Comparing VE Locomotion Interfaces

Mary C. Whitton*
University of North Carolina at Chapel Hill

Joseph V. Cohn†
U.S. Naval Research Laboratory

Jeff Feasel
University of North Carolina at Chapel Hill

Paul Zimmons
University of North Carolina at Chapel Hill

Sharif Razzaque
University of North Carolina at Chapel Hill

Sarah J. Poulton‡
Hollins University
Roanoke, VA

Brandi McLeod+
University of Maryland at College Park

Frederick P. Brooks, Jr.
University of North Carolina at Chapel Hill

ABSTRACT

To compare and evaluate locomotion interfaces for users who are (virtually) moving on foot in VEs, we performed a study to characterize task behavior and task performance with different visual and locomotion interfaces. In both a computer-generated environment and a corresponding real environment, study participants walked to targets on walls and stopped as close to them as they could without making contact.

In each of five experimental conditions participants used a combination of one of three locomotion interfaces (really walking, walking-in-place, and joystick flying), and one of three visual conditions (head-mounted display, unrestricted natural vision, or field-of-view-restricted natural vision). We identified metrics and collected data that captured task performance and the underlying kinematics of the task.

Our results show: 1) Over 95% of the variance in simple motion paths is captured in three critical values: peak velocity; when, in the course of a motion, the peak velocity occurs; and peak deceleration. 2) Correlations of those critical value data for the conditions taken pairwise suggest a coarse ordering of locomotion interfaces by "naturalness." 3) Task performance varies with interface condition, but correlations of that value for conditions taken pairwise do not cluster by naturalness. 4) The perceptual variable, τ (also known as the time-to-contact) calculated at the point of peak deceleration has higher correlation with task performance than τ calculated at peak velocity.

CR Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality.
Keywords: Locomotion, velocity profile, tau, time-to-collision, motor control, vision, perception-action.

*{whitton|feasel|zimmons|sharif|brooks}@cs.unc.edu
†cohn@itd.nrl.navy.mil
‡spoulton@hollins.edu
+brandi_mcleod@hotmail.com

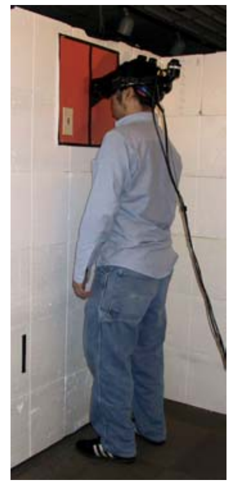

Figure 1: The approach wall and stop task

1 INTRODUCTION

One of the unsolved problems in virtual environment systems research is building a locomotion interface that enables users on foot to move through virtual spaces much larger than the real space enclosing the VE system. We have embarked on a series of studies to compare and evaluate locomotion interfaces for users who are on foot. Although there are a number of ingenious mechanical locomotion interfaces (several are described in [1] and [2]), we focused on three: really walking (as a standard), simple walking-in-place, and flying with a joystick (or gamepad).

The goals of the work reported here are to develop metrics for characterizing and comparing users' movements under the different interface conditions, to correlate these metrics with task performance metrics, and to begin to investigate metrics that relate interface condition to performance. From this study, and those that follow, we hope to develop a model and guidelines to advise system builders on the choice of locomotion interface. Specifically, our results will be used to advise the development of to-be-fielded systems for training dismounted warfighters (i.e., infantry soldiers and marines). The final evaluation metric for the interfaces and systems will be the amount that real-world skills improve as a result of training in an immersive virtual environment.

Most studies on locomotion are restricted to either a virtual or a real environment. Our facility supports both. Exploiting this, we developed methods to characterize the output of a locomotion interface, i.e. the user's path of motion through the virtual environment, and we then used those characterizations to compare different interface conditions to each other and to natural motion in the real environment. We derived metrics from the motion paths, measured performance on a simple task (walking up to a wall and stopping, Figure 1), and explored how motor control of this simple action, captured in the perceptual variable τ, differs among interfaces.

New Knowledge. Principal-component analysis of motion path data revealed that the first three principal components, respectively, are primarily related to 1) maximum velocity; 2) when, during the motion, maximum velocity occurs (percent time elapsed from the start of motion to the time of onset of deceleration, referred to here as *percent time*), and 3) maximum deceleration. Repeated-measures analysis of variance showed that interface condition had a significant effect ($p < 0.05$) on variables peak velocity and peak deceleration, as well as on task performance and on τ measured at peak velocity and at peak deceleration.

To further understand the effects of condition, we performed both comparisons and correlations of data for the 10 possible pairs of conditions taken two at a time. We'll refer to these as *pairwise comparisons* and *pairwise correlations*. Pairwise comparisons of means planned *a priori* revealed 53 of 60 to be significant; this data was unilluminating. Rank ordering of the results of the pairwise correlations of the motion-path-related data (peak velocity, peak deceleration, and τ computed at those points) suggests a clustering and ordering of interfaces by how much they are like locomotion in the real world. User questionnaire data and task performance data only partially supported this, with the differences likely attributable to the quality of our walking-in-place interface. Pairwise correlations of the task performance metric and τ measured at 1) peak deceleration and 2) peak velocity do not follow the clustering and ordering of the motion-path related variables.

The data show that τ measured at peak deceleration correlates more strongly with task performance than with τ measured at peak velocity. This suggests the use of τ calculated at peak deceleration in further work investigating motor control, motion paths, and performance.

The results, as yet, are insufficient to establish strong design guidelines; they do, however, provide a framework for further studies to develop and test models of locomotion interface efficacy.

2 Previous Research

2.1 Comparing Locomotion Interfaces

Metric: Positional Accuracy. Iwata and Yoshida [3] compared the ability of users to reproduce paths through a VE comparing the performance of users of their Torus Treadmill and to those who used a joystick locomotion interface. In the first study, users walked a straight path to a target cone. In the second study the users went to a first target cone, made a turn, and then walked to a second cone. In both studies the users walked along the path with the cones visible, and were then asked to walk to the same path again without cones. Accuracy was measured as how close the users came to the cone positions as they walked the paths without seeing the cones. When walking in a straight line the users overshot the target by nearly equal amounts in both locomotion conditions. In the two-cone condition, the results showed significantly larger total error for users of the joystick than of the Torus Treadmill.

Metric: Cognition. Zanbaka, et al. [4] studied the effect of four locomotion interfaces on cognition in an immersive virtual environment. Their four conditions varied in locomotion control, viewpoint control, and display device. Data collected were the results of a cognition questionnaire (CQ), sketch maps, and responses to the Steed-Usoh-Slater (SUS) Presence Questionnaire. The CQ probed three categories: Knowledge, the recall or recognition of specific information; Understanding and Application, comprehension and application of information; and Higher Mental Processes, information analysis, synthesis, and evaluation.

Post hoc analysis investigating trends in the Understanding and Application scores and Higher Mental Processes scores showed significantly higher performance for the really walking condition over joystick locomotion. Similar significant differences in performance on Understanding and Application scores were found between the between real walking condition and a condition where the user viewed the environment on a monitor and controlled motion and viewpoint with a joystick. Overall, the research results provided evidence that there are cognitive benefits attributable to physically and naturally walking in a virtual environment when the application involves problem solving and interpretation of material.

Metric: Multisensory Realism. There is significant evidence that the level of realism of sensory immersion, interpreted as the degree to which one or more sensory modalities are stimulated in ways corresponding to the real world, plays a role in supporting the ability to perform complex actions within a Virtual Environment [5]. Grant & Magee [6] demonstrated a critical difference in internalizing spatial information when users are allowed to explore a large-scale VE using either a joystick or a simple walking interface. Results favor the interface that enabled the more natural locomotion. Other studies, which have assessed the utility of VEs to train spatial navigation and wayfinding skills ([7]), have shown that in addition to providing adequate proprioceptive stimulation, e.g. natural walking, VEs must also provide adequate visual information.

Metric: Sense of Presence. In [8] , Slater, Usoh, and Steed proposed that sense of presence is a function of both 1) the match between the sensory input provided by the VE and the "internal representation systems…employed by the participant" and 2) the match between proprioceptive cues and visual feedback, i.e. realistic visuals and realistic motion cues. Two studies have compared the efficacy of locomotion techniques by measuring the level of presence evoked in users of different locomotion techniques while holding visual condition constant.

Slater, Steed, and Usoh [8] compared walking-in-place and on-off flying in an environment that evoked a strong reaction in the users, a visual cliff. Usoh, et al. [9] report on a follow-up study which included both of the original conditions, walking-in-place and pushbutton flying, and added the condition of really walking. Both studies investigated whether users would experience higher presence in a VE if they moved though the environment using a locomotion technique that provided proprioceptive stimuli similar to natural walking.

Post-experience presence questionnaires for the first study showed that walking-in-place produces higher levels of presence than moving by pushing a button, provided that users identify with their avatar. The second study showed that both really walking and walking-in-place conditions yielded significantly higher levels of presence than did push-button flying; a strong trend suggested that really walking produced higher sense of presence than walking-in-place.

2.2 Analysis of Motions from Tracker Logs

A motion path is a sequence of position-time samples of a tracked point on the user's body. From such a path it is easy to derive a velocity profile (velocity vs. time) and measures such as maximum velocity, peak deceleration, and τ.

We analyzed motions in our study in essentially the same way as reported in Mason, et al. [10]. Although their research is focused on reaching movements, their reported data preparation techniques are particularly relevant. Mason used a 3 DoF tracker to collect motion paths of reaching hands. The raw tracker data were interpolated and smoothed with a low pass Butterworth filter before the velocity profile and other measures were computed. Their paths were algorithmically truncated to a consistent starting condition (when velocity reached 5 mm/sec, in their case). We employed nearly identical techniques.

2.3 Tau and Motion Control

Basing their work on the Gibsonian notion of treating visual input as an optic flow field [11], Lee and Reddish [12] suggest that the onset of certain motions is controlled by the visual information individuals receive. They propose that the control mechanism is based on an optically defined parameter, time-to-collision, τ. In perceptual terms, τ is the expansion of visual information on the retina; operationally, it is computed as user's distance-from-collision divided by user's velocity (i.e., x/x') (c.f. [13] for detailed derivations). Lee and Reddish's work demonstrates this relationship between τ (computed as x/x') and motor behavior for a single system, and Schoner [14] demonstrates that it holds in a number of other previously described experimental systems.

In practice, τ is typically reported at a critical value within a motion path, e.g., at peak velocity [15] or at the point in the path where τ varies least [16]. The value of τ at peak velocity, i.e., the onset of deceleration for simple motions, has been widely explored as an element in motor control strategies [15]. Research has also identified and examined control situations in which τ remains nearly constant across experimental conditions [15, 16],. Such cases are important because of the notion that features of movement that are invariant across conditions reflect the nature of the neural control of motor actions [17].

In this work, movement control is assumed to depend primarily on information provided visually. Previous work [12, 16, 18, 19] suggests that τ, the perceived time-to-contact, may serve as an environmentally specified (i.e., situation specific) variable for controlling movements in a dynamic setting. Here, we investigate the notion that τ captures the relationship between motion path metrics, task performance, and interface type, providing another means to quantify the impact of interface design on performance.

2.4 Interface Efficacy and Naturalness

Osgood's notion of identical elements proposes that the closer the training environment is to the actual one, the more likely it is that the training environment will prove effective [20]. Prior work cited in section 2.1 supports the notion that the more natural or realistic the locomotion interface, the more effective it is as evaluated on some metric. These ideas suggest that we examine our data to see if it supports the idea of a "naturalness" ordering of locomotion interfaces. Using motion-path-derived variables for real and VR conditions we can examine how much the users' motions under different conditions are like their motions in the real environment. We can perform similar comparison for task performance and we can extend the scope of our comparisons to motor control mechanisms by comparing values of τ under different interface conditions.

3 User Study

The Academic Affairs Institutional Review Board at the University of North Carolina at Chapel Hill approved the user study reported here.

User motion in the real world is the standard against which we compare various locomotion interfaces. Based on our assumption that the user's head will be tracked in any future fielded VE-based training systems, and to maximize the chance that our metrics can be used in future field studies, we based our metrics on motion parameters derived from only head-tracking data.

3.1 Conditions and Task

Each of the five conditions we studied included one of three locomotion interfaces and one of three visual interfaces. Of the three locomotion interfaces, one was natural (really walking), and two were artificial (joystick flying and walking-in-place).
Similarly, of the three visual interfaces, one was natural (unobstructed natural vision) and two were artificial (field-of-view-restricted natural vision and a head-mounted display (HMD)).

We collected data for three VE conditions: really walking in virtual reality (VRW), joystick flying (JS), and walking-in-place (WIP). We collected data for two conditions where users could see naturally: really walking with unhindered vision (Real) and really walking with field-of-view-restricted vision (Cowl). The field-of-view of the Cowl and the HMD are the same.

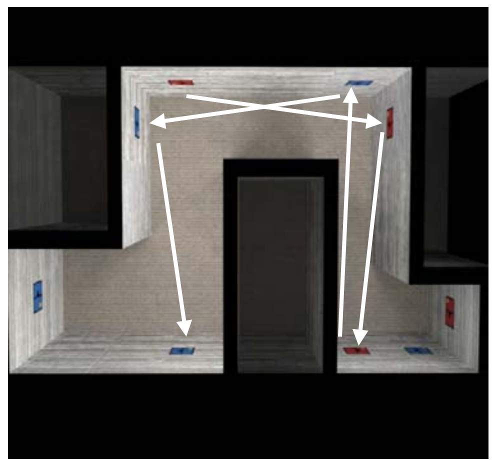

Figure 2: The virtual maze viewed from overhead. The arrows show the approaches to the five targets used in the study.

Each participant navigated a maze in each condition. The order of the conditions was determined by a modified Latin Square. As they moved through the maze, the participants saw targets on the walls. We instructed participants to walk up to the target and stop as close to it as they could without touching the wall. When participants felt they were as close to the target plane as possible, they signaled by pressing a button and then proceeded to the next target. Experimenters noted if the participants bumped the wall, but this information was not used in the analysis presented here.

3.2 Equipment and Software

Our environments were modeled using 3D Studio Max™; our custom virtual environment application was developed using Visual C++ 6.0 on Windows™ XP, the WildMagic™ game engine by Magic Software, Inc., and the VRPN library for communication with peripherals. The study application was run on a dual-Xeon 1.7GHz PC with 1 GB of RAM and an nVidia GeForce™4 Ti 4600 graphics card. For the VR conditions, the participants wore a Virtual Research Systems V8 HMD with 640x480 tri-color pixel resolution in each eye and a horizontal field-of-view of 47 degrees. The head was tracked with a

3rdTech™ HiBall 6DoF optical tracker with a 22' x 30' tracked area. The tracker sensor was mounted on the HMD for the VRW, JS, and WIP conditions. For the Cowl condition the tracker was mounted on a modified V8 HMD shell that restricted FOV. In the Real condition, the tracker was mounted to a simple headband.

Participants carried a Logitech® Cordless Rumblepad as a button input device and for locomotion in the joystick flying condition. Midpoint in the joystick's 256-level output range was set to a speed comparable to normal walking. In the JS and WIP conditions, the Rumblepad vibrated when the user collided with a wall. A Crossbow™ Solid State 2G accelerometer mounted on the HMD provided input data for in the walking-in-place interface. The acceleration was compared to a threshold to identify footfalls, and the footfalls in turn produced the forward motion of an average stride. The direction of motion for both the JS and WIP conditions was the view direction.

3.2.1 Matching Real and Virtual Environments

We refer to the environment as a maze, though it was simply a corridor with four turns (Figure 2). There were targets on the walls at several locations. The real maze was constructed from ReddiForm™ Styrofoam blocks. The walls were 1.8 m (6 ft) tall. The 45 cm square targets have a vertical stripe to indicate the center of the target, and an arrow pointing in the direction of the next target. During a pilot of the study it became apparent that our stark VE included no familiar objects users could use to judge size and scale. In response, we added a light switch cover plate to each target. The virtual maze and targets matched the real maze and targets (Figure 3).

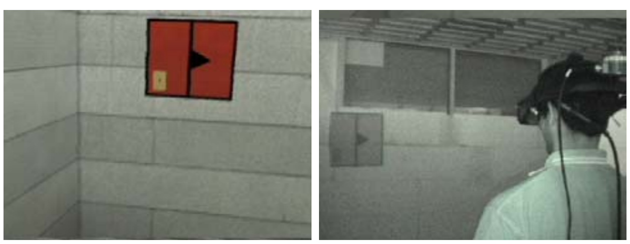

Figure 3: First-person (left) and third-person views of user approaching a target

3.2.2 Data Collection and Preprocessing

Motion Paths. We updated frames and logged the 6DoF head-pose data at 160 Hz.[1] We extracted 3DoF motion paths from the pose data.

Path Preprocessing. To provide a consistent starting point for each target approach, we algorithmically truncated each path so that it began at the same distance from the target plane.

Because the motion paths in each target approach are essentially perpendicular to the target plane, we projected the 3DoF points onto such a line. This not only reduced the dimensionality of the data but eliminated positional variations caused by side-to-side and up-and-down head movements characteristic of walking.

Filtering with a low-pass Butterworth filter eliminated head-bobbing movements from the paths and eliminated any high-frequency motions caused by tracker jitter. The cutoff frequency was empirically defined to eliminate ripples in the path data that would become exaggerated in later differentiation steps.

[1] Our system runs with sync-to-vertical-refresh turned off. No users reported seeing image "tearing"; we attribute this to the relatively slow switching time of the LCDs in the HMD.

Principal-component analysis. Observations of velocity profiles for our five conditions in an earlier, exploratory study (Figure 4) suggested that curve height, skew, and steepness-of-deceleration differentiate the curves for the different conditions. In the current study we sought quantitative confirmation of our observations using principal-component analysis (PCA).

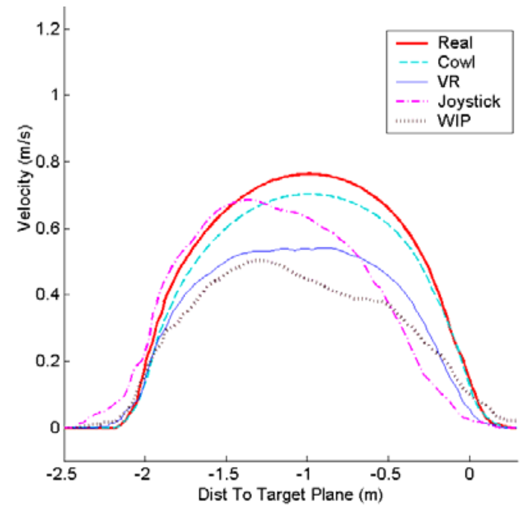

Figure 4: These curves from an exploratory study show the mean values of the paths for all subjects for each condition for one target. We observed that peak velocity differs between conditions and that the peak velocity for joystick and WIP are skewed left. Less apparent is the change in steepness in deceleration.

Quoting from the Wikipedia web site, "In statistics, **principal components analysis (PCA)** is a technique that can be used to simplify a dataset; more formally it is a linear transformation that chooses a new coordinate system for the data set such that the greatest variance by any projection of the data set comes to lie on the first axis (then called the first principal component), the second greatest variance on the second axis, and so on. PCA can be used for reducing dimensionality in a dataset while retaining those characteristics of the dataset that contribute most to its variance by eliminating the later principal components (by a more or less heuristic decision)." We use PCA in exactly this way to reduce the dimensionality of the path vector data.

The input to the PCA is a set of feature vectors derived from the path vectors, one feature vector for each path vector. To generate the feature vectors for the velocity profiles, we first used finite difference techniques to generate position-velocity data from the position-time path vectors. The velocity data were then resampled in distance to produce a 50-element feature vector. Fifty was chosen empirically as a tradeoff between PCA computational cost and resolution. We repeated this process for each path vector, creating a set of 50-element feature vectors that were the input to the PCA.

4 Results

We collected data on approaches and stopping at five wall-mounted targets, under each of five interface conditions. We derived six dependent variables to quantify the impact of interface on performance: three motion path metrics—peak velocity, percent of time to peak velocity (*percent time*),[2] and peak

[2] We normalized time as the distances between targets differ.

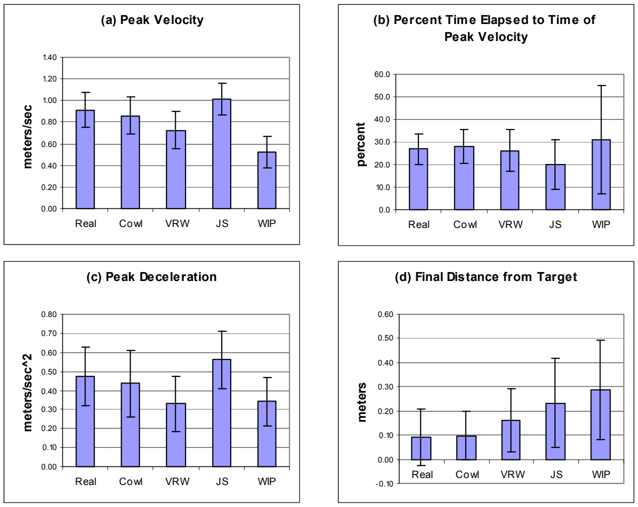

Figure 5: Mean and +/- 1 σ of four dependent variables

deceleration; task performance—final distance to target; and time-to-collision, τ, at peak velocity and at peak deceleration.

4.1 Study participants

Participants were recruited from among students at UNC-Chapel Hill and were paid for participating in two-hour sessions. All subjects were able to walk unassisted and passed a screening for health and susceptibility to motion sickness.

Thirty-two participants (21 male, 11 female) ranged in age from 18 to 42, with a mean age of 24. Video-game-playing experience was bi-modal, with 11 participants reporting less than one hour/week and 15 reporting over 10 hours/week. The participants were generally naïve VE users: only six had been in an immersive environment more than twice before, and none had experienced a VE more recently than a month before.

Users signed informed consent forms and filled out demographic questionnaires before entering the virtual environment. Post-VE experience questionnaires asked them to both rate and rank their experiences with the five interfaces. Open-ended questions elicited additional qualitative data.

4.2 Description of Quantitative Metrics

4.2.1 Motion-Path Derived Metrics

Principal components. The first three PCs account for over 95% of the variance in the data. Our velocity profile feature vectors have meaning when graphed. Varying each of the first three PCs of a velocity profile showed them, respectively, to most noticeably affect the curve shapes as follows: the height of the curve (peak velocity), the skew of the peak of the curve (the percent time of peak velocity), and how steeply the curve falls (peak deceleration). The analysis confirmed what our observation of earlier data suggested: these three characteristics are the major factors defining the shape of velocity profiles for different locomotion conditions. The result of the PCA helped us identify which discrete values derived from the motion paths we would use to compare paths, and allowed us to reduce the dimensionality of the motion path description from (in our work) 500 element path vectors to 3 values.

Motion Path Variables. Based on the outcome of the PCA, we examined the data for peak velocity, percent time to peak velocity, and peak deceleration for each of the five conditions. Figures 5 (a), (b), and (c) show the means and standard deviations for these data.

4.2.2 Task performance: final distance from target

In the present work, task performance is defined as the absolute value of the distance between the user's final position and the plane of the target at which they are stopping. Figure 5 (d) shows the mean and standard deviation for this metric.

4.2.3 Time-to-collision: τ

Two values of τ, taken at the times of peak velocity and peak deceleration, were computed from the data. The process used was nearly identical to that presented in [12]. Figure 6 shows the means and standard deviations for τ calculated at peak velocity and peak deceleration.

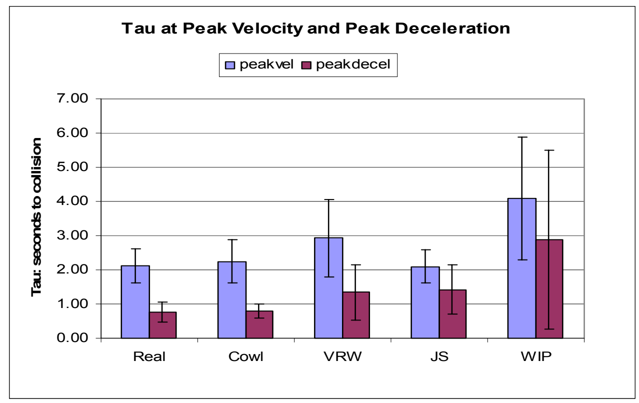

Figure 6: Mean and +/- 1 σ for τ at critical motion path variables.

4.3 Statistical Analysis

The overall statistics are based on a 5 x 5 (Targets x Interface-Type) repeated measures analysis-of-variance (ANOVA) run on each of the six dependent variables: peak velocity, percent time to peak velocity, peak deceleration, final distance, τ at peak velocity, and τ at peak deceleration. Since there were missing data, statistical analyses were performed as six separate within-subjects ANOVAs[3], with corresponding *a priori* contrasts, rather than using a single MANOVA procedure. The results of the ANOVAs were adjusted using the Benjamini-Hochberg method, which controls for false detection rates resulting from multiple hypotheses testing [21]. Main effects of Target and Interface Type were considered, as were planned pairwise comparisons. Although we report main effects for Targets, we do not discuss Targets further in this paper.

4.3.1 ANOVA results

Each of our ANOVA procedures compared the means of one of the dependent variables for three cases: differences in means across all targets and all conditions (Overall), across five Targets, and across five Conditions.

Motion Path Metrics. The ANOVA for peak velocity showed Overall significance ($p<0.0001$), with a significant main effect for

[3] Our analysis used an ANOVA process that accounted for fact that each subject did the task in each condition (subjects repeated). The mathematics for this model are approximate, not exact, so the probabilities are found using the Chi-Square rather than the F statistic. This model does not return a measure of the variance attributable to Target and Condition.

both Targets ($p<0.0001$) and Interface ($p<0.0001$). The ANOVA for percent time showed Overall significance ($p<0.002$) and a significant main effect for Targets ($p<0.0057$). Percent time did not show a main effect for Interface ($p=0.1583$). The ANOVA for peak deceleration also showed Overall significance ($p<0.0003$) and significant main effects for Targets ($p<0.005$) and Interface ($p<0.0001$).

Task Performance. The ANOVA for task performance, measured as the final stopping distance relative to each target, showed Overall significance ($p<0.002$), with a significant main effect for both Targets ($p<0.003$) and Interface ($p<0.001$).

Time-to-collision, τ. The ANOVA for time-to-collision measured at peak velocity showed Overall significance ($p<0.0002$), with a significant main effect for both Targets ($p<0.0001$) and Interface ($p<0.0001$). The ANOVA for time-to-collision measured at peak deceleration showed overall significance ($p<0.002$, with a significant main effect for both Targets ($p<0.03$) and Interface ($p<0.0001$).

4.3.2 Planned pairwise comparisons

A series of ten *a priori* pairwise comparisons were calculated for each of the six variables, using the Real condition (really walking and unobstructed natural vision) as the baseline. We used Benjamini-Hochberg's method to control for inflated error rates in these comparisons. Of the 60 comparisons, only 7 were not significant. We had to explore further.

4.3.3 Correlation of τ and task performance

To determine which of our two measures of τ might better describe motion control strategies used for locomotion in the different interface conditions, we computed correlations between τ and the task performance measure for both τ calculated at peak velocity and peak deceleration. Figure 7 shows the correlation coefficients for τ at these two critical points and final distance.

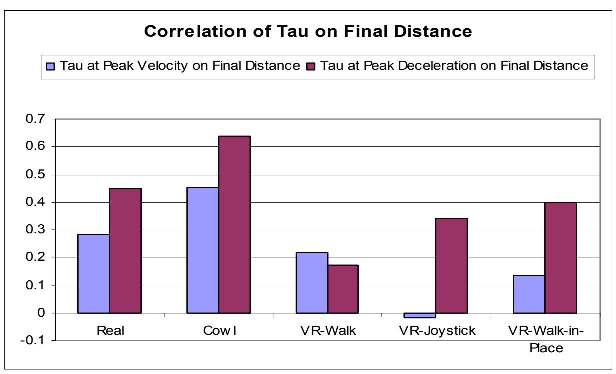

Figure 7: Correlation coefficients of τ on final distance at two critical points.

4.3.4 Pairwise correlations

To better understand the relationships of the conditions, we performed pairwise correlations on the five dependent variables that show a main effect for Interface Type. The ordered results for the Peak Velocity, Peak Deceleration, and Tau at those points are shown in Figure 8. The order of the pairs on the Y-axis is nearly identical on the four charts; the ordering of the pairs that include the Real condition is consistent: Real-Cowl>Real-VRW>Real-WIP>Real-JS. The ordering for the final distance correlations (not shown) is somewhat different overall, including, for the pairs including Real, the order is: Real-VRW>Real-Cowl>Real-JS>Real-WIP.

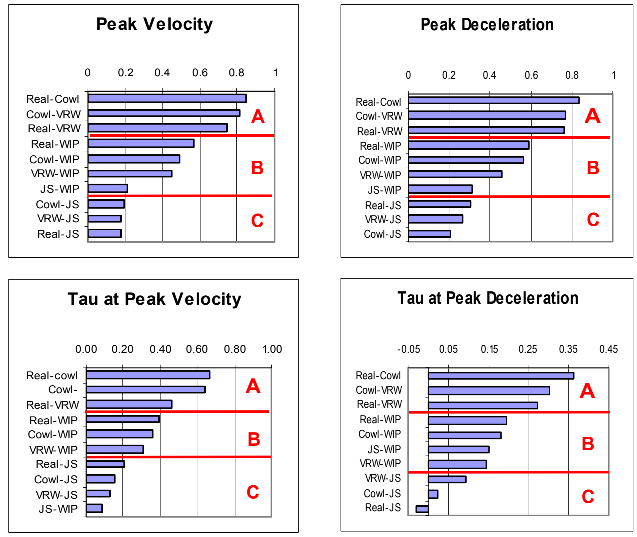

Figure 8: Correlations on dependent variables by conditions taken pairwise. Red horizontal lines denote bins that suggest a coarse ordering of interfaces.

4.4 User Experience

Users rated which interface condition (the combination of locomotion technique and vision condition) they thought they performed the task best with, and the one they thought they did the worst with. Twenty of the 32 participants indicated that their best performance was in the Real condition; the other participants were widely spread over the other four conditions, with each being selected between 2 and 4 times. Users commented:

> *The headband [Real] seemed easiest because I had full peripheral vision...*
>
> *The blinders/cowl [Cowl] and headband [Real] didn't change my sense of location and sense of body so I performed best with them.*
>
> *I don't think the lack of peripheral vision influenced my performance much. It did, however, influence the way I moved, I think.*

Comparing the three VR conditions, one participant commented:

> *Really walking in VR allowed me to choose exactly how far I want to move; so I could move slowly until I was just a couple of centimeters away from the target. This "analog control" made it easier for me. I also thought that really walking in VR was easier than the gamepad [JS] because the movements were more natural.*

The joystick was rated as worst or next to worst by 20 of the 32 participants. One participant, who also reported game use of 5-10 hours/week during at least part of the previous year, noted its limitations:

> *[The joystick (JS)] was not good, but better than walking-in-place [WIP]. The joystick allowed far more sensitive adjustments once I got close, but no sidestepping ability.*

Of the 18 people who reported more than 5 hours/week of video game usage at some time during the last year, only 3 rated joystick as the interface with which they performed best.

Twenty-five of 32 participants reported they performed worst in the walking-in-place condition. The open-ended comments and experimenter observations support this result. We attribute this to the quality of our WIP interface. One user commented:

> *I felt I had the least control here with my speed and turning. Maybe I wasn't stomping enough, but I felt I couldn't move myself around the way I wanted.*

Participants were also asked to rank the interfaces according to how well they performed while using them, from best (1) to worst (5). The modal values for those responses are reported in the middle row of Table 1. The bottom row is the number of responses in that mode.

Table 1: Modal values for ranking of conditions by best to worst performance and instances of that ranking (of the 32 total).

Real	Cowl	VRW	JS	WIP
1	2	3	4	5
22 of 32	15 of 32	17 of 32	16 of 32	24 of 32

5 Discussion and Conclusions

5.1 Walking-in-Place Interface

The interpretation of our results is complicated by the quality of our walking-in-place interface. The WIP technique as implemented at the time of this study was difficult to use. Heavy stomping was required for some participants to trigger a footstep; turning was difficult. We see the effect of these difficulties in the relatively large variances for percent time of peak velocity and final distance. Footsteps, when they were recognized, resulted in a constant size movement in the direction of gaze. This put an unnatural constraint on Final Distance in this condition: the users reached a point from which they could get no closer to the target without colliding with it.

While our current version of walking-in-place lacks sensitivity, the version we used previously in Usoh et al. [9] had unnatural stopping and starting delays. The poor showing of these two simple WIP locomotion interfaces means that system designers proposing to use WIP need to (1) engineer them very carefully, and (2) validate them against real walking. The Gaiter system, [22], represents a significant effort to develop a WIP system allowing both natural motion and natural exertion.

5.2 Discussion of results

A significant challenge to the development of any type of human - computer interface lies in the quantification of performance increments or decrements when one interface design is chosen over another. We have chosen to work on an interface for which design requirements are only beginning to be developed, and have created a framework within which metrics can be validated using a task whose complexity can, over time, be amplified.

We developed metrics in three categories: properties of the motion paths, task performance, and, an exploratory investigation of time-to-collision, τ, as a measure of interface efficacy. As suggested by previous work, as we examined our results we looked for patterns in the ordering of the interfaces.

Motion Path Metrics. The ANOVAs showed that two of the three motion path metrics, peak velocity and peak deceleration, show a main effect for Interface Type. If you order results of pairwise correlations for these data (Figure 8) the order of the pairs along the Y-axis is nearly identical for each case. Conservatively, these orderings show similar clusters with interfaces with real locomotion (walking) interfaces and natural vision grouped at one end, and conditions with less natural locomotion interfaces (walking-in-place and joystick) and computer-generated visual input (HMD) at the other end. The mixed interface, real walking with HMD visuals, falls in-between.

Task Performance Metric. Our performance metric was Final Distance to Target and the ANOVA shows a main effect for Interface. For this variable, smaller is better. Figure 5 shows an ordering of the interfaces that groups as that described above.

User Experience. User questionnaire responses, Table 1, also support the three bin ordering.

Time-to-Collision, τ. Whereas exploring performance and motion metrics provides insight into which interface may be most effective, it does not explain the mechanism through which differences in the metrics arise. Such understanding would provide general principles to help developers make better design decisions. Arguably, the manner in which visual information is provided, and the manner in which the interface enables the user to act upon such information, are critical. Consequently, a model that captures a relationship between the two might prove useful, especially if it could be shown to relate to overall performance in some fashion.

Previous research [23] suggests that the human nervous system may plan simple braking maneuvers by defining the point at which braking starts or at the point of maximum deceleration. Moreover, there is evidence that the time-to-contact variable, τ, an indirect measure of control strategy, captures this relationship when calculated at such critical values and correlated to performance metrics [15, 24]. As Figure 7 suggests, τ calculated at peak deceleration correlates better—and captures more variability—than τ calculated at peak velocity and should be used in future research relating motor control and performance.

However, final distance and τ at peak deceleration did not correlate particularly well (Figure 7). We speculate that since humans brake to avoid collisions, exploring number of collisions as a task performance metric might prove illuminating. We did not formally collect or analyze collision data in this work.

Table 2: Bin into which each pairwise correlation falls when correlations are coarsely grouped high (bin A) to low (bin C). Data are from Peak Deceleration in Figure 8.

	Visual Condition	Real	Re-strict-ed FOV	HMD	HMD	HMD
Loco--motion Condition		Real	Cowl	VR-Walk	VR-WIP	VR-JS
Walk	Real					
Walk	Cowl	A				
Walk	VR-Walk	A	A			
Walk-in-place	VR-WIP	B	B	B		
Joystick	VR-JS	C	C	C	C	

5.3 Observations

The data presented here are not so conclusive as to warrant basing design guidelines on them. However, a number of interesting observations can be made from the data in Table 2, which recasts the data for Peak Velocity from Figure 8 in a way that shows which pairwise correlations fall into which bin.

We observed the following:

Visual Interface—Field-of-View appears to have no effect in the task used in this study. Holding locomotion condition constant (looking at rows), correlations for each visual condition fall into the same bin. This is believable as this particular task required only looking straight ahead.

Visual Interface—Real or HMD appears to have no effect. Again, holding the locomotion interface constant, we observe that the data for the Cowl (real vision) and HMD (computer generated visuals) fall into the same bin. Again, this is understandable as the task was designed so as not to require high visual acuity.

Locomotion Interface. Holding visual condition constant (looking at columns), we observed that locomotion interface does have an effect on size of the correlation values. In each column, the more "real" locomotion methods are in the A bin (high correlations) than the less natural ones in bins B and C (lower correlations. These data indicate that the motions with walking-in-place and joystick locomotion do not correlate well with (i.e., are not like) motions when walking naturally. Since we do not yet understand the potential impact of a locomotion interface dissimilar too real walking, this observation should serve as a caution to developers: carefully consider whether there might be unintended consequences of adopting walking-in-place or joystick interfaces.

6 Acknowledgements

Support for this research is, or has been, provided by the Office of Naval Research and the NIH National Institute for Biomedical Imaging and Bioengineering and National Center for Research Resources.

This work would not have been completed without the many significant contributions of Luv Kohli, Matt McCallus, Christopher Oates, and Christopher VanderKnyff. We also thank William Becker of Strategic Analysis, Inc. for thoughtful advice that helped frame this research.

References

[1] J. M. Hollerbach, "Locomotion Interfaces," in *Handbook of Virtual Environments: Design, Implementation, and Applications*, K. M. Stanney, Ed., Mahwah, N.J.: Lawrence Erlbaum Asso., 2002, pp. 239-254.

[2] R. P. Darken, T. Allard, and L. B. Achille, "Spatial Orientation and Wayfinding in Large-Scale Virtual Spaces II," *Presence: Teleoperators & Virtual Environments*, vol. 8, pp. 3-6, 1999.

[3] H. Iwata and Y. Yoshida, "Path Reproduction Tests Using a Torus Treadmill," *Presence*, vol. 8, pp. 587-597, 1999.

[4] C. Zanbaka, B. Lok, S. Babu, D. Xiao, A. Ulinski, and L. F. Hodges, "Effects of travel technique on cognition in virtual environments," *Proceedings of Virtual Reality, 2004*. IEEE , 27-31 March 2004.Pages:149 - 286.

[5] S. S. Chance, F. Gaunet, A. C. Beall, and J. M. Loomis, "Locomotion Mode Affects the Updating of Objects Encountered During Travel: The contribution of vestiular and proprioceptive inputs to path integration," *Presence*, vol. 7, pp. 168-178, 1998.

[6] S. C. Grant and L. E. Magee, "Contributions of proprioception to navigation in virtual environments," *Human Factors*, vol. 40, pp. 489-497, 1998.

[7] B. G. Witmer, W. J. Sadowski, and N. M. Finkelstein, "VE-based training strategies for acquiring survey knowledge," *Presence: Teleoperators and Virtual Environments*, vol. 11, pp. 1-18, 2002.

[8] M. Slater, M. Usoh, and A. Steed, "Taking Steps: The Influence of a Walking Technique on Presence in Virtual Reality," *ACM Transactions on Computer-Human Interaction*, vol. 2, pp. 201-219, 1995.

[9] Usoh, Martin, K. Arthur, M. Whitton, A. Steed, M. Slater, and F. Brooks, "Walking>Virtual Walking>Flying, in Virtual Environments " *Proceedings of SIGGRAPH'99 (*Los Angeles , CA, August 11-13), *Computer Graphics Annual Conference Series, 1999*, pp. 359-364.

[10] A. H. Mason, M. A. Walji, E. J. Lee, and C. L. MacKenzie, "Reaching Movements to Augmented and Graphic Objects in Virtual Environments," *Proceedings of Conference on Human Factors in Computing Systems* (CHI 2001), Seattle, WA, 2001.

[11] J. J. Gibson, *The ecological approach to visual perception*. Hillsdale, NJ: Lawrence Erbaum Associates, 1979/1986.

[12] D. N. Lee and P. E. Reddish, "Plummeting gannets: a paradigm of ecological optics," *Nature*, vol. 293, pp. 293-294, 1981.

[13] Bessemer, D., Boldovici, J, & Bolton, A. (2002) *The Elements Of Training Evaluation*, Army Research Institute for the Behavioral and Social Sciences, Alexandria, VA.

[14] G. Schoner, "Dynamic theory of action-perception patterns: The time before contact paradigm," *Human Movement Science*, vol. 13, pp. 415-439, 1994.

[15] P. Foo, J. A. S. Kelso, and G. C. de Guzman, "Functional stabilization of unstable fixed points: Human pole balancing using time to balance information," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 26, pp. 1281-1297, 2000.

[16] D. N. Lee, "A theory of visual control of braking based on information about time-to-collision," *Perception*, vol. 5, pp. 437-459, 1976.

[17] R. A. Schmidt, *Motor Control and Learning*. Champaign, IL: Human Kinetics Books, 1988.

[18] Lee, D. N. (1980). Visuo-motor coordination in space-time. In G. E.Stelmach & J. Requin (Eds.), *Tutorials in in motor behavior* (pp. 281-285), Amsterdam: North-Holland.

[19] Turvey, M. T., & Carello, C. (1986). "The ecological approach to perceiving-acting." *Ada Psychologica,* **63,** 133-155.

[20] Osgood, C.E. (1949) "The Similarity Paradox in Human Learning: A Resolution," *Psychology Review* 36: 132-143.

[21] Benjamini, Y., Hochberg, Y. (1995). "Controlling the False Discovery Rate: a Practical and Powerful Approach to Multiple Testing ", *Journal of the Royal Statistical Society B,* **57** 289-300.

[22] Templeman, J., "Performance based design of a new virtual locomotion control," The Capability of Virtual Reality to Meet Military Requirements: Papers presented at the RTO Human Factors and Medicine Panel (HFM) Workshop, Orlando, Florida, 1997.

[23] Cohn, J.V. & Kelso, J.A.S. (1999). "On getting your finger to the right place at the right time." Proceedings of *1999 North American Society for Psychology of Sport and Physical Activity Conference*, St. Petersburg, FL

[24] Treffher, P. J., & Kelso, J. A. S. (1995). "Functional stabilization of unstable fixed points." In B. Bardy, R. Bootsma, & Y. Guiard (Eds.), *Studies in perception and action*, pp. 83-86. Hillsdale, NJ: Erlbaum.