

Towards Virtual Reality Infinite Walking: Dynamic Saccadic Redirection

QI SUN, Stony Brook University, NVIDIA, and Adobe Research

ANJUL PATNEY, NVIDIA

LI-YI WEI, Adobe Research

OMER SHAPIRA, NVIDIA

JINGWAN LU and PAUL ASENTE, Adobe Research

SUWEN ZHU, Stony Brook University

MORGAN MCGUIRE and DAVID LUEBKE, NVIDIA

ARIE KAUFMAN, Stony Brook University




Fig. 1. *Triggering and harnessing temporary blindness via saccades for room-scale redirected walking in VR.* Our system renders a virtual environment into a pair of HMD views while tracking the user eye gaze. (a) shows a HMD left eye rendering for the viewer with overlaid visualizations of tracked eye gaze (green circle) and view frustum (lower left corner). When saccades (rapid eye movements) and head rotations are detected, our system rotates the virtual environments to redirect the users (b). Such rotations are visible during normal viewing conditions, but can be imperceptible during eye or head movements. (c) photographs our experimental setup with a Vive HMD augmented with SMI gaze tracking. Superimposed are the top view of the recorded movements of the physical path in a $3.5\text{m} \times 3.5\text{m}$ real room and the virtual path in a much larger $6.4\text{m} \times 6.4\text{m}$ synthetic space. Scene courtesy of NOT_Lonely (Vitaly).

Redirected walking techniques can enhance the immersion and visual-vestibular comfort of virtual reality (VR) navigation, but are often limited by the size, shape, and content of the physical environments.

We propose a redirected walking technique that can apply to small physical environments with static or dynamic obstacles. Via a head- and eye-tracking VR headset, our method detects saccadic suppression and redirects the users during the resulting temporary blindness. Our dynamic path planning runs in real-time on a GPU, and thus can avoid static and dynamic obstacles, including walls, furniture, and other VR users sharing the same

physical space. To further enhance saccadic redirection, we propose subtle gaze direction methods tailored for VR perception.

We demonstrate that saccades can significantly increase the rotation gains during redirection without introducing visual distortions or simulator sickness. This allows our method to apply to large open virtual spaces and small physical environments for room-scale VR. We evaluate our system via numerical simulations and real user studies.

CCS Concepts: • **Human-centered computing** → **Virtual reality**; • **Computing methodologies** → **Perception**;

Additional Key Words and Phrases: virtual reality, redirected walking, human perception, saccade

ACM Reference Format:

Qi Sun, Anjul Patney, Li-Yi Wei, Omer Shapira, Jingwan Lu, Paul Asente, Suwen Zhu, Morgan McGuire, David Luebke, and Arie Kaufman. 2018. Towards Virtual Reality Infinite Walking: Dynamic Saccadic Redirection. *ACM Trans. Graph.* 37, 4, Article 67 (August 2018), 13 pages. <https://doi.org/10.1145/3197517.3201294>

Authors' addresses: Qi Sun, Stony Brook University, NVIDIA, Adobe Research; Anjul Patney, NVIDIA; Li-Yi Wei, Adobe Research; Omer Shapira, NVIDIA; Jingwan Lu; Paul Asente, Adobe Research; Suwen Zhu, Stony Brook University; Morgan McGuire; David Luebke, NVIDIA; Arie Kaufman, Stony Brook University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2018/8-ART67 \$15.00 <https://doi.org/10.1145/3197517.3201294>

1 INTRODUCTION

Room-scale virtual reality (VR) increases presence and decreases discomfort caused by visual-vestibular inconsistency by allowing the user to walk freely in a physical space [Usoh et al. 1999]. However, a direct one-to-one mapping from virtual to physical space is impractical for most applications. Today's room-scale experiences either constrain the virtual space through scenario design or frequently interrupt the users and break their sense of presence by requiring them to walk back to the center of the physical room or consciously teleport in the virtual world. A major challenge for VR is embedding a large virtual space within a small, irregular, multi-user physical space while minimizing interruptions. The ideal solution would create the perception of infinite walking in the virtual space within a small, finite physical space.

Treadmills or other physical devices can address the infinite walking problem, but are undesirable for general applications because they are expensive, bulky, and can compromise the user's balance, while also preventing free user movements such as kneeling and jumping. The current state-of-the-art techniques for solving the mapping problem using only a head-mounted display (HMD) are redirected walking [Razzaque et al. 2001, 2002; Steinicke et al. 2010] and warping [Dong et al. 2017; Sun et al. 2016]. These methods create a distorted mapping of the virtual environment by applying to the world subtle rigid-body and nonlinear transformations, respectively. These magnify the effective physical space, but state-of-the-art methods still require an unoccluded space of 36 m² to be simultaneously imperceptible and effective [Azmandian et al. 2015]. This is a significant step towards practical room-scale VR for unconstrained scenarios, but it is still too large to accommodate many home and office rooms. We believe the main cause is the perceptually-imposed limitation of traditional redirection systems that cannot respond to the real-time user and environmental changes.

We present a novel, *dynamic* solution to the infinite walking problem. It is the first to be demonstrated as effective for physical areas as small as 12.25 m². This significant advance beyond previous results meets for the first time the standard for practicality: these bounds match the recommended consumer HMD room-scale installation bounds, e.g., for HTC Vive and Oculus Rift. Our key innovation is redirecting the user much more aggressively, yet still imperceptibly, by tracking rapid eye movements called saccades using a HMD equipped with internal gaze-tracking cameras, and incorporating guided navigation and planning based on the scenario.

Saccades are rapid eye movements, during which viewers are momentarily blind in a phenomenon called *saccadic suppression*. Saccades occur frequently, but our high-level visual system prevents conscious awareness of the blindness. The visual system also essentially recalibrates its orientation after a saccade on the assumption that the world itself has not changed [Hopp and Fuchs 2004]. We exploit that assumption to change the virtual world imperceptibly and avoid predicted future collisions with physical objects. Our method retains faithful visual and vestibular experiences across a broader range of virtual and physical spaces than previous methods. To further enhance the effectiveness of the technique, we also employ subtle gaze directions to opportunistically trigger additional

saccades, and a content-aware path planner to adapt to dynamic environmental changes. Our main contributions are:

- An end-to-end redirected walking system based on saccadic suppression, effective for consumer room-scale VR;
- A real-time path planning algorithm that automatically avoids static and dynamic obstacles by responding to individuals' eye movements – our optimization links user behavior and physical changes, considers possibilities of near future through real-time sampling, and finds the best numerical solution for online camera manipulation;
- The use of subtle gaze direction (SGD) methods in VR to induce more saccades for the system to exploit;
- Validation through simulations and real redirected walking scenarios with game-like tasks, such as search and retrieval.

2 RELATED WORK

2.1 Redirected Interaction in VR

Redirected interaction, such as walking [Dong et al. 2017; Hodgson and Bachmann 2013; Razzaque 2005; Razzaque et al. 2001; Sun et al. 2016] and touching [Azmandian et al. 2016c; Cheng et al. 2017], has received recent attention in the graphics and HCI community as a technique that uses mapping and rendering methods to enhance presence. It works by modifying what the user sees while they are physically interacting with their surroundings [Azmandian et al. 2017]. Due to the dominance of vision over other senses, the user perceives the physical interaction as being consistent to the visual stimulus. This way, physical interactions can be redirected. In particular, redirected walking can influence the user's walking path in an imperceptible fashion, simulating larger virtual environments within smaller physical ones and avoiding walls and obstacles.

Researchers have proposed two primary methods of redirected walking: those that work by dynamically scaling user motion and head rotation for the virtual camera [Azmandian et al. 2017; Razzaque et al. 2001, 2002; Steinicke et al. 2010] due to sensory conflicts in virtual environments [Steinicke et al. 2008], and those that work by warping the virtual scene [Dong et al. 2017; Sun et al. 2016].

Notwithstanding the specific technique, contemporary redirected techniques assume that users are aware of the environment at all times. The techniques do not consider perceptual masking effects like saccades, blinks, and other perceptual suppressions. In this paper, we enhance redirected interaction by detecting masking effects and amplifying redirection during the events without introducing virtual scene warping. Concurrent work by Langbehn et al. [2018] conducts perceptual experiments to measure translation and rotation thresholds during eye blinks to facilitate redirected walking. In comparison, our method considers rotations but not translations during eye saccades, in conjunction with subtle gaze direction and GPU path planning for real-time performance.

2.2 Gaze-contingent Rendering in VR

Gaze-contingent graphics is a widely studied area with several applications in medicine, optometry, vision science, and computer graphics [Duchowski et al. 2004; Reder 1973]. However, due to the increasing availability of high-quality eye trackers [Vincent and

Brannan 2017] as well as growing research into potential applications, gaze-contingent rendering has also gained popularity in virtual and augmented reality. When used for foveated rendering, gaze-contingent rendering helps improve visual quality without performance compromises [Albert et al. 2017; Duchowski and Çöltekin 2007; Guenter et al. 2012; Levoy and Whitaker 1990; Luebke and Hallen 2001; Patney 2017; Patney et al. 2016]. When used to simulate high-dynamic range [Jacobs et al. 2015] or to provide focus cues [Duchowski et al. 2014], gaze-contingent graphics enable new experiences on contemporary displays. Finally, eye-tracking is a useful interaction tool for virtual environments [Pfeuffer et al. 2017]. Thus, eye-tracking support is foreseen in the next generation commodity VR/AR devices. Our system employs eye-tracking to determine occurrences of perceptual suppression for VR redirected walking.

2.3 Saccadic and Blink Suppression

A saccade is the rapid eye movement that occurs when we change fixation points. During normal viewing, saccades occur several times a second, contain extremely fast motion (up to $900^\circ/\text{sec}$), and are long (20–200 ms) compared to VR frame durations [Bahill et al. 1975], although the speed at which they can be detected varies depending upon the chosen algorithms [Andersson et al. 2017]. Saccades are among many behaviors that trigger temporary perceptual suppression. Others include masking by patterns, tactile saccades [Ziat et al. 2010], and blinks [Ridder III and Tomlinson 1997]. While our system for redirected walking could potentially extend to any of these, we explicitly evaluate it under saccades in this paper.

Saccadic suppression (a.k.a. *saccadic omission*) of perception occurs before, during, and after each saccadic eye motion [Burr et al. 1994]. While the exact mechanism behind it is an area of active research [Burr et al. 1994; Diamond et al. 2000; Ibbotson and Cloherty 2009], the characteristics are well-known [Matin 1974; McConkie and Loschky 2002; Ross et al. 2001]. Our system exploits the particular documented phenomenon – suppression of *image displacement* [Bridgeman et al. 1975; Li and Matin 1990].

A key property of visual saccades is that they are ballistic in nature [Bahill et al. 1975] and their velocity profile and landing position can often be predicted mid-flight [Arabadzhiyska et al. 2017; Han et al. 2013]. This, in addition to saccadic suppression lasting for a short period after the saccade itself completes, suggests that detecting saccades and altering rendering based on the detection should be fairly tolerant of current VR eye-tracking-to-photon latency of around 35 ms [Albert et al. 2017]. Recent work established reorientation and repositioning thresholds for VR during saccades [Bolte and Lappe 2015] and blinks [Langbehn et al. 2016]. We leverage those established perceptual thresholds to build and evaluate a redirected walking system.

2.4 Subtle Gaze Direction

Subtle gaze direction (SGD) uses image-space modulation to direct a viewer's gaze to a specific target [Bailey et al. 2009]. When applied in peripheral regions these can direct attention without affecting net perception of the scene. Previous work used SGD to trigger

controlled saccades to enhance visual search performance [McNamara et al. 2008, 2009] and as a narrative tool [McNamara et al. 2012]. Recent work suggests that SGD can drive user gaze in VR experiences as well [Grogorick et al. 2017; Sridharan et al. 2015]. We integrate SGD into our system to dynamically and subtly increase the frequency of saccades, which we then exploit as opportunities for imperceptible transformation of the world.

3 PILOT STUDY OF VISUAL SACCADDES

The efficacy of redirection during saccadic suppression depends on several factors, including frequency and duration of saccades, perceptual tolerance of image displacement during saccadic suppression, and the eye-tracking-to-display latency of the system.

To quantify these, we have conducted a short pilot study with six participants using an HTC Vive HMD with integrated SMI eye-tracking. They were instructed to walk a pre-defined path in the small “Van Gogh room” scene and search for six fixed task objects. We recorded their gaze orientations (Figures 2e and 2f) and used the method of adjustment to identify the angular rotation redirections. Specifically, we tuned the rotation angles up/down until the participants could/could not recognize the difference between saccadic redirection, head-only redirection, and walking without redirection by answering “Yes, I noticed something in the camera orientation” or “No, I do not. They are all normal and the same”.

We determined no participant could detect camera rotation less than $12.6^\circ/\text{sec}$ (0.14° at 90 frames per second) when their gaze velocity was above $180^\circ/\text{sec}$. We increase redirection for longer saccades linearly, which is consistent with previous perceptual experiments [Bolte and Lappe 2015; Li and Matin 1990]. Bolte and Lappe [2015] have shown that “participants are more sensitive to scene rotations orthogonal to the saccade”. However, since our overall system computes across multiple frames (Section 5.2), saccade directions may change within this period. To guarantee imperceptibility, we choose a conservative gain threshold assuming orthogonal saccades.

We then augmented the data from our experiment with captured head and gaze orientation recorded from a participant playing commercial VR arcade games *NVIDIA VR Funhouse* (Funhouse), and horror defense game *The Brookhaven Experiment* (Brookhaven), for 10 minutes each (Figure 2). While less controlled as experimental settings, these represent the state of the art for VR presence, rendering quality, and entertainment tasks. They are more realistic and less biased for evaluating the potential for redirected walking than our specially-constructed lab scenario. For each frame in the collected data, we used our previously measured gaze thresholds to predict the maximum imperceptible redirection.

Over one-minute intervals, the proportion of redirected frames varied between 2.43% and 22.58% in Funhouse, and between 10.25% and 22.02% in Brookhaven. The average proportion of frames with redirection was approximately 11.40% for Funhouse, and approximately 15.16% for Brookhaven, which can sufficiently provide $1.4^\circ/\text{sec}$ and $1.9^\circ/\text{sec}$ angular gains. We conclude that the frequency and distribution of redirection depend on the content, yet contain significant extra gains due to saccadic suppression.




Fig. 2. Saccade analysis for VR applications. We recorded head and gaze data for a user playing two VR games and a simple scene with VR walking to estimate the potential benefits from saccadic redirection. We plot five seconds of angular gaze velocity for these applications, showing frames that we detected as saccades (above the $180^\circ/\text{sec}$ threshold visualized in green lines). Section 2.4 describes our pilot study setup and analysis using (e). Scene (e) courtesy of ruslans3d.

4 METHOD

Reorientation is a technique that modifies the user’s virtual camera to decrease the likelihood of exiting the physical play area. Since minor changes in the virtual camera during head rotation are generally imperceptible, this helps provide richer experiences without the user noticing the redirection. Our system also reorients the virtual camera, but it does so not only during head rotations, but also during, and slightly after, eye saccades. Similar to the case with head rotation, small changes to the camera orientation during saccades are imperceptible, and hence offer opportunities for introducing more frequent and greater amounts of redirection.

Our redirected walking method consists of the following three parts:

Saccade detection Use gaze tracking to detect saccades and identify opportunities to reorient the virtual camera for redirection (Section 4.1).

Dynamic path planning Use the saccade detection thresholds and the physical space around the user to dynamically determine the best virtual camera orientation for redirection (Sections 4.2 and 4.3).

Subtle gaze direction (SGD) Render temporally-modulated stimuli in a user’s visual periphery to induce visual saccades (Section 4.4).

Algorithm 1. Overview of our approach. We perform saccade-aware redirection and dynamic path planning before each frame. We begin by detecting saccade to determine the acceptable perceptual thresholds for virtual camera redirection. We then run our path planning optimization, amortized over several frames. After redirection, we apply image/object-space SGD to scene rendering.

```

1: PathPlanningState = Ready
2:  $\Delta\theta = 0$ 
3: function RENDERREDIRECTED( $t, \mathbf{M}(t)$ )
4:    $E_{curr} = \text{GETLATESTYEPOSE}$ 
5:    $H_{curr} = \text{GETLATESTHEADPOSE}$ 
6:    $G_{curr} = \text{COMBINEHEADGAZE}(H_{curr}, E_{curr})$ 
7:    $\Delta_g = \text{MEASUREANGLE}(G_{curr}, G_{prev})$ 
8:    $\Delta_h = \text{MEASUREANGLE}(H_{curr}, H_{prev})$ 
9:    $\Gamma_g = 0$ 
10:   $\Gamma_h = 0$ 
11:  > Detect saccades
12:   $\Delta_t = \text{GETFRAMEDELTA TIME}$ 
13:  if  $\Delta_g > 180 \cdot \Delta_t$  then
14:    > Gaze gain threshold
15:     $\Gamma_g = 12.6 \cdot \Delta_t$ 
16:  end if
17:  > Path planning (amortized over 2-5 frames)
18:  if PathPlanningState is Ready then
19:    Initialize optimization by sampling  $\mathcal{S}$  using Equation (3)
20:    PathPlanningState = Running
21:  else if PathPlanningState is Running then
22:    Perform iterations of planning optimizer (Equation (8))
23:    if Optimization is done then
24:      Update redirection angle  $\Delta\theta$  (Equation (8))
25:      PathPlanningState = Ready
26:    end if
27:  end if
28:  > Perform redirection
29:  if  $\Delta\theta > 0$  then
30:    > Head rotation gain [Steinicke et al. 2010]
31:    if ( $\text{sgn}(\Delta\theta) = \text{sgn}(\Delta_h)$ ) then  $\lambda = 0.49$  else  $\lambda = -0.2$ 
32:     $\Gamma_h = \lambda \cdot \Delta_h$ 
33:     $\Delta\theta_t = \text{sgn}(\Delta\theta) \cdot \min(\|\Gamma_h\| + \|\Gamma_g\|, \|\Delta\theta\|)$ 
34:    > From Equations (1) and (2)
35:     $\mathbf{M}(t+1) \leftarrow \mathbf{R}(\Delta\theta_t)\mathbf{M}(t)$ 
36:     $\Delta\theta = \Delta\theta - \Delta\theta_t$ 
37:  end if
38:  > Subtle gaze direction (SGD) and rendering
39:  if SGDMode is ObjectSpace then
40:    Modulate material luminance of selected objects
41:  end if
42:  Draw current frame
43:  if SGDMode is ImageSpace then
44:    Modulate luminance of selected peripheral pixels
45:  end if
46:  Display rendered frame
47:   $G_{prev} = G_{curr}$ 
48:   $H_{prev} = H_{curr}$ 
49: end function

```

Algorithm 1 summarizes the steps that constitute each frame of our approach. During each frame, we first use current and previous gaze orientation to detect visual saccades, identifying the opportunity for redirected walking. We then update our dynamic path planning algorithm, which we amortize over 2–5 frames to maintain real-time performance. After its final iteration, our path planning algorithm returns a direction and magnitude of desired redirection. If the current frame is a candidate for redirection, either due to an ongoing saccade or head rotation, we modify the camera viewpoint in the direction of desired redirection by a magnitude subject to our perceptual limits. Finally, while rendering the frame, we add subtle gaze direction – temporally-modulated stimuli in a user’s peripheral vision to imperceptibly encourage visual saccades. We can apply SGD stimuli in either object space or image space.

4.1 Saccade Detection for Camera Reorientation

Saccade detection. Once calibrated, our high-speed eye-tracker is relatively noise-free. Thus we use a simple heuristic to determine whether users are currently making visual saccades. At the beginning of each frame, we use the previous two gaze samples to estimate the current angular velocity of the user’s gaze. If the angular velocity is greater than $180^\circ/\text{sec}$, we conclude that a saccade is either currently ongoing or has recently finished.

In our implementation we use the average position of the user’s left and right gaze locations. This helps reduce noise in detecting location and in estimating velocity. More robust detection (e.g., *Hidden Markov Model* or *Hidden Markov Model* [Andersson et al. 2017]) are potential future research for lower-quality tracking devices.

Due to the latency of contemporary eye-trackers as well as VR rendering and display pipelines, saccade detection generally lags actual saccades by tens of milliseconds. However, since the duration of visual saccades ranges from 20–200 ms and saccadic suppression lasts for 100 ms after a saccade begins [McConkie and Loschky 2002; Ross et al. 2001], we find that our detection is relatively tolerant of tracking and rendering latency, especially for saccades with large angular amplitude. Our pilot studies as described in Section 2.4 indicated that the empirically-determined threshold of $180^\circ/\text{sec}$ accounts for this tolerance.

Camera reorientation thresholds. When saccades are detected within a frame, we slightly re-orient the virtual camera by up to $0.14^\circ/\text{frame}$ as described in Section 2.4. If we respect this threshold, our path planning algorithm can successfully perform redirections with meaningful direction and magnitude without alerting the user. Saccadic redirection can be combined with conventional head-only reorientation. For the latter, we use the previously studied angular gain threshold within $[-20\%, 49\%]$ [Steinicke et al. 2010] pre-calibrated within this range for individual users as some may have lower detection thresholds than others [Grechkin et al. 2016]. Although rotation during head movement allows more redirection, large head rotations are less frequent than large saccades, so we expect an overall improvement by using both for redirected walking. The saccadic detection threshold $180^\circ/\text{sec}$ and gain speed $12.6^\circ/\text{sec}$ were set through our pilot study (Section 2.4).

4.2 Dynamic Path Planning

The saccade-guided camera manipulation and subtle gaze direction (SGD) facilitate VR redirected walking. However, to guide users away from both stationary and moving obstacles, the system must dynamically compute the virtual camera orientation in each frame. Existing off-line mapping approaches [Dong et al. 2017; Sun et al. 2016] require slow pre-processing, incompatible with saccadic actions that happen dynamically and unpredictably in real time. We would also like to avoid any visual distortion caused by virtual scene warping and rely only on larger, rigid transformation gains enabled by saccadic suppression. Thus, we present a real-time dynamic path planning approach driven by perceptual factors (such as SGD), scene properties (e.g. floor layouts and scene object placements), and GPU parallelization.

Formulation. For a given frame t and a 2D virtual position $\mathbf{x} = (x, y)$, we model the corresponding physical position $\mathbf{u} = (u, v)$ using an affine transformation $\mathbf{M}$ between the virtual and physical spaces:

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{M}(t) (\mathbf{x} - \mathbf{x}_c(t)) + \mathbf{x}_c(t) \quad (1)$$

$$\mathbf{M} = [\mathbf{R}|\mathbf{T}]$$

where $\mathbf{x}_c(t)$ is the user’s current virtual space position. This formulation interprets $\mathbf{x}$ and $\mathbf{u}$ as the next virtual and real user positions to allow optimization for the near future, such as avoiding obstacles.

The goal of the real-time path planner is to find the next frame’s optimal translation $\mathbf{T}(t+1)$ and rotation $\mathbf{R}(t+1)$ components so that the redirected walking path during saccades can guide users away from boundaries and obstacles. In our initial investigations we have found $\mathbf{R}$ to be much more effective than $\mathbf{T}$ with saccades and head rotations, so we set $\mathbf{T}(t) = 0$ to reduce the real-time, multidimensional computation workload:

$$\mathbf{M}(t+1) \leftarrow \begin{bmatrix} \cos(\Delta\theta(t)) & -\sin(\Delta\theta(t)) \\ \sin(\Delta\theta(t)) & \cos(\Delta\theta(t)) \end{bmatrix} \mathbf{M}(t) \quad (2)$$

where $\Delta\theta$ is the redirection angle to optimize for (Section 4.3).

Dynamic sampling. Inspired by [Dong et al. 2017; Sun et al. 2016], we perform optimization via virtual scene samples. However, instead of global uniform sampling, we dynamically allocate the sample set $\mathcal{S}$ locally, adapting to the user’s position and orientation to enhance optimization quality and speed. Specifically, we design an importance-based real-time sampling mechanism emphasizing areas that are (1) close to the user’s current position and (2) visible and within the user’s current camera frustum, to predict possibilities in the nearer future, as exemplified in Figure 3. To achieve fast results, we created a closed-form formulation for the intuition above. The importance is computed in the polar coordinates $(r(\mathbf{x}), \theta(\mathbf{x}))$ of the virtual space with $\mathbf{x}$ as the origin:

$$I(\mathbf{x}) = (-\text{erf}(\alpha_0^r r(\mathbf{x}) + \alpha_1^r) + \alpha_2^r) \times \left(\exp\left(-\frac{(\cos(\theta(\mathbf{x}) - \theta_c) - 1)^2}{\alpha_0^a}\right) + \alpha_1^a \right) + \alpha^o \quad (3)$$

where $\text{erf}(x) = \frac{1}{\sqrt{\pi}} \int_{-x}^x e^{-t^2} dt$ is the error function, θ_c is the user’s current virtual camera direction, $\alpha_{i \in \{0,1,2\}}^r$ and $\alpha_{i \in \{0,1\}}^a$ are parameters fitting to a given space size, and α^o is added to avoid zero




Fig. 3. *Illustration of Equations (3) and (8).* Suppose the user is standing at the center of the room while facing right, as shown in (a). The room contains a static obstacle (table on the top) and a dynamic obstacle (person on the right). (b) shows the sampling importance heatmap from Equation (3). (c) plots the corresponding samples. Their color gradients represent energy values from Equation (8). Energies are high for samples outside the physical space and close to the moving human obstacle. Energies are low around the table, because it is far from the user's current position and orientation. The dark blue rectangle in (a) and (c) shows the available physical space.

samples at low importance areas. The importance value is higher at areas close to the user's current position (smaller r) and orientation (θ closer to θ_c). This is illustrated in the heatmap and the corresponding sample set $\mathcal{S}$ in Figure 3. To obtain uniform sampling parameters, we numerically normalize the virtual space to a 1×1 unit. In this space, we use $\alpha_0^r = 30$, $\alpha_1^r = -3$, $\alpha_2^r = 1.15$, $\alpha_0^a = 0.01$, $\alpha_1^a = 0.1$, $\alpha^o = 0.01$ in our experiments. Implementation details of performing the sampling are described in Section 5.2.

Based on $\mathcal{S}$, we propose the following energy terms that guide users away from physical boundaries and obstacles, keep the redirection from being noticeable by the users, and respond to dynamic user saccades and environment changes in real time.

Static boundary avoidance. Similar to [Dong et al. 2017; Sun et al. 2016], the redirection should automatically help users avoid static physical boundaries like walls. We adapt the soft barrier function from [Dong et al. 2017]:

$$E_B(t, \Delta\theta) = \sum_{l_i} \sum_{\mathbf{x}(t) \in \mathcal{S}} w_b(\mathbf{x}(t)) \left(d(\mathbf{u}, l_i) + \sqrt{d(\mathbf{u}, l_i)^2 + \epsilon} \right)^{-1} \quad (4)$$

$$\mathbf{u} \triangleq \mathbf{u}(\mathbf{x}(t), t + 1)$$

where l_i is the i -th edge of the physical boundary polygon, $d(\mathbf{u}, l_i)$ is the distance between user's real position $\mathbf{u}$ and boundary edge l_i , and $\mathbf{u}(\mathbf{x}, t + 1)$ is a function of $\Delta\theta$ (Equations (1) and (2)). The term $w_b(\mathbf{x})$ weighs $\mathbf{x}$'s importance for boundary avoidance. Intuitively, w_b should emphasize the virtual samples closer to current user's virtual position $\mathbf{x}_c$, since the user will more likely reach those points. We fit w_b as an exponential function of the distance $d(\mathbf{x}, \mathbf{x}_c)$ between $\mathbf{x}$ and $\mathbf{x}_c$:

$$w_b(\mathbf{x}) = \exp(-d(\mathbf{x}, \mathbf{x}_c)^2 / \alpha_0^b) + \alpha_1^b, \quad (5)$$

where α_0^b is used to ensure that the weights $w_b(\mathbf{x})$ are appropriate for the size of the virtual space and α_1^b is used to avoid zero weights. We use $\alpha_0^b = 0.01$, $\alpha_1^b = 0.002$ in our experiments. We further calculated w_b from $\mathcal{S}$, which prioritizes virtual regions that are closer to the current user position and orientation (Equation (3)). Note that

Equation (4) represents physical boundaries as polygon clusters and thus can handle non-convex or curved shapes via polygonization.

Moving obstacle avoidance. One major limitation of previous redirected walking approaches is the inability to handle dynamically moving obstacles like other people in the same physical room [Azmandian et al. 2017]. Our dynamic sampling and GPU accelerated redirection planning let our redirection respond to such real-time physical environment changes.

To analytically model obstacles and obtain high gradients at barrier edges, we use a weighted error function instead of the Gaussian barrier function in [Sun et al. 2016] to guide users away from obstacles:

$$E_O(t, \Delta\theta) = \sum_{\mathbf{o} \in \mathbf{O}} \sum_{\mathbf{x} \in \mathcal{S}} w_o(\mathbf{x}, \mathbf{u}^o) \operatorname{erf} \left(\alpha_0^m(r^o) \|\mathbf{u}(\mathbf{x}, t) - \mathbf{u}^o\|^2 + \alpha_1^m \right) \quad (6)$$

where $\mathbf{O}$ is the set of obstacles, $\{\mathbf{u}^o\}$ and $\{r^o\}$ are the dynamic position and radius of each obstacle $\mathbf{o}$, and the linear parameters α_0^m and α_1^m are used to fit the sizes of the obstacles with regard to the erf function. We set $\alpha_0^m < 0$ so that E_O is lower for $\mathbf{u}(\mathbf{x}, t)$ further away from $\mathbf{u}^o$. The obstacle avoidance parameters $\alpha_{0,1}^m$ should adapt to the obstacle sizes to properly guide users away from potential collision. Specifically, we let $\alpha_0^m(r^o) = \frac{-1}{r^o}$, $\alpha_1^m = 2$. Since dynamic obstacles tend to be smaller than wall boundaries, for efficiency and to reduce potential interference with Equation (4), we consider the obstacles only when users are nearby:

$$w_o(\mathbf{x}, \mathbf{u}^o) = \begin{cases} 1 - \frac{d(\mathbf{x}, \mathbf{x}_c)}{2r^o} & d(\mathbf{u}, \mathbf{u}^o) < 2r^o, \|\theta(\mathbf{x}) - \theta_c\| < 15^\circ \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

where $\mathbf{u} = \mathbf{u}(\mathbf{x}, t, \Delta\theta)$ is the redirected physical position of $\mathbf{x}$ at the current time t .

4.3 Real-time Optimization and Redirection

Given the energy terms above and a given time frame t , the optimal redirected mapping is defined as

$$\arg \min_{\Delta\theta} E(t, \Delta\theta) = E_B(t, \Delta\theta) + wE_O(t, \Delta\theta). \quad (8)$$

We set $w = 500$ in our experiments. The visualization of the object among each sample in $\mathcal{S}$ can also be seen from Figure 3c.

Dynamic path planning. Our system applies only rigid rotation from the optimized $\Delta\theta(t)$ during saccades and head rotations. Not having a distortion energy term makes it simpler to optimize than warping-based methods [Dong et al. 2017; Sun et al. 2016].

Note that the perceptually unnoticeable angular gain from saccade suppression is limited to $[-\Delta\theta_{\max}, \Delta\theta_{\max}]$, where $\Delta\theta_{\max}$ is $12.6^\circ/\text{sec}$ in Section 2.4. To match this constraint while obtaining real-time performance responding to users' dynamic saccadic actions, we implement the optimization as a GPU-based line searching method; details and performance comparison are shown in Section 5.2 and Table 1. It is based on the iterative cubic + quadratic zoom searching method with Wolfe condition [Nocedal and Wright 2006]. With the optimized $\Delta\theta$, we redirect the virtual camera when saccades and/or head rotations are detected.




Fig. 4. *Subtle Gaze Direction (SGD) stimuli used in our study.* This example illustrates the stimuli used in our implementation of subtle-gaze direction. The green inset shows an example of image-space SGD stimulus, and the magenta inset shows an example of object-space SGD stimulus. The blue circle indicates the user gaze. Scene courtesy of Barking Dog.

4.4 Subtle Gaze Direction for Saccades

The frequency and extent of visual saccades vary with user, content, and task. However, they directly influence the opportunity for saccadic suppression. Thus, in order to improve the effectiveness of saccadic redirected walking, we would like to increase the occurrence of saccades. In the spirit of using distractors to improve redirected walking [Chen and Fuchs 2017; Peck et al. 2010] without introducing noticeable content change, we propose to utilize subtle gaze direction (SGD) [Bailey et al. 2009] to encourage saccades.

Instead of guiding users to look at particular objects or regions, as is the goal of conventional SGD, our primary goal is to encourage larger and more frequent saccades. Hence, we place SGD stimuli of temporally-varying luminance modulations at a user's peripheral vision, as inspired by Grogorick et al [2017]. The radius of our stimulus is 3.5° with a smooth Gaussian fall-off.

Following Sridharan and Bailey [2015], we prioritize SGD target locations to overlay objects and image features that are already visually salient. We can select these locations in two different ways, which we call *image-space SGD* and *object-space SGD*. Both are shown in Figure 4.

Image-space SGD finds salient peripheral pixels in the rendered image of the current frame. Using visual contrast as the saliency measure, we implement image-space SGD by selecting regions with high local contrast to ensure GPU efficiency. To further speed up the search, we down-sample the image via MIPMAP. Section 5.1 describes details of our implementation. Our preliminary studies suggested that image-space SGD stimuli in a walking experience are either too hard to perceive or too prominent and hence are undesirably distracting. We believe this is because existing SGD mechanisms for either stationary desktop [Bailey et al. 2009] or relatively static AR [Booth et al. 2013] and VR [Grogorick et al. 2017] scenarios may not suffice for highly dynamic redirected walking.

Algorithm 2. *Image-space SGD.* Our image-space SGD approach searches for high-contrast regions in a down-sampled version of the currently rendered frame. We use the center of the tile with highest contrast as the center of our stimulus.

```

1:  $I$ : current frame (rendered, but not displayed)
2: function IMAGE_SPACE_SGD( $I$ )
3:   Compute MIPMAPs for  $I$ 
4:   Select the 5th MIPMAP image  $I_5$ 
5:   Compute the local Weber contrast for each  $3 \times 3$  tile in  $I_5$ 
6:   Find peripheral pixel  $p_5^{max} \in I_5$  with max local contrast
7:   Locate the tile  $t_{max}$  in  $I$  corresponding to  $p_5^{max}$ 
8:   Perform SGD modulation centered at  $t_{max}$ 
9: end function

```

Thus, we also implement object-space SGD, a method that performs SGD modulation directly on the textures or materials of chosen virtual objects, so the users will perceive them as actual scene motion or appearance modulations instead of rendering artifacts. Our object-space SGD approach is straightforward. For each frame, we find scene objects that belong to a manually chosen set (e.g. targets of our task), and modulate the color of their diffuse material. To ensure subtlety of SGD, we can choose to apply SGD only to objects that lie in a user's peripheral vision, or to those that are close to the user's current virtual position. Sampled stimuli are shown in Figure 4. Note that in our pipeline, object-space SGD works by modifying materials *before* we begin drawing a frame, while image-space SGD works by modifying pixels *after* drawing a frame. Since image-space and object-space SGD approaches are orthogonal, they can be combined for evaluation.

5 IMPLEMENTATION

Our system is implemented using an eye-tracked HMD – an HTC Vive augmented with an SMI eye tracker with 250Hz update and 6.5ms response latency, driven by a desktop computer with one NVIDIA Titan Xp GPU, an Intel i7-7700K CPU, and 32GB RAM. For implementing our redirected walking methods in a real-time VR rendering environment, we used the Unity Pro engine, the redirected walking toolkit [Azmandian et al. 2016c], ShaderLab, and DirectX HLSL pixel and compute shaders.

5.1 Subtle Gaze Direction

Image-space SGD. Our image-space SGD approach involves applying temporal modulations to pixels in a user's visual periphery. To improve its effectiveness, we use a content-aware approach that prioritizes high-contrast image regions for stimulus placement. Searching for pixels with high local contrast can be an expensive per-frame computation. For acceleration, we compute the contrast on a down-sampled version of the current frame, which we obtain by generating MIPMAPs for the current frame. After estimating and finding the region with maximum local contrast, we generate the SGD stimulus by modulating the luminance of a Gaussian-shaped region around the center of the high-contrast region. Algorithm 2 provides an overview of this approach.

Object-space SGD. We perform object-space SGD as luminance modulations on the diffuse textures or materials of specific scene

objects. In general, we would like to select a salient object as the target of SGD. For our study, we simply chose SGD objects from the set of target objects used in our search task, while restricting the set to only those objects that are close to the user's virtual viewpoint.

5.2 GPU-Based Sampling and Line Search

Performing summation for importance-based samplings over all virtual space areas is slow on the CPU. For fast parallel processing, we distribute the importance sampling task in each local virtual space area into threads in the GPU, each of which performs sampling independently and adds the importance values atomically. Then the overall sampling budget, which depends on GPU capability, is distributed to each thread based on their local value. The portion is computed by dividing the sum of all areas. In our experiments, the budget was set as 500. This significantly reduces the sampling time, to less than 5ms. This step takes only 1 frame.

In the line searching step, we adapt a searching approach with strong Wolfe condition [Nocedal and Wright 2006] to find the optimal redirection angle $\Delta\theta$ by minimizing Equation (8). Since the computation of objective Equation (8) and its derivatives of each sample in S are independent of each other, we also parallelize the computation of each virtual space sample as a thread in the GPU with atomic operation. The parallelization reduces the computation time to $< 5ms$ per iteration. However, line search is an iterative process, which multiplies the computation time of the objective and derivative calculation. To leverage the high VR rendering refresh rate (90FPS for HTC Vive), we distribute the iterations into multiple consecutive frames. In the final system, we perform 2 iterations per frame. This amortizes the path planning optimization over 2–5 frames to maintain real-time performance.

6 EVALUATION

We evaluated our method with two user studies (Sections 6.2 and 6.3) and several simulations (Section 6.4). The study participants were randomly chosen from internal and external volunteers. One of them was aware of the research, but not the study hypothesis. The study was conducted in a much larger physical space, with a subset of that space designated as the bounds for our redirected walking method. This ensured participant safety without worst-case stimuli, provided a continuous experience, simulated a challenging small room, and facilitated error measurement whenever a participant strayed outside the bounds.

The results show that our method provides significant improvements to redirected walking in VR. We also examine the impact of the three key aspects of our method, saccadic redirection, dynamic path planning, and the use of SGD. While our user studies help to understand the practical effectiveness of our method and identify possible VR sickness, our simulations evaluate our method across a much broader set of conditions with a controlled, consistent set of synthetic inputs and recorded walk-throughs.

6.1 Measurement

In the studies, we record participants' real/virtual-world positions as well as head orientations and gaze positions. We then visualize the virtual and physical path of each trial, and compute the *error*



Fig. 5. Error measure for user studies. We compute the error ϵ in Equation (9) as the total area (shown striped) that is out of bounds or within obstacles.

area for each path—the area outside the physical space or inside an obstacle, as shown in Figure 5. The measure combines the effect of path length with how far each position is from the boundary. With equal path length, a redirected walking technique is more effective by bringing users back after shorter excursions than guiding them far away from boundaries.

To quantify the effectiveness of the redirection, we compare the error area for the virtual path (without redirection) to the area for the physical path. Smaller ratios indicate more effective redirection. Specifically, we define the effectiveness of the redirected walk as the saving ratio ξ , defined as:

$$\epsilon(p(t \rightarrow \mathbf{u})) = \int \min_{l_i}^{\oplus} d_s(l_i, \mathbf{u}(t)) dt \quad (9)$$

$$\xi = \frac{1 - \epsilon(p_r)/\epsilon(p_v)}{\int h(t) dt} \quad (10)$$

where p is a given physical path that maps a given time t to a physical position $\mathbf{u}(t)$; p_v and p_r are the paths without and with redirection respectively, as visualized in Figure 1c. $\min^{\oplus}$ finds the minimum non-negative signed distance d_s (positive/negative for outside/inside the real space domain) between exterior-or-interior boundary segment l and real user position $\mathbf{u}$, and $h(t)$ is the user's head rotation angle at time frame t . ϵ is the total area that is out of bounds or within obstacles. The saving ratio ξ shows how much a redirected path can reduce the error cost compared with the original virtual path overlaid on the real environment. Since we used the savings from head-only gain as the baseline, we normalized ξ by the total head rotations, as users may have a different number of head rotations for multiple trials with different virtual paths.

6.2 User Study: Impact of Saccades

Overview. In our first user study, we evaluate whether the use of saccades with and without traditional image-space SGD [Grogoric et al. 2017] can improve the effectiveness of a redirected walking system. We instructed participants to perform a typical target-retrieval task. Each participant's goal was to search and count all instances of a specific target object in a VR environment.

Task and Stimuli. The study consisted of three experiments, all using our dynamic path planning algorithm as the redirection method. Each user did each experiment once.

- (1) Non-saccadic redirected walking, with head rotation gain only (NON-SACCADE);
- (2) Saccadic redirected walking (SACCADE);
- (3) Saccadic redirected walking with image-space SGD from Sections 4.4 and 5.1 (IMAGE-SGD-I).




Fig. 6. *The effects of saccades for redirected walking.* We plot the average saving error ratios ξ (Equation (9)) and 95% confidence intervals over all users in Section 6.2. Notice that saccades can provide greater error saving for redirected walking than head-rotation alone. The confidence error bar indicates that the exact gains (thus variance) likely vary across users and within experiences.

Figure 4 shows a screenshot of the VR environment and task stimuli used in this study. Each participant started from the same corner in the virtual room and was instructed to find as many instances of randomly colored, positioned, and scaled balls as possible. Each trial lasted 80 seconds. The sizes of the virtual and physical room were 3.1m×3.1m and 2.0m×2.0m respectively. To encourage walking during retrieval, we dynamically control the transparency of target objects based on their distances to the current user position $\mathbf{x}_c$. Specifically, for the i -th target at position $\mathbf{x}_t$, we update its material alpha (a_i) at time t as

$$a_i(t) = \exp(-\|\mathbf{x}_c(t) - \mathbf{x}_t\|^2 / 0.05). \quad (11)$$

Prior studies have used fog for a similar purpose [Hodgson and Bachmann 2013]. While it is a good alternative, we opt for object transparency so that the overall environment is consistently visible at all times.

At the end of each trial, each participant was asked to complete the Kennedy Lane SSQ [Kennedy et al. 1993] for simulator sickness. After the 3 trials, the participant was asked, “Did you notice any camera modulation or difference among all trials?”

Participants. 9 users (3 female, 33.3%) participated in the study. The average age was 26.7 ($SD = 1.66$). The median of self-reported experiences with VR was 4, with 1 being least familiar, and 5 being most familiar. We adopted a within-subject design. The order of the three experiments were counterbalanced across participants. Subjects were not informed of the study hypothesis. Between successive trials, a mandatory 3-minute break was enforced.

Results. We statistically analyze the recorded error measures among the SACCAD, NON-SACCAD and IMAGE-SGD-I experiments from the study.

Saving ratio The introduction of extra rotation during saccade enables more opportunities to perform stronger angular manipulation, thus smaller physical space usage. To numerically evaluate the capability we compare NON-SACCAD to SACCAD. The average saving ratio ξ was $2.01e-3$ ($SD = 1.95e-3$) for NON-SACCAD, and $3.39e-3$ ($SD = 1.98e-3$) for SACCAD, as shown in Figure 6. There was a significant main effect of SACCAD on ξ ($F_{1,8} = 15.01, p < 0.005$).

Saccadic angular gains To evaluate the impact of SGD, we calculated the sum of all saccadic angular gains between SACCAD and IMAGE-SGD-I. The total saccadic redirected angle across




Fig. 7. *Simulator sickness from Section 6.2 and Section 6.3.* The dashed lines show the minimum sickness rates from [Dong et al. 2017]. Our system is shown to reduce sickness compared to a warping-based approach.

all users was 163.82° ($SD = 28.79^\circ$) for SACCAD, and 148.63° ($SD = 22.99^\circ$) for IMAGE-SGD-I. Single factor repeated measures ANOVA did not show a significant main effect of SGD on the saccadic angular gains ($F_{1,8} = 3.306, p = 0.107$).

Subjective feedback Nausea and oculomotor levels are reported below 2 by all users, except for one user who reported to have often experienced VR perceptual anomalies including general discomfort, nausea, and vertigo, as shown in the left half of Figure 7. All users answered “no” to the post-trial question, indicating that the saccadic redirection was perceptually unnoticeable.

Discussion. The ξ between SACCAD and NON-SACCAD indicates that SACCAD can greatly help redirected walking by reducing errors by 68.7% on average. It is better in performance than NON-SACCAD and in comfort than a recent warping-based redirected walking method (Figure 7).

Figure 8 compares redirection methods, with 8a and 8d demonstrating that saccadic redirection reduces the chance of hitting physical boundaries, allowing much larger differences between virtual and physical environments.

Image-space gaze direction cues did not trigger saccades for all study subjects in the search tasks. From the saccadic angular gains results we can conclude that while gaze direction cues in general can help trigger saccades in desktop and VR displays (based on our initial development with sitting/static setups), previously reported image-space methods [Bailey et al. 2009; Grogork et al. 2017] may not be as effective in the highly dynamic redirected walking scenario, especially when it involves search and retrieval tasks (e.g., in real VR games). This observation was also derived from our post-interview with users: most reported that they were focused on the task object retrieval while constantly moving. They paid much less attention to, or ignored, the detailed image content, which changes rapidly and contains the image-space SGD stimuli. The result and discovery inspired us to explore a task-matching, object-space SGD variant that we used for a follow-up user study.

6.3 User Study: Image-space SGD Vs. Object-space SGD

Overview. As described above, for highly dynamic redirected walking applications, image-space SGD stimuli were often not as effective as they are in relatively static scenarios like image viewing [Bailey et al. 2009] or searching while being seated [Grogork et al. 2017]. We conducted a second study with object-space SGD as




Fig. 8. Path planning comparisons under different types of obstacles and redirection methods. All experiments share the same physical room (visualized as the dark blue rectangles). Using a randomly sampled user’s historical virtual head and gaze data from Figure 1c, we simulate our dynamic path planning with no obstacle in (a), with a static obstacle (the red circle) in (b), with a dynamic obstacle (the red curve) in (c), using head-rotation without saccade in (d), and using traditional S2C redirection in (e). Both (d) and (e) have no obstacle and can be compared to (a). All but (a) are simulated paths. The saturation gradients in (c) stand for time propagation.

described in Section 4.4. Using a similar setup as in Section 6.2, we evaluated the relative effectiveness of image-space and object-space SGD in increasing the frequency of saccades.

For image-space SGD, we applied SGD stimuli using the algorithm in Section 5.1; for object-space SGD we simply modulated the target objects’ luminance. The study consisted of three experiments:

- (1) Image-space SGD only (IMAGE-SGD-II, Figure 4);
- (2) Object-space SGD only (OBJ-SGD, Figure 4);
- (3) Both object-space and image-space SGD (DUAL-SGD).

Participants. Another 9 users (2 female, 22.2%) participated in the study. The average age was 26.7 ($SD = 2.24$). The median of self-reported experiences with VR was 3, with 1 being least familiar, and 5 being most familiar. The order of the three experiments were counterbalanced across participants. A mandatory 3-minute break was enforced between successive trials.

Results. We compared the effect from different SGD approaches.

Saccadic angular gain The total saccadic angle gain across all users was 145.13° ($SD = 21.83^\circ$) for IMAGE-SGD-II, 156.78° ($SD = 23.41^\circ$) for OBJ-SGD, and 167.48° ($SD = 22.56^\circ$) for DUAL-SGD. There was a significant main effect of SGD method on the total redirected angles ($F_{2,16} = 6.417, p < 0.05$). Pair-wise comparison with Holm correction showed the differences between DUAL-SGD and the other two experiments were significant ($p < 0.05$ for both experiments), but not between IMAGE-SGD-II and OBJ-SGD ($p = 0.168$).

Subjective feedback No users noticed any camera modulation. All users reported nausea below 2 and oculomotor below 3, as shown in the right half of Figure 7.

Discussion. Compared with traditional image space SGD, the object-plus-image space SGD achieved better results. This shows that in a highly dynamic redirected walking VR scenario, the impact of image-space SGD becomes weaker. However, having the task objects with similar flickering appearances to image SGD might trigger more saccades since users were looking for such stimuli. Task-dependent SGD design can be an interesting direction for future, more exhaustive studies.

The user perception of saccadic redirection was similar to the first study. Saccadic redirection using the parameters we selected in Section 2.4 was imperceptible in our VR exploration and object-retrieval task. Further, since we used head-pose redirection from Steinicke et al. [2010] in conjunction with saccadic redirection, we can infer no perceptual impact of the two working together.

6.4 Simulation: Redirection Methods

In addition to user studies, we conducted simulations to evaluate our method over a wider set of conditions but using a consistent set of input virtual paths and head orientations for fair comparison. During each study trial, we recorded virtual user position $\mathbf{x}$, head rotation angles, and gaze point of regard in each time frame t . We use recorded rather than procedurally generated user paths for better realism. Although saccadic redirection was enabled while recording, for simulation we used only users’ virtual paths, which are solely dependent on object placement and the individuals’ virtual movements, to avoid bias toward or against any particular redirection approach, such as Steer-to-Center (S2C). The path planners then return the corresponding $\Delta\theta$ values (0 for methods not considering eye/head rotation), allowing us to update $\mathbf{M}(t+1)$ and to get the simulated physical position $\mathbf{u}(t+1)$ at the recorded $\mathbf{x}(t+1)$ using Equation (1). With this mechanism, we can simulate different physical paths with different path planners and/or angular gains, based on the same virtual path as another trial. Error measure analysis (Equation (10)) can also be performed on the new physical path. When simulating virtual spaces with difference sizes, by assuming the same walking speeds, we can rescale the recorded virtual coordinates and insert extra time frames by interpolation.

Dynamic path planning versus S2C. Measuring path planning approaches is sensitive to a specific user’s virtual traveling path for each trial. To obtain a fair comparison, we simulate S2C redirection results with a same user movement history, as described in Section 6.4.

With all 18 users from Sections 6.2 and 6.3, the average ξ was $3.06e-3$ ($SD = 1.52e-3$) for the dynamic path planning condition, and $0.75e-3$ ($SD = 2.12e-3$) for the corresponding simulated S2C




Fig. 9. *The effectiveness of dynamic path planning.* Here we show the average saving ratios ξ and 95% confidence intervals over all trials from IMAGE-SGD-I and IMAGE-SGD-II in Section 6.2, and corresponding simulated S2C redirection with identical input virtual path. It can be seen that the redirection with our dynamic planning approach shows stronger error saving than S2C redirection (both with saccadic redirection and SGD).

condition, as shown in Figure 9. Saccadic suppression had a significant main effect on ξ ($F_{1,17} = 26.12, p < 0.005$). Because eye actions such as saccades occur frequently and uniformly, as shown in Figure 2, it allows stronger and more uniform rotation gains. Figure 8a and Figure 8e compare results with the same user history.

Obstacle and multi-user collision avoidance. Traditional redirection planning approaches such as S2C [Azmandian et al. 2016b; Hodgson and Bachmann 2013] handle convex-shaped laboratory spaces like rectangular rooms. However, in consumer use-cases, the physical rooms often include static obstacles like furniture and may even contain other people. Consequently, practical VR play areas are non-convex and often dynamic. In such cases, content-unaware methods are highly likely to cause collisions, as seen by example in the supplementary video. In contrast, our dynamic technique and a real-time implementation can respond to physical-world changes, guiding users away from boundaries and obstacles.

To simulate multi-user scenarios, we use the recorded physical paths from IMAGE-SGD-I and II as moving obstacle positions, and then use our dynamic path planner to simulate new walking paths with their corresponding virtual space records as input. The dynamic planner reduces the error ϵ from obstacles by 94.2% on average ($SD = 3.9\%$). The overall average ξ is $2.82e-3$ ($SD = 1.82e-3$) for the simulation, which is lower than the original ($3.06e-3$). However, ANOVA did not show a significant main effect ($F_{1,17} = 1.055, p = 0.319$). This means that our method may introduce extra boundary errors by avoiding moving obstacles, but this is not statistically significant.

Figures 8b and 8c show additional simulated paths for redirection around static and dynamic obstacles. The supplemental video also contains a non-simulated example, where, since our current setup cannot track multiple users, the “moving obstacle” is just another person instructed to walk along a predetermined path.

Dynamic path planning versus static scene warping. The static scene warping methods in [Dong et al. 2017; Sun et al. 2016] depend on significant occlusions in the virtual scene to drive unnoticeable geometric warping. These methods can thus cause visible artifacts or scaling for open virtual spaces. Our dynamic path planning method can handle both open and occluded virtual spaces, since it does not rely on any scene appearance. Figure 10 shows a comparison. Moreover, unlike Sun et al. [2016] and Dong et al. [2017], our planning



(a) static warping (b) dynamic path planning

Fig. 10. *Static scene warping versus dynamic path planning.* Prior static warping methods such as [Dong et al. 2017; Sun et al. 2016] rely on sufficient occlusions and may cause unnatural distortions or translation/rotation gains for sufficiently open spaces as shown in (a). Our method, in contrast, does not cause scene distortions or noticeable translation/rotation gains (b). Scene courtesy of Tirgames.

Table 1. *Performance comparison between our GPU-based path planner (Section 4.2) and a CPU implementation.* The GPU time consumption already includes memory transferring between GPU and CPU. The three parts are dynamic sampling (Figure 3), the computation of the cost function Equation (8) and its derivatives.

platform \ step	derivative	function	sampling
GPU	0.005	0.004	0.0045
CPU	0.025	0.016	0.74

approach runs in real-time, so it can also redirect the user to fit physical environmental changes.

6.5 Performance

Table 1 compares our GPU-based sampling and optimization with a corresponding CPU implementation. It shows that we are able to achieve a significant speedup compared to the CPU, enabling real-time dynamic path planning without latency. Combined with our amortization approach from Section 5.2, we are able to run our overall system including eye tracking, dynamic path planning, and rendering at 80-85 FPS depending on rendering complexity.

7 APPLICATIONS

Beyond redirected walking with greater perceptual comfort and visual quality, our system can benefit other applications:

Cinematic VR. Although users can freely explore in virtual scenes, directors who produce immersive stories may intend to redirect the user to a certain part of the scene. Our path planning approach (Section 4.2) can adapt to story-based objectives to achieve this.

Home entertainment. Our method lets multiple users explore the same or different virtual scenes while sharing one physical room. Home entertainment applications often contain multiple game players in the same room. It could encourage game industry development towards the VR platform by avoiding unnatural motion controllers (e.g., mouse/keyboard or gamepad) and enabling practical features such as inter-user collaboration and competition.

Education. In architectural design education or virtual museum navigation scenarios, users should be guided to follow an ideal path to increase exposure or avoid getting lost. Our redirection approach can be adjusted to guide users towards pre-defined virtual paths.




Fig. 11. Simulation results of different physical spaces. (a) shows the result in Figure 1c with 12.25 m² physical and 40.96 m² virtual spaces. (b) shows a simulation result for a physical room of 47.61 m². The simulated user can walk through a virtual space of 4.9× larger area. (c) shows a simulation result for a physical room of 134.56 m². The virtual space can be 11.1× larger in area. As a comparison, in the user experiments Section 6.2, the physical room size was 3.61 m², the virtual space can be 2.0× in area. The trend is plotted in Figure 11d.

8 LIMITATIONS AND FUTURE WORK

In this paper we showed that rotation-based redirection during saccades is effective in both room-scale and large-scale VR (Figure 11) and that our GPU implementation allows real-time path planning (Table 1). The real-time performance also allows the planning to avoid moving obstacles and changing geometry. However, recent researches on robotics and artificial intelligence fields may be adapted to the redirection planning approach for faster and more robust response to the dynamic environmental changes.

Limiting our redirection transformations to rotational gain simplified the planning optimization, enabling real-time performance. We plan to investigate whether translational gain can be incorporated into the optimization while maintaining real-time performance.

There are many opportunities for enhancing the redirection system, including saccade prediction [Arabadzhiyska et al. 2017; Han et al. 2013] to compensate for tracking latency, learning [Gatys

et al. 2017] to enhance gaze guidance, redirection during blink suppression [Langbehn et al. 2018; Ridder III and Tomlinson 1997] to provide more opportunities for redirection, and additional forms of distractors [Chen and Fuchs 2017; Peck et al. 2010] to encourage more eye movement.

Compared to warping based methods [Dong et al. 2017; Sun et al. 2016], our rigid-transformation based redirection allows exploring open virtual spaces without distracting visual distortions. However, Suma et al. [2013] showed that warping provides chances to overlay virtual objects onto physical obstacles in applications like mixed reality. To further enhance the visual and tactile consistency between the virtual and physical environments, we plan to investigate incorporating limited degrees of warping [Azmandian et al. 2016a].

Our system works for room-scale physical environments. However, the space saving benefit from the saccadic gain increases greatly as the available physical area grows. Figure 11 shows the comparison and trend; it would be interesting to investigate whether further gains could come from tuning the system for larger areas.

ACKNOWLEDGMENTS

We would like to thank Michael Stengel for early discussions, Rachel Albert for conducting initial vision studies, Mavey Ma for video dubbing, Jonghyun Kim, Joohwan Kim, Erik Lindholm, Alexander Majercik, and Ward Lopes for helping our GTC live demo, and the anonymous reviewers for their valuable suggestions. This project is partially supported by National Science Foundation grants CNS1302246, NRT1633299, CNS1650499, and a gift from Adobe.

REFERENCES

- Rachel Albert, Anjul Patney, David Luebke, and Joohwan Kim. 2017. Latency Requirements for Foveated Rendering in Virtual Reality. *ACM Trans. Appl. Percept.* 14, 4, Article 25 (2017), 13 pages.
- Richard Andersson, Linnea Larsson, Kenneth Holmqvist, Martin Stridh, and Marcus Nyström. 2017. One algorithm to rule them all? An evaluation and discussion of ten eye movement event-detection algorithms. *Behavior Research Methods* 49, 2 (2017), 616–637.
- Elena Arabadzhiyska, Okan Tarhan Tursun, Karol Myszkowski, Hans-Peter Seidel, and Piotr Didyk. 2017. Saccade Landing Position Prediction for Gaze-contingent Rendering. *ACM Trans. Graph.* 36, 4, Article 50 (2017), 12 pages.
- Mahdi Azmandian, Timofey Grechkin, Mark Bolas, and Evan Suma. 2015. Physical Space Requirements for Redirected Walking: How Size and Shape Affect Performance. In *Eurographics Symposium on Virtual Environments (2015)*. The Eurographics Association, Kyoto, Japan, 93–100.
- Mahdi Azmandian, Timofey Grechkin, Mark Bolas, and Evan Suma. 2016a. Automated Path Prediction for Redirected Walking Using Navigation Meshes. In *IEEE Symposium on 3D User Interfaces*. 63–66.
- Mahdi Azmandian, Timofey Grechkin, Mark Bolas, and Evan Suma. 2016b. The redirected walking toolkit: a unified development platform for exploring large virtual environments. In *2016 IEEE 2nd Workshop on Everyday Virtual Reality (WEVR)*. 9–14.
- Mahdi Azmandian, Timofey Grechkin, and Evan Suma Rosenberg. 2017. An evaluation of strategies for two-user redirected walking in shared physical spaces. In *2017 IEEE Virtual Reality (VR)*. 91–98.
- Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. 2016c. A Demonstration of Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces (ISS '16)*. ACM, New York, NY, USA, 501–504.
- A. Terry Bahill, Michael R. Clark, and Lawrence Stark. 1975. The main sequence, a tool for studying human eye movements. *Mathematical Biosciences* 24, 3-4 (1975), 191–204.
- Reynold Bailey, Ann McNamara, Nisha Sudarsanam, and Cindy Grimm. 2009. Subtle Gaze Direction. *ACM Trans. Graph.* 28, 4, Article 100 (2009), 14 pages.
- Benjamin Bolte and Markus Lappe. 2015. Subliminal Reorientation and Repositioning in Immersive Virtual Environments using Saccadic Suppression. *IEEE TVCG* 21, 4 (2015), 545–552.

- Thomas Booth, Srinivas Sridharan, Ann McNamara, Cindy Grimm, and Reynold Bailey. 2013. Guiding Attention in Controlled Real-world Environments. In *SAP '13*, 75–82.
- Bruce Bridgeman, Derek Hendry, and Lawrence Stark. 1975. Failure to detect displacement of the visual world during saccadic eye movements. *Vision Research* 15, 6 (1975), 719–722.
- David C. Burr, M. Concetta Morrone, John Ross, and others. 1994. Selective suppression of the magnocellular visual pathway during saccadic eye movements. *Nature* 371, 6497 (1994), 511–513.
- Haiwei Chen and Henry Fuchs. 2017. Supporting Free Walking in a Large Virtual Environment: Imperceptible Redirected Walking with an Immersive Distractor. In *Proceedings of the Computer Graphics International Conference (CGI '17)*. ACM, New York, NY, USA, Article 22, 6 pages.
- Lung-Pan Cheng, Eyal Ofek, Christian Holz, Hrvoje Benko, and Andrew D. Wilson. 2017. Sparse Haptic Proxy: Touch Feedback in Virtual Environments Using a General Passive Prop. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 3718–3728.
- Mark R. Diamond, John Ross, and Maria C. Morrone. 2000. Extraretinal control of saccadic suppression. *The Journal of Neuroscience* 20, 9 (2000), 3449–3455.
- Zhi-Chao Dong, Xiao-Ming Fu, Chi Zhang, Kang Wu, and Ligang Liu. 2017. Smooth Assembled Mappings for Large-scale Real Walking. *ACM Trans. Graph.* 36, 6, Article 211 (2017), 13 pages.
- Andrew T. Duchowski and Arzu Çöltekin. 2007. Foveated gaze-contingent displays for peripheral LOD management, 3D visualization, and stereo imaging. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 3, 4 (2007), 6.
- Andrew T. Duchowski, Nathan Cournia, and Hunter Murphy. 2004. Gaze-contingent displays: A review. *CyberPsychology & Behavior* 7, 6 (2004), 621–634.
- Andrew T. Duchowski, Donald H. House, Jordan Gestring, Rui I. Wang, Krzysztof Krejtz, Izabela Krejtz, Radosław Mantiuk, and Bartosz Bazyluk. 2014. Reducing Visual Discomfort of 3D Stereoscopic Displays with Gaze-contingent Depth-of-field. In *Proceedings of the ACM Symposium on Applied Perception (SAP '14)*. ACM, New York, NY, USA, 39–46.
- Leon A. Gatys, Matthias Kümmerer, Thomas Wallis, and Matthias Bethge. 2017. Guiding human gaze with convolutional neural networks. *ArXiv e-prints* (2017). arXiv:cs.CV/1712.06492
- Timofey Grechkin, Jerald Thomas, Mahdi Azmandian, Mark Bolas, and Evan Suma. 2016. Revisiting Detection Thresholds for Redirected Walking: Combining Translation and Curvature Gains. In *Proceedings of the ACM Symposium on Applied Perception (SAP '16)*. ACM, New York, NY, USA, 113–120.
- Steve Grogoric, Michael Stengel, Elmar Eisemann, and Marcus Magnor. 2017. Subtle Gaze Guidance for Immersive Environments. In *SAP '17*. Article 4, 7 pages.
- Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D Graphics. *ACM Trans. Graph.* 31, 6 (2012), 164:1–164:10.
- Peng Han, Daniel R. Saunders, Russell L. Woods, and Gang Luo. 2013. Trajectory prediction of saccadic eye movements using a compressed exponential model. *Journal of vision* 13, 8 (2013), 27–27.
- Eric Hodgson and Eric Bachmann. 2013. Comparing Four Approaches to Generalized Redirected Walking: Simulation and Live User Data. *TVCG* 19, 4 (2013), 634–643.
- J. Johanna Hopp and Albert F. Fuchs. 2004. The characteristics and neuronal substrate of saccadic eye movement plasticity. *Progress in Neurobiology* 72, 1 (2004), 27–53.
- Michael R. Ibbotson and Shaun L. Cloherty. 2009. Visual Perception: Saccadic Omission - Suppression or Temporal Masking? *Current Biology* 19, 12 (2009), R493–R496.
- David E. Jacobs, Orazio Gallo, Emily A. Cooper, Kari Pulli, and Marc Levoy. 2015. Simulating the Visual Experience of Very Bright and Very Dark Scenes. *ACM Trans. Graph.* 34, 3, Article 25 (2015), 15 pages.
- Robert S. Kennedy, Norman E. Lane, Kevin S. Berbaum, and Michael G. Lilienthal. 1993. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The International Journal of Aviation Psychology* 3, 3 (1993), 203–220.
- Eike Langbehn, Gerd Bruder, and Frank Steinicke. 2016. Subliminal Reorientation and Repositioning in Virtual Reality During Eye Blinks. In *Proc. of Spatial User Interaction*. 213–213.
- Eike Langbehn, Frank Steinicke, Markus Lappe, Gregory F. Welch, and Gerd Bruder. 2018. In the Blink of an Eye: Leveraging Blink-Induced Suppression for Imperceptible Position and Orientation Redirection in Virtual Reality. *ACM Trans. Graph.* 37, 4, Article 66 (2018).
- Marc Levoy and Ross Whitaker. 1990. Gaze-directed volume rendering. *ACM SIGGRAPH Computer Graphics* 24, 2 (1990), 217–223.
- Wenxun Li and Leonard Martin. 1990. The influence of saccade length on the saccadic suppression of displacement detection. *Perception & Psychophysics* 48, 5 (1990), 453–458.
- David Luebke and Benjamin Hallen. 2001. Perceptually driven simplification for interactive rendering. In *Rendering Techniques 2001*. Springer, 223–234.
- Ethel Martin. 1974. Saccadic suppression: a review and an analysis. *Psychological bulletin* 81, 12 (1974), 899.
- George W. McConkie and Lester C. Loschky. 2002. Perception onset time during fixations in free viewing. *Behavior Research Methods, Instruments, & Computers* 34, 4 (2002), 481–490.
- Ann McNamara, Reynold Bailey, and Cindy Grimm. 2008. Improving Search Task Performance Using Subtle Gaze Direction. In *Proceedings of the 5th Symposium on Applied Perception in Graphics and Visualization (APGV '08)*. ACM, New York, NY, USA, 51–56.
- Ann McNamara, Reynold Bailey, and Cindy Grimm. 2009. Search Task Performance Using Subtle Gaze Direction with the Presence of Distractions. *ACM Trans. Appl. Percept.* 6, 3 (2009), 17:1–17:19.
- Ann McNamara, Thomas Booth, Srinivas Sridharan, Stephen Caffey, Cindy Grimm, and Reynold Bailey. 2012. Directing Gaze in Narrative Art. In *Proceedings of the ACM Symposium on Applied Perception (SAP '12)*. ACM, New York, NY, USA, 63–70.
- J. Nocedal and S. Wright. 2006. *Numerical Optimization*. Springer New York.
- Anjul Patney. 2017. Perceptual Insights into Foveated Virtual Reality. In *NVIDIA GPU Technology Conference 2017 Talks*. San Jose, California.
- Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards Foveated Rendering for Gaze-tracked Virtual Reality. *ACM Trans. Graph.* 35, 6 (2016), 179:1–179:12.
- Tabitha C. Peck, Henry Fuchs, and Mary C. Whitton. 2010. Improved redirection with distractors: A large-scale-real-walking locomotion interface and its effect on navigation in virtual environments. In *Virtual Reality Conference*. IEEE, 35–38.
- Ken Pfeuffer, Benedikt Mayer, Diako Mardanbegi, and Hans Gellersen. 2017. Gaze + Pinch Interaction in Virtual Reality. In *Proc. of Spatial User Interaction (SUI '17)*. 99–108.
- Sharif Razzaque. 2005. *Redirected Walking*. Ph.D. Dissertation. Chapel Hill, NC, USA. Advisor(s) Brooks, Jr., Fredrick P. AAI3190299.
- Sharif Razzaque, Zachariah Kohn, and Mary C. Whitton. 2001. Redirected walking. In *Proc. of Eurographics*, Vol. 9. Manchester, UK, 105–106.
- Sharif Razzaque, David Swapp, Mel Slater, Mary C. Whitton, and Anthony Steed. 2002. Redirected Walking in Place. In *EGVE '02*. 123–130.
- Stephen M. Reder. 1973. On-line monitoring of eye-position signals in contingent and noncontingent paradigms. *Behavior Research Methods* 5, 2 (1973), 218–228.
- William H. Ridder III and Alan Tomlinson. 1997. A comparison of saccadic and blink suppression in normal observers. *Vision Research* 37, 22 (1997), 3171–3179.
- John Ross, M. Concetta Morrone, Michael E. Goldberg, and David C. Burr. 2001. Changes in visual perception at the time of saccades. *Trends in Neurosciences* 24, 2 (2001), 113–121.
- Srinivas Sridharan and Reynold Bailey. 2015. Automatic Target Prediction and Subtle Gaze Guidance for Improved Spatial Information Recall. In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception (SAP '15)*. 99–106.
- Srinivas Sridharan, James Pieszala, and Reynold Bailey. 2015. Depth-based Subtle Gaze Guidance in Virtual Reality Environments. In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception (SAP '15)*. ACM, New York, NY, USA, 132–132.
- Frank Steinicke, Gerd Bruder, Jason Jerald, Harald Frenz, and Markus Lappe. 2008. Analyses of Human Sensitivity to Redirected Walking. In *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology (VRST '08)*. ACM, New York, NY, USA, 149–156.
- Frank Steinicke, Gerd Bruder, Jason Jerald, Harald Frenz, and Markus Lappe. 2010. Estimation of Detection Thresholds for Redirected Walking Techniques. *IEEE TVCG* 16, 1 (2010), 17–27.
- Evan A. Suma, D. Krum, and Mark Bolas. 2013. Redirected Walking in Mixed Reality Training Applications. In *Human Walking in Virtual Environments: Perception, Technology, and Applications*. Springer, 319–331.
- Qi Sun, Li-Yi Wei, and Arie Kaufman. 2016. Mapping Virtual and Physical Reality. *ACM Trans. Graph.* 35, 4, Article 64 (2016), 12 pages.
- Martin Usoh, Kevin Arthur, Mary C. Whitton, Rui Bastos, Anthony Steed, Mel Slater, and Frederick P. Brooks, Jr. 1999. Walking > Walking-in-place > Flying, in Virtual Environments. In *SIGGRAPH '99*. 359–364.
- Peter Vincent and Ritchie Brannan. 2017. S7797 Tobii Eye Tracked Foveated Rendering for VR and Desktop. (2017).
- Mounia Ziat, Vincent Hayward, C.Élaine Chapman, Marc O. Ernst, and Charles Lenay. 2010. Tactile suppression of displacement. *Experimental Brain Research* 206, 3 (2010), 299–310.