

Predicting Destination using Head Orientation and Gaze Direction During Locomotion in VR

Jonathan Gandrud^{*} and Victoria Interrante
Department of Computer Science and Engineering, University of Minnesota




Figure 1: Eye gaze, head orientation, and head position are used to predict the participant's future path direction after a decision waypoint.

Abstract

This paper reports preliminary investigations into the extent to which future directional intention might be reliably inferred from head pose and eye gaze during locomotion. Such findings could help inform the more effective implementation of realistic detailed animation for dynamic virtual agents in interactive first-person crowd simulations in VR, as well as the design of more efficient predictive controllers for redirected walking. In three different studies, with a total of 19 participants, we placed people at the base of a T-shaped virtual hallway environment and collected head position, head orientation, and gaze direction data as they set out to perform a hidden target search task across two rooms situated at right angles to the end of the hallway. Subjects wore an nVisorST50 HMD equipped with an Arrington Research ViewPoint eye tracker; positional data were tracked using a 12-camera Vicon MX40 motion capture system. The hidden target search task was used to blind participants to the actual focus of our study, which was to gain insight into how effectively head position, head orientation and gaze direction data might predict people's eventual choice of which room to search first. Our results suggest that eye gaze data does have the potential to provide additional predictive value over the use of 6DOF head tracked data alone, despite the relatively limited field-of-view of the display we used.

Keywords: gaze direction, locomotion, virtual environments

Concepts: • Human-centered computing ~ Interaction paradigms; Virtual reality;

*e-mail: gandr027@umn.edu, interrante@cs.umn.edu.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

SAP '16, July 22-23, 2016, Anaheim, CA, USA
© 2016 ACM. ISBN 978-1-4503-4383-1/16/07 \$15.00
DOI: <http://dx.doi.org/10.1145/2931002.2931010>

1 Introduction

Head tracking is an essential component of every virtual reality system, required for the presentation of perspectively correct imagery. Eye tracking requires extra effort and cost, in the form of additional hardware and software as well as the inconvenience of a cumbersome calibration process, yet gaze information, if it were available, might offer multiple advantages in diverse VR applications. In this paper, we specifically consider the question of how gaze information might help us to infer where someone is intending to walk, particularly in the case where a decision between two alternative walking directions needs to be made.

One important potential application for such information is in the design of more effective controllers for redirected walking. Redirected walking was introduced by Razzaque *et al.* [2001] as a method for enabling users to experience the illusion of physically walking through a larger virtual space within the confines of a smaller physical area. Redirection is achieved by judiciously introducing subtle discrepancies into the typical 1-1 mapping between a user's actual locomotor actions in the physical space and the observed effect of those actions in the virtual world. For example, a very simple redirection controller might allow a user to move forward at a speed that is 7x faster in the virtual world than in the real world [Interrante *et al.* 2007] or to turn around by 360° in the virtual world while turning only 180° in reality [Williams *et al.* 2007]. While classical redirection controllers rely on instantaneous measures of attributes such as head position, head orientation, and velocity to inform the implementation of simple heuristics such as "steer-to-center" or "steer-to-orbit" [Hodgson and Bachman 2013], many in the research community have recognized the enhanced potential for more effective redirection when a person's future locomotion path is either known in advance or can be inferred [e.g. Nitzche *et al.* 2004, Nescher and Kunz 2012, Goldfeather and Interrante 2012, Zmuda *et al.* 2013, Zank and Kunz 2015].

A second important application for eye gaze information is in the realistic and effective animation of autonomous virtual agents. Our typical experience of the everyday world is as a shared space, so much so that architects routinely incorporate human figures into their drawings of planned spaces, both to provide a sense of scale and to "bring the drawing to life" [Anderson 2002]. Beyond

the numerous challenges in automating realistic locomotor behavior at the macro level [Karamouzas *et al.* 2014], correctly modeling detailed behavior such as gaze direction can be important not only to the visual realism of the avatar animation but also to the potential effectiveness of the avatar’s non-verbal communication [Ruhland *et al.* 2014].

2 Related Work

Real-world studies of gaze behavior during locomotion have mainly focused on the role of gaze in maintaining postural and locomotor stability. Early studies found that the head is generally oriented in the direction of travel, that head reorientation occurs in near temporal coincidence with gaze realignments when subjects turned by 30°–60°, and that, when the route is obstacle-free, gaze is primarily directed either straight ahead or at the goal [Hollands *et al.* 2002, Patla 2004], results that echo observations of eye gaze behavior while driving [Land and Lee 1994]. However, later work began to raise some concerns about the extent to which these earlier findings might have been unduly influenced by the design of the experimental setup [Pelz *et al.* 2009], and more recent work suggests that gaze direction in fact anticipates head orientation, which in turn anticipates the reorientation of other body segments [Bernardin *et al.* 2012, Wilkie *et al.* 2010]. None of these studies, however, have involved a decision-making component.

Classical research in gaze behavior during decision-making [Shimojo *et al.* 2003, Glaholt and Reingold 2009] suggests to us that by observing people’s eye gaze as they walk, we may be able to infer their future locomotor intent significantly in advance of the point at which that locomotor decision needs to be instantiated. As summarized in a recent review [Orquin and Loose 2013], when presented with an alternative force-choice task in an image-viewing setting, people reliably look longer and more often at the alternative they eventually choose. In addition, there is a slight bias towards eventually choosing the alternative they look at first, and peoples’ last fixation before their decision tends to be on the chosen alternative. Somewhat surprisingly, given that people are of course free to look wherever they want, studies conducted in natural (as opposed to picture-viewing) contexts have found that most gaze is goal-directed and only moderately affected by salience [Tatler *et al.* 2011].

While there has been an abundance of prior research in gaze behavior, little of that work has focused on assessing the potential of using gaze to dynamically predict where people intend to walk while proceeding towards a potential wayfinding decision point. Most relevant to the implementation of appropriate gaze behavior for signaling the future directional intent of autonomous agents, Nummenmaa *et al.* [2009] found that when seated observers were asked to watch a video of an approaching avatar, displayed on a 20” computer monitor, and to press a key to indicate which way they would choose to pass, while the character either looked constantly to one side, or approached with a straightforward gaze that suddenly shifted to one side, subjects consistently chose to pass on the side the avatar was looking away from. This suggests that people can infer future locomotor intent from the gaze behavior of animated agents, at least when implicitly prompted to do so, and underlines the potential importance of implementing correct agent gaze behavior in situations where collision avoidance is a concern.

Most closely related to the work we describe here is a very recent paper by Zank and Kunz [2016], which focuses on the potential

for using gaze direction in a redirected walking controller. Following a similar but slightly different experimental design as ours, they find that predictive models based on gaze direction typically allow the earlier correct prediction of a user’s eventual turn direction than models that rely on positional data alone, though at a small cost of more frequent late incorrect predictions. Our work complements these findings by providing additional data obtained under slightly different conditions, contributing to the development of a comprehensive, integrated understanding of the potential advantages of collecting eye gaze data for use in conjunction with head position and head orientation information to infer a person’s most likely future heading direction after reaching a distant decision point.

3 Our Experiment

In this study, we measure the head pose, head orientation and eye gaze direction of participants as they navigate down the length of a virtual hallway (figures 1 & 2). The participants are directed to approach a doorway at the end of the virtual hallway and, as a distractor task, asked to count the number of circles that appear in a virtual picture frame that hangs on the wall within the rooms beyond those doors. Our true intention is to determine which of the two virtual doorways the participant will visit first, using the head tracking and eye gaze direction data. We performed two pilot studies, followed by a final experiment.




Figure 2. (left) Downward view of the virtual hallway. Participant starts at the narrow end and walks over to one of the doorways. (middle) View of the virtual hallway, as seen by the participant. (right) View inside the right doorway.

3.1 Participants

The first pilot study had six participants (5 male and 1 female), ages ranging from their 20’s to their 60’s. The second pilot study had six participants (4 male and 2 female), in their 20’s and 30’s. This study had seven participants (all male), in their 20’s and 30’s. Participants were recruited from our University and local community, and were compensated with \$10 gift cards.

3.2 VR Equipment & Software

Participants wore an nVisorST50 optical see through head-mounted display equipped with an Arrington Research ViewPoint eye tracker (see figure 2). During the experiment, the participant’s view of the real world was blocked so that they saw only the virtual scene. The display field of view was 40° horizontal, and separate SXGA (1280 x 1024) images were presented to each eye. Positional data were tracked using a 12-camera Vicon MX40+ motion capture system. Head position (XY coordinates) and orientation (rotational yaw) were recorded using the Vicon DataStream SDK 1.4 (x64). The eye tracking system was calibrated using the ViewPoint EyeTracker software from Arrington Research, and eye tracking data were collected using a C-language DLL provided by Arrington. The Unreal Engine 4

was used for creation of the virtual hallway and to run the virtual reality simulation. We wrote a C-language DLL wrapper to retrieve frames of eye tracking data from the Arrington DLL. We also wrote a C# program to retrieve frames of head tracking data from the Vicon motion tracking system. For each frame of head tracking data, the C# program also retrieved a data frame of eye tracking data. Head tracking and eye tracking data were thereby recovered synchronously and saved to a log.




Figure 3. Our nVisor ST50 Head-Mounted Display with Arrington eye tracking camera and retroreflective motion tracking dots.

3.3 Method Overview & Hypothesis

We collected head position, head orientation and gaze direction data as participants walked from the narrow end of a T-shaped virtual hallway towards one of two virtual doorways at the end of the hallway (see figure 2). The goal of our experiment was to assess the extent to which the information provided by each of these data streams forecasts the participant’s eventual choice of which room to enter first. To blind participants to the true purpose of our experiment, we devised a cover story which was that we were seeking to understand where people look when they search for targets in an image, under natural viewing conditions.

At the start of each trial, participants stood at the end of a long hallway and their task was to enter each of the rooms at the end of the hallway and report the number of circles in each image (see figure 4). Between every trial, the experimenter changed the images in the virtual picture frames with a keystroke. We performed six trials for each of the seven participants in this study.



Figure 4. Examples of the photos that appeared in the virtual picture frames. The participant is asked to count the circles in each photo (2 circles on left, 1 circle on right).

Prior to this reported experiment, we conducted two pilot studies, with six participants each. In the first pilot study, we recorded binary information about the eye gaze direction (i.e. whether it was on the left side or the right side of the display’s midline). While the results were encouraging, we recognized that more complete data would help us make more nuanced predictions of the destination, so we elected to measure the actual magnitude of eye gaze direction (in radians) in later studies. In the second pilot study, we collected six trials of data from each of six different participants. From this study we determined that it would be necessary to account for slippage in the head mounted display

(HMD) as the participant walks around. We then experimented with various mechanical and procedural measures to try to prevent HMD slippage but determined that such measures could not be used if they impeded participants’ perceived freedom of head movement. Therefore, for our final experiment, we introduced improvements to the experimental procedure both to reduce the likelihood of slippage between the calibration and experimentation phases (e.g. conducting the full eye tracking calibration process from a standing position, so that the participant does not need to stand up from a chair and risk moving the HMD) and to allow us to estimate and correct for any slippage that does occur. Preventing slippage in the HMD is important to ensuring the accuracy of the eye tracking measurements.

3.3.1 Experimental Protocol

The experimenter helps the participant put on the HMD securely (but to the participant’s comfort) and position the eye tracking camera and IR light (see figure 3) so that the eye appears on-camera (as viewed using the ViewPoint EyeTracker software that comes with the Arrington eye tracker). The participant is asked to move their right eye to the 4 compass directions, so that the experimenter can observe whether tracking is ever lost. The camera position is adjusted, if necessary. The process of adjusting the camera is frequently time-consuming. The experimenter then helps the participant prepare for eye tracking calibration by leading the participant to a computer screen which is positioned at head height. The HMD’s eye cover plate is off at this time. The experimenter instructs the participant to look through the HMD, and to center the right eye on the computer screen. The participant is told to position the right eye so that the eye does not see a pair of red tape strips affixed to the right and left edges of the computer monitor, but DOES see a pair of green vertical lines displayed as a bitmap image along the inner edges of the monitor, on the left and right sides. This procedure is used in order to center the participant on the computer screen, enabling definition of a coordinate system to track eye coordinates in the range [0, 1]. This [0, 1] range was verified in pilot testing by the experimenter, and in later data analysis we reviewed a plot of each participant’s eye movements to verify that the movements were within the [0, 1] range. Eye tracker calibration is done using the Arrington software. The participant views a sequence of green squares that appear at all vertices of a grid, all across the computer monitor. After calibration is complete, the participant leaves the computer screen. The computer screen is not used at any further points of the experiment. The HMD and eye tracker devices are connected to our computer systems using a bundle of cables. The cables are long enough to afford the participant full access to all points of the virtual hallway. The participant is told to hold the cable bundle in one hand, so that the cable bundle does not tug on the head during walking. The participant is led to the starting point, which is demarcated with tape on the floor of the physical lab space. This tape in the physical lab space corresponds to the far end of the virtual T-shaped hallway. The virtual hallway model is launched using the Unreal Engine software. The virtual scene is now visible to the participant on the HMD. The physical cover plate is affixed to the HMD so that the participant no longer sees the physical lab in front of the HMD, and can focus on viewing the virtual hallway model.

From this position in the virtual hallway, the participant can see three sconces hanging on the virtual wall (one on the left wall, one in the middle, and one the right wall – see figure 2). Before each

trial, we collect a data log while the participant gazes at the middle sconce, and use the median head yaw (i.e. the h term in Equation 3) from these data logs to compute a *display slippage correction factor*, γ , to compensate for any slippage that might have occurred in the HMD during the trials since eye tracking calibration (see next section for details). We also collect data logs while the participant gazes at the other two sconces, in order to avoid biasing the participant. The order in which they are asked to gaze at the three sconces is randomized.

For each trial:

- a) Experimenter collects a separate data log while the participant gazes at each of the middle, left and right sconces, in a randomized order.
- b) Experimenter initiates the data collection program and invites the participant to start walking when ready.
- c) Participant walks over to a doorway and tells the experimenter how many circles appear in the picture. The picture only becomes visible after the participant enters the doorway, thereby forcing the participant choose a door.
- d) Participant walks to the other doorway and verbally reports the number of circles.
- e) Experimenter stops the data collection program.
- f) Participant returns to the starting point.

3.4 Data Analysis

We wish to predict which of the two virtual doors the participant will approach first in the virtual hallway. We consider the following inputs:

- 1) Head direction vector
- 2) Gaze direction vector
- 3) Participant's current XY position in the hallway
- 4) XY position of each doorway's midpoint

Head direction is given as a yaw rotation, in radians, by the Vicon tracking system. The gaze direction is defined as: head direction + eye gaze direction relative to head + γ . γ is the *display correction factor* that compensates for HMD slippage between trials. The Y axis of our Vicon tracking system runs along the length of the virtual hallway. We measure the participant's walking progress using the Y axis coordinate. We partition the Y axis into segments of 100 mm and report the proportions of predictions within each segment (see figures 7, 8 & 9). The dependent variable is a prediction of the destination: {left door, right door, indeterminate}. For the data samples within each segment, we predict the destination door using each of the following three methods.

- 1) Head Direction Method
- 2) Gaze Direction Method
- 3) Position Method (relative to virtual hallway midline)

Head direction is used to predict the outcome as follows: We draw a vector from the participant's current position to the right door, and a vector from the current position to the left door. We find the angles β_1 and β_2 , between the head direction and these other two directions, as shown in figure 5. The smaller angle indicates the prediction. In figure 5, we predict the right door, since $\beta_2 < \beta_1$. An indeterminate prediction occurs if the difference between β_1 and β_2 is less than a threshold (see the Data Cleanup section of this paper for details).

Position is used to predict the outcome as follows: If the participant is to the right of the midline that divides the virtual hallway, then we predict the right door. If the participant is to the left of the midline, then we predict the left door. Otherwise we issue a prediction of 'indeterminate'. The midline is parallel to the Y axis, and slightly offset from the Y axis.




Figure 5 (left). Geometry that is used by the Head Direction and Gaze Direction prediction methods.

Figure 6 (right). Trigonometry that is used to convert the eye tracking measurement, a , to an angle, α , (in radians).

Figure 6 illustrates the conversion from eye tracking coordinates, a , in a range $[0, 1.0]$ to an angle, α , in radians. The *display slippage correction factor* is computed for each trial. The *average display correction factor* across the 42 trials is small: -0.00721326 radians. The PC screen (shown as the horizontal line in figure 6) is present during the one-time calibration process, but not thereafter. We use the one-time measurements taken from the PC screen, along with the following trigonometry, to convert the eye tracking measurement into an angle.

$$d = \frac{z}{\tan(\theta)} \quad (1)$$

$$\alpha = \text{atan}\left(\frac{a}{d}\right) \quad (2)$$

$$\gamma = -\frac{\pi}{2} - h \quad (3)$$

$$\bar{\alpha} = \alpha + \gamma \quad (4)$$

In these equations: h is the median head rotation (yaw) while the participant views the middle sconce, before the trial begins; d is the distance between the participant's right eye and the computer screen during calibration; α is the angle (in radians) of the right eye's gaze; z is half the width of the computer monitor screen (in mm); θ is half of the horizontal field of view of the HMD (as found on the nVisor data sheet); the $-\pi/2$ term corresponds to the axis of walking (i.e. the Y axis), in the coordinates of our Vicon tracking system; γ is the correction factor that we add to the participant's gaze direction to account for slippage in the HMD; a is the gaze space coordinate read directly from the Arrington software, in the range $[0, 1.0]$; and $\bar{\alpha}$ is the eye angle, corrected to account for any slippage of the HMD on the participant's head, which may have occurred subsequent to eye tracking calibration. Head direction (yaw) is provided directly by the Vicon tracking system (in radians). Gaze direction, in world coordinates, is the sum of 1) head direction 2) the *display slip correction factor*, γ , and 3) the gaze direction relative to the head, α .

3.5 Data cleanup & Validation

Before analyzing our data, we subjected it to a careful clean-up process. First, we filtered the data stream for outliers due to glitches in tracking. We detected bad values in the head direction data by looking for a sharp gradient in the data stream. Through this process we removed 9 data points (out of a total set of over 25,000). Additionally, we determined a noise threshold in the eye

tracking data of 0.01 radians. We measured this noise threshold by having the experimenter stare straight down the hallway for a few seconds, without moving, and collecting a data log. This noise threshold accounts for the “indeterminate” outcome for the Head Direction and Gaze Direction prediction methods. If the difference between β_1 and β_2 is less than that noise threshold, we issue an indeterminate prediction.

We wrote software to replay movies of experimental trials and verify the gaze direction by drawing the gaze point as a moving sphere in the scene. This process helped us to understand participants’ eye movements. We also reviewed plots of the raw eye tracking data, verifying that the values fell within the range [0, 1]. Further, we reviewed plots of the head directions and gaze directions for each of the 42 trials, ensuring that the radian measurements looked reasonable. Additionally, the experimenter collected 7 validation trials on himself, each with a known structure. For example, in one validation trial, he intentionally pointed his head at the left door for the entire length of the hallway while fixating his eyes at the right door. He verified that the Head Direction prediction method predicted the left door, whereas the Gaze Direction prediction method predicted the right door at all points along the walking path. The predictions of each of these 7 validation trials were as expected.

4 Results

We report the following views of the data:

- 1) Predictions of the destination door, reported for every 100 mm segment of the Y axis (see Figures 7, 8 & 9)
- 2) Plots that show the walking path of selected trials with erroneous predictions, where each data point is color-coded according to the destination that the data point predicts (see figures 10, 11 & 12)
- 3) Head and gaze directions averaged over each 100 mm segment of the Y axis (see figures 13 & 14)
- 4) Plots that show the averaged XY positions across all participants for trials that ended at the left virtual doorway and for trials that ended at the right virtual doorway (see figure 15)
- 5) Plots that show the averaged head direction and averaged gaze direction for each 100 mm segment along the Y axis (see figures 16 & 17)

5 Discussion

The plots in figures 7, 8 and 9 depict the mixture of prediction outcomes using each of the 1) head direction, 2) gaze direction and 3) position methods. We stop making predictions at $Y = -1700$ mm because the doorways lie just beyond that point. For each prediction method, the average predictions improve as the trial proceeds (i.e. as the participant advances from 2500 mm to -1700 mm). Notice that there are some incorrect predictions for each of the methods, even in the final bucket. Figure 10 shows the two trials that account for these incorrect predictions in the Head Direction prediction method. The final bucket (which occurs between -1600 mm and -1700 mm, just above the “ -1.7 m” tic mark) is shown in figure 10 to contain numerous data points that predict the right door (and are therefore colored green), whereas we would expect these points to predict the left door (and be colored blue). These two trials were the only trials that contained incorrect head direction predictions in the last segment. In the




Figure 7. The predictions for each 100 mm segment of the Y axis, using the Head Direction prediction method. Walking proceeds from 2500 mm to -1700 mm.

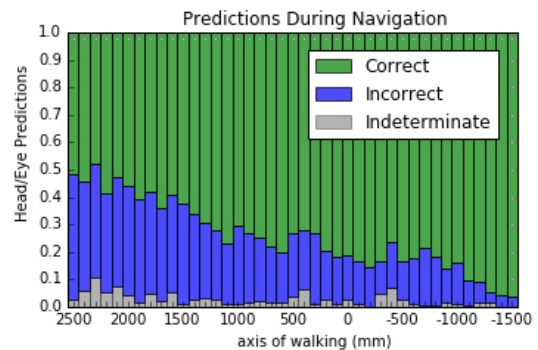


Figure 8. The predictions for each 100 mm segment of the Y axis, using the Gaze Direction prediction method.

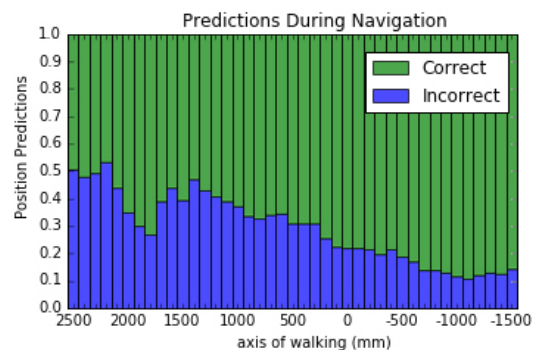


Figure 9. The predictions for each 100 mm segment of the Y axis, using the Position prediction method.

other 40 trials, the last bucket was fully predictive of the left door. Figure 11 shows two trials where the gaze direction prediction method yielded incorrect predictions at a late stage. Figure 12 shows two trials where the position prediction method yielded late incorrect predictions.

The plots in figures 13 and 14 depict the overall trends in head direction and gaze direction. Figure 13 shows head directions (green) and gaze directions (blue) for trials that ended at the right virtual doorway. Figure 14 shows the same, but for trials that ended at the left virtual doorway. The horizontal axis corresponds to 100mm segments of the Y Axis (the axis of walking). The vertical axis corresponds to the average direction of all data points within that segment. The plots for the 30 trials that ended in the right doorway show that, on average, within a segment, the gaze direction turns toward the destination door with a greater




Figure 10. Experimental trials for two different participants, where the Head Direction method results in a late incorrect prediction. These trials are the only two trials in which there were incorrect predictions in the final 100 mm segment, which occupies $[-1.6\text{m to } -1.7\text{m}]$ along the Y axis (indicated by the tic marks). Walking proceeds from Start to P2 (left door).




Figure 11. Experimental trials for two different participants, where the Gaze Direction method produces notable incorrect predictions. Walking proceeds from Start to P1 (right door).




Figure 12. Experimental trials for two different participants, where the Position Method produces a late incorrect prediction. Walking proceeds from Start to P1 (right door).

magnitude the head direction. The plots for the 12 trials that ended in the left doorway show somewhat more variability, but tell the same story. The participant walks from $Y = 2500\text{ mm}$ (on the right side of figures 13 & 14) and stops walking just beyond $Y = -1700\text{ mm}$ (on the left side of figures 13 & 14). This effect is further illustrated in figures 16 and 17, which diagrammatically show the head and gaze directions at 100 mm segments along the walking path. Notice the difference between the averaged head directions and gaze directions. The gaze directions tend to point to the destination door with a greater magnitude at an earlier point

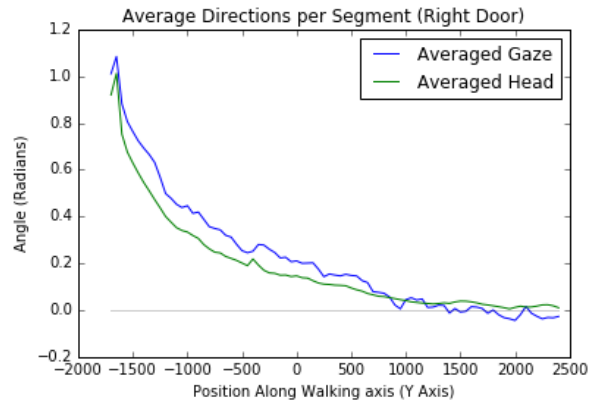


Figure 13. Head and gaze directions for each segment of 100 mm along the Y Axis, averaged over the 30 total trials where the participant walked to the right door first. Walking starts slightly above 2500 mm (on the right) and ends slightly below -1700 mm (on the left).

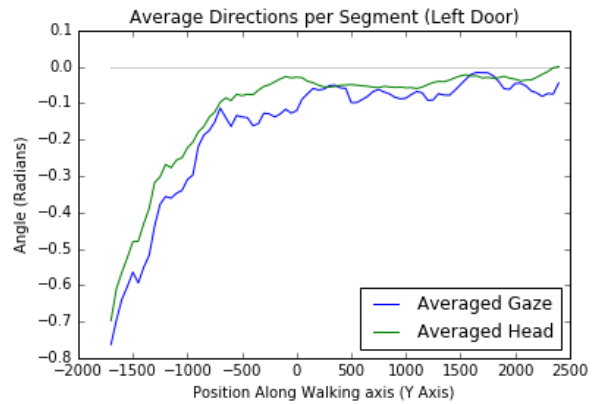


Figure 14. Head and gaze directions for each segment of 100 mm along the Y Axis, averaged over the 12 total trials where the participant walked to the left door first. Walking starts slightly above 2500 mm (on the right) and ends slightly below -1700 mm (on the left).

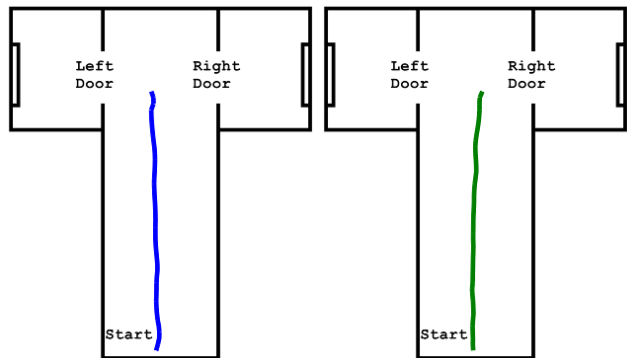


Figure 15. Averaged XY positions for the trials that ended at the left and right doors, respectively.

than the head directions. This effect is especially pronounced in the 30 trials that end at the right door, likely due to the greater quantity of data (compared to trials that ended at the left door, of which there were only 12). The gaze directions “fan-out” more noticeably than the head directions.

The diagram on the left of figure 15 shows the averaged XY positions of the participant for the 12 trials where the participant chose the left door. The diagram on the right shows the average XY positions for the 30 trials where the participant chose the right door. These averaged plots begin at the start of data collection and run until -1700 mm along the Y axis, just before the participant reaches the doorways.




Figure 16. The average head directions and gaze directions are drawn at each 100 mm segment along the walking path, up to the -1700 mm mark. Directions are averaged over the 30 trials that end at the right door.




Figure 17. The average head directions and gaze directions are drawn at each 100 mm segment along the walking path, up to the -1700 mm mark. Directions are averaged over the 12 trials that end at the left door.




Figure 18 (left). Logistic Regression model for gaze data.

Figure 19 (right). Logistic Regression model for head data.

Using a 3-way ANOVA (4 subject x 42 distance interval x 2 turn directions) to analyze the gaze data of the subset of subjects who chose each door at least some of the time, we found a statistically significant main effect on gaze angle of the ultimate door choice: $\{F(1,795) = 451.72, p < 0.01\}$. Considering a normalized gaze direction of 0 = pointing directly at the left door and 1 = pointing directly at the right door, the population marginal mean of gaze direction was 0.3706, $se = 0.0088$ when the left door was chosen and 0.6361, $se = 0.0088$ when the right door was chosen. This

indicates that relative gaze direction (considered in aggregate over all of the 42 distance intervals) has the potential to be diagnostic of door choice.

Similarly, using a 3-way ANOVA (4 subject x 42 distance interval x 2 turn directions) to analyze the head orientation data of the four subjects who chose each door at least some of the time, we found a statistically significant main effect on head orientation of the ultimate door choice: $\{F(1,795) = 783.03, p < 0.01\}$. Considering a normalized head direction of 0 = pointing directly at the left door and 1 = pointing directly at the right door, the population marginal mean of head direction was 0.3965, $se = 0.0060$ when the left door was chosen and 0.6362, $se = 0.0060$ when the right door was chosen. This indicates that relative head direction (considered in aggregate over all 42 distance intervals) has the potential to be diagnostic of door choice.

The population marginal means of relative gaze orientation over the 30 trials in which people eventually chose the right door were significantly different between the start of the trial (where the mean gaze orientation was 0.5054, $se = 0.0421$) and at distances beyond 2.7m from the start (1.4m from the end, where the mean gaze orientation was 0.7555, $se = 0.0421$). Likewise, the population marginal means of relative head orientation were significantly different between the start of the trial (where the mean head orientation was 0.5375, $se = 0.0257$) and at distances beyond 2.8m from the start (1.3m from the end, where the mean head orientation was 0.6957, $se = 0.0257$). We observe that gaze orientation turned significantly towards the right at a slightly earlier point along the hallway than head orientation did, suggesting that gaze angle has the potential to be a slightly earlier predictor of eventual door choice than head angle. In addition, we note that the population marginal mean of gaze orientation was slightly closer to 1 at the end of these trials (0.8778) than the population marginal mean of head orientation was (0.8441), suggesting that as people start to turn their head to the right in anticipation of choosing to enter the right door, they also turn their eyes even further in that same direction, on average.

Across the trials where people eventually chose the left door, there were no significant differences in the population marginal means of relative gaze orientation between any points along the hallway; however, the population marginal means of relative head orientation were significantly different between the start of the trial (where the mean head orientation was 0.5060, $se = 0.0342$) and at distances beyond 3.7m (0.4m from the end, where the mean head orientation was 0.3149, $se = 0.0342$). This supports the notion that head orientation may be a more reliable predictor of the eventual direction of travel than gaze orientation in some cases.

Finally, we used logistic regression to fit a predictor $F_g(x)$ to the relative gaze orientation data (figure 18) and a predictor $F_h(x)$ to the relative head orientation data (figure 19). From these plots, which show the relative gaze or head orientation along the horizontal axis and the probability of choosing the right vs left door along the vertical axis, with the actual data plotted in blue and the fitted function plotted in red, we can see that there is a significantly larger range over which the gaze data cannot yield a reliable prediction, relative to the head orientation data, and we also see that the gaze data cannot reliably predict a choice of the left door. The broader distribution of blue circles across the range from 0 to 1 in figure 18 relative to figure 19 shows how the gaze

direction tended to extend more nearly towards the doors than did the head direction, and the fact that this broad distribution occurs for both eventual door choices reflects the tendency, in at least some cases, for people to fixate each of the doors before making a final decision.

Conclusions

Overall, we obtained a greater proportion of correct predictions (81.66% vs. 76.19% for the eventual choice of left door and 80.71% vs. 77.46% for the eventual choice of right door) using head orientation than gaze orientation. At the same time, over all of the trials that ended in a choice of the left door, the average relative gaze orientation was closer than the average relative head orientation to the left door (0.3502 vs. 0.3811), and the same was true for the trials that ended in a choice of the right door (0.6808 for gaze vs 0.6448 for head). This suggests that both head orientation and gaze orientation have the potential to be useful in predicting a person's future direction of locomotion. Devising a more robust predictor that incorporates both of these sources of information could be an important direction for future work.

Acknowledgements

This work was supported by the National Science Foundation through grant CRI-1305401 Virtual Reality Infrastructure and Technology Development to Support Architectural Education and Basic Research in Immersive Design, Embodied Interaction, Spatial Cognition. The first of the initial pilot studies was conducted in partnership with Deepa Dongapure. We are grateful to Peng Liu for his assistance with the environment modeling and body tracking, and to all of our participants for their patient and dedicated efforts.

References

- ANDERSON, A.T. 2002. On the human figure in architectural representation. *Journal of Architecture Education*, 55, 4, 238–246.
- BERNADIN, D., KADONE, H., BENNEQUIN, D., SUGAR, T., ZAOU, M., BERTHOZ, A. 2012. Gaze anticipation during human locomotion. *Experimental Brain Research*, 223, 1, 65–78.
- GLAHOLT, M.G. AND REINGOLD, E.M. 2009. The time course of gaze bias in visual decision tasks. *Visual Cognition*, 17, 8, 1228–1243.
- GOLDFEATHER, J. AND INTERRANTE, V. 2012. Adaptive redirected walking in a virtual world. *IEEE Virtual Reality Workshop on Perceptual Illusions in Virtual Environments*, 17–20.
- HODGSON, E. AND BACHMANN, E. 2013. Comparing four approaches to generalized redirected walking: simulation and live user data. *IEEE Transactions on Visualization and Computer Graphics*, 19, 4, 634–643.
- HOLLANDS, M.A., PATLA, A.E., VICKERS, J.N. 2002 “Look where you’re going!”: gaze behavior associated with maintaining and changing the direction of locomotion. *Exp. Brain Research*, 142, 2, 221–230.
- INTERRANTE, V., RIES, B., ANDERSON, L. 2007. Seven league boots: a new metaphor for augmented locomotion through moderately large scale immersive virtual environments. *IEEE Symposium on 3D User Interfaces*, 167–170.
- KARAMOUZAS, I., SKINNER, R. AND GUY, S.J. 2014. Universal power law governing pedestrian interactions. *Physical Review Letters*, 113, 23, 238701.
- LAND, M.F. AND LEE, D.N. 1994. Where we look when we steer. *Nature*, 369, 742–744.
- NESSCHER, T. AND KUNZ, A. 2012. Analysis of short term path prediction of human locomotion for augmented and virtual reality applications. *International Conference on Cyberworlds*, 15–22.
- NITZSCHE, N., HANEBECK, U.D., SCHMIDT, G. 2004 Motion compression for telepresent walking in large target environments. *Presence: Teleoperators and Virtual Environments*, 13, 1, 44–60.
- NUMMENMAA, L., HYÖNÄ J., HIETANEN J.K. 2009. I’ll walk this way: eyes reveal the direction of locomotion and make passersby look and go the other way. *Psychological Science*, 20, 12, 1454–1458.
- ORQUIN, J.L. AND LOOSE, S.M. 2013. Attention and choice: a review on eye movements in decision making. *Acta Psychologica*, 144, 1, 190–206.
- PATLA, A. 2004. Gaze behaviors during adaptive human locomotion: insights into how vision is used to regulate locomotion. Chapter 17 in *Optic Flow and Beyond*, Vaina, L.M., Beardsley, S.A., Rushton, S.K., eds., Springer, 383–399.
- PELZ, J., PURINGTON, J. HERBERT, A. 2009. Travel gaze? Re-examining gaze behavior during locomotion. *Journal of Vision*, 9, 8, 422.
- RAZZAQUE, S., KOHN, Z., WHITTON, M.C. 2001 Redirected Walking. *Eurographics (short paper presentation)*.
- RUHLAND, K., ANDRIST, S., BADLER, J., PETERS, C., BADLER, N., GLEICHER, M., MUTLU, B., McDONNELL, R. 2014. Look me in the eyes: a survey of eye and gaze animation for virtual agents and artificial systems. *Eurographics State of the Art Report*, 69–91.
- SHIMOJO, S., SIMION, C., SHIMOJO, E., SCHEIER, C. 2003. Gaze bias both reflects and influences preference. *Nature Neurosci.*, 6, 12, 1317–1322.
- TATLER, B., HAYHOE, M.M., LAND, M.F., BALLARD D.H. 2011. Eye guidance in natural vision: reinterpreting salience. *Journal of Vision*, 11, 5, 5, 1–23.
- WILKIE, R.M., KOUNTOURIOTIS, G.K., MERAT, N., WANN, J.P. 2010. Using vision to control locomotion: looking where you want to go. *Experimental Brain Research*, 204, 4, 539–547.
- WILLIAMS, B., NARASIMHAM, G., RUMP, B., MCNAMARA, T.P., CARR, T.H., RIESER, J.J., BODENHEIMER, B. 2007. Exploring large virtual environments with an HMD when physical space is limited. *ACM/SIGGRAPH Symposium on Applied Perception in Graphics and Visualization*, 41–48.
- ZANK, M. AND KUNZ A. 2015. Using locomotion models for estimating walking targets in immersive virtual environments. *International Conference on Cyberworlds*, 229–236.
- ZANK, M. AND KUNZ A. 2016. Eye tracking for locomotion prediction in redirected walking. *IEEE Symposium on 3D User Interfaces*, 49–58.
- ZMUDA, M.A., WONER, J.L., BACHMANN, E.R., HODGSON, E. 2013. Optimizing constrained-environment redirected walking instructions using search techniques. *IEEE Transactions on Visualization and Computer Graphics*, 19, 11, 1872–1884.