay of 2 1/2 frames. A saccade was indicated if the filtered velocity exceeded a threshold of 60 deg/s. If the velocity exceeded 900 deg/s at the first or second sample point, the saccade detector was not triggered as such a rapid acceleration is physiologically unlikely. Screening and removal of these later cases made the system more robust against noise and other artifacts (e.g. undetected eyeblinks).

The two algorithms were tested on the sample data that we collected. For both algorithms, the false alarm rate was unacceptably high even if the threshold was increased. A saccade length of 50 ms was set as a lower bound limit for making a trans-saccadic change. Every detected saccade that was shorter in duration than this threshold was categorized as overestimated in the sense that an initiated change would not occur before the saccade finished. Even with a threshold of 60 deg/sec for the 5-point differentiator, saccades with a duration shorter than 33.32, which were either really short saccades or noise, were occasionally detected as regular long-duration saccades (see figure 4). The ratio of correctly detected saccades to false alarms for the 5-point differentiator with a threshold of 60 deg/sec was 2.8:1 and was 3.8:1 for the 2-point differentiator with a threshold of 220 deg/s on the second frame. Raising the thresholds lowered the number of false alarms but also increased the number of missed suitable saccades.

To reduce the number of overestimated saccades we combined the two algorithms by adding the 5-point differentiator threshold to the 2-point differentiator algorithm. In addition to this and in order to guarantee that an update was not triggered at the saccade peak velocity, the acceleration is estimated over the previous three frames of data in order to verify that it is continually increasing. A change was only triggered if all three (2-point differentiator threshold: 200 deg/s, 5-point differentiator threshold: 60 deg/s, accelerating eye-movement) conditions were satisfied.

A comparison of the resulting algorithm and the two base algorithms can be seen in figure 4. Based on the correctly detected and overestimated saccades, the false alert rate with the new hybrid algorithm is as low as 16.2%. A low false detection rate should be the highest priority. Even though this hybrid algorithm is more conservative than the base algorithms, the rate of missing suitable saccades remains low enough to allow a reasonable update frequency.

4. Detectability of Trans-Saccadic Image Translation

In the second experiment, the sensitivity of subjects to translation of complex real world images during a saccade was investigated and compared with sensitivity to translation in the absence of a saccade. The goal was to determine whether

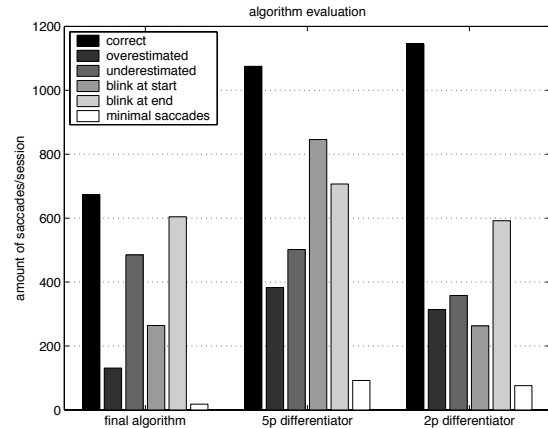


Figure 4: The figure shows the amount of saccades of each type for each class summarized from all sessions (12 sessions x 150s = 1800s).

graphics updates that would normally be perceptually disruptive could be hidden in the trans-saccadic intervals.

For the setup of this experiment the subjects freely viewed a complex scene. After a specified number of saccades, the scene was changed either during the saccade or during a fixation. The change was a translation of the whole scene to the left or the right. Three different sizes of translation (0.4, 0.8 and 1.2 degree of viewing angle) were performed and sensitivity evaluated.

4.1. System Setup

The system setup for this experiment differed merely in the output device from the system in Section 3. Due to delays associated with input resampling by the DLP projector, the Vision Station could not be used in this experiment. Instead a large screen projection display generated by a Barco 808 CRT projector with a resolution of 1024x768 and a refresh rate of 120 Hz was used.

4.2. Subjects

Six normal subjects including two of the authors participated in this study. One subject was female and five were male. Age ranged from 22 to 38. Three subjects were already familiar with the eye tracking system and the task. The other three subjects were complete naive to the whole procedure.

4.3. Material

Instead of the rendered 3D scenes used in the previous experiment, images of complex natural scenes were shown to the subjects. The images were recorded with a Nikon Coolpix5400 digital camera at a resolution of 5.1 megapixels and afterwards down-sampled to the screen resolution of

1056x792. The scene size had to be bigger than the actual screen-size of 1024x768 to be able to perform a translation of the picture without shifting the border of the screen, which could have affected the detection performance.

4.4. Calibration

The calibration procedure was changed to a single center point calibration mode supported by the system. This procedure was faster and more comfortable for the subjects. We confirmed that the accuracy of this procedure was sufficient given that we did not need to perform high accuracy point of regard measurements.

4.5. Methods

Participants examined a set of 45 images in a random order. Each image was shown for 10 seconds. The trials were divided into five categories. Three of them differed in size of the saccade-contingent image change, one had no change at all and the last category had changes that were not synchronized with an eye movement. Ten images were randomly assigned to each category of saccade-contingent displays and ten images were assigned to the no change category. To compare the detectability of the saccade-contingent changes with equivalent changes in the absence of a saccade, five images were assigned to trials where saccades were triggered after a specific time rather than during a saccade. These changes were expected to be easily detected if the subjects did not coincidentally perform a blink or a saccade.

Subjects were asked to study the scenes and memorize them in order to perform a memory test afterwards. The purpose of this task was to ensure that subjects actively viewed the scenes. After the scene was shown for 10 seconds a black screen was displayed and the subjects were asked to indicate whether there was a change in the last scene or not. After entering the decision by pressing a button on the keyboard, the next scene was shown. The frequency with which the changes occurred was not communicated to the subjects. Each scene was changed once per trial. An example of the type of change that would occur was shown to the subjects prior to the experiment.

The changes were a translation either to the right or to the left triggered at the third saccade that was detected in the viewing period for each image. The direction of translation depended on the internal scene number. Odd numbers shifted the image to the left, even numbers to the right. However, as the scenes were shown in random order neither the direction nor the category was predictable. The stimulus displacement was a horizontal shift of 0.4, 0.8 or 1.2 degree in terms of visual angle from the eye position of the subject. There is some evidence from experiments with small target displays that displacements that are small with respect to the size of the saccade are less detectable [MJA03, MC96]. It

was interesting to investigate whether this was also true for translations of the entire visual environment.

4.6. Results and Discussion




Figure 5: Relation between saccade length and change detection. Saccade length is measured in frames (8.3 ms) starting from the detection of the saccade.

It was found that saccades that met the detection criterion occurred in average of 5.9 times per picture. Depending on the scene and subject this number varied between 2 and 16 saccades. However, in some cases the participant made insufficient saccades to trigger the scheduled change. Out of 210 trials where a change was scheduled (30 displacements x 7 subjects) data for 188 displacement trials was obtained. The change detector was limited to only respond to horizontal saccades in order to limit possible noise introduced by the eyelids while performing vertical eye-movements.

An analysis of false alarm and miss rates during the time-triggered (i.e. non-saccade-contingent) scene changes was performed. Responses could include 2 types of errors: detecting a change during the control condition where no change occurred (false alarm) and failing to detect when the image was changed after a specified time interval (miss). Subjects acted conservatively with respect to change detection and triggered no false alarms during the no-change condition. In the timed-change condition 5 out of 35 changes (14.4 percent) were missed. After analysis of the data it was found that on four of these timed trials the change occurred during a coincidental saccade and one on a coincidental blink occurred. Thus, it appears that suppression during a coincidental saccade or blink explains why these usually obvious changes were missed.

The primary aim of the second experiment was to evaluate if it was possible to reliably detect saccades and change the display without the subject noticing it. If the subject detected the change we wanted to know if it was due to a small saccade or other possible reasons.

	detected	not detected
0.4 deg	6.25	93.75
0.6 deg	14.29	85.71
1.2 deg	17.65	82.35

Table 1: Detection ratio as a function of translation size

In total we obtained 188 trials with saccade-contingent scene changes. Over all the subjects, 86 of these changes were detected and 102 were not detected. The detection ratio ranged from 32.3% to 67.6% in individual subjects. Only one subject in one session detected more changes than they missed. The average detection level was about 44.5% and for most subjects clearly under chance level of 50%.

The detection rate was higher than we expected based on results from previous studies [TSHB02, MJA03, Irw92] from which detection probability could be inferred. To try to account for this, the saccade size was compared with the detection rate for saccades in each size category. It was possible that on some trials the saccade length was too short to perform the change. The majority of changes were triggered with 33.2 ms remaining in the intra-saccadic interval. The total saccade length in this category was predominately about 50 ms. Based on our total system delay measurements this is approximately equivalent to the time needed to finish a display change after the saccade started. The average time left after the detection of the saccade was 38 ms for the detected changes and 57.4 ms for undetected changes. This indicates that the duration of the saccade was a critical factor since it was crucial that the saccade duration exceeds the end-to-end latency.

Since highest spatial frequency information in a scene is not readily visible any sooner than 6 ms after the peak of the overshoot of the eye movement, it is supposed that perception begins at about the time the eye stops rotating. Also to acquire a detection rate below 0.10, eye velocities above 40 deg/s are required [ML02]. Thus, we conclude that most changes that were detected were the result of short saccade durations. Detection rate decreased if the saccade length increased.

To evaluate if the translation magnitude affected the detection of a change, the 3 different translation conditions were analyzed. To eliminate detections based on short saccade length, only saccades with duration of at least 66 ms were considered. With a translation of 0.4 degree, only 6.2 percent of changes were detected while 93.8 percent were not seen. Large translations were easier to see and the detection rate for 1.2 deg translations was 17.6 percent while just 82.4 percent of the changes stayed undetected, see table 1.

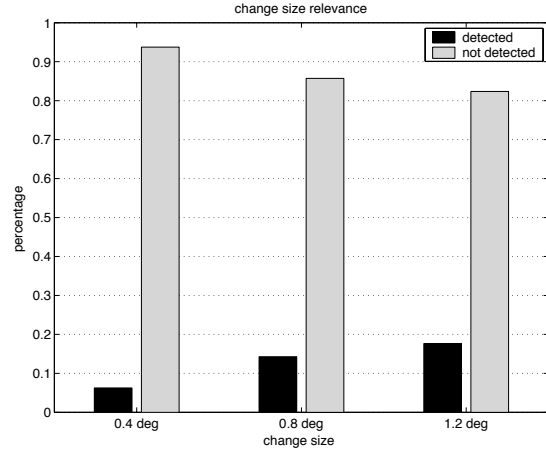


Figure 6: Correct detection and miss rates for changes during saccades that were longer than 58.3 ms.

5. Conclusion and Discussion

We have shown that just like changes in simple displays on computer monitors, scene changes in realistic images presented in virtual environments can be suppressed during saccadic eye movements. Our system was similar to a normal virtual-environment display without any special hardware except for the video-based eye tracker. Our aim was to evaluate if saccadic suppression could be used to perceptually hide graphics updates in virtual environments. Other applications of gaze-contingent displays rely on high accuracy that can be obtained with a precision video-based eye-tracking system. One problem with these types of systems is their temporal resolution. The limited sampling rate means that the estimate of the size of the saccade and prediction of whether a graphic update will be perceptible must be based upon only a few samples from the beginning of a saccade.

Our algorithm was conservatively designed and was able to reliably detect large saccades with a false alarm rate of 20% although it missed some long-duration saccades. It was found that updates were not perceptible for saccades bigger than 15 degree or longer than 66 ms in duration. As end-to-end system delay was about 51 ms and could vary up to 59 ms, for an imperceptible change one has to ensure that the saccade duration exceeded these values. Even when detectible, the visual disturbance was barely noticeable.

6. Future Work

After the detection of a saccade there must be sufficient time to complete the change before the end of the eye movement. In the current implementation, this limits us to triggering on fairly large saccades, which impacts the maximum update rate. One possibility could be to encourage the user to make big saccades, but this method limits the potential in

most applications. In order to increase the opportunity for saccade-contingent scene changes in virtual environments, the saccade detector needs to be enhanced. Latency may be reduced somewhat with high speed cameras or higher bandwidth systems such as infrared limbus trackers. But using these different system implies other disadvantages for example lower accuracy (limbus) or a limited range (double Purkinje).

With an accurate gaze tracker it would be possible to perform accurate point-of-gaze measurements and would offer the ability to perform saccade-contingent scene changes. Our technique could be combined with other gaze-contingent techniques such as multi-resolution displays.

The technique described in this paper could be used in a distributed virtual environment that is partly controlled or responsive to eye-movements. Errors in scene presentations due to network delay could be corrected while the user performs a saccade.

7. Acknowledgements

Robert Allison gratefully acknowledges the support of CRESTech and NSERC (Canada). Rainer Herpers gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft (DFG) grant HE 2967/2-1, which partially funded the presented work.

References

- [Azu93] AZUMA R.: Tracking requirements for augmented reality. *Communications of the ACM* 36, 7 (7 1993). 1
- [CH97] CONNER B., HOLDEN L.: Providing a low latency user experience in a high latency application. In *Symposium on Interactive 3D Graphics* (1997), pp. 45–48. 1
- [DBS98] DEUBEL H., BRIDGEMAN B., SCHNEIDER W. X.: Immediate post-saccadic information mediates space constancy. *Vision Research* 38 (1998), 3147–3159. 3
- [Duc03] DUCHOWSKI A. T.: *Eye Tracking Methodology: Theory and Practice*. Springer, 2003. 2
- [HRJA03] HOGUE A., ROBINSON M., JENKIN M. R., ALLISON R. S.: A vision-based tracking system or fully immersive displays. In *ACM International Conference Proceedings Series: Proceedings of the Eurographics Workshop on Virtual Environments* (2003), pp. 179–187. 1
- [Irw92] IRWIN D. E.: Memory for position and identity across eye movements. *Journal of Experimental Psychology: Learning, Memory and Cognition* 18, 2 (1992), 307–317. 7
- [LZ99] LEIGH J. R., ZEE D. S.: *The Neurology of Eye Movements*. Oxford University Press, 1999. 2, 4
- [Mat97] MATIJASEVIC M.: A review of networked multi-user virtual environments, 1997. 1
- [MC96] MCCONKIE G. W., CURRIE C. B.: Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology: Human Perception and Performance* 22, 3 (1996), 563–581. 2, 3, 6
- [MJA03] MACASKILL M. R., JONES R. D., ANDERSON T. J.: Saccadic suppression of displacement: Effects of illumination and background manipulation. *Perception* 32 (2003), 463–474. 2, 3, 6, 7
- [ML02] MCCONKIE G. W., LOSCHKY L. C.: Perception onset time during fixations in free viewing. *Behavior Research Methods, Instruments and Computers* 34, 4 (2002), 481–490. 7
- [ORC99] O'REGAN J., RENSINK R., CLARK J.: Change-blindness as a result of "mudsplashes". *Nature* 34 (1999), 398. 2
- [RLZ*02] ROBINSON M., LAURENCE J., ZACHER J., HOGUE A., ALLISON R., JENKIN M., HARRIS L. R., STUERZLINGER W.: Ivy: The immersive virtual environment at york. 7th Annual Immersive Projection Technology (IPT) Symposium May 23-24, Orlando. 2
- [RTW89] ROBINSON R. M., THOMAS M. L., WETZEL P. A.: *Eye Tracker Development On The Fiber Optic Helmet Mounted Display*. Tech. rep., Air Force Human Resources Laboratory, Operations Training and University of Dayton, 1989. 2
- [SL97] SIMONS D. J., LEVIN D. T.: Change blindness. *Trends in Cognitive sciences* 1, 261–267 (1997). 2
- [SL98] SIMONS D. J., LEVIN D. T.: Failure to detect changes to people in a real-world interaction. *Psychonomic Bulletin and Review* 5 (1998), 644–649. 2
- [TSHB02] TRIESCH J., SULLIVAN B. T., HAYHOE M. M., BALLARD D. H.: Saccade contingent updating in virtual reality, 2002. 2, 7
- [WBV*99] WELCH G., BISHOP G., VICCI L., BRUMACK S., KELLER K., COLUCCI D.: *The HiBall Tracker: High Performance Wider-Area Tracking for Virtual and Augmented Environments*. Tech. rep., Universtiy College London, 1999. 1
- [Yar67] YARBUS A. L.: *Eye movements and vision*. New York: Plenum (1967). 4