Compensating for Eye Tracker Camera Movement

Susan M. Kolakowski* Jeff B. Pelz†
Visual Perception Laboratory, Carlson Center for Imaging Science,
Rochester Institute of Technology, Rochester, NY 14623 USA

Abstract

An algorithm was developed to improve prediction of eye position from video-based eye tracker data. Eye trackers that determine eye position relying on images of pupil and corneal reflection positions typically make poor differentiation between changes in eye position and movements of the camera relative to the subject's head. The common method employed by video-based eye trackers to determine eye position involves calculation of the vector difference between the center of the pupil and the center of the corneal reflection under the assumption that the centers of the pupil and the corneal reflection change in unison when the camera moves with respect to the head. This assumption was tested and is shown to increase prediction error. Also, predicting the corneal reflection center is inherently less precise than that of the pupil due to the reflection's small size. Typical approaches thus generate eye positions that can only be as robust as the relatively noisy corneal reflection data. An algorithm has been developed to more effectively account for camera movements with respect to the head as well as reduce the noise in the final eye position prediction. This algorithm was tested and is shown to be particularly robust during the common situation when sharp eye movements occur intermixed with smooth head-to-camera changes.

CR Categories: D.1.1 [Programming Techniques]: Applicative (Functional) Programming; F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems; G.4 [Mathematical Software];

Keywords: eye tracking, camera compensation, algorithm, noise

1 Introduction

Eye trackers have proved to be powerful tools in understanding a broad range of behaviors. From Delabarre's [1898] work in the late 1800's to today's systems, eye trackers have allowed objective measures of task performance. Eye movements are unique in that they provide an external key to behavior and cognition. Observers are typically unaware of their eye movements and, except in extraordinary circumstances, do not exert conscious control over the more than 100,000 eye movements executed daily. The ability to examine behavior on a timescale of hundreds of milliseconds has yielded many insights. Extending that capability into complex tasks in the natural world is allowing a new class of experiments for examining truly natural behaviors.

*e-mail: smk8165@cis.rit.edu
†e-mail: pelz@cis.rit.edu

ETRA 2006, San Diego, California, 27–29 March 2006.
© 2006 ACM 1-59593-305-0/06/0003 $5.00

Land et. al. [1999] used a head-mounted eye tracker to study their subjects performing the "everyday" activity of making a cup of tea. Studying subjects during natural tasks such as this provides insight into how people perform everyday over-learned activities. Pelz and Canosa [2001] also studied subjects during natural tasks with the use of a portable eye tracker. One of these tasks involved hand-washing which elicited complex eye movements not observable during less complex tasks more commonly studied. Previous eye trackers tended to restrict their use to stabilized laboratory configurations. Researchers using such eye trackers have often substituted pictures or other artificial setups for natural environments. Henderson and Ferreira [2004] argue that results from scene depiction studies may not generalize to the real world environment and that the use of pictorial scene depictions introduces artifacts.

Eye trackers have evolved from crude devices that were often painful to use and highly imposing [Delabarre 1898]. By contrast, modern lightweight systems allow complex, natural movements of the eyes, head, and body [Babcock and Pelz 2004]. The majority of current eye tracker systems are image-based where video images of the eye are used to compute the point of gaze in an observer's field of view. Video-based eye trackers illuminate the eye with an infrared source (typically IR LEDs). Most visible wavelengths are absorbed in the pigment epithelium, but incident radiation in the deep red and near IR are reflected. The retina is a retroreflector; incident illumination is reflected back along the incident path. If the optical axis of a camera imaging the eye is coaxial with the IR illuminator, the retroreflected light back illuminates the pupil. The result is a 'bright-pupil' image. If the axis of the eye camera is not coincident with the illuminator, the reflected light does not enter the eye camera, resulting in a 'dark-pupil' image (see Figure 1).

Figure 1: Dark-pupil eye image frame.

Both bright- and dark-pupil systems process a digitized video stream to locate the pupil. By setting a threshold value at an appropriate level, the pupil region can be isolated. Locating the pupil center is robust and relatively low in noise because the pupil subtends a relatively large area. Noise in the centroid measurement is limited because of the size; the inevitable video noise at the 'edges' of the pupil region tends to average out along the circumference.

Under ideal conditions any change of relative position of the pupil within the eye camera's field-of-view would represent an eye movement. Through a calibration, the pupil's position could then be transformed to gaze position in an appropriate reference frame. That ideal condition, however, cannot be achieved without rigid constraints on observer motion. Even dental bite bars allow enough residual movement to cause artifacts in apparent eye motion due to

movements of the eye camera relative to the head. To correctly determine that a subject's eye is gazing, the eye tracker must compensate for any movements of the eye tracking camera with respect to the subject's head. Currently, video-based eye trackers compensate for this camera movement by tracking both the *corneal reflection* (CR) and the pupil. If the camera moves with respect to the head, the pupil and CR images tend to move in the same direction and are usually assumed to have moved the same relative distance.

The eye camera captures the virtual image of the iris formed by the cornea, and the first-surface reflection from the surface of the cornea. A specular reflection on the nearly spherical cornea will move only a portion of the offset of the pupil image. If the camera moves with respect to the head, the pupil and CR images tend to move together. A common method used to compensate for camera movement in many video-based eye trackers is referred to as the "Pupil Minus Corneal Reflection (P-CR)" technique. This technique uses the vector difference between the center of the pupil and the center of the corneal reflection to determine eye position. This method compensates for camera movement under the assumption that when the camera is translated with respect to the eye, the pupil and corneal reflection are translated by the same amount. This assumption, which says that a camera movement would not affect the P-CR vector, was tested with results reported in Section 5.

The corneal reflection region in the eye tracker's eye image is much smaller than the pupil region (see Figure 1). As a consequence, the spatial noise inherent in all video systems will be more significant to the estimation of the center of the CR than to the estimation of the center of the pupil. In addition, any variation in the surface of the cornea or tear layer can degrade the CR signal. The pupil region contains approximately 25 times the number of pixels in the corneal reflection region. In sum, the CR signal can be much higher in noise than the pupil signal. The vector difference in the P-CR calculation will be at least as noisy as the corneal reflection data. In practice there is a tradeoff between noise in the P-CR offset signal and the temporal response; the operator can select the number of video fields (16.7 msec per field) over which to average. Our goal was to develop an algorithm that maintained the noise level of the pupil signal without adversely affecting the mean temporal response. Successful systems will compensate for movements of the eye camera with respect to the head, while maintaining the ability to detect small, rapid eye movements that would otherwise be obscured by CR noise.

Karmali and Shelhamer [2004] attacked the problem of compensating for camera movement primarily through digital image processing techniques. In their preferred solution, they create eyelid templates for comparison with each eye image frame via cross-correlation to determine the amount of camera translation. Xie et. al [1998] exploit the fact that head movements are much slower than eye movements in their "head movement compensation" algorithm. The authors use a Kalman filter within their algorithm that tracks the center of the eye as a reference point via digital image processing to determine camera movement with respect to the head. These techniques require additional image processing, do not fully compensate for camera movement, and suffer from artifacts due to any errors and/or noise in localizing the landmarks used to compensate for camera/head movements.

A new algorithm to calculate eye position from video-based eye trackers was implemented using the method of computing and extracting movements of the camera with respect to the head. Once isolated, the relative motion (or camera movement) can be low-pass filtered. The calculated camera position data are smoothed before being subtracted from the output pupil position data because these movements are less frequent and slower than eye movements. This method compensates for camera movement through the use of the corneal reflection data and yet is as accurate and low in noise as pupil-only data.

The following Section provides an overview of the eye tracking system. Section 3 explains the theory behind the new algorithm. The implemented algorithm is described in Section 4. Section 5 covers the methods used to test the algorithm and discusses the results obtained from these tests. Finally, Section 6 consists of conclusions and possible future extensions of this work.

2 Overview

The algorithm was developed on data from the RIT wearable eye tracker [Babcock and Pelz 2004], a customized dark-pupil, head-mounted portable eye tracker. The ISCAN 726/426 Pupil/Corneal Reflection Analysis System was utilized. This video-based eye tracker determines eye position based on calibration of the vector difference of the pupil and corneal reflection positions. The system delivers floating point data in terms of pixel position with a resolution of 3/10 of a pixel horizontally and 1/10 of a pixel vertically. The video signal is divided into a 512 horizontal x 256 vertical pixel matrix such that the effective operational matrix is 1500 horizontal by 2000 vertical eye position data points [ISC 2001]. The vector difference between the pupil and corneal reflection positions within the eye image are calibrated to the scene image to produce the final point of regard data.

The final point of regard data output by the eye tracker may be the result of any combination of eye movements, camera movements (with respect to the head) and noise. It is desirable to separate eye movements from camera movements when analysing eye tracking data.

The eye movements to be studied via monocular eye tracking experiments consist of saccades, smooth pursuit, optokinesis and vestibular-ocular reflex. Saccades are rapid movements made to shift the point of gaze. In contrast, smooth pursuit is employed to track a moving object. Optokinesis is similar to smooth pursuit in that it is a smooth movement invoked to stabilize an image of a moving target on the retina except it is involuntarily elicited through head or body movements. Vestibular-ocular reflex, also invoked by head or body movements, stabilizes the image of a fixated object as the head or body moves relative to the object. Additionally, fixations, which stabilize the eye for higher acuity at a given point, are often of importance to the eye tracking researcher. These are all described in greater detail in [Carpenter 1988].

Head movements made during tasks may be related to the visual behaviors (such as rotation of the head to capture a new field of view) or unrelated (such as nodding). In the case of a wearable eye tracker, as the head rotates to incorporate a new field of view, the eye tracker rotates with it such that the subsequent point of gaze is determined by the new scene visible to the eye tracker's scene camera. If the camera and head do not move in exact synchrony, the movement of the camera with respect to the head must be taken into account. Furthermore, at any point in time a head-mounted eye tracker may shift on the head creating another type of camera movement that needs to be taken into account. These camera movements are much slower than a saccade.

In this paper, a technique to acquire more accurate eye movements from the final eye position data produced by the eye tracker will be discussed. It will also be shown how this technique may be used to decrease the noise in the data.

3 Theory

In order to create an algorithm that will distinguish eye movements with respect to the head from camera movements with respect to the head, differences between these two movements as "seen" by the eye tracker had to be considered. We will refer to these movements as *eye movements* and *camera movements* respectively. During eye movements, the eye rotates within the socket such that the center of the pupil with respect to the scene camera moves a greater distance than does the corneal reflection (see Figure 2). On the other hand, during a camera movement, the difference in displacement of the center of the pupil and the corneal reflection is much less noticeable (see Figure 3).

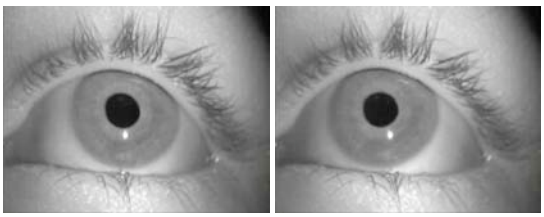

(a) Before (b) After

Figure 2: Eye images obtained before and after an eye movement.

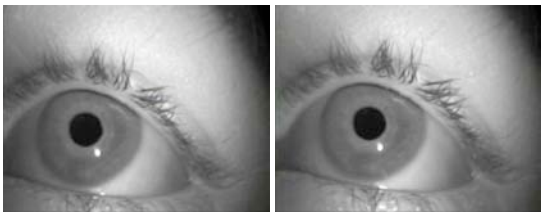

(a) Before (b) After

Figure 3: Eye images obtained before and after a camera movement.

As a camera moves parallel to features being imaged on a planar surface, the imaged features would move the same distance. This simplistic scenario does not apply to the case of a camera imaging the eye because the pupil is located within the optics of the eye and the corneal reflection is situated on the curved surface of the cornea. In reality, the center of pupil obtained by the eye tracker is the center of the virtual image of the pupil created by the optics of the human eye. As the camera is translated, light rays enter the optics of the eye at varying angles such that the virtual image of the pupil actually moves less than the amount the camera has moved. Additionally, the infrared source imaging the eye moves with the camera. This illumination will hit the first surface of the eye at different angles as its translated in front of the eye such that the corneal reflection also does not move the same amount as the camera. The virtual image of the pupil located within the eye ball moves a greater distance than does the corneal reflection. Therefore, each time the camera moves, the P-CR distance changes a small amount such that the assumed eye position does not remain stationary as it should in the absence of eye movements. Knowledge of the relationship between the pupil and corneal reflection displacements and how they differ for the two types of movements measured allows us to obtain and separate eye position and camera position data using the pupil and corneal reflection data. These eye and camera position data will be output in terms of the amount the center of the pupil has moved due to the corresponding type of movement.

Abbreviation	Data Represented
P_{track}	Pupil data output by eye tracker
P_{eye}	Pupil data from eye movements
P_{cam}	Pupil data from camera movements
CR_{track}	Corneal Reflection data output by eye tracker
CR_{eye}	Corneal Reflection data from eye movements
CR_{cam}	Corneal Reflection data from camera movements

Table 1: Variable Abbreviations.

4 Algorithm

The foundation for this camera extraction technique is two gain values: *cam_gain* and *eye_gain* (see Equations 1 and 2). These two values represent the fraction of a one unit pupil movement that the corneal reflection moves during a specific type of movement. In other words, an *eye_gain* of one-half would mean that the corneal reflection moves half the amount that the pupil moves during an eye movement. The abbreviations used in the equations in this paper are summarized in Table 1.

$$cam\_gain = \frac{\Delta CR_{cam}}{\Delta P_{cam}} \tag{1}$$

$$eye\_gain = \frac{\Delta CR_{eye}}{\Delta P_{eye}} \tag{2}$$

These gain values were calculated for multiple subjects (see Section 5) by performing a linear regression least squares fit to the pupil and corneal reflection data for each subject. In order to perform this least squares fit, the first value in the pupil and corneal reflection data arrays was subtracted from the entire corresponding array such that each array started at zero (allowing us to ignore the Δ's in Equations 1 and 2). For this reason, subjects were asked to fixate at a point in the center of their field of view at the beginning of each trial used for gain determination. The linear relationship was supported by empirical data showing that the gains remained constant across small and large movements; to further support this, the R^2 statistic for the regressions are shown in Table 2.

The actual pupil and corneal reflection data output by the eye tracker include the combined affect of both camera and eye movements (see Equations 3 and 4). The P_{eye} data are desired to determine eye position with respect to the world and the P_{cam} data are desired to map the eye position back to the scene camera.

$$P_{track} = P_{eye} + P_{cam} \tag{3}$$

$$CR_{track} = CR_{eye} + CR_{cam} \tag{4}$$

There are now four equations and four unknowns: P_{eye}, P_{cam}, CR_{eye} and CR_{cam}. To solve for P_{cam} in terms of the known variables, we start with a rearrangement of Equation 3. First, both sides of the equation are scaled by the *eye_gain* and a substitution is made using Equation 2. Then CR_{cam} is subtracted from both sides and substitutions using Equations 1 and 4 are made. The equation is rearranged to produce the final equation for P_{cam}:

$$P_{cam} = \frac{CR_{track} - P_{track} \cdot eye\_gain}{cam\_gain - eye\_gain} \tag{5}$$

These camera position data, in terms of amount the pupil has moved due to the camera movement, are a function of the pupil and corneal reflection data obtained by the eye tracker and the camera and eye gain values. The pupil position during eye movements can then be extracted, through rearrangement of Equation 3, by simply subtracting this camera position data from the pupil position data output by the eye tracker.

Now there are two separate arrays of data - one for the camera position and one for the eye position - that are both based on the pupil and corneal reflection eye tracking data and that can be altered separately. Therefore, since camera movements occur at a slower rate than eye movements, the camera position data may be smoothed to reduce the added noise introduced by the corneal reflection data. Looking at Equation 3, as P_{cam} is smoothed, the amount of noise in P_{eye} will approach the amount of noise in P_{track}. For our preliminary results presented in Section 5, we have chosen to first apply a median filter and then a gaussian filter to smooth the camera data.

5 Algorithm Testing and Results

The algorithm described in Section 4 was tested using data collected by the RIT wearable eye tracker [Babcock and Pelz 2004] and output by the ISCAN 726/426 Pupil/Corneal Reflection Analysis System, both described in Section 2. These eye tracking data were obtained by having five subjects perform three sets of trials. The trials consisted of looking through calibration points presented to them on a projection screen 3 meters in front of them. Each calibration point was an "X" character which subtended a viewing angle of approximately 0.5 degrees. A fixation slide of one centered point was displayed between each trial. A calibration slide of nine points as shown in Figure 4 was used for the first and third sets of trials.

Figure 4: Calibration slide presented during first and third sets of trials.

For the first set of trials, subjects were asked to keep their heads and the camera stationary while they looked through the nine calibration points requiring horizontal and vertical eye movements. Data from these trials were used to determine each subject's horizontal and vertical eye gain. A 14-second segment of data from one subject for this trial is shown in Figure 5. During this segment, the subject looked across the top horizontal line of calibration points in Figure 4. The difference between the amount of noise in the Pupil data versus the noise in the CR data can be seen in a zoomed in view of this plot in Figure 6.

The second set of trials were used to determine each subject's horizontal and vertical camera gains. The subjects were told to keep their head still and fixate on a single calibration point in the center

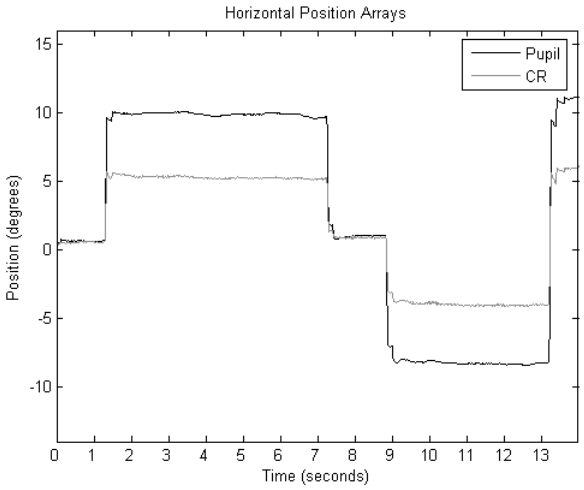

Figure 5: Sample data from first set of trials. Subject made only eye movements

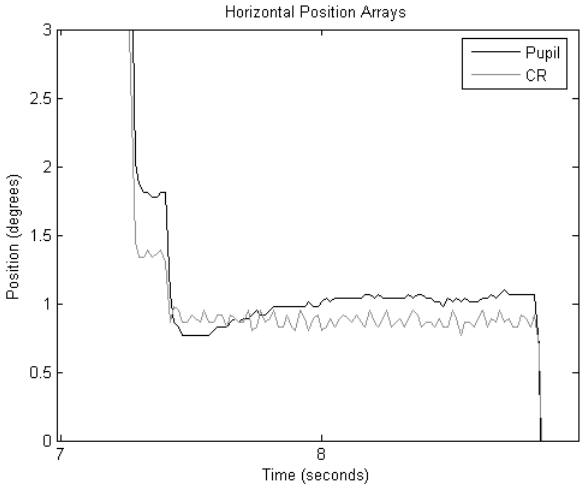

Figure 6: Close-up of eye tracker Pupil and CR data during an eye movement trial.

of the screen. With their eye and head still, subjects were asked to move the head-mounted eye tracker in both horizontal and vertical directions. Subjects moved the eye tracker a small amount such that their eye remained in the eye camera's field of view throughout the trial. A 39-second segment of data of one subject from this camera movement trial is shown in Figure 7. During this segment, the subject shifted the eye tracker back and forth on the nose making both small and large camera movements. As displayed in this figure, the pupil and corneal reflection do not move the same amount during a camera movement. In fact, the corneal reflection moves less than the pupil when the camera moves.

The last set of trials consisted of both eye and camera movements. Each subject was asked to look through the nine calibration points while moving the eye tracker around as before. A nine-second segment of Pupil and CR data from this trial is shown in Figure 8. During this segment, the subject moved the camera while looking through the center horizontal row of calibration points. Data from these trials were used to determine the success of the algorithm.

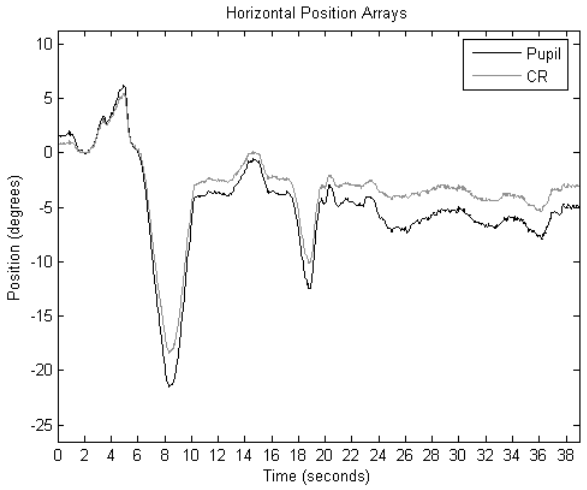

Figure 7: Sample data from second set of trials. Subject made only camera movements.

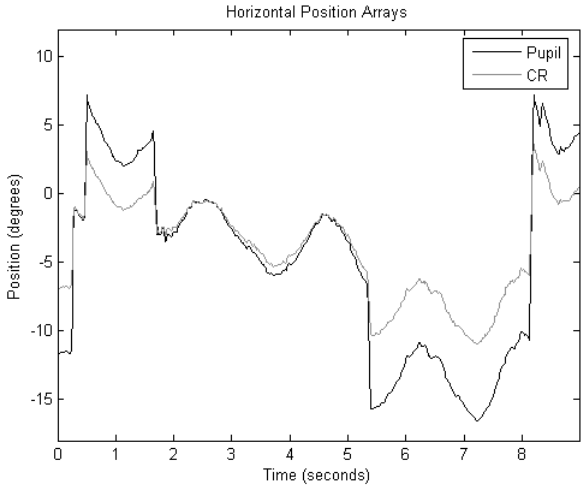

Figure 8: Sample data from third set of trials. Subject made eye and camera movements.

5.1 Gain Results

Horizontal and vertical eye gains were determined for each subject using data from the eye movement trials. As described in Section 4, the gain values were calculated using a linear regression model. Each subject's horizontal and vertical camera gains were likewise determined based on data from the camera movement trials. The resulting gain values as well as the standard deviation and average of these values across subjects is shown in Table 2. Additionally the R^2 statistic for each regression is shown. The vertical eye movement gain for subject CJL has been omitted due to poor CR vertical data during the eye movements trials.

5.2 Algorithm Results

The algorithm described in Section 4 was applied to the data collected. Equation 5 was used to calculate the horizontal and vertical

	Eye Gain		Camera Gain	
Subject	Horizontal	Vertical	Horizontal	Vertical
ABC	0.5161	0.4372	0.8143	0.9238
R^2	0.9878	0.9650	0.9768	0.9914
AEF	0.4974	0.3064	0.7936	0.8805
R^2	0.9947	0.9347	0.9667	0.9990
CJL	0.5456	–	0.8876	0.7864
R^2	0.9893	–	0.9992	0.9957
JBP	0.5277	0.3323	0.8697	0.8965
R^2	0.9975	0.9577	0.9880	0.9851
JLS	0.4862	0.3715	0.8966	0.9111
R^2	0.9993	0.9837	0.9972	0.9971
mean	0.5146	0.3618	0.8524	0.8797
σ	0.0236	0.0569	0.0458	0.0546

Table 2: Gain results for five subjects including mean and standard deviation, σ, for each gain. Top number in each row represents the gain value, below it is the corresponding R^2.

Camera position arrays using the mean gain values presented in Table 2. The Camera arrays were then smoothed with a median filter of width 7 fields followed by a gaussian filter with a standard deviation of 4 fields. The horizontal and vertical Eye arrays were then computed by subtracting the corresponding smoothed Camera array from the corresponding Pupil array (Equation 3). The results presented here are from a section of subject ABC's data. These data were chosen to show the strength of the algorithm to be used on various subjects because this subject's gain values overall were the furthest from the means.

Results of the algorithm applied to the data displayed in Figure 8 are shown in Figures 9 and 10. Figure 9 shows a section of the Pupil and CR data arrays output by the eye tracker along with the Camera and Eye arrays calculated by our algorithm. Note that the Camera array in this plot is below zero for the length of the trial. This is accurate and due to the cumulative camera displacement after the previous trials and accounts for the Eye position data remaining around 0 degrees when the subject is fixated on the center point while the P-CR array is fluctuating about -0.75 degrees (see Figure 10). Figure 10 compares our Eye position array to a scaled version of the eye tracker's P-CR array. The Eye array output by our algorithm is at the same scale as the Pupil array whereas the original P-CR array is at approximately half that scale. We chose to have our Eye array at the same scale as the Pupil array for comparison to eye trackers that only track the pupil. The P-CR array was scaled in Figure 10 in order to show a side-by-side comparison.

Aside from comparing our Eye array to the eye tracker's output, it is also important to note the usefulness of the calculated Camera array. This Camera array, whose computation is not built into the eye tracker's functionality, explains what is happening to all cameras attached to the head gear. For the RIT wearable eye tracker [Babcock and Pelz 2004], these cameras include the camera imaging the eye and the camera imaging the scene. Therefore, this Camera array will be useful in a calibration routine by describing not only the motion of the eye image but the motion of the scene image as well.

5.2.1 Noise Reduction

Figure 11 shows a zoomed in view of the comparison of our Eye array with the P-CR array for the first trial (see Figure 5). This figure also shows the comparison between the Pupil and CR arrays for the same time segment. This provides an example of how smoothing the Camera array before computing the Eye array can result in

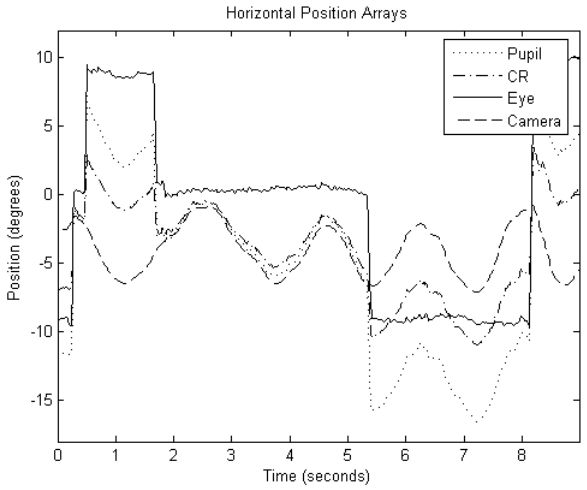

Figure 9: Sample data results from third set of trials.

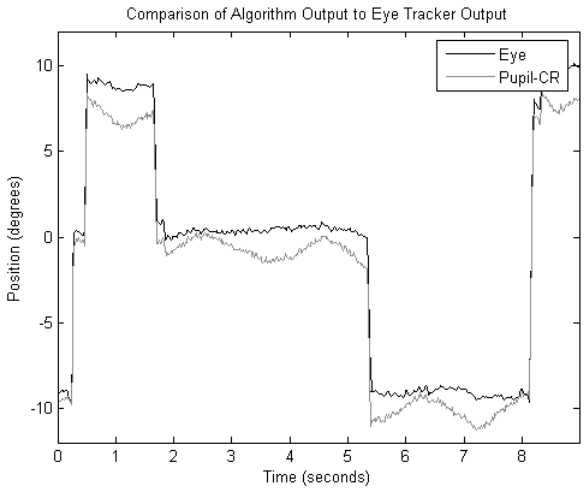

Figure 10: Sample comparison of raw eye tracker output to our algorithm output. The P-CR data was output by the eye tracker and the Eye position was output by our algorithm.

less noise artifacts in the final eye position output while maintaining small eye movements. The noise characteristics of our algorithm's Eye array closely resemble the noise in the Pupil array whereas the noise inherent in the P-CR array more closely resembles the significantly noisier CR array.

The difference between the amount of noise in the Eye and P-CR arrays is not as noticeable for the data from the third set of trials (shown in the previous section) because the difference between the noise in the Pupil and CR arrays for this trial was also not as noticeable. A side-by-side comparison of the Pupil and Corneal Reflection data versus the Eye and P-CR data during the third type of trial is shown in Figure 12.

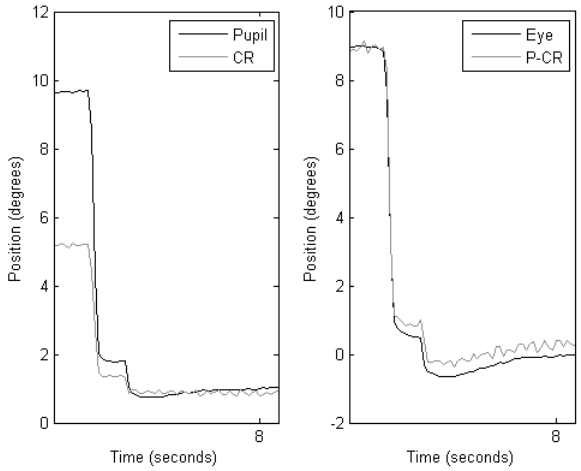

Figure 11: Comparison of noise relationship between eye tracker's Pupil and CR data to the noise relationship between algorithm's output (Eye) and eye tracker's output (P-CR) from eye movements trial.

6 Conclusion

Calculating the position of the eye tracker's camera serves multiple uses. Knowledge of the camera position at every field for which pupil and corneal reflection data are obtained allows for the correction of the Pupil array in the presence of camera motion. Additionally, since camera motion is smooth and typically much slower than eye movements, the Camera array can be smoothed such that the final output Eye array contains a comparable amount of noise to that in the Pupil array. The current P-CR technique used by video-based eye trackers produces an output that can only be as low in noise as the corneal reflection data. With our algorithm, the corneal reflection data are important to determining the camera position but do not contribute to the noise in the final output.

Aside from its increased susceptibility to noise, the P-CR technique is flawed due to the fact that the Pupil minus CR vector difference is affected by camera motion. To remove the affect of the camera movement on this vector difference, the CR array needs to be scaled before being subtracted from the Pupil array. Another scale factor is necessary in order for final array to be on the same scale as the Pupil array. These two scaling values are based on the relationships between the pupil and corneal reflection during eye and camera movements. Applying these scaling factors (deemed the eye and camera gains) to the Pupil and CR arrays allows determination of individual arrays representing the motion of the eye tracker's camera and the movement of the subject's eye (both with respect to the subject's head).

The success of this algorithm is dependent on the distance between the subject's eye and the camera imaging the eye. As the eye camera is moved further from the subject's eye (as in remote stationary eye tracking systems), the camera gain approaches the eye gain and consequently eye and camera movements become less distinguishable.

This algorithm can be customized for each subject by having the subject perform eye and camera movements, as in our trials (see Section 5), before the experimental task but this is not necessary. The average gain values may be generalized for all subjects and used as constant parameters within the algorithm. The results

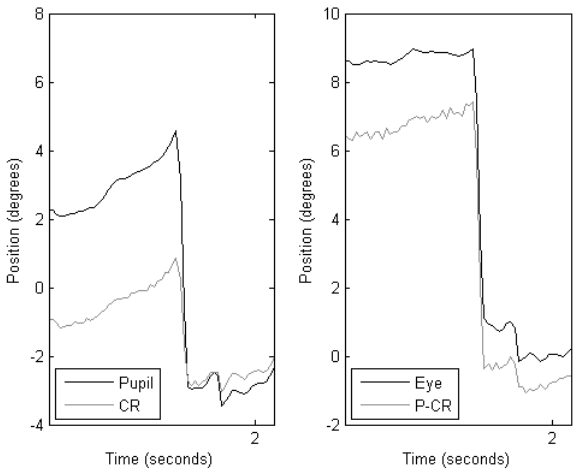

Figure 12: Comparison of noise relationship between eye tracker's Pupil and CR data to the noise relationship between algorithm's output (Eye) and eye tracker's output (P-CR) from eye and camera movements trial.

shown in this paper were produced using the average gain values. This was done to show the robustness of this algorithm to be applied to multiple subjects without acquiring further parameters during individual subject calibration. Studying more subjects to gather gain values from a larger population is planned.

The final Eye array produced by this algorithm requires a new calibration routine to produce data for point of regard (POR) in scene because it is not on the same scale as the P-CR array (which is calibrated to POR via the ISCAN Analysis System). The new calibration routine will make use of the additional information provided by the camera position data. The Camera array provides knowledge of how any camera attached to the eye tracker's head gear has moved with respect to the subject's eye. Therefore, for eye trackers using head-mounted scene cameras that may move during a task, the Camera array can be used to map the eye position back to scene image coordinates. This is important because even perfect compensation for eye camera movement will cause errors in gaze position in scene image coordinates if movement of the scene camera is not considered. This calibration routine is the next step in our goal to improve eye tracking data.

Future extensions of this work will include exploration of smoothing options to apply to the Camera array. Collection of data to investigate the noise characteristics of the system as well as the size and frequency of typical camera movements during various tasks has commenced to support this effort.

Acknowledgements

The authors would like to acknowledge Christopher Louten for his help with data collection and Mitchell Rosen for all of his valuable comments and suggestions. We appreciate the participation of our subjects. This research was funded in part by a grant from the National Science Foundation.

References

BABCOCK, J. S., AND PELZ, J. B. 2004. Building a lightweight eyetracking headgear. In *ETRA'2004: Proceedings of the Eye tracking research & applications symposium on Eye tracking research & applications*, ACM Press, New York, NY, USA, 109–114.

CARPENTER, R. H. S. 1988. *Movements of the Eyes (2nd ed.)*. Pion Limited, London.

DELABARRE, E. B. 1898. A method of recording eye-movements. *American Journal of Psychology 9*, 4, 572–574.

HENDERSON, J. M., AND FERREIRA, F. 2004. *The Interface of Language, Vision, and Action: Eye Movements and the Visual World*. Pyschology Press.

ISCAN, INC. 2001. *RK-726PCI Pupil/Corneal Reflection Tracking System*, January.

KARMALI, F., AND SHELHAMER, M. 2004. Automatic detection of camera translation in eye video recordings using multiple methods. In *Proceedings of the 26th Annual International Conference of the IEEE EMBS*, vol. 1, 1525–1528.

LAND, M., MENNIE, N., AND RUSTED, J. 1999. The roles of vision and eye movements in the control of activities of daily living. *Perception 28*, 1311–1328.

PELZ, J. B., AND CANOSA, R. 2001. Oculomotor behavior and perceptual strategies in complex tasks. *Vision Research 41*, 3587–3596.

XIE, X., SUDHAKAR, R., AND ZHUANG, H. 1998. A cascaded scheme for eye tracking and head movement compensation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans 28*, 4, 487–490.