Validating Simulation-Based Evaluation of Redirected Walking Systems

Mahdi Azmandian, Rhys Yahata, Timofey Grechkin, *Member, IEEE*,
Jerald Thomas, *Student Member, IEEE*, and Evan Suma Rosenberg, *Senior Member, IEEE*

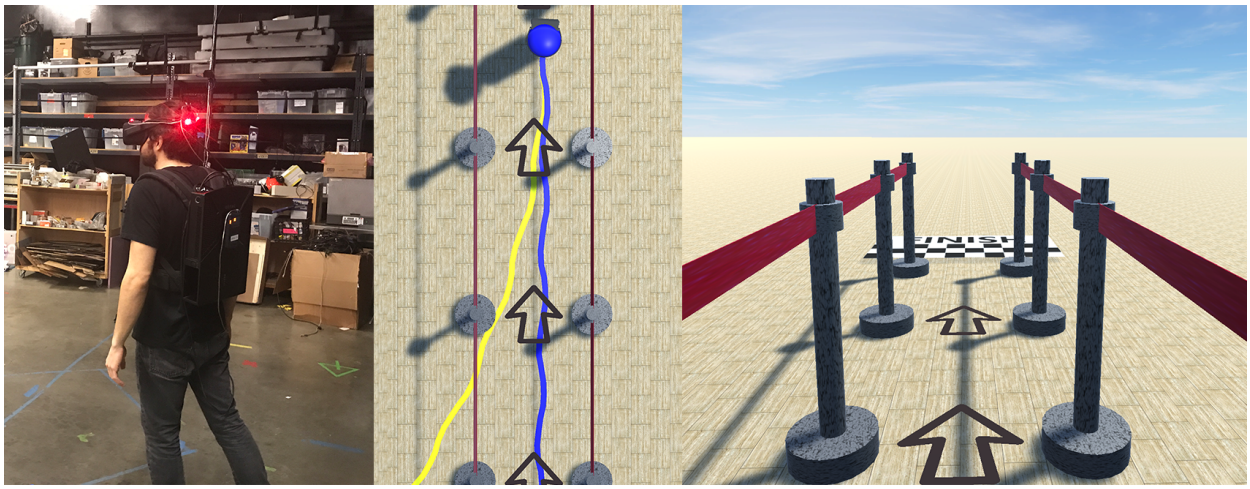

Fig. 1. Snapshot of the user study setup. Participant wearing virtual reality equipment is shown on the left, the user's perspective in the virtual world on the right, and in the center a top view showing the real (yellow) and virtual (blue) trajectories influenced by redirection.

Abstract—Developing effective strategies for redirected walking requires extensive evaluations across a variety of factors that influence performance. Because these large-scale experiments are often not practical with user studies, researchers have instead utilized simulations to systematically test different algorithm parameters, physical space configurations, and virtual walking paths. Although simulation offers an efficient way to evaluate redirected walking algorithms, it remains an open question whether this evaluation methodology is ecologically valid. In this paper, we investigate the interaction between locomotion behavior and redirection gains at a micro-level (across small path segments) and macro-level (across an entire experience). This examination involves analyzing data from real users and comparing algorithm performance metrics with a simulated user model. The results identify specific properties of user locomotion behavior that influence the application of redirected walking gains and resets. Overall, we found that the simulation provided a conservative estimate of the average performance with real users and observed that performance trends when comparing two redirected walking algorithms were preserved. In general, these results indicate that simulation is an empirically valid evaluation methodology for redirected walking algorithms.

Index Terms—Virtual reality, redirected walking, locomotion, simulation

1 Introduction

Exploring arbitrarily large virtual environments within a limited physical space is a critical challenge for virtual reality systems. A promising solution for this problem is redirected walking, which involves decoupling the user's virtual path from the real world trajectory [31]. The key principle behind this approach is to leverage unnoticeable perceptual manipulations such as imperceptibly small visual rotations (*rotation gain* and *curvature gain*) and translations (*translation gain*) to redirect the user away from the boundaries of the physical space. Redirection provides the benefits of unconstrained physical walking in virtual environments, such as an enhanced sense of presence [42], efficient navigation [32] [38], and improved cognitive maps of the environment [33] at a limited cost of some cognitive load on the user [8] and without interfering with navigation and spatial cognition [15] [39].

In order to effectively utilize redirected walking in a virtual reality experience, a system needs to employ a *redirection strategy* that determines the real-time gain levels that are applied over time [1]. The performance of a redirection strategy may be influenced by numerous parameters such as the size and shape of the tracked physical space, the structure of the virtual environment, and the user's locomotion behavior [3]. Additionally, other factors such as users' perceptual thresholds for redirection gains and the selected wall-contact resolution technique (also known as *reorientation* or *resetting*) can also influence performance. Thus, systematic evaluation of redirected walking strategies is often non-trivial and requires comparison of numerous experimental conditions over many trials. Therefore, researchers frequently make use of simulation to determine parameters for newly developed algorithms

- *Mahdi Azmandian was with the Institute for Creative Technologies, University of Southern California.*
- *Rhys Yahata is with the Institute for Creative Technologies, University of Southern California.*
- *Timofey Grechkin was with the Institute for Creative Technologies, University of Southern California.*
- *Jerald Thomas is with the Department of Computer Science & Engineering, University of Minnesota. E-mail: thoma891@umn.edu*
- *Evan Suma Rosenberg is with the Department of Computer Science & Engineering, University of Minnesota. E-mail: suma@umn.edu*

Manuscript received 6 Sept. 2021; revised 3 Dec. 2021; accepted 7 Jan. 2022. Date of publication 17 Feb. 2022; date of current version 29 Mar. 2022. Digital Object Identifier no. 10.1109/TVCG.2022.3150466

and to compare the effectiveness of different redirection strategies. Although simulation is not a substitute for human studies, it is especially useful for conducting large-scale evaluations that are otherwise impractical or impossible to perform with live users.

Simulations provide an efficient and economical methodology for the evaluation of redirected walking systems. One major advantage is that simulations allow performance analyses to be conducted rapidly and efficiently without the overhead and cost of user studies. This typically occurs during development to determine ideal parameters or test different variants of a novel algorithm, as well as afterwards to compare new strategies with previously developed approaches from the literature. Furthermore, simulation allows the evaluation with lengthy repeated trials that would be unreasonable to expect from human study participants and avoids introducing confounding factors that may arise due to random variability between individual users. Finally, large or complex physical space configurations can be simulated if the researcher or developer does not have immediate access to a suitable motion tracking space in the real world. While simulations are a useful for evaluating redirected walking strategies, this is predicated on the notion that simulation-based analysis provides a reasonable estimate of performance with live users.

Simulations are therefore first and foremost an analysis tool, a method by which we evaluate algorithms. The merits of simulations may not directly be evident to virtual reality practitioners, but the insights gained from their use can indirectly benefit all users. For example, researchers have reported scientific results from simulation experiments (e.g., [3, 5]). Furthermore, the efficient and economical nature of simulation-based analysis accelerates the overall process of research, development, and testing of new approaches. However, there has been limited prior work to empirically validate simulation as an evaluation methodology for redirected walking systems. This work, therefore, aims to investigate the validity of this approach by comparing performance metrics computed through simulation with redirected walking results from live users.

In this paper, we investigate how real users' locomotion behavior is affected by redirected walking and conduct empirical comparisons with simulated user data. To achieve this, we analyze the interaction between locomotion and redirection gains at the micro and macro level. From the micro-scale inspection, we understand specific characteristics of locomotion behavior that cause deviations from simulation while applying translation gains, curvature gains, rotation gains, and resets. The macro-inspection evaluates the overall redirected walking system performance across an entire experience, and we observe that simulations provide conservative estimates for the performance of average users. Comparative trends across the different strategies were also preserved, which provides evidence that simulation-based evaluation is a valid methodology for evaluating redirected walking algorithms.

2 Background and Related Work

Initially proposed by Razzaque [31], redirected walking allows users to explore virtual environments that are larger than the real space they walk in. This is accomplished by altering the mapping between the real and virtual coordinate systems. Redirected walking exploits the fact that human vision will generally override the vestibular and proprioceptive systems when these sensory signals are not congruent, so long as the magnitude of the conflict is kept within manageable levels [21]. There have been a multitude of developments in redirected walking research since its inception, and this section aims at only covering those most relevant to this work. A comprehensive literature review has been conducted by Nilsson et al. [26].

Redirection gains are the set of techniques used to perform the alteration of the mapping between real and virtual coordinate systems. Translation and rotation gains directly scale the user's translations and rotation respectively, thus increasing or decreasing the amount of virtual displacement and/or turning perceived by the user. Curvature gain causes the user to walk in a curved path by applying imperceptible virtual rotations when they are moving. Other techniques have been proposed, such as bending and strafing gains, but are not yet widely used by the community [19, 34, 46]. The amount of a gain that can be applied without users noticing, known as the detection threshold, has been estimated through empirical experimentation (e.g., [12, 17, 25, 35]). It is commonly accepted that when applying redirected walking, the gains should be set to a value within the established limits to avoid detection by the user.

Because gain values are capped by the detection thresholds, it is not always possible to prevent the user from leaving the boundary of the physical tracked space. One potential solution is to temporarily exceed the gain thresholds [6, 22, 27, 37]. This will likely cause users to notice the redirection and possibly induce negative reactions such as cybersickness. Therefore, it is usually preferred to use techniques that incorporate reorientation events [29]. These events stop the user before they cross a boundary and require them to perform a reorientation task. Typically, the virtual experience remains paused until the reorientation task is complete, which can be disruptive. However, it is possible to incorporate reorientation events into the experience in a contextually relevant way that does not break the user's sense of presence (e.g., [11, 28]). The reorientation task results in users physically facing a specific direction, such as the center of the physical space.

Resets are a type of reorientation event in which the task involves applying rotation gains while a user turns in place [43]. The two common types of resets are face-center and 2:1-turn, both of which require the user to rotate 360 degrees in the virtual world. The major difference between the two is that face-center resets apply rotation gain such that the user ends up physically facing the center of the tracking space [43], whereas 2:1-turn resets scale a 180 degree real rotation to a 360 degree virtual turn [29]. Recently, other types of resets have also been explored to accommodate for more complex physical space geometries [41]. In general, reorientation events interrupt the flow of an experience and should be used as a fail-safe to keep users from leaving the physical space. For this reason, the number of resets encountered is the primary metric used to evaluate the performance of redirected walking algorithms.

2.1 Redirected Walking Algorithms

The primary goal of redirected walking is to maximize the use of the physical tracked space. This is achieved by the selection of gains, generally on a frame by frame basis. If these gains are selected in an intelligent manner, the resulting physical path for a given virtual path will have fewer resets compared to a naive implementation of redirection. To achieve this, researchers have defined two categories of redirected walking algorithms: reactive and predictive The main difference between the two is whether or not the algorithm has access to information regarding the future virtual movements of the user. Reactive algorithms have no such information and have to choose gain values based only on the current and previous states of the system. Predictive algorithms have some varying amount of information about the user's future movements and can therefore apply gains strategically to reduce the occurrence of resets. Generally speaking, predictive algorithms outperform reactive algorithms, but reactive algorithms are easier to implement and can be used on a greater variety of virtual environment layouts.

Traditionally, Steer-to-Center (S2C) and Steer-to-Orbit (S2O) are the two most commonly used reactive algorithms [30]. S2C aims to redirect the user towards the center of the physical space, whereas S2O tries to steer the user into a circular path that has a center at the center of the physical space. Studies have shown that in most cases S2C outperforms S2O, although S2O does perform well in specific situations [3, 14]. These algorithms work on a simple heuristic (choosing gains that steer the user to a specific physical point) and are not suitable for more complex physical environments that may contain obstacles such as furniture. Two classes of reactive algorithms have been recently introduced to handle such situations. The first utilizes artificial potential fields to create a more complex heuristic [6, 9, 10, 22, 41], and the second relies on machine learning to optimize for a particular physical environment [20, 36]. These methods have been shown to outperform S2C in complex physical environments. However, thus far all reactive algorithms perform similarly when the physical environment rectangular and free of obstacles. For this reason, and to stay in line with

the majority of previous literature, this work makes use of the S2C algorithm in a square and obstacle-free physical environment.

Predictive approaches require some level of knowledge about the user's future movements. Planned path algorithms such as COPPER assume that the complete linear path of the user can be provided before the experience begins [1]. Although these are the most restrictive type of redirected walking algorithms, they can produce the best results for tightly choreographed experiences. Dynamic predictive algorithms can be deployed in a wider range of virtual environments (e.g. a maze or an office complex), but still require sufficient structure to create a model that predicts where the user is likely to move [24, 44, 47]. If users deviate from a planned path or do not behave as the prediction model suggests, a new trajectory must be calculated which may result in sub-optimal performance. Researchers have investigated methods for countering "drift" during walking tasks while gains are applied [2, 23].

2.2 Redirected Walking Simulation

One significant challenge in the literature has been gauging the performance of redirected walking algorithms. This is due to the fact that performance depends on a variety of interacting factors including user behavior, physical space dimensions, the structure of the virtual environment, and the type of virtual path. Internal parameters such as perceptual thresholds and the method used to recover from would-be collisions can add more complexity and makes comparative evaluation of such algorithms a non-trivial problem. The main way to overcome these factors is to increase the number of study participants and number of trials each participant performs. This may often require experimental designs that are impractical or even impossible to implement, which has motivated numerous researchers to evaluate redirected walking algorithms using simulation instead of live human subjects (e.g. [3, 5, 6, 14, 16, 22, 40, 41, 44, 47]).

Ultimately, the goal of the redirected walking algorithm is to continuously manipulate the mapping between the coordinate systems of the real world and the virtual environment such that the user can naturally move through a larger virtual environment in a smaller physical space. When simulations are employed to compare the efficacy of two or more algorithms, the user locomotion element is replaced with synthetic movements generated by the system. Redirection gains are then applied to the movements of the "simulated user." Therefore, a "good" simulation would produce outcomes that are sufficiently similar to those expected from live users, such that the performance metrics computed when evaluating different redirected walking strategies can be considered valid and generalizable.

The main criticism of simulation based redirected walking experiments is the lack of proof to show that results from a simulation based experiment can be transferred to real world experiences. The crux of this argument usually concerns how the simulated user, specifically its locomotion, is modeled. We have identified three main ways in which human walking can be modeled for simulation based experiments. The first is to generate a highly realistic model of human locomotion. This is a non-trivial problem affecting research in many fields including kinematics, cybernetics, and computer animation (e.g., [7, 13, 45]). Representing locomotion in this way requires complex modeling of the biomechanics of human movement and the ability to generate a broad variety of locomotion behaviors, but would produce outcomes that closely resemble real user movements at a micro-scale (i.e., at a granular frame-by-frame and throughout smaller segments of the virtual path). However, the performance measures for evaluating redirected walking algorithms, such as the total number of resets, are generally computed at the macro-scale (i.e., across the entire path taken in a VR experience). Therefore, complex biomechanical modeling or the laborious task of recording several user's locomotion of a given path may not be necessary if a simpler simulation model can reliably produce the macro-level behaviors and metrics relevant for performance evaluation.

The second approach defines rules for computing the orientation and position of the simulated user. For example, the simulated user can be programmed to translate and rotate at predetermined speeds and precisely follow a given virtual path. Due to the simplicity of this method, it is used in the vast majority of simulated redirected walking experiments. The third method was recently introduced by Bachmann et al. and uses pre-recorded users' tracking data as control for the simulated user [6]. This data-driven method could potentially incorporate the idiosyncrasies of natural human locomotion without requiring complex biomechanical modeling. However, it would only be possible to implement virtual paths for which you already have a sufficient quantity of recorded tracking data, and gathering this data can itself be a large undertaking.

3 User Study Design

This section describes our study that incorporates trials from both the micro- and macro-scale experiments. Real user locomotion data for this study was gathered through the virtual reality experience and compared to data generated from the simulation framework.

Participants. Thirty participants (18 male and 12 female) were recruited for our study using online postings (Craigslist volunteers section, Facebook posts and Twitter). Participants were between 18 and 70 years old with a mean age of 38.8 years and a median age of 40 years and were required to have normal or corrected normal vision. We did not pre-screen participants for gaming or virtual reality experience. Participants were paid USD $25 for their efforts. This study was reviewed and approved by the University's Institutional Review Board (IRB).

Apparatus. Participants wore an Oculus Rift CV1 head-mounted display (HMD) with Ausdom wireless stereo headphones connected to a Zotac VR GO Backpack PC that was strapped to their shoulders. The CV1 HMD has a 1080×1200 per eye resolution, 90Hz refresh rate, 110 degree nominal field of view (FOV), and an internal 9 degree of freedom (9DOF) inertial sensor. The participants were tracked in an 8 meter by 8 meter space with a PhaseSpace Impulse X2 motion capture system. To enable 6DOF head tracking by the Impulse X2 system, we attached 6 non-coplanar active LEDs to the CV1. At run-time, head orientation was tracked using the CV1's lower latency internal inertial sensor, while positional information was supplied by the PhaseSpace system. We also performed periodic orientation drift correction between trials using orientation data from the Impulse X2 system. Sound cues were provided using a pair of wireless "over the ear" Ausdom headphones. The virtual reality experience was designed and rendered using the Unity game engine.

Procedure. After signing the consent form, participants completed a pre-experiment Simulator Sickness Scale questionnaire [18]. Participants were then escorted to the physical space and introduced to the hardware used for the study. After putting on the equipment (Oculus HMD, Zotac Backpack, and Ausdom Headphones), they were loaded into an open-ended virtual environment with beige flooring and a blue sky (Figure 1). Users verbally verified they could clearly see and hear the visuals and test music playing and then proceeded with the tutorial. Instructions were explained through a voice-over heard on the headphones that guided them through all the elements of the experimental task. Participants were then asked to verbally confirm they had understood the tutorial elements and offered a chance to ask questions about the experience elements. Participants would then immediately begin the first phase of the study. After the first phase, the equipment would be taken off, allowing participants to take a 5-minute break (or longer if requested). Then the second phase would begin with the equipment being mounted again and resuming the virtual experience. With the second phase completed, the equipment was taken off and users asked to complete a post-experiment Simulator Sickness Scale questionnaire. At the end, they were also asked to provide basic demographics information, such as age and gender. Each of the two phases took 8-12 minutes depending on the participants movement speed. On average, the entire study took approximately 50 minutes including breaks.

Trials. The study task consisted of 32 interleaved trials, comprising 24 for the micro-level inspection, 4 for the macro-level inspection, and 4 for a counter-deviation experiment that is reported separately in [1]. These trials were randomized and split into two 16-trial phases, each containing the same number trials from each type.

Fig. 2. Sample illustration of locomotion behavior causing deviations during translation gain. The predicted user motion and its expected effect on the target waypoint (red) is contrasted with the actual results (black).

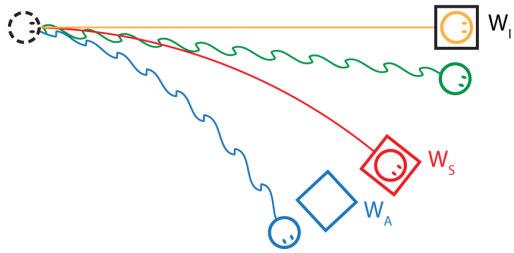

Fig. 3. Illustration of trajectories and end waypoints for calculating manipulation/deviation measures. The virtual trajectories for the simulated (yellow) and real user (green) are shown along with their corresponding real trajectories (red and blue). W_I is the end waypoint's initial pose at the beginning of a trial. W_A and W_S are the end waypoint poses for the real and simulated users, respectively.

Task. At the beginning of each trial, the participant was guided to the initial position and orientation. This was achieved by showing a directional visual cue pointing towards the trial's starting platform. The platform was placed on the floor, shaped as an arrow enclosed in a circle, pointing in the trials' starting direction. Once the user arrived at the platform and aligned their orientation with the arrow, the platform would be replaced with a starting line, and the first two segments of the trial's virtual path would be loaded. The path appeared as a 1 meter-wide walkway designated by stanchions with ropes. As the participant would begin to walk along the path, more segments of the path would appear, with previously cleared portions incrementally disappearing. This would done to prevent confusion caused by the virtual path potentially crossing itself since it was created by random generation. It is important to note that the user would always see the upcoming two path segments and the last cleared path segments at all times so that the the user would never be waiting on a segment to appear to make their navigation choices. Three elements were used to help with guiding the user along the correct direction: 1) arrows placed along the virtual path on the floor, 2) a 10 meter trail indicating the user's recent trajectory and 3) a "wrong way" sign appearing when the user faced more than 135 degrees away from the correct direction. To prevent collisions with the boundary, proximity cues were implemented using an alert sound and a red screen tint that gradually increased inversely proportional to the distance from boundary.

In reset trials (belonging to the micro-inspection) and path trials (belonging to the macro-inspection), occasionally a reset would be triggered. To initiate the reset, an audio cue was played, and as the scene faded to black, a stop sign would appear with the caption "Turn in Place." After sufficiently rotating in place, the sign would disappear and the environment would reappear, allowing the user to resume progression. Resets were triggered at the 3 meter mark in reset trials, and for path trials, whenever the user's distance to a boundary dropped below 1 meter. As an extra measure, to alert the user of her proximity to hard boundary limits, an alert tone would fade in and the screen would gradually tint to red with an audio and visual intensity proportional to boundary proximity. The end of the virtual path was delineated with a finish line on the floor similar to the starting line. Once the end of the path was reached, depending on the trials remaining, the user would either see a directional cue to the next platform or sign indicating that it was time to take off the headset.

Simulated User. To perform simulated experiments, we used a modified version of the evaluation platform included in the open-source Redirected Walking Toolkit [4]. A walking user was simulated by an autonomous agent programmed to traverse the virtual path by walking toward the next waypoint with a constant linear velocity of 1 m/s while maintaining its heading toward the waypoint (i.e. attempting to walk on a straight line in the virtual environment). Upon reaching a waypoint, the simulated user stopped and turned in place with angular velocity of 90 deg/s to face the next waypoint. Because physical boundaries prevent real users from exceeding physical space limits, simulated users must also be prevented from violating boundary constraints. Resets were therefore initiated using a safety trigger that was placed 1m from each side of the physical space boundary. Upon notification of a reset, the simulated user would decelerate at a constant rate, taking 0.5 seconds to come to a complete stop before beginning the reset.

For this study, no noise was introduced to the simulated user's translation and rotation. This guaranteed the simulated user would walk along the virtual path defined by the series of waypoints. Note that the simulated user is deterministic, therefore only one trial was run for each condition of the experiment.

4 Experiment 1: Micro-Scale Inspection

The first experiment examines elements of human locomotion that may influence how gains are applied. The simulated user abstracts a real person's movements by making the assumption that people walk in perfectly straight lines, only face the target direction while walking, and rotate around their local y-axis without deviating from the initial pivot point. Our goal with this part of the experiment is to isolate characteristics of human locomotion that contribute to the deviation from the expected results of the simulation. In order to control and accurately study the impact locomotion behavior has on redirection, we test each redirection technique in isolation. The redirection techniques that we studied were the three basic gains (translation, rotation, and curvature) and overt reorientations (resets).

To examine how translation gains are affected, the user walks along a 5 meter straight path with a fixed gain applied. In this case, we determine whether small swaying movements cause the user to deviate from the expected outcome (see Figure 2). We also investigate whether intermediary motions cancel each other out, causing the overall manipulation to only depend on the user's net displacement. Similar to the translation gain case, we will examine curvature gain by having the user walk along a 5 meter straight virtual path with a fixed curvature gain applied. For this inspection, we will try to explain variance in manipulations by the sum distance traveled and the user's average lateral distance from the virtual path. For rotation gain, the user traverses a 6 meter L-shaped path with gains only being applied when the user is transitioning from one leg of the path to the other. More specifically, rotation gain is enabled when the user completes 2/3rds of the first leg and disabled after completing the first 1/3rd of the second leg, effectively applying rotation gains within 1 meter before and after the turn. Finally, to examine resets, the user traverses a 6 meter straight path that triggers a reset half way through. In this case, we compare the result from the simulated user reaction model for resets to how a real user responds, similarly identifying the main aspects explaining deviations.

4.1 Manipulation and Deviation Measures

To analyze the effect of redirection across users, we define two sets of measures. *Manipulation* measures aim to gauge the overall effect of gains in comparison to their absence. *Deviation* measures compare how the effect of gains for a real user differ from that of a simulated user. We use Figure 3 as a sample trial for illustrating how these measures are calculated. The virtual trajectories of the simulated and real users are shown in yellow and green respectively, as they walk towards the square-shaped target end waypoint in the path. The corresponding real trajectories for the simulated and real users are shown in red and blue respectively, along with the resulting pose of the end waypoint.

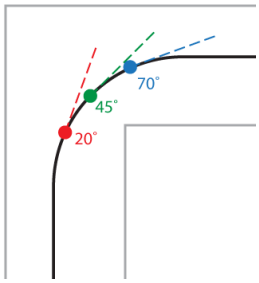

Fig. 4. Illustration of key points in the user's trajectory for estimating *turn point error* and *turn radius*. The first time the user's virtual angle with the target direction reaches 70, 45, and 20 degrees is marked.

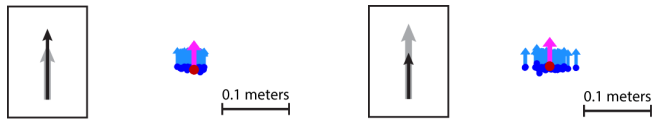

(a) $g_t < 1$ (b) $g_t > 1$

Fig. 5. End waypoint pose visualization for translation gain trials. The real pose of the end waypoints are shown for the real user trials (blue), depicting how they shift away from the expected end waypoint pose of the simulation (red). We observe an overall symmetric variance in the deviations for these trials. The icon to the left communicates the trial type and condition. The grey arrow indicates the virtual path traversed, while the black arrow shows the real path influenced by the gain.

To perform a fair and accurate comparison of the effect of gains across trials, we need to formulate measures that express the influence of redirection. To achieve this, we compare the pose of the end waypoint in the virtual world to its corresponding pose in the real world by the end of the trial (the blue and black squares in Figure 3). When no gains are applied, these waypoints will always be aligned, regardless of the user's end pose. However, after redirection, the user's motions interacting with the gains will cause the end waypoint's pose to translate and rotate relative to its initial pose at the beginning of the trial. When the user reaches the end of the trial, the overall influence of redirection can be interpreted by the changes in the reference pose. Our reference for no gains applied is therefore W_I, which is essentially the end waypoint's *initial* pose at the beginning of the trial. The *actual waypoint* W_A and the *simulation waypoint* W_S are the end waypoint poses for the real and simulated users respectively. To measure the overall manipulation, we compare W_A and W_I, and to measure deviation, we compare W_A and W_S. The distance between poses are compared to calculate positional manipulation/deviation, and the magnitude of angular difference between the poses is calculated to obtain rotational manipulation/deviation.

4.2 Hypotheses

For the micro-scale inspection, we hypothesize that the user's locomotion behavior determines the variance in overall manipulation. But more importantly, we hypothesize that the variance in manipulations among users along with the deviations from observed redirection results can be explained by measurable characteristics of locomotion behavior. This means that for each gain, factors can be identified that would explain variance. We suspect that for same gains, values throughout the trajectory would matter, for some, the overall change would only matter (effects of intermediate values would cancel out), and some variables would have little effect on the outcome.

4.3 Independent and Dependent Variables

The independent variables for micro-inspection are essentially the explanatory factors for the variance in redirection results. The following lists the independent variables for each gain type:

(i) Translation Gain
  (a) **translation gain** (amount of constant translation gain applied during trial)
  (b) **net translation** (amount of overall displacement between start and end of trial)

(ii) Curvature Gain
  (a) **curvature gain** (amount of constant curvature gain applied during trial)
  (b) **sum distance traveled** (sum of magnitude of user displacement between frames from start to end of trial)
  (c) **average lateral distance from virtual path** (distance between user and expected virtual path, averaged throughout the trial)

(iii) Rotation Gain
  (a) **rotation gain** (amount of constant rotation gain applied during trial)
  (b) **average lateral distance from virtual path** (distance between user and expected virtual path)
  (c) **turn point error** (distance between between turn waypoint and user's position when reaching a 45 degree angle with the target direction, indicating how far off the user was from the expected turn point when half of the 90 degree rotation was completed, Figure 4)
  (d) **turn radius** (approximated user turn radius measured by inspecting the two points along the path where the user's angle with the target direction reached 70 and 20 degrees, Figure 4).
  (e) **net turn angle** (overall change in rotation before and after turn area (the 1-meter proximity of the turn waypoint)

(iv) Resets
  (a) **displacement in reset** (overall change in position during the trial)
  (b) **mid-reset point error** (distance between mid-reset point (point where the reset task's overall virtual rotation reached 180 degrees) and the simulated user's reset task stopping point)
  (c) **end-reset point error** (distance between end-reset point (point where the reset task's overall virtual rotation reached 360 degrees) and the simulated user's reset task stopping point)
  (d) **mid-reset point distance to trigger** (distance between mid-reset point and where the user was when the reset was triggered)
  (e) **end-reset point distance to trigger** (distance between the end-reset point and where the user was when the reset was triggered)
  (f) **reset turn diameter** (approximated diameter of the circular trajectory of the user turning in place as measured by accounting for mid-reset and end-reset positions)

The dependent variables included the amounts of positional manipulation, positional deviation, and rotational manipulation. Note that rotational manipulation is only measured when relevant (translation gain does not introduce rotational manipulation, and a 2:1-turn reset affects orientation uniformly for all users). Additionally, since rotational deviation can be derived from rotational manipulation (by subtracting the simulated user's rotational manipulation from the user's rotational deviation), modeling rotational deviation would be redundant and is therefore not included.

4.4 Trials

The 24 trials in this category consisted of 10 walk trials, 12 turn trials, and 2 reset trials. The 10 walk trial conditions were $g_t = 1$, $g_t = 1.2$, $g_t = 0.86$, $g_c = 7.63$, $g_c = -7.63$, each being tested twice. The 12 turn trial conditions were in the form of {+90, -90} degree turns with a rotation gain of $g_r = \{0.8, 1, 1.49\}$, each being tested twice. The gain levels were selected based on internal pilot testing and to be within the upper and lower bounds of perceptual sensitivity observed in prior studies (e.g., [35]). The 2 reset trials both used a 2:1-Turn reset.

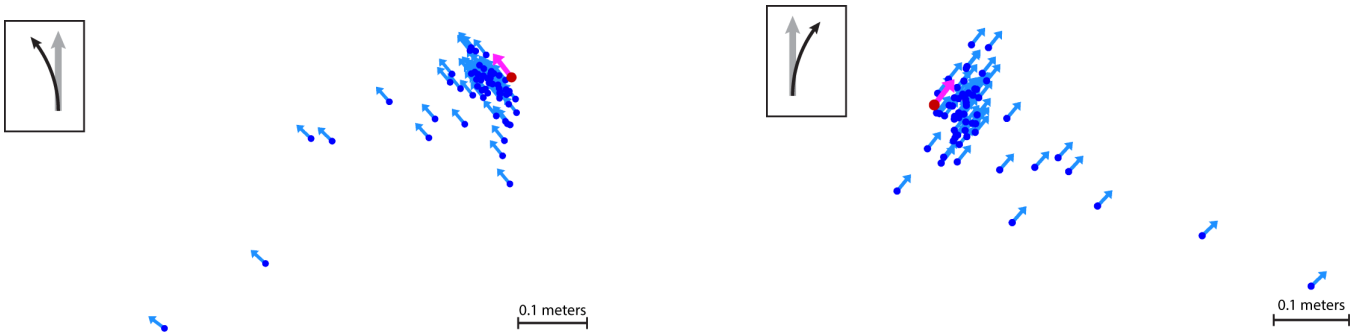

(a) $g_c < 1$ (b) $g_c > 1$

Fig. 6. End waypoint pose visualization for curvature gain trials. The real pose of the end waypoints are shown for the real user trials (blue), depicting how they shift away from the expected end waypoint pose of the simulation (red). We observe an over-application of curvature gain for real users in contrast to simulation. The icon to the left communicates the trial condition. The grey arrow indicates the virtual path traversed, while the black arrow shows the real path influenced by the gain.

4.5 Results

4.5.1 Translation Gain

Figure 5 compares end waypoint poses for all translation gain trials for real and simulated users. To examine how positional manipulation is affected by applied translation gain and net displacement, we constructed a multiple linear regression model (TG1). The model had a good fit ($R^2 = 0.99$). We treated translation gain as a two-level factor and net displacement as a continuous predictor variable. The model shows that both variables were significant predictors of positional manipulation (see Table 1). Positional manipulation increased by $0.19 \pm 0.01m$ for each meter of net displacement. On average, positional manipulation for human participants was 0.81m (SE = 0.00016) for $g_t = 0.86$ and 1.03m (SE = 0.00024) for $g_t = 1.26$.

We also constructed a multiple linear regression model to explore how positional deviation is affected by translation gain and net displacement as main effects (TG2). The model had a good fit ($R^2 = 0.99$). The models shows that both variables significantly affect positional deviation. Positional deviation increased by $0.18 \pm 0.01m$ for each meter of net displacement. The average positional deviation was 0.009m (SE = 0.00069) for $g_t = 0.86$ and 0.012m (SE = 0.0011) for $g_t = 1.26$.

Discussion. The results indicate that effect of translation gain can be explained by the net displacement, which implies that the shape of the virtual trajectory is not a factor and the gains essentially cancel each other out throughout the path. Although it can be argued that translation gain has a small impact on deviation when applied in isolation, it can however play a part in increasing deviation when combined with other gains. For instance, the user's lateral motions along the virtual path can cause intermediary deviations that can shift the pivot point of applied rotations, which can in turn magnify deviations.

4.5.2 Curvature Gain

Figure 6 compares end waypoint poses for curvature gain trials for real and simulated users. To examine how positional manipulation is affected by sum distance traveled and average lateral distance from the virtual path, we constructed a multiple linear regression model (CG1). We treated both sum distance traveled and average lateral distance from the virtual path as continuous predictor variables. The model shows that both variables were significant predictors (see Table 1). Positional manipulation increased by $0.32 \pm 0.01m$ for each meter of sum distance traveled and decreased by $0.17 \pm 0.01m$ for each meter of average lateral distance from the virtual path. On average, positional manipulation for human participants was 1.68m (SE = 0.011) for $g_c = -7.63$ and 1.69m (SE = 0.013) for $g_c = 7.63$.

We constructed a multiple linear regression model to explore how positional deviation is affected by these variables as main effects (CG2). The model shows that both were significant predictors. Positional deviation increased by $0.32 \pm 0.01m$ for each meter of sum distance traveled and decreased by $0.41 \pm 0.01m$ for each meter of average lateral distance from the virtual path. On average, positional deviation for human participants was 0.078m (SE = 0.011) for $g_c = -7.63$ and 0.082m (SE = 0.013) for $g_c = 7.63$.

Analysis	Factor	Regression	Sig.
TG1	translation gain	$F(1,117) = 2256535$	$< .01$
	displacement	$F(1,117) = 364.31$	$< .01$
TG2	translation gain	$F(1,117) = 136.16$	$< .01$
	displacement	$F(1,117) = 364.31$	$< .01$
CG1	distance traveled	$F(1,117) = 1523.00$	$< .01$
	lateral distance	$F(1,117) = 5.61$	.02
CG2	distance traveled	$F(1,117) = 1444.31$	$< .01$
	lateral distance	$F(1,117) = 31.64$	$< .01$
CG3	distance traveled	$F(1,117) = 563.35$	$< .01$
	lateral distance	$F(1,117) = 0.003$	.10

Table 1. Statistical results for the multiple linear regression analyses of translation gain and curvature gain trials.

We also constructed a multiple linear regression model to explore how rotational manipulation is affected by these variables as main effects (CG3). The model shows that sum distance traveled was a significant predictor of rotational manipulation, however average lateral distance from the virtual path was not significant. Rotational manipulation increased by $7.65 \pm 0.01°$ for each meter of sum distance traveled. On average, rotational manipulation for human participants was 39.37° (SE = 0.21) for $g_c = -7.63$ and 39.63° (SE = 0.33) for $g_c = 7.63$. For rotational deviation the average was 1.30° (SE = 0.21) for $g_c = -7.63$ and 1.56° (SE = 0.33) for $g_c = 7.63$.

Discussion. Based on the results, we can see that the sum distance traveled is the main factor for explaining how much curvature gain is applied. Therefore, a real user's excess lateral motions, in comparison to the simulated user, cause an increase in injected rotation gain. The user's lateral distance from the virtual path, which can be caused by both lateral motions and simply not walking along the middle of the path, affect the pivot point for applying the rotations of curvature gain. The variance in pivot point influences the positional measures, but since the amount of curvature gain applied is not affected, the rotational measures are not influenced.

4.5.3 Rotation Gain

Figure 7 compares end waypoint poses for rotation gain trials for real and simulated users. For each trial, we measured the user's virtual position when the angle between their orientation and the target turn

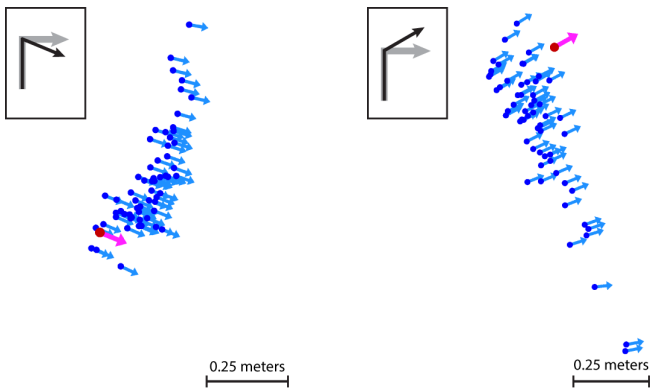

(a) $g_r < 1$ (b) $g_r > 1$

Fig. 7. End waypoint pose visualization for rotation gain trials. The real pose of the end waypoints are shown for the real user trials (blue), depicting how they shift away from the expected end waypoint pose of the simulation (red). We observe an under-application of rotation gain for real users in contrast to simulation. The icon to the left communicates the trial condition. The grey arrow indicates the virtual path traversed, while the black arrow shows the real path influenced by the gain.

orientation reached 70, 45 and 20 degrees. The distance between between the turn waypoint and user's position at 45 degrees was defined as the *turn point error*, indicating how far off the user was from the expected turn point when half of the 90 degree rotation was completed. The distance between the 70 degree and 20 degree marks was also used to approximate the *turn radius*. The *net turn angle* was also calculated as the overall change in rotation before and after the turn area (the 1 meter proximity of the turn waypoint).

To examine how positional manipulation is affected by rotation gain, turn point error, turn radius, net turn angle, and average lateral distance from the virtual path, we constructed a multiple linear regression model (RG1). We treated rotation gain as a two-level factor and all other factors as continuous predictor variables. The model shows that all variables except turn radius were significant predictors (see Table 2). Positional manipulation increased by $0.01 \pm 0.01m$ for each degree of net turn angle, decreased by $0.66 \pm 0.01m$ for each meter of average lateral distance from the virtual path, and increased by $0.08 \pm 0.01m$ for each meter of turn point error. On average, positional manipulation for human participants was 0.98m (SE = 0.020) for $g_r = 0.8$ and 1.26m (SE = 0.034) for $g_r = 1.49$.

We also constructed a multiple linear regression model to explore how positional deviation is affected by these variables as main effects (RG2). The model shows that all variables except turn radius were significant. Positional deviation decreased by $0.009 \pm 0.001m$ for each degree of net turn angle, increased by $0.57 \pm 0.01m$ for each meter of average lateral distance from the virtual path, and increased by $0.14 \pm 0.01m$ for each meter of turn point error. On average, positional deviation for human participants was 0.23m (SE = 0.019) for $g_r = 0.8$ and 0.32m (SE = 0.029) for $g_r = 1.49$.

We also constructed a multiple linear regression model to explore how rotational manipulation is affected by these variables as main effects (RG3). The model shows that all were significant. Rotational manipulation increased by $0.31 \pm 0.01°$ for each degree of net turn angle, decreased by $2.53 \pm 0.01°$ for each meter of average lateral distance from the virtual path, increased by $0.81 \pm 0.01°$ for each meter of turn point error and increased by $0.55 \pm 0.01°$ for each meter of turn radius. On average, rotational manipulation for human participants was 20.18° (SE = 0.41) for $g_r = 0.8$ and 25.11° (SE = 0.70) for $g_r = 1.49$. The average rotational deviation was 2.57° (SE = 0.36) for $g_c = -7.63$ and 4.57° (SE = 0.66) for $g_c = 7.63$.

Discussion. Overall, among the three gain types, we see that rotation gain has the largest impact on deviations. The amount of applied rotation gain can primarily be explained by the net turn angle. Therefore, similar to translation gain, rotation gains can cancel each other out when the user turns side to side, and only the overall change in heading is what determines the amount of rotation that is injected. For rotation gain trials, we observe that overall less rotation is injected compared to simulation, which is the opposite case in contrast to curvature gain. Upon further inspection, we noticed that this is due to the fact that users begin slightly facing towards the waypoint after the current target waypoint, even before reaching a 1 meter distance from their current target waypoint. Also, in some cases, users wouldn't fully turn towards the next waypoint. As a result, the net turn angle would be less than expected; thus, less rotation gain would be injected.

Factor	RG1	RG2	RG3
rotation gain	1211.73**	10.63*	4096.17**
turn point error	4.32*	16.97**	4.49*
turn radius	3.31	0.46	7.70**
net turn angle	705.80**	418.47**	3728.88**
lateral distance	34.71**	30.92**	5.20*

Table 2. Statistical results for the multiple linear regression analyses of rotation gain trials. Scores are $F(1,111)$, $^{*}p < .05$, $^{**}p < .01$.

Similar to curvature gain, pivot points heavily influence positional deviations for rotational gains, and this is also captured by the other factors (average lateral distance from the virtual path and turn point error). Upon further inspection, we noticed users tend to cut corners, essentially completing a turn segment without getting very close to the turn waypoint. This causes the rotation to be injected along points that differ from the turn waypoint, which is where the simulated user performs the turn. In the future, this turning behavior can potentially be modelled in simulated users to better resemble the average user, which should reduce the overall deviations along turns.

It is important to point out the fundamental difference between injecting translations (via translation gains) and rotations (via curvature and rotation gains). Only the vector of net translation influences the translation gain applied, regardless of the origin and where the translation took place. For rotations on the other hand, the specific place where the rotation is injected (the pivot point) matters, and variance in this pivot point manifests in positional deviations. Furthermore, with injecting rotations, small variances in rotation can result in positional deviations of great magnitude. Therefore, injecting rotations play a greater part in causing deviations. The magnitude of rotations involved with rotation gain is substantially greater than curvature gains, making rotation gain responsible for the majority of deviation as well.

4.5.4 Resets

Figure 8 compares end waypoint poses for rotation gain trials for real and simulated users. For each trial we measured the *reset displacement*, which is the overall change in position during the reset task. We also marked the user's *mid-reset* and *end-reset* positions in the real world, signifying when the reset task's overall virtual rotation had reached 180 and 360 degrees, respectively. These positions were compared to the point where the reset was triggered. We also measured *mid-reset* and *end-reset point error* by measuring the distance of these points to the simulated user's reset task stopping point. The reset turn diameter was also approximated by the distance between the mid-reset and end-reset positions.

A multiple linear regression model was calculated to predict positional manipulation based on displacement in reset, mid-reset point distance to trigger, end-reset point distance to trigger, and reset turn diameter as main effects (RST1). The model shows that all variables except reset displacement were significant (see Table 3). The average positional manipulation was 4.92m (SE = 0.071). It increased by $1.20 \pm 0.20m$ for every meter of reset turn diameter, decreased by $0.56 \pm 0.18m$ for every meter in mid reset point distance, and decreased by $1.53 \pm 0.57m$ for every meter in end reset point distance.

Similarly, a multiple linear regression model was calculated to predict positional deviation based on main effects of these variables (RS2).

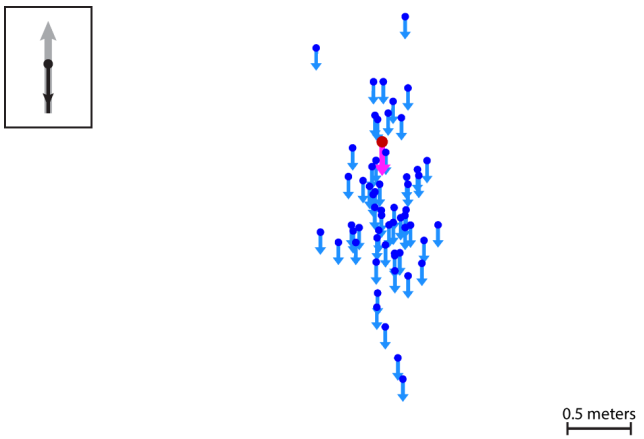

Fig. 8. End waypoint pose visualization for reset trials. The real pose of the end waypoints are shown for the real user trials (blue), depicting how they shift away from the expected end waypoint pose of the simulation (red). We observe translational deviations in both forward and backward directions. The icon to the left communicates the trial condition. The grey arrow indicates the virtual path traversed, while the black arrow shows the real path influenced by the reset.

The model shows that mid-reset point error and displacement in reset were significant factors. The average positional deviation was 0.66m (SE = 0.047). It increased by $1.42 \pm 0.16m$ for every meter of mid-reset error and decreased by $0.36 \pm 0.10m$ for every meter of displacement.

Discussion. Reset trials are somewhat comparable to rotation gain trials in that they both involve injecting rotation gains. However, in a 2:1-turn reset, the injected rotation is always 180 degrees, meaning the rotational measures are the same across users. On the other hand, the broad variations in how users react to resets substantially influence the pivot points for injecting the rotation, and thus results in variation in positional measure that is much greater than rotation gain trials. There is not a consistent overshooting or undershooting as we observed with curvature and rotation gain respectively, but rather early and late reactions can result in a decrease or increase in positional manipulation. It is also important to note that users begin rotating before coming to a full stop, which has a similar effect to an early reaction to reset. This explains the majority of end poses appearing lower than the simulated end pose in Figure 8. In general, the variation in pivot points for injecting rotation during the reset and their distance from the expected pivot point from the simulation explain the positional deviation, which is captured by the explaining factors.

The measured reset displacement is an important factor in determining the appropriate *reset buffer distance*, which dictates at what proximity to a boundary a reset must be triggered to keep users safe. Although the median reset displacement for real users is $0.76m$ (similar to $0.78m$ for the simulated users), this value greatly varies across users, which in turn explains variation in positional deviation. The 10th percentile for reset displacement was $0.46m$ and the 90th percentile was $1.21m$. Based on the data, a $1.21m$ reset buffer distance may be suitable. However, given the relatively high average age for our study population, a lower value such as 1 meter would probable be sufficient for the average user. Furthermore, the reset buffer distance can generally be adjusted for each user with a calibration process to both ensure safety for users with slower reactions and make better use of the physical space for those with quicker reflexes.

5 Experiment 2: Macro-Scale Inspection

The second experiment examines how human locomotion behavior influences and affects the overall performance of redirection. Although our simulation may not perfectly replicate exactly how a real person would traverse a virtual path segment, it may be able to provide a useful high-level performance metric for estimating how different redirected walking approaches can be expected to perform on a potential path or environment layout. This examination will explore whether user locomotion behavior causes deviations that will either ultimately impact the efficacy of redirection or cancel out over time. Additionally, this part of the study will compare the results for real users with the simulation in order to determine two other features. First, the comparison will identify and estimate systematic error of our simulation platform. Furthermore, the juxtaposition of real and simulated data will help us determine if relative performance trends are preserved. In other words, if the simulation expects a specific configuration to out-perform another, will this prediction will also hold true for real users?

For this examination, the user traverses a random, procedurally generated virtual path made up of straight segments connected by randomly selected 90 degree left or right turns. Each path comprised 20 waypoints over total walking distance of 100 meters. The length of the straight portions of the path were uniformly distributed between 2 to 8 meters, inclusively. Gains are applied along the path according to the chosen redirection strategy and resets are used to ensure that the user remains in the physical space. The performance metric used to compare the results is the reset count, which is a critical factor in measuring the efficacy of redirected walking. Additionally, it captures the overall impact of locomotion manipulation variability across users.

Analysis	Factor	Regression	Sig.
RST1	reset displacement	$F(1,54) = 0.52$	.81
	mid-reset point distance	$F(1,54) = 9.96$	$< .01$
	end-reset point distance	$F(1,54) = 7.22$	$< .01$
	reset turn diameter	$F(1,54) = 37.78$	$< .01$
RST2	reset displacement	$F(1,54) = 13.23$	$< .01$
	mid-reset point error	$F(1,54) = 81.22$	$< .01$
	end-reset point error	$F(1,54) = 2.17$	.15
	reset turn diameter	$F(1,54) = 2.47$	.12

Table 3. Results for the multiple linear regression analyses of reset trials.

5.1 Hypotheses

For the macro-scale inspection, we hypothesize that results from simulation will preserve the performance trends observed with real user data. In this experiment, this will be investigated by testing if one strategy outperforms another based on real user data, and then evaluating whether with this relationship also holds true within the simulation platform. More generally, we are interested in evaluating whether the average overall redirection performance for real users at a macro-scale is similar to results when using simulated users.

5.2 Independent and Dependent Variables

This experiment examined the following independent variables:

(a) **user** (the set of real users is contrasted against a single deterministic simulated user)
(b) **redirection strategy** (S2C and S2O both paired with a the face-center reset form the two strategies under study)
(c) **virtual path** (two procedurally generated random paths with 90-degree turns)

The performance factor measured as the dependent variable was the total number of resets triggered in each trial.

5.3 Trials

The 4 trials in this category were 2 different procedurally generated, random virtual paths, each tested once with the steer-to-center strategy and once with steer-to-orbit. In all case, a face-center reset was set to be triggered 1 meter from the boundary. The paths were constructed from 20 waypoints (not including the initial point), with random 90 or -90 degree turns, each placed between 2 to 8 meters (sampled uniformly) apart, for an expected length of 100 meters.

5.4 Results

The reset count for each redirection strategy in each path is shown for real users vs. the simulated user in Figure 9. For path 1 with S2C, using a one-sample t-test comparing the expected reset count of 11 from the simulated user compared to real users (M=11.1 , SD=1.37) we failed to reject the null hypothesis ($t(29) = 0.39, p = 0.693$). For path 1 with S2O, a one-sample t-test comparing the expected reset count of 15 from the simulated user compared to real users (M=14.3 , SD=1.17) was significant ($t(29) = -3.25, p = 0.003$). For path 2 with S2C, a one-sample t-test comparing the expected reset count of 12 from the simulated user compared to real users (M=11.23 , SD=1.00) was also significant ($t(29) = -4.17, p < 0.001$). Finally, for path 2 with S2O, a one-sample t-test comparing the expected reset count of 15 from the simulated user compared to real users (M=15.2 , SD=1.65) was not significant ($t(29) = 0.88, p = 0.38$).

Discussion. The results from our study indicate an overall similarity between redirection performance measured via simulation and the average performance for real users. Overall, we can argue that results from simulation can provide a conservative estimate for average performance with real users. More importantly, results from the simulation preserved the trends present in real user data, specifically S2C outperforming S2O in both paths across real and simulated users.

The similarity in performance between real and simulated users may seem to contradict our previous observations that deviations rapidly manifest in short trials. However, deviation inherently does not entail an increase or decrease in performance, but rather implies that a real user's path is becoming out of sync with the simulated user. The diverging trajectories can potentially lead to a real user observing a reset earlier or later than the simulated user, with neither of the two scenarios being necessarily more likely than the other.

Redirection performing better for some users in comparison to our simulation can potentially be explained by asymmetric gains. S2C and S2O are greedy strategies that take the user's motions at each frame as an opportunity to get closer to a target objective (the user facing the center or orbiting around the center). This means real users who exhibit more motion than a simulated user (such as excess lateral motions and head motions) have more chances to steer themselves, potentially improving the performance of the redirection strategy. In contrast to using symmetric gains, an asymmetric gain approach uses motions only to improve the user's state, and will stop injecting motions if necessary. Therefore, extra motions may not necessarily be conducive to the success of a planned redirection strategy that uses fixed gains.

While the performance of redirection for some users can be better than our simulation, in other cases we notice a decline in performance. We believe that this could be attributed to resets being triggered in situations where they may not be absolutely necessary. The most common example of this would be a user brushing against a reset trigger due to side-to-side lateral motions while walking almost parallel to the physical space boundary. Another example would be back-to-back resets triggering when a user struggles with both performing the reset task and resuming progression in the correct direction. These situations can be avoided by developing more sophisticated reset techniques.

6 Limitations and Future Work

The results presented in this paper have several limitations. The experiments were designed to evaluate the efficacy of simulations without addition of "noise" that mimics motion tracking jitter or postural sway from natural walking. The simulation also used a fixed walking speed, and the effects of velocity changes on deviations are not yet known. Future studies to investigate these factors at the micro-scale would be valuable. The micro-scale inspection can also be extended by examining how combined gains (such as simultaneous translation and curvature) can manifest deviations. Furthermore, trials with asymmetric gains can help with identifying alternative locomotion behaviors that cause deviations and help to further explain the behavior of strategies such as S2C. Another avenue also worth investigating is how the application of gains influences the virtual trajectory. We would like to understand whether an observer in the virtual world can tell if gains are being applied to a user or not, just by observing the user's walking behavior.

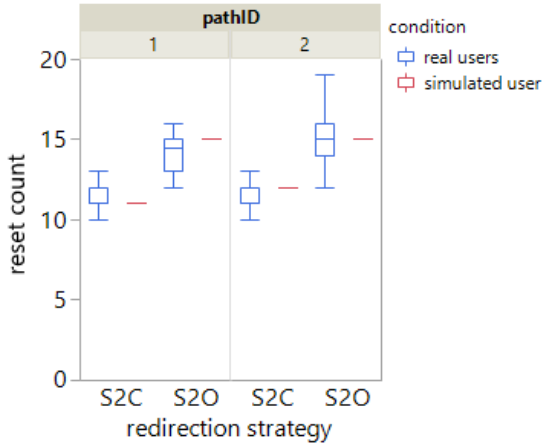

Fig. 9. Reset counts for macro inspection trials across two test virtual paths and two redirection strategies. Since the simulated user is deterministic, no variations are present in the corresponding data.

The macro-scale inspection focused on two common reactive strategies (S2C and S2O), and there have been a variety of more advanced reactive and predictive strategies introduced in recent years. Although our results generally affirm the validity of simulation-based evaluation, more research is needed to compare simulation and live users using a wider variety of redirected walking algorithms, such as approaches based on artificial potential fields [6, 9, 10, 22, 41] and machine learning [20, 36]. It would also be valuable to conduct studies using more complex physical space configurations, such as those with non-convex boundaries or interior obstacles, and different types of virtual reality experiences that use redirected walking, such as those that integrate passive haptics [40].

7 Conclusion

In this work, we investigated how redirection gains interact with user locomotion behavior to understand how the outcome of redirection can vary across real and simulated users. By looking at the gains applied at a micro-scale, we learned that the main sources of deviation are lateral movements causing increased curvature gains, premature and gradual turning along with shortcuts influencing rotation gain pivot points, and variations in reaction to resets, with the most deviation being caused by the latter two. Looking at gains applied at a macro-scale, we observed similar average performance between users and our simulated model, with simulated user results offering a conservative estimate for the average performance of real users. We observed that though deviations cause real trajectories to diverge, they did not seem to directly decrease or increase performance.

The investigation of how deviations manifest and affect performance sheds light on the differences we can expect from redirection when applied with real users versus simulation. These findings open the door to performing simulation-based studies that can help with understanding redirection performance factors, testing and developing novel strategies and algorithms, and also performing cost-benefit analyses for a custom virtual reality configuration. We envision researchers and developers performing large-scale experiments and prototyping novel approaches using simulation, rapidly advancing the field of redirected walking.

Acknowledgments

This work was sponsored by the U.S. Army Research Laboratory (ARL) under contract number W911NF-14-D-0005. Statements and opinions expressed do not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

REFERENCES

[1] M. Azmandian. *Design and Evaluation of Adaptive Redirected Walking Systems*. PhD thesis, University of Southern California, 2018.

[2] M. Azmandian, M. Bolas, and E. Suma. Countering user deviation during redirected walking. In *Proceedings of the ACM Symposium on Applied Perception*, pp. 129–129. ACM, 2014.

[3] M. Azmandian, T. Grechkin, M. Bolas, and E. Suma. Physical space requirements for redirected walking: how size and shape affect performance. In *Proceedings of the 25th International Conference on Artificial Reality and Telexistence and 20th Eurographics Symposium on Virtual Environments*, pp. 93–100. Eurographics Association, 2015.

[4] M. Azmandian, T. Grechkin, M. Bolas, and E. Suma. The redirected walking toolkit: A unified development and deployment platform for exploring large virtual environments. In *Everyday VR Workshop, IEEE VR*, 2016.

[5] M. Azmandian, T. Grechkin, and E. S. Rosenberg. An evaluation of strategies for two-user redirected walking in shared physical spaces. In *Virtual Reality (VR), 2017 IEEE*, pp. 91–98. IEEE, 2017.

[6] E. R. Bachmann, E. Hodgson, C. Hoffbauer, and J. Messinger. Multi-user redirected walking and resetting using artificial potential fields. *IEEE Transactions on Visualization and Computer Graphics*, 25(5):2022–2031, 2019.

[7] R. Boulic, N. M. Thalmann, and D. Thalmann. A global human walking model with real-time kinematic personification. *The visual computer*, 6(6):344–358, 1990.

[8] G. Bruder, P. Lubas, and F. Steinicke. Cognitive resource demands of redirected walking. *IEEE transactions on visualization and computer graphics*, 21(4):539–544, 2015.

[9] T. Dong, X. Chen, Y. Song, W. Ying, and J. Fan. Dynamic artificial potential fields for multi-user redirected walking. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 146–154. IEEE, 2020.

[10] T. Dong, Y. Shen, T. Gao, and J. Fan. Dynamic density-based redirected walking towards multi-user virtual environments. In *IEEE Virtual Reality and 3D User Interfaces*, pp. 626–634, 2021.

[11] T. Grechkin, M. Azmandian, M. Bolas, and E. Suma. Towards context-sensitive reorientation for real walking in virtual reality. In *2015 IEEE Virtual Reality (VR)*, pp. 185–186. IEEE, 2015.

[12] T. Grechkin, J. Thomas, M. Azmandian, M. Bolas, and E. Suma. Revisiting detection thresholds for redirected walking: combining translation and curvature gains. In *Proceedings of the ACM Symposium on Applied Perception*, pp. 113–120. ACM, 2016.

[13] M. Günther and H. Ruder. Synthesis of two-dimensional human walking: a test of the λ-model. *Biological cybernetics*, 89(2):89–106, 2003.

[14] E. Hodgson and E. Bachmann. Comparing four approaches to generalized redirected walking: Simulation and live user data. *IEEE transactions on visualization and computer graphics*, 19(4):634–643, 2013.

[15] E. Hodgson, E. Bachmann, and D. Waller. Redirected walking to explore virtual environments: Assessing the potential for spatial interference. *ACM Transactions on Applied Perception (TAP)*, 8(4):22, 2011.

[16] J. E. Holm. *Collision Prediction and Prevention in a Simultaneous Multi-User Immersive Virtual Environment*. PhD thesis, Miami University, 2012.

[17] C. Hutton, S. Ziccardi, J. Medina, and E. Suma Rosenberg. Individualized calibration of rotation gain thresholds for redirected walking. In *ICAT-EGVE*, 2018.

[18] R. S. Kennedy, N. E. Lane, K. S. Berbaum, and M. G. Lilienthal. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology*, 3(3):203–220, 1993.

[19] E. Langbehn, P. Lubos, G. Bruder, and F. Steinicke. Bending the curve: Sensitivity to bending of curved paths and application in room-scale vr. *IEEE transactions on visualization and computer graphics*, 23(4):1389–1398, 2017.

[20] D.-Y. Lee, Y.-H. Cho, and I.-K. Lee. Real-time optimal planning for redirected walking using deep q-learning. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 63–71. IEEE, 2019.

[21] J. Lishman and D. Lee. The autonomy of visual kinaesthesis. *Perception*, 2(3):287–294, 1973.

[22] J. Messinger, E. Hodgson, and E. R. Bachmann. Effects of tracking area shape and size on artificial potential field redirected walking. In *IEEE Conference on Virtual Reality and 3D User Interfaces*, 2019.

[23] R. A. Montano-Murillo, P. I. Cornelio-Martinez, S. Subramanian, and D. Martinez-Plasencia. Drift-correction techniques for scale-adaptive vr navigation. In *ACM Symposium on User Interface Software and Technology*, p. 1123–1135, 2019.

[24] T. Nescher, Y.-Y. Huang, and A. Kunz. Planning redirection techniques for optimal free walking experience using model predictive control. In *3D User Interfaces (3DUI), 2014 IEEE Symposium on*, pp. 111–118. IEEE, 2014.

[25] C. T. Neth, J. L. Souman, D. Engel, U. Kloos, H. H. Bulthoff, and B. J. Mohler. Velocity-dependent dynamic curvature gain for redirected walking. *IEEE transactions on visualization and computer graphics*, 18(7):1041–1052, 2012.

[26] N. C. Nilsson, T. Peck, G. Bruder, E. Hodgson, S. Serafin, M. Whitton, F. Steinicke, and E. S. Rosenberg. 15 years of research on redirected walking in immersive virtual environments. *IEEE Computer Graphics and Applications*, 38(2):44–56, 2018.

[27] N. Nitzsche, U. D. Hanebeck, and G. Schmidt. Motion compression for telepresent walking in large target environments. *Presence: Teleoperators & Virtual Environments*, 13(1):44–60, 2004.

[28] T. C. Peck, H. Fuchs, and M. C. Whitton. Evaluation of reorientation techniques and distractors for walking in large virtual environments. *IEEE transactions on visualization and computer graphics*, 15(3):383–394, 2009.

[29] T. C. Peck, H. Fuchs, and M. C. Whitton. Improved redirection with distractors: A large-scale-real-walking locomotion interface and its effect on navigation in virtual environments. In *Virtual Reality Conference (VR), 2010 IEEE*, pp. 35–38. IEEE, 2010.

[30] S. Razzaque. *Redirected walking*. University of North Carolina at Chapel Hill, 2005.

[31] S. Razzaque, Z. Kohn, and M. C. Whitton. Redirected walking. In *Proceedings of EUROGRAPHICS*, vol. 9, pp. 105–106. Citeseer, 2001.

[32] R. A. Ruddle and S. Lessels. The benefits of using a walking interface to navigate virtual environments. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 16(1):5, 2009.

[33] R. A. Ruddle, E. Volkova, and H. H. Bülthoff. Walking improves your cognitive map in environments that are large-scale and large in extent. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 18(2):10, 2011.

[34] A. Serrano, D. Martin, D. Gutierrez, K. Myszkowski, and B. Masia. Imperceptible manipulation of lateral camera motion for improved virtual reality applications. *ACM Transactions on Graphics*, 39(6), 2020.

[35] F. Steinicke, G. Bruder, J. Jerald, H. Frenz, and M. Lappe. Estimation of detection thresholds for redirected walking techniques. *IEEE transactions on visualization and computer graphics*, 16(1):17–27, 2010.

[36] R. R. Strauss, R. Ramanujan, A. Becker, and T. C. Peck. A steering algorithm for redirected walking using reinforcement learning. *IEEE transactions on visualization and computer graphics*, 26(5):1955–1963, 2020.

[37] J. Su. Motion compression for telepresence locomotion. *Presence: Teleoperators and Virtual Environments*, 16(4):385–398, 2007.

[38] E. Suma, S. Finkelstein, M. Reid, S. Babu, A. Ulinski, and L. F. Hodges. Evaluation of the cognitive effects of travel technique in complex real and virtual environments. *IEEE Transactions on Visualization and Computer Graphics*, 16(4):690–702, 2010.

[39] E. A. Suma, D. M. Krum, S. Finkelstein, and M. Bolas. Effects of redirection on spatial orientation in real and virtual environments. In *3D User Interfaces (3DUI), 2011 IEEE Symposium on*, pp. 35–38. IEEE, 2011.

[40] J. Thomas, C. Hutton Pospick, and E. Suma Rosenberg. Towards physically interactive virtual environments: Reactive alignment with redirected walking. In *26th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–10, 2020.

[41] J. Thomas and E. S. Rosenberg. A general reactive algorithm for redirected walking using artificial potential functions. In *IEEE Conference on Virtual Reality and 3D User Interfaces*, 2019.

[42] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks Jr. Walking¿ walking-in-place¿ flying, in virtual environments. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pp. 359–364. ACM Press/Addison-Wesley Publishing Co., 1999.

[43] B. Williams, G. Narasimham, B. Rump, T. P. McNamara, T. H. Carr, J. Rieser, and B. Bodenheimer. Exploring large virtual environments with an hmd when physical space is limited. In *Proceedings of the 4th symposium on Applied perception in graphics and visualization*, pp. 41–48. ACM, 2007.

[44] N. L. Williams, A. Bera, and D. Manocha. Redirected walking in static

and dynamic scenes using visibility polygons. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4267–4277, 2021.

[45] Q. Yang, J. Qin, and S. Law. A three-dimensional human walking model. *Journal of Sound and Vibration*, 357:437–456, 2015.

[46] C. You, E. Suma Rosenberg, and J. Thomas. Strafing gain: A novel redirected walking technique. In *ACM Symposium on Spatial User Interaction*, p. 26. ACM, 2019.

[47] M. A. Zmuda, J. L. Wonser, E. R. Bachmann, and E. Hodgson. Optimizing constrained-environment redirected walking instructions using search techniques. *IEEE transactions on visualization and computer graphics*, 19(11):1872–1884, 2013.