Walking in the Wild: Safe and Natural Redirected Walking in Open Physical Spaces

Xinda Liu[1], Qinyu Zhang[1], Guoqiang Yang[1], Yunchen Li[1], Jian Wu[2]*, Guohua Geng[1], and Lili Wang[2]

[1]School of Information Science and Technology, Northwest University, Xi'an, China

[2]State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, China

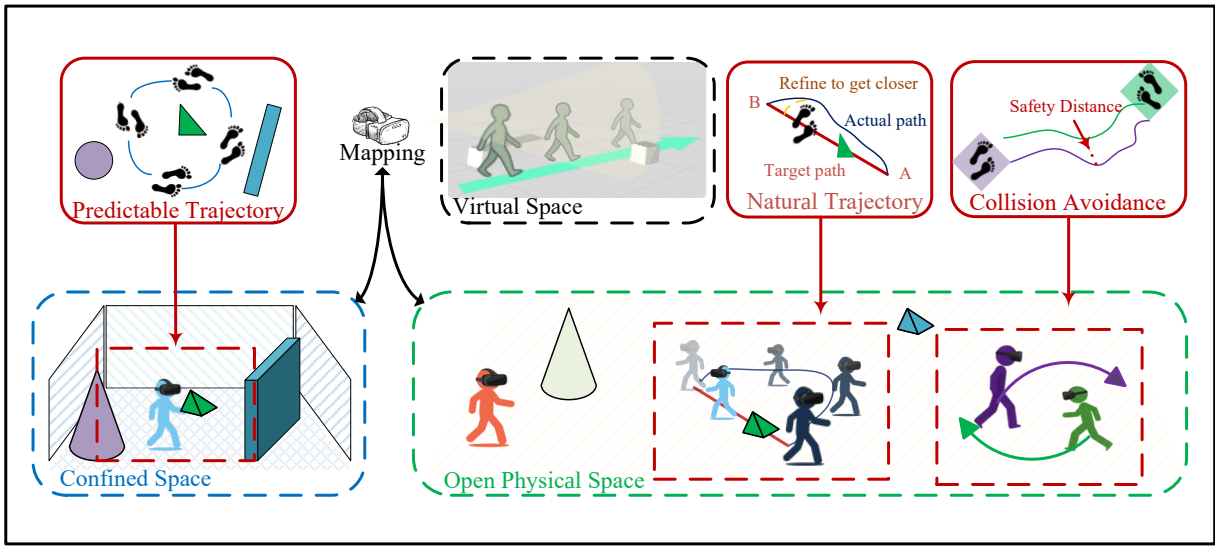

Figure 1: The comparison of classic Redirected Walking (RDW) and our proposed Open Physical Spaces RDW (OPS-RDW). The left side depicts classic RDW within a confined area, where obstacles are fixed and physical walking paths are well-defined and predictable. In contrast, the right side presents the OPS-RDW scenario, which involves dynamic and unpredictable obstacles, potential multi-user interactions, and without a predefined boundary model. The goal of OPS-RDW is to maintain natural walking trajectories consistent with the intended navigation direction, while the redirection algorithm ensures real-time avoidance of dynamic obstacles and guarantees user safety.

Abstract

Redirected Walking (RDW) enables continuous locomotion in virtual environments (VEs) within limited physical spaces. However, classic RDW methods rely on static settings with tight boundaries, which can reduce their applicability in large shared physical spaces where boundary constraints are not dominant, especially under dynamic and multi-user conditions. To overcome this, we introduce a new task: **Redirected Walking in Open Physical Spaces (OPS-RDW)**, which allows users to navigate expansive VEs safely and naturally despite dynamic obstacles and without requiring a fixed boundary model. We further propose **Dynamic Control and Redirection with Safety Constraints (DyCoRe)**, which formulates OPS-RDW as a constrained optimization problem. DyCoRe uses Dynamic Control Barrier Functions to model real-time collision avoidance constraints and solves an online Quadratic Programming problem to compute optimal velocities that minimize path deviation while ensuring safety. These velocities are mapped into real-time redirection gains, guiding users along natural paths while reducing collision risk. A relaxation mechanism is incorporated to handle infeasible scenarios. Extensive simulations suggest consistent improvements in safety and obstacle clearance over state-of-the-art methods. User studies demonstrate that DyCoRe significantly improves navigation continuity, reduces the number of resets, and tends to reduce perceived discomfort compared to classic RDW strategies. DyCoRe provides an efficient and learning-free solution for safe and natural VR locomotion in open physical spaces with dynamic obstacles and multiple users.

Index Terms: Redirected Walking (RDW), Open Physical Spaces, Dynamic Obstacles, Safety Constraints.

1 Introduction

Virtual Reality (VR) provides a deeply immersive virtual environment experience, but the constraints of physical space often limit its full potential, especially when users wish to freely explore large-scale virtual worlds. Redirected Walking (RDW) technology overcomes this limitation by dynamically adjusting the user's walking direction and trajectory in the physical space, enabling them to walk through expansive virtual environments within a confined physical space. RDW technology has been widely applied in various fields, including virtual museums [6, 42], immersive entertainment

* Corresponding author. X. Liu, Q. Zhang, G. Yang, Y. Li, and G. Geng are with the School of Information Science and Technology, Northwest University, Xi'an, China. E-mail: liuxinda@nwu.edu.cn, zhangqinyu@stumail.nwu.edu.cn, yangguoqiang@stumail.nwu.edu.cn, liyunchen@stumail.nwu.edu.cn, gengguohua@nwu.edu.cn. J. Wu and L. Wang are with the State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, China. E-mail: lanayawj@buaa.edu.cn, wanglily@buaa.edu.cn.

[19, 46], and emergency-response training [23, 13], providing rich virtual experiences in space-constrained environments and advancing the application of VR technology across multiple scenarios.

Most existing RDW methods have been tested primarily in controlled environments, where the physical space is typically static and the distribution of obstacles is fixed. However, the increasing demand for greater immersion and more complex user interactions in VR has made isolated and constrained environments insufficient for supporting users in freely and naturally exploring expansive virtual worlds [44, 33]. This trend is already visible in location-based VR venues, where users walk within large-scale tracked areas to explore much larger virtual environments [28]. As deployments move to larger and more complex open physical spaces, the performance of existing methods in ensuring safe and natural walking remains underexplored in dynamic, complex, and multi-user shared tracked spaces. In particular, users must not only interact with surrounding objects but also avoid collisions with other users in these large shared tracked spaces, considerably increasing the difficulty of RDW methods [35, 21, 6] in handling such complex environments. These methods typically rely on static or controlled environments, where the movement of users and obstacles is more structured or controlled, making it difficult to adapt to the complexity and variability of real-world, multi-user settings.

In this work, we introduce a new task—Open Physical Spaces Redirected Walking (OPS-RDW), which addresses RDW in large-scale open physical spaces with dynamic obstacles, including co-located users. The core objective of this task is to enable users to engage in VR activities within expansive and dynamic open physical spaces, facilitating natural and free walking in these complex environments. An overview of the OPS-RDW task is illustrated in Figure 1. Specifically, this task must address several key challenges: ensuring real-time avoidance of dynamic obstacles without requiring a fixed or predefined boundary model, guaranteeing the user's safety under dynamic, multi-user conditions, and maintaining the naturalness of the user's walking trajectory while ensuring alignment with the intended navigation direction despite environmental changes.

To this end, we propose a novel method, **Dynamic Control and Redirection with Safety Constraints (DyCoRe)**. Our insight is to approach the OPS-RDW task as a dynamic safety constraint problem, which is addressed through a staged closed-loop process. First, we utilize Dynamic Control Barrier Functions (D-CBF) to define real-time safety constraints between the user and surrounding dynamic agents. Next, we apply online Quadratic Programming (QP) optimization to minimize the deviation between the user's walking path and the desired navigation direction, while ensuring the safety constraints are satisfied. Finally, the optimized velocity is mapped to redirection gains, guiding the user along a safe and natural walking path, with a relaxation mechanism ensuring feasible navigation even in dense or dynamic scenarios. Simulations provide supporting evidence, while a co-located user study shows that DyCoRe significantly improves navigation continuity and reduces resets, and tends to reduce perceived discomfort. By enforcing real-time safety constraints, DyCoRe reduces disruptions in complex, dynamic scenarios.

Our main contributions are summarized as follows:

(1) We introduce the OPS-RDW, which aims to enable users to engage in VR activities within shared tracked spaces, addressing key challenges such as real-time avoidance of dynamic obstacles and ensuring the safety and naturalness of walking paths in complex environments.

(2) We propose DyCoRe, a method for OPS-RDW that uses D-CBF for safety constraints and QP optimization to minimize path deviation. The optimized velocity is mapped to redirection gains, guiding the user along a safe, natural path, with a relaxation mechanism for feasibility in dense or unpredictable environments.

(3) We validate DyCoRe through extensive experiments, showing a 21.58% reduction in collisions with dynamic obstacles and a 20.12% increase in average distance from obstacles, demonstrating its superior performance in complex, multi-user environments.

2 Related Works

Research on RDW can be broadly divided into three categories: classic geometric methods, perception-driven approaches, and learning-based or multi-user extensions [32, 25]. In the following, we review works in each category and discuss their limitations for large shared tracked spaces with dynamic agents, without explicit boundary models.

2.1 Classic RDW methods.

RDW initially focused on geometric gains, with strategies such as Steer-to-Center (S2C), Steer-to-Multiple-Targets (S2MT), and Steer-to-Orbit (S2O) [35]. These approaches assume a static, enclosed, and fully known physical environment, redirecting users by applying rotation, translation, or curvature gains within perceptual thresholds. Comparative studies have evaluated and extended these strategies, showing that S2C is generally more robust, whereas S2O performs better in long straight-line walking scenarios. Variants of S2C have further been proposed to mitigate performance degradation when the angle between the walking and center directions becomes too large [47, 12]. Progress in threshold research has quantified the imperceptible ranges of rotation, translation, and curvature gains, establishing parameter foundations such as the relationship between curvature radius, walking speed, and field of view, as well as the independence of gain combinations [9, 31]. Nevertheless, these methods are intrinsically tied to static, pre-modeled indoor settings, and their effectiveness diminishes markedly in dynamic or unstructured tracked spaces. Recent surveys emphasize that the core assumptions of traditional RDW do not generalize to real-world scenarios, underscoring the necessity of more adaptive paradigms [1, 27, 41].

2.2 Perception-Driven RDW.

In recent years, research has increasingly focused on incorporating human factors and perception-driven models into RDW [45, 38]. These approaches dynamically adjust parameters based on individual thresholds [24, 18], eye movements [7, 45], and physiological signals [45], aiming to optimize redirection intensity and timing without compromising user comfort. Early work explored the suppression window induced by blinking or saccades, applying larger rotations or translations when the user is not directly looking ahead, thereby concealing redirection [38, 20, 3]. More recent studies have integrated saccade prediction, attention masking, and dynamic gaze modeling into online control systems [14, 15]. Gaze speed, blink rate, and gaze stability have been identified as reliable predictors of rotation-gain sensitivity, enabling real-time, user-specific adaptations. However, excessive gains can disrupt spatial memory and path recall, necessitating a balance between efficiency and cognitive load [4]. Despite the increasing sophistication of these perceptual models, many still rely on pre-sampling or offline calibration of individual users. When environmental conditions change unexpectedly, both thresholds and attention patterns shift with context. This makes it difficult for existing models to adapt in real-time to such dynamic changes.

2.3 Learning-Based RDW.

With the rapid advancement of machine learning, particularly reinforcement learning, RDW research has increasingly explored data-driven strategies to optimize user redirection. Lee et al. [21] introduced Steer-to-Optimal-Target, which selects the optimal redirection target from multiple candidates using Deep Q-Learning to evaluate and score alternatives. This approach was later extended to

support multiple users in Multiuser-Steer-to-Optimal-Target [22], incorporating pre-reset actions and additional refinements to improve multi-user coordination. Similarly, Strauss et al. [37] leveraged Proximal Policy Optimization (PPO) to train deep neural networks that predict rotation, translation, and curvature gains directly from a user's real-time position and orientation. Chang et al. [5] further advanced this paradigm by integrating a center-based translation gain strategy with a novel Turn-to-Furthest reset technique, demonstrating enhanced redirection performance. Beyond reinforcement learning, recent works have also explored physics- and alignment-inspired methods. Artificial Potential Function (APF)-based controllers, such as TAPF [40] and subsequent multi-user adaptations, use repulsive and attractive forces to guide users away from obstacles and towards desirable virtual areas. Alignment-based approaches, including REA [39] and ARC [43, 44], focus on synchronizing physical and virtual environments via geometric or visibility-based alignment to facilitate safer and more predictable walking. While these methods demonstrate improved redirection in constrained or semi-structured spaces, they often struggle with efficiency and obstacle avoidance in large, tracked, and dynamic environments.

2.4 Multi-User RDW.

The development of RDW for multi-user and tracked spaces has advanced rapidly in recent years. Artificial Potential Field-based approaches utilize repulsive and attractive forces for real-time obstacle avoidance, with early work demonstrating reduced collisions and fewer resets in multi-user room scenarios [2, 30]. Recent advancements in this area include predictive APF and alignment-based control methods, which reduce virtual-physical conflicts by aligning environmental or visible domain information, thereby facilitating more natural walking in non-rectangular spaces [8, 43, 44, 10]. In the context of tracked spaces, researchers have integrated cutting-edge exploration techniques with APF, allowing users to expand tracking areas while preserving natural walking in irregular physical spaces [29]. Moreover, joint optimization strategies that address occlusion consistency and feasible domains have been introduced, leveraging physical-virtual visibility polygons and obstacle consistency constraints to enhance safety and immersion, particularly in augmented reality or semi-open environments [44]. The evolution of multi-user interaction in RDW reached a significant milestone with the introduction of AVF-RDW in 2025, which combines vector fields with interaction constraints to support physical interactions between multiple users, such as handshakes and object transfers [36]. Additionally, spatial partitioning approaches have been proposed, dynamically dividing physical areas for multi-user environments to minimize interference [16]. However, these methods remain highly dependent on pre-existing environmental models, typically requiring prior mapping or stable boundary tracking. Their performance often deteriorates significantly in dynamic public spaces, where factors such as human movement patterns and boundary changes introduce substantial unpredictability. Many of these techniques still necessitate resets or predefined paths to ensure safety. Consequently, there is a pressing need for a solution that can operate efficiently and naturally within large, unmodeled, and dynamic tracked spaces.

In summary, because most prior RDW approaches still rely on predefined environment models or structured layouts, their robustness degrades in large tracked spaces dominated by dynamic agents; thus, we propose DyCoRe, which prioritizes real-time safety constraints without explicitly remodeling the environment or prescribing paths.

3 Problem Formulation

To formalize our proposed task, Open Physical Spaces Redirected Walking (OPS-RDW), we first outline the classic Redirected Walking (RDW) formulation and then extend it to tracked, dynamic, multi-user spaces.

3.1 Classic Redirected Walking

RDW aims to enable users of VR systems, who are constrained within a limited physical space, to walk naturally within expansive virtual environments. This is achieved by applying subtle geometric transformations to the user's visual input. Formally, let the trajectory of the user in the virtual environment at time t be denoted as $V(t) \in \mathbb{R}^n$, and the corresponding trajectory in the physical environment as $P(t) \in \mathbb{R}^n$. The relationship between these trajectories is described by a mapping function:

$$P(t) = f(V(t)), \tag{1}$$

where $f(\cdot)$ represents a mapping function, and it is required to be continuously differentiable and locally injective. These conditions ensure the smoothness and reversibility of the physical motion, which are crucial for maintaining a natural and uninterrupted user experience in the physical space. In practice, such mapping functions are typically implemented through gain models, including translation gains g_t, rotation gains g_r, and curvature gains g_c, which geometrically adjust the user's perceived motion. These adjustments are designed to remain imperceptible to the user, ensuring that the transformations do not disrupt the natural walking experience.

3.2 Open Physical Spaces Redirected Walking

OPS-RDW extends Classic RDW to dynamic, complex, and multi-user shared tracked spaces. We define the OPS-RDW task as a spatial mapping problem under dynamic constraints, with the mapping function given by:

$$P(t) = f(V(t), E(t)), \tag{2}$$

where $E(t)$ represents the state of the physical environment at time t, including the distribution of dynamic obstacles and the trajectories of other users. This mapping function f not only depends on the virtual trajectory $V(t)$ but must also respond in real-time to changes in the environment. Moreover, the user's physical trajectory must remain within a safe and feasible region:

$$P(t) \in F(E(t)), \tag{3}$$

where $F(E(t))$ represents the safe, feasible region determined by the environmental state at time t. Additionally, the deviation between the physical and virtual trajectories must be kept within perceptual thresholds to ensure that the user's movement feels natural. In summary, the mathematical formulation of OPS-RDW can be conceptualized as solving the mapping function f under dynamic environmental constraints, with the objective of maximizing the natural consistency between virtual and physical motion while ensuring safety.

4 Dynamic Control and Redirection with Safety Constraints

Building on the problem formulation of OPS-RDW, our goal is to realize the mapping function f such that the physical trajectory $P(t)$ satisfies both the safety constraint $P(t) \in F(E(t))$ and the naturalness constraint $\|P(t) - V(t)\|$ within the perceptual threshold. To this end, we propose **Dynamic Control and Redirection with Safety Constraints (DyCoRe)**, which approaches OPS-RDW as a dynamic safety-constrained optimization problem. Specifically, we employ Dynamic Control Barrier Functions (D-CBF) to encode safety constraints imposed by dynamic obstacles, and solve an online Quadratic Programming (QP) problem to minimize deviation from the intended navigation direction under these constraints. The

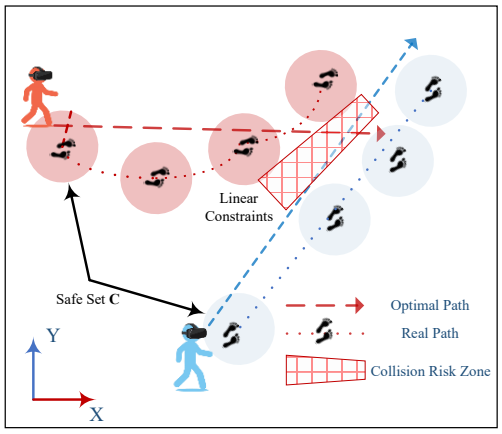

Figure 2: Illustration of dynamic safety constraints with D-CBF. The method ensures that each user's safety set is maintained without intersecting with others, thus preventing collisions in the open physical space.

optimized motion commands are then translated into redirection gains, ensuring a safe and natural walking experience. The following subsections detail the three main components of DyCoRe: dynamic safety constraints with D-CBF, online optimization with QP, and redirection gain mapping.

4.1 Dynamic Safety Constraints with D-CBF

In OPS-RDW, users are immersed in VR environments with fuzzy physical boundaries and multiple co-existing users. To guarantee safe locomotion, we employ D-CBFs to transform user–obstacle interactions into linear constraints on user velocity, as shown in Figure 2.

Let the user position at time t be $p_u(t) \in \mathbb{R}^2$, and the position of obstacle o_i be $p_{o_i}(t) \in \mathbb{R}^2$. The minimum safety distance is denoted by $d_{\text{safe}} > 0$. We define the safety function as:

$$h_i(p_u, p_{o_i}) = \|p_u - p_{o_i}\|^2 - d_{\text{safe}}^2, \tag{4}$$

with the associated safe set:

$$C = \{(p_u, p_{o_i}) \mid h_i(p_u, p_{o_i}) \geq 0, \forall i\}. \tag{5}$$

To ensure that the safe set C effectively constrains the system state, the safety function is required to satisfy the Zeroing CBF condition:

$$\dot{h}_i(p_u, p_{o_i}, u) + \alpha h_i(p_u, p_{o_i}) \geq 0, \tag{6}$$

where $u \in \mathbb{R}^2$ is the user's virtual control velocity and $\alpha > 0$ is a hyperparameter regulating the system's responsiveness when approaching the safety boundary. This inequality ensures that if the system state starts inside the safe set C (i.e., $h\_i \geq 0$), its future evolution will not drive it outside, thereby guaranteeing safety at all times. The time derivative of h_i is

$$\dot{h}_i = 2(p_u - p_{o_i})^\top (u - v_{o_i}), \tag{7}$$

where $v_{o_i} \in \mathbb{R}^2$ denotes the obstacle velocity. Substituting this into the Zeroing CBF condition yields a linear inequality in u:

$$a_i^\top u \geq b_i, \tag{8}$$

where $a_i = 2(p_u - p_{o_i})$, $b_i = 2(p_u - p_{o_i})^\top v_{o_i} - \alpha h_i(p_u, p_{o_i})$. This constraint defines a convex polytope for the feasible set of u, which can be directly embedded into a QP solver to explicitly enforce safety during optimization. The parameter α controls constraint conservativeness: larger values impose stricter restrictions near obstacles, while smaller values allow greater freedom at the cost of increased risk.

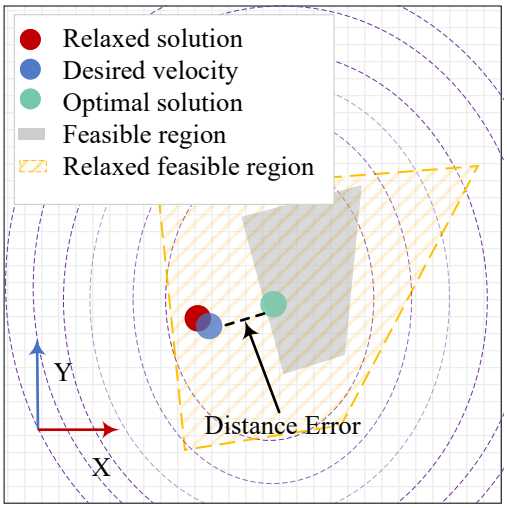

Figure 3: Illustration of QP-based online optimization with relaxation. The desired velocity may fall outside the feasible region; the optimal solution is projected inside, and when infeasible, a relaxed solution is obtained within the relaxed feasible region.

4.2 Online Optimization with QP

QP allows for the effective integration of dynamic safety constraints derived from D-CBF. This ensures that the system adheres to obstacle avoidance requirements throughout the optimization process. In OPS-RDW, however, it is essential to not only consider safety but also ensure the naturalness of the path. Specifically, the user's trajectory should closely follow the desired path to enhance immersion and reduce motion sickness.

Based on Equ. 8, a straightforward solution would be to minimize the Euclidean distance between the control speed u and the desired speed u_{des} using QP, while ensuring the satisfaction of safety constraints. The optimal speed u^* is formulated as:

$$u^* = \arg\min_u \|u - u_{\text{des}}\|_2^2 \quad \text{s.t.} \quad a_i^\top u \geq b_i, \tag{9}$$

where$\|.\|_2^2$ denotes euclidean distance. This formulation transforms the path planning problem into a standard convex optimization problem, which can be efficiently solved in real time using a solver. However, in extreme cases, such as when the user is very close to obstacles or when the goal direction conflicts with obstacles, the linear constraints may conflict with the objective, leading to an infeasible QP solution. we introduce a non-negative relaxation variable $\delta \geq 0$, which allows for the violation of certain constraints when necessary, with a penalty applied based on the degree of violation. The modified constraints are:

$$a(s)^\top u + \delta \geq b(s), \quad \delta \geq 0, \tag{10}$$

where s denotes the system state vector. The optimization objective is then extended as follows:

$$(u^*, \delta) = \arg\min_{u,\delta} \|u - u_{\text{des}}\|_2^2 + \lambda \delta^2, \tag{11}$$

where $\lambda > 0$ is the penalty coefficient, controlling the tolerance for violating safety constraints; δ represents the distance by which the safety boundary is violated, with smaller values indicating greater

safety. The final optimization problem is:

$$(u^*, \delta) = \arg\min_{u,\delta} \begin{cases} \|u - u_{\text{des}}\|_2^2 + \lambda \delta^2 \\ \text{s.t.} \quad a(p_u, p_{oi})^\top u + \delta \geq b(p_u, p_{oi}), \\ \quad \delta \geq 0. \end{cases} \tag{12}$$

To illustrate this relaxation mechanism, we visualize the relationship between the desired velocity, the optimal solution within the feasible region, and the relaxed solution when constraints become infeasible in Figure 3.

In practice, the QP controller is embedded into the update loop of each frame. After receiving the user and obstacle states, the required components for the QP problem are constructed, and the optimal control velocity u^* is computed using an efficient solver. This control velocity u^* both satisfies the obstacle avoidance requirement and closely follows the desired path, forming the basis for generating redirection gains in the system.

4.3 Redirection Gain Mapping

User motion is inherently voluntary in VR system, preventing direct control over physical speed or trajectory. Consequently, the ideal velocity u^* serves only as an intention signal rather than a direct command. Redirection Gain Mapping treats u^* as the desired velocity and dynamically adjusts user motion by combining three types of gains g_t, g_r, g_c based on the real-time user velocity v_{real}, thereby guiding users along the intended path while preserving natural motion perception.

The scaling gain g_t modulates the user's step length to align virtual locomotion with the desired speed while maintaining safety constraints:

$$g_t = \frac{\|u^*\|}{\|v_{\text{real}}\| + \varepsilon}, \tag{13}$$

where $\|u^*\|$ is the magnitude of the desired virtual velocity, $\|v_{\text{real}}\|$ is the user's actual speed, and ε is a small constant to prevent division by zero. Curvature gain g_c introduces subtle path deviations to create a turning illusion, encouraging the user to naturally correct their heading. We first compute the direction difference:

$$\theta_{\text{real}} = \arg(v_{\text{real}}), \quad \theta_{\text{cbf}} = \arg(u^*), \quad \Delta\theta = \theta_{\text{cbf}} - \theta_{\text{real}}, \tag{14}$$

and then map this deviation to curvature gain using a linear approximation:

$$g_c = 1 + k \cdot \Delta\theta, \tag{15}$$

where k is an empirically chosen coefficient controlling directional adjustment sensitivity. The rotation gain g_r adjusts the virtual head rotation velocity to synchronize with the user's actual head motion, ensuring alignment between the user's viewpoint and the planned navigation direction:

$$g_r = \frac{\omega_{\text{cbf}}}{\omega_{\text{real}} + \varepsilon}, \tag{16}$$

where ω_{real} is the user's measured head angular velocity, and ω_{cbf} represents the desired virtual angular velocity, derived from the trajectory planner or the rate of navigation direction change; if not explicitly designed, ω_{cbf} can be set to a fixed value or smoothly propagated from the previous frame. The three gains are applied in a coordinated manner and updated in real time based on QP optimization outputs. Each gain is low-pass filtered to suppress high-frequency variations, ensuring that the resulting user trajectories are smooth, safe, and perceptually natural. By modulating perceived motion rather than directly controlling it, this mechanism enables effective path guidance while preserving immersive and voluntary navigation.

DyCoRe formulates OPS-RDW as a dynamic, safety-constrained optimization problem, combining D-CBF-based obstacle avoidance with online QP to generate control velocities that balance safety and naturalness. These velocities are then translated into real-time redirection gains, subtly guiding user motion while preserving immersive and voluntary locomotion. Together, this method ensures safe, smooth, and perceptually natural VR navigation in complex, dynamic environments.

5 Simulation Study

We first conduct simulation experiments against seven state-of-the-art RDW techniques to examine the behavior and robustness of the proposed DyCoRe method. Simulation allows us to explore performance trends and variability across diverse navigation paths and dynamic obstacle configurations under controlled conditions.

5.1 Experimental Setup

We simulate user locomotion using the motion model of Razzaque et al. [34]. The simulation defines a constrained physical workspace of 15×15 m^2 to bound the virtual walker's movement. This large-scale finite environment serves as a testbed to evaluate the algorithm's performance in scenarios theoretically applicable to large tracked spaces. The virtual walker moves at a constant speed of 1.0 m/s and initially faces the positive y-axis. We procedurally generate 50 random virtual navigation paths, each consisting of 10–15 waypoints. Inter-point distances are uniformly sampled from 5 to 15 m, and steering angles are constrained within $[-90^\circ, +90^\circ]$, resulting in an average path length of approximately 100 m.

Reset handling. A reset is triggered whenever the virtual walker approaches within 1.0 m of the simulated workspace boundary or the distance to a dynamic obstacle drops below 0.5 m. Upon such an event, the system pauses the virtual walker and executes a standard Stop-and-Turn strategy to reorient the user toward a safe direction. This protocol is applied identically across all evaluated methods. While all resets are recorded, only segments involving dynamic collisions are retained for subsequent performance evaluation.

Dynamic obstacles. Dynamic obstacles are introduced to emulate other moving users. Obstacles are randomly generated along the environment boundary and traverse the workspace in a straight line towards a target on the opposite side at a constant speed of 0.9 m/s. Obstacles are continuously spawned such that at least one dynamic obstacle is present at all times.

Evaluation and baselines. We evaluate performance using three metrics: Reset Count (RC), Virtual Distance Between Resets (VDBR), and Time Elapsed Between Resets (TEBR). DyCoRe is compared against seven representative RDW methods, including S2C, S2O, T-APF, ZigZag, a DeepLearning-based approach, MessAPF, and PassiveAPF. All methods are implemented within the OpenRDW v1.0 framework [26] and evaluated on the same 50 paths with identical obstacle schedules using fixed random seeds. Simulations are executed on a Windows workstation equipped with an Intel Core i7-14700HX CPU and 64 GB RAM.

5.2 Results

In this simulation study, all methods are evaluated under identical path and obstacle conditions, resulting in a paired experimental design for all simulation metrics.

Statistical Analysis. As all methods are evaluated under a paired experimental design, we employ non-parametric repeated-measures statistical analyses. For each metric, overall differences among the eight methods are assessed using a Friedman test, followed by post-hoc Wilcoxon signed-rank tests for pairwise comparisons between DyCoRe and each baseline. Holm–Bonferroni correction is applied to account for multiple comparisons, and results are considered statistically significant only when the corrected $p < 0.05$.

Reset Count (RC). Figure 4 (a) illustrates the distribution of reset counts across all evaluated methods in the simulation study.

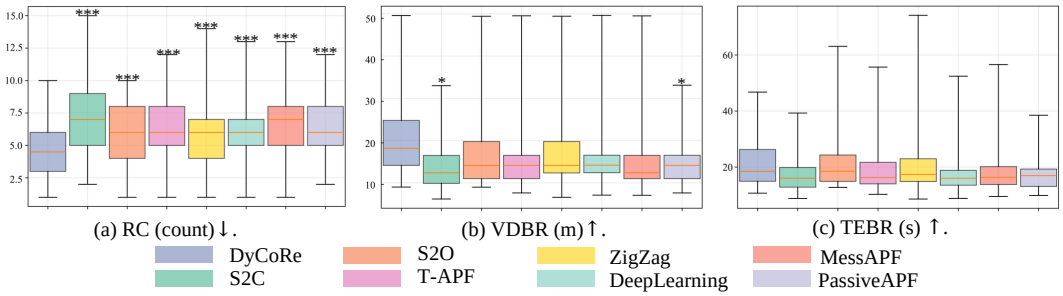

Figure 4: Simulation results across eight RDW methods. **(a)** Reset Count (RC), **(b)** Virtual Distance Between Resets (VDBR), and **(c)** Time Elapsed Between Resets (TEBR). ↓ denotes lower is better, ↑ denotes higher is better. ∗ indicates statistically significant differences between DyCoRe and each baseline after correction for multiple comparisons; non-significant comparisons are omitted for clarity.

DyCoRe yields the lowest median reset count with consistently reduced dispersion, indicating fewer interruptions and more stable navigation under dynamic conditions. Statistical testing confirms a significant effect of the redirection method on reset frequency, with DyCoRe incurring fewer resets than all baseline approaches after correction for multiple comparisons. Together, these results indicate that DyCoRe provides a clear robustness advantage in mitigating reset events under dynamic disturbances, which is further examined through real user experiments in the following section.

Virtual Distance Between Resets (VDBR). Figure 4 (b) presents the virtual distance traveled between consecutive resets. DyCoRe attains a higher median VDBR compared to most baseline methods, indicating longer uninterrupted navigation in many simulation trials. Statistical testing reveals that this improvement is significant for a subset of baselines, while differences with others do not reach statistical significance after correction for multiple comparisons. Overall, these results suggest that DyCoRe tends to improve spatial continuity in dynamic environments, although the magnitude of this benefit varies across comparison methods.

Time Elapsed Between Resets (TEBR). Figure 4 (c) summarizes the time elapsed between consecutive resets. DyCoRe exhibits competitive median TEBR values relative to baseline methods, with reduced variability in several trials. However, after correction for multiple comparisons, differences in TEBR do not reach statistical significance across baselines. This indicates that, in simulation, DyCoRe does not consistently extend the temporal duration between resets, and that temporal robustness may be more sensitive to scenario-specific factors than spatial or event-based measures. As discussed in the following section, the practical implications of these trends are further examined through real user experiments.

6 User Study

We conducted a controlled virtual walking experiment with real users to evaluate the applicability and user experience of the proposed DyCoRe method, which provides direct evidence beyond simulation-based analysis. The experiment aims to evaluate the comprehensive performance of the method in terms of navigational efficiency, robustness to dynamic obstacles, and user comfort, and compare it with multiple typical redirection strategies to ensure the comprehensiveness and reproducibility of the conclusions.

6.1 Participants

We recruited 30 participants, 15 male and 15 female, aged 22 to 28 years old, mean 25.1 years old. This study was approved by the Medical Ethics Committee of the Northwest University (IRB Protocol No. 241113089, approved on 2024-11-13), and all participants provided written informed consent. Participants reported normal vision, no history of severe motion sickness, and basic VR experience, but no prior training in redirected walking. Participants were organized into pairs to facilitate the study of dynamic obstacle interaction. Within each trial, distinct roles were assigned: one participant acted as the redirected walker, while the other acted as a dynamic obstacle walker. We used a circular role rotation to assign the tracked obstacle: participant n served as the obstacle for participant $n+1$, and participant 30 served as the obstacle for participant 1. This design yielded one redirected walker session per participant, covering all conditions. Anonymized identifiers were used throughout the study to protect privacy and ensure ethical compliance in data collection. An experimenter continuously monitored the sessions to ensure participant safety.

6.2 Experiment Design

The user study followed a paired, within-subject design in which two participants simultaneously walked in a shared physical space under four redirected walking conditions: DyCoRe, S2C, S2O, and T-APF.

Experimental conditions and design. We compare the proposed method with representative reaction-based redirected walking algorithms, including S2C, S2O, and T-APF, consistent with the scope of this study. In each session, the redirected walker navigated the virtual environment subject to the specific redirection algorithm, while the dynamic obstacle walker acted as a dynamic obstacle encountered in the shared physical space. We fixed the obstacle trajectory across the four methods to ensure comparability, generating a random path per pair for the dynamic obstacle walker and reusing it unchanged across methods. To mitigate potential learning effects, a short familiarization session was provided before the formal experiment to allow participants to become accustomed to the virtual walking task and redirection effects. In addition, all participants followed the same sequence of virtual waypoints in each condition to ensure consistent task difficulty across methods. The order of the four conditions was fixed across participants due to logistical constraints. These measures were adopted to reduce potential learning effects while maintaining consistency across experimental conditions.

Apparatus and environment. The experiment was conducted in a large, flat outdoor open space to allow unconstrained walking and to ensure participant safety by avoiding external obstacles, as shown in Figure 6. The experimental platform was built on OpenRDW v1.0 [26]. A virtual boundary of 15×15 m^2 was configured in software, while the available physical area was substantially larger than this boundary. We used two Pico 4 Ultra head-mounted displays (HMDs) with a binocular resolution of 2160×2160 pixels per eye. Although both participants wore the same HMD model, the devices served different roles. For the redirected walker, the HMD rendered the immersive virtual environment and executed the redirection algorithms. For the dynamic obstacle walker, the HMD rendered a minimal virtual environment containing visual guidance cues to direct the participant along the pre-defined path; the participant's pose was wirelessly streamed in real time to the walker's

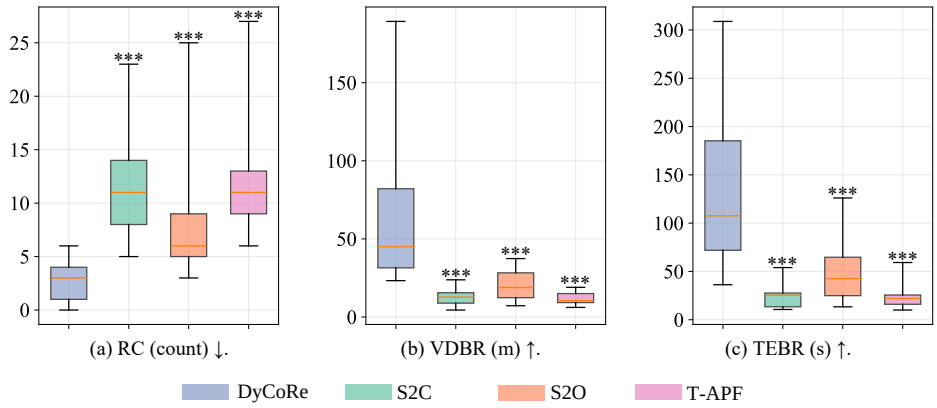

Figure 5: User study results across four RDW methods. **(a)** Reset Count (RC, times), **(b)** Virtual Distance Between Resets (VDBR, meter), and **(c)** Time Elapsed Between Resets (TEBR, second). ↓ denotes lower is better, ↑ denotes higher is better. $*$ indicates statistically significant differences between DyCoRe and each baseline after correction for multiple comparisons.

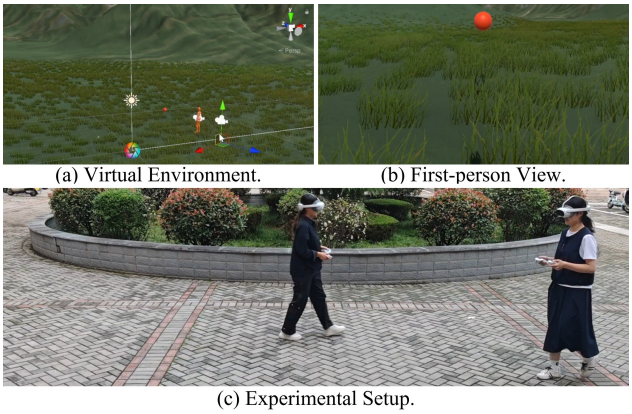

(a) Virtual Environment. (b) First-person View.

(c) Experimental Setup.

Figure 6: Overview of the experimental environment and physical setup. (a) The global view of the virtual environment constructed in Unity. (b) The participant's first-person perspective showing the target waypoint. (c) The physical experimental site in a large outdoor open space, where two participants navigate the shared area simultaneously.

system to represent the dynamic obstacle state in the shared virtual environment. A networked architecture consisting of two laptop computers and a dedicated wireless router supported the multi-user system: each HMD was paired with a laptop serving as a rendering workstation, and inter-device synchronization was established via a dedicated Wi-Fi 6 router, forming a low-latency WLAN. To align both devices to a unified global reference frame, we employed a physical initialization protocol. Before each session, the two HMDs were placed at a predefined floor landmark and oriented in the same cardinal direction. Upon launching the application, each HMD's internal tracking initialized its virtual origin $(0,0,0)$ and forward vector from this physical pose, effectively aligning the two local coordinate systems and enabling spatially consistent interaction and collision detection. The application was developed in Unity and ran at a locked frame rate of 60 Hz. At each frame update, the system synchronized both users' global coordinates in real time. A collision event was defined purely by proximity: a collision was recorded when the Euclidean distance between the two users' planar coordinates (x,z) fell below a safety threshold of 0.5 m.

Virtual task and procedure. The virtual task required participants to navigate through a sequence of target waypoints in an open virtual environment. The virtual scene was designed as a visually open, large-scale flat grassland without explicit boundary cues related to physical space limits and to focus the evaluation on redirection behavior under dynamic obstacles. Target locations were indicated by a virtual crystal placed at the current goal position. Participants were instructed to walk toward the crystal using natural walking. When a participant reached and touched the crystal, the current target disappeared and the next target crystal appeared at a different location in the virtual environment. Each trial consisted of a sequence of 15 to 20 waypoints. Before the formal experiment, participants completed a short training session in the same virtual environment. This familiarization phase allowed participants to become accustomed to the virtual walking task and the redirected walking experience, thereby reducing initial discomfort and learning effects during data collection.

Measures. We collected both objective locomotion metrics and subjective user feedback to evaluate the performance and user experience of different redirected walking methods. **RC** measures the number of times a participant was required to reset their walking direction during a trial due to safety constraints. A lower reset count indicates fewer interruptions to continuous walking and is commonly used as an indicator of redirection efficiency. We computed the **VDBR** for each participant. This metric reflects how long users can walk continuously without interruption. Larger distances between resets indicate smoother redirection and improved locomotion continuity. We also recorded the **TEBR** during each trial. This measure captures the temporal continuity of walking and complements the distance-based metric. After completing each experimental condition, participants completed the **Simulator Sickness Questionnaire (SSQ)**. The SSQ provides a subjective measure of discomfort related to simulator sickness, with lower scores indicating better user comfort.

6.3 Results

In the user study, we compare the DyCoRe, S2C, S2O and T-APF strategies across three key performance metrics under dynamic obstacle conditions: RC, VDBR, and TEBR. Statistical analysis using Friedman tests revealed significant overall differences among the four methods across all metrics, as illustrated in Figure 5.

Overall statistical analysis. We conducted non-parametric statistical analyses on each objective performance metric to examine overall differences among the four RDW methods, given the within-subject experimental design and the non-normal distribution of the collected data. Friedman tests were applied separately to RC, VDBR, and TEBR. The results indicate significant overall differences among the four methods across all three metrics, with large effect sizes as measured by Kendall's coefficient of concordance. Based on these significant overall effects, we subsequently per-

Table 1: Overall statistical results of the user study based on Friedman tests.

Metric	Test	$\chi^2(3)$	p	Kendall's W
RC	Friedman	53.34	< 0.001	0.71
VDBR	Friedman	48.79	< 0.001	0.65
TEBR	Friedman	48.22	< 0.001	0.64
SSQ	Friedman	15.53	< 0.01	0.21

formed pairwise Wilcoxon signed-rank tests to compare DyCoRe with each baseline method, applying Holm–Bonferroni correction to control for multiple comparisons. An overview of the overall statistical results is summarized in Table 1.

Comparing DyCoRe with S2C. Pairwise comparisons using Wilcoxon signed-rank tests with Holm–Bonferroni correction revealed that DyCoRe achieved significantly better performance than S2C across RC, VDBR, and TEBR (all $p < 0.001$). The DyCoRe shows a significantly lower number of resets compared to S2C in the left panel. The median reset count of S2C is nearly three times higher, and its distribution shows a wider variance, indicating more frequent interruptions in navigation and a less robust response to dynamic disturbances. In contrast, DyCoRe exhibits tighter bounds and a lower maximum, suggesting greater stability. The center further supports DyCoRe's advantage in terms of spatial performance. The average virtual distance between resets for DyCoRe is substantially higher than that of S2C. The median value for S2C remains clustered at the lower end, while DyCoRe not only achieves a higher median but also exhibits a broader, upward-shifted range, demonstrating that it enables more extended travel without requiring recovery or reinitialization. The right panel further indicates DyCoRe's improved temporal continuity. It achieves a significantly longer average time between resets, with a wider interquartile range and higher maximum value compared to S2C. This indicates that DyCoRe is capable of sustaining effective operation for longer durations under dynamic conditions, reducing the frequency of disruptions over time.

In summary, the results highlight DyCoRe's advantages over S2C. DyCoRe consistently achieves fewer resets, longer travel distances, and extended uninterrupted operation times, indicating its enhanced robustness, efficiency, and reliability in dynamic environments.

Comparing DyCoRe with S2O. Pairwise Wilcoxon signed-rank tests with Holm–Bonferroni correction showed that DyCoRe performed better than S2O across RC, VDBR, and TEBR ($p < 0.001$ for all metrics). DyCoRe exhibits a lower median reset count, with a narrower IQR and a lower maximum value compared to S2O. This suggests that DyCoRe is better at maintaining continuous operation, requiring fewer resets during navigation, even under frequent dynamic disturbances. The more compact distribution of reset counts for DyCoRe reflects more consistent navigation behavior across participants. The middle panel illustrates the difference between DyCoRe and S2O in terms of spatial continuity. DyCoRe achieves a substantially higher median and upper quartile for VDBR than S2O, indicating that DyCoRe can sustain longer uninterrupted travel between resets. In contrast, S2O's distribution is tightly clustered around shorter distances, reflecting its limited spatial continuity. This reflects improved spatial continuity for DyCoRe compared to S2O. The right panel illustrates the temporal robustness of the two strategies, represented by TEBR. DyCoRe shows a significantly higher TEBR than S2O. This indicates improved temporal continuity for DyCoRe under dynamic conditions. The wider distribution of TEBR for DyCoRe reflects more sustained operation across varied conditions compared to S2O.

Overall, the results indicate that DyCoRe achieves better performance than S2O across the evaluated metrics under dynamic obstacle conditions. DyCoRe achieves fewer resets, longer uninterrupted walking distances, and extended periods of continuous operation across participants.

Comparing DyCoRe with T-APF. Pairwise Wilcoxon signed-rank tests with Holm–Bonferroni correction indicated that DyCoRe performed better than T-APF across all three metrics (all $p < 0.001$). As shown in the left panel, DyCoRe results in a substantially lower reset count compared to T-APF. The median number of resets for DyCoRe is significantly lower, and the overall distribution is more compact, with a much smaller maximum value. This indicates that DyCoRe enables more stable navigation behavior, with fewer interruptions triggered by dynamic obstacles. In contrast, T-APF exhibits a higher median and a wide variance in reset counts, suggesting frequent failures and poor adaptability in the presence of moving agents.

The middle panel highlights the spatial advantage of DyCoRe. The VDBR for DyCoRe is markedly higher, both in terms of median and maximum values. DyCoRe is able to maintain extended travel distances between resets, reflecting greater path continuity and resilience. On the other hand, T-APF displays a narrow and low-valued distribution, with almost all resets occurring within short travel segments. This suggests limited capability in sustaining navigation without interruption.

The right panel further indicates DyCoRe's improved temporal stability. The TEBR for DyCoRe is significantly higher than that of T-APF, with both a broader interquartile range and a maximum exceeding 300. This indicates that DyCoRe can operate continuously for much longer durations before resets are triggered. In contrast, T-APF suffers from short operational intervals and limited temporal endurance, making it less suitable for scenarios that require sustained, robust control under dynamic conditions.

Collectively, these results support DyCoRe's advantages over T-APF. By achieving fewer resets, longer uninterrupted paths, and greater operational durations, DyCoRe provides improved robustness and continuity in the tested dynamic obstacle settings.

Simulator sickness questionnaire (SSQ) scores. The SSQ score is a measure of the severity of symptoms experienced due to simulator sickness, with higher scores indicating more severe symptoms. We calculated the weighted SSQ Total Score for each participant using the standard formula derived by Kennedy et al. [17] Statistical analysis was performed on these Total Scores to determine the significance of the differences between conditions. Figure 7 illustrates the comparison of SSQ across different strategies: DyCoRe, S2O, S2C, and T-APF. The SSQ comprises 16 specific symptoms, including general discomfort, fatigue, headache, eye strain, nausea, and various types of dizziness, providing a comprehensive assessment of participants' subjective discomfort during the experiment.

Figure 7 shows that the DyCoRe strategy, represented by the blue line, consistently yields lower scores across most SSQ items compared to the other strategies. In particular, DyCoRe shows consistently lower scores in items such as Headache, Nausea, Eye Fatigue, and the dizziness-related items, indicating reduced discomfort compared to the baseline strategies. Moreover, the DyCoRe strategy also shows lower scores in cognitively related symptoms, such as Blurred Vision and Difficulty Concentrating, further supporting its potential to enhance user comfort. In summary, the SSQ results suggest that DyCoRe may improve subjective comfort by reducing symptoms across multiple SSQ items.

6.4 Discussion

The experimental results from both the simulation and user studies provide consistent evidence that the proposed DyCoRe constraint-based field strategy improves performance compared to a wide range of baseline methods, including S2O, S2C, T-APF, ZigZag,

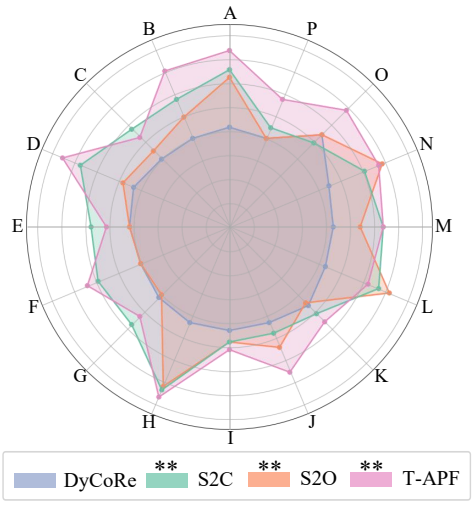

Figure 7: Radar chart of mean SSQ item scores for four RDW methods: DyCoRe, S2C, S2O, T-APF. Larger radii indicate more severe symptoms. Items A–P: A: General Discomfort; B: Fatigue; C: Headache; D: Eye Fatigue; E: Difficulty Focusing; F: Increased Saliva Secretion; G: Sweating; H: Nausea; I: Difficulty Concentrating; J: Feeling Heavy in the Head; K: Blurred Vision; L: Dizziness (Eyes Open); M: Dizziness (Eyes Closed); N: Dizziness; O: Vomit; P: Hiccups. $*$ indicates statistically significant differences between DyCoRe and each baseline after correction for multiple comparisons.

DeepLearning, MessAPF, PassiveAPF, and the user-tested counterparts S2O and S2C.

In simulation, DyCoRe achieved the lowest reset counts and generally longer VDBR under dynamic obstacles. While TEBR showed a descriptive upward trend, it was not statistically reliable after multiple-comparison correction. Overall, DyCoRe exhibited more stable behavior than T-APF and ZigZag, which showed higher variability and failure rates in complex motion patterns.

The user study results provide further corroboration of these findings. In all comparative evaluations involving S2O, S2C, and T-APF, DyCoRe consistently exhibited superior spatial continuity and temporal robustness. Participants using DyCoRe experienced fewer resets, covered longer distances without interruption, and spent more time navigating successfully between dynamic collisions. The statistical distributions of reset intervals were shifted toward favorable values and showed lower variance, indicating not only better average performance but also more predictable and stable behavior across runs. Moreover, DyCoRe outperformed traditional potential-field methods such as T-APF and learning-based approaches such as DeepLearning without relying on heavy offline training or extensive model tuning. Its interpretable structure and lightweight computational requirements make it a practical choice for real-world applications with strict resource constraints or stringent safety and reliability demands.

Despite these advantages, several limitations should be acknowledged. Although DyCoRe performed well in the tested dynamic scenarios, its adaptability to highly cluttered or adversarial environments with intelligent moving obstacles remains to be validated. The current implementation also does not explicitly incorporate learning-based feedback or global planning, which may further enhance performance in large-scale or partially observable settings.

In conclusion, results from both simulation and real-user experiments demonstrate that DyCoRe is a robust, efficient, and stable navigation strategy for dynamic environments, offering practical advantages over existing solutions and opening new possibilities for safer and more reliable autonomous systems.

7 Limitations and Conclusion

7.1 Limitations

While DyCoRe demonstrates notable advantages in open physical spaces, several limitations remain. First, DyCoRe's scalability in highly dense environments may be limited. Our experiments were conducted under moderate obstacle densities; in more crowded settings with highly dynamic or uncooperative obstacle behaviors, performance may degrade unless more advanced risk modeling or explicit crowd-interaction mechanisms are incorporated. Second, DyCoRe does not incorporate online learning or automatic parameter adaptation, and thus relies on fixed, hand-tuned parameters. While this design promotes stability and predictability, it may reduce robustness across varying environmental conditions, such as changes in obstacle velocity, density, or scale. Third, although our user study provides insights under controlled real-time interaction, further validation is needed for uncontrolled real-world deployments, where environmental noise, system latency, and operational uncertainties may introduce additional challenges. In addition, fixed condition order may introduce learning effects. DyCoRe assumes a managed, co-located tracked space with low-latency state sharing; untracked pedestrians and large-area network deployments are not considered. Moreover, our simulation baselines use OpenRDW v1.0 implementations that may differ from original descriptions; therefore, simulation-based comparisons should be interpreted with this toolkit dependency in mind [11].

7.2 Conclusion

In this work, we propose a novel RDW task, open physical spaces RDW (OPS-RDW) to address the challenges of enabling safe and natural VR locomotion in dynamic, large-scale tracked environments. To tackle this, we also propose the Dynamic Control and Redirection with Safety Constraints (DyCoRe) method that integrates Dynamic Control Barrier Functions with online optimization to balance safety and natural walking in OPS-RDW scenarios. Unlike traditional artificial potential field methods and modern data-driven approaches, DyCoRe leverages geometric constraints and dynamic risk evaluation to generate smooth and reliable trajectories without relying on learning-based policies or global maps.

We evaluated DyCoRe through a multi-tiered approach: while preliminary simulations characterized broader performance trends, a formal user study ($N = 30$) provided rigorous empirical validation, confirming our findings with high statistical significance. In real-world trials, DyCoRe not only achieved substantially higher VDBR and TEBR compared to state-of-the-art baselines ($p < 0.001$) but also exhibited a more stable and concentrated performance distribution. The synergy of objective locomotion metrics and lower SSQ scores confirms that DyCoRe provides a more reliable, consistent, and comfortable redirection experience.

Ultimately, DyCoRe offers a robust and scalable solution for VR navigation in open physical spaces. Its ability to ensure real-time safety and continuity without complex prior knowledge sets it apart from existing methods. As immersive environments evolve, DyCoRe provides a promising foundation for multi-user applications and unpredictable dynamic scenarios.

8 Acknowledgements

This work is partially supported by National Natural Science Foundation of China through Project (62571051), by Open Project Program of State Key Laboratory of Virtual Reality Technology and Systems, Beihang University (No. VRLAB2024C02), by General Projects of the Shaanxi Provincial Department of Science and Technology (2025JC-YBQN-801).

REFERENCES

[1] M. Azmandian, R. Yahata, T. Grechkin, and E. S. Rosenberg. Adaptive redirection: A context-aware redirected walking meta-strategy. *IEEE Transactions on Visualization and Computer Graphics*, 28(5):2277–2287, 2022. 2

[2] 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. 3

[3] B. Bolte and M. Lappe. Subliminal reorientation and repositioning in immersive virtual environments using saccadic suppression. *IEEE transactions on visualization and computer graphics*, 21(4):545–552, 2015. 2

[4] G. Bruder, P. Lubos, and F. Steinicke. Cognitive resource demands of redirected walking. *IEEE transactions on visualization and computer graphics*, 21(4):539–544, 2015. 2

[5] Y. Chang, K. Matsumoto, T. Narumi, T. Tanikawa, and M. Hirose. Redirection controller using reinforcement learning. *IEEE Access*, 9:145083–145097, 2021. 3

[6] J.-J. Chen, H.-C. Hung, Y.-R. Sun, and J.-H. Chuang. Apf-s2t: Steering to target redirection walking based on artificial potential fields. *IEEE Transactions on Visualization and Computer Graphics*, 30(5):2464–2473, 2024. 1, 2

[7] K. Davis, T. Hayase, I. Humer, B. Woodard, and C. Eckhardt. A quantitative analysis of redirected walking in virtual reality using saccadic eye movements. In *International Symposium on Visual Computing*, pp. 205–216. Springer, 2022. 2

[8] 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. 3

[9] 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, 2016. 2

[10] C. Hirt, N. Isaak, C. Holz, and A. Kunz. Predictive multiuser redirected walking using artificial potential fields. *Frontiers in Virtual Reality*, 5:1259429, 2024. 3

[11] C. Hirt, Y. Kompis, C. Holz, and A. Kunz. The chaotic behavior of redirection–revisiting simulations in redirected walking. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 524–533. IEEE, 2022. 9

[12] E. Hodgson, E. Bachmann, and D. Waller. Steering immersed users of virtual environments: Assessing the potential for spatial interference. *ACM: Transactions on Applied Perception*, 8:1–22, 2011. 2

[13] Y. Hoshikawa, K. Fujita, K. Takashima, M. Fjeld, and Y. Kitamura. Redirecteddoors+: Door-opening redirection with dynamic haptics in room-scale vr. *IEEE Transactions on Visualization and Computer Graphics*, 30(5):2276–2286, 2024. 2

[14] G. K. Illahi, M. Siekkinen, T. Kämäräinen, and A. Ylä-Jääski. Real-time gaze prediction in virtual reality. In *Proceedings of the 14th international workshop on immersive mixed and virtual environment systems*, pp. 12–18, 2022. 2

[15] S.-B. Jeon, J. Jung, J. Park, and I.-K. Lee. F-rdw: Redirected walking with forecasting future position. *IEEE Transactions on Visualization and Computer Graphics*, 31(4):1970–1984, 2024. 2

[16] S.-B. Jeon, S.-U. Kwon, J.-Y. Hwang, Y.-H. Cho, H. Kim, J. Park, and I.-K. Lee. Dynamic optimal space partitioning for redirected walking in multi-user environment. *ACM Transactions on Graphics (TOG)*, 41(4):1–14, 2022. 3

[17] 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. 8

[18] D. Kim, S. Kim, J.-e. Shin, B. Yoon, J. Kim, J. Lee, and W. Woo. The effects of spatial configuration on relative translation gain thresholds in redirected walking. *Virtual Reality*, 27(2):1233–1250, 2023. 2

[19] H. Kim, S.-B. Jeon, and I.-K. Lee. Locomotion techniques for dynamic environments: Effects on spatial knowledge and user experiences. *IEEE transactions on visualization and computer graphics*, 30(5):2184–2194, 2024. 2

[20] E. Langbehn, F. Steinicke, M. Lappe, G. F. Welch, and G. Bruder. In the blink of an eye: leveraging blink-induced suppression for imperceptible position and orientation redirection in virtual reality. *ACM Trans. Graph.*, 37(4), July 2018. 2

[21] 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. 2

[22] D.-Y. Lee, Y.-H. Cho, D.-H. Min, and I.-K. Lee. Optimal planning for redirected walking based on reinforcement learning in multi-user environment with irregularly shaped physical space. In *2020 IEEE conference on virtual reality and 3D user interfaces (VR)*, pp. 155–163. IEEE, 2020. 3

[23] H. J. Lee, H. Kim, and I.-K. Lee. Multimodal turn in place: A comparative analysis of visual and auditory reset uis in redirected walking. *IEEE Transactions on Visualization and Computer Graphics*, 2025. 2

[24] J. Lee, S. Hwang, A. Ataya, and S. Kim. Effect of optical flow and user vr familiarity on curvature gain thresholds for redirected walking. *Virtual Reality*, 28(1):35, 2024. 2

[25] Y.-J. Li, F. Steinicke, and M. Wang. A comprehensive review of redirected walking techniques: Taxonomy, methods, and future directions. *Journal of Computer Science and Technology*, 37(3):561–583, 2022. 2

[26] Y.-J. Li, M. Wang, F. Steinicke, and Q. Zhao. Openrdw: A redirected walking library and benchmark with multi-user, learning-based functionalities and state-of-the-art algorithms. In *2021 IEEE International symposium on mixed and augmented reality (ISMAR)*, pp. 21–30. IEEE, 2021. 5, 6

[27] J.-H. Liu, Y.-F. Ren, Q. W. Gan, K. Huang, F. X. Y. Chen, E.-X. Luo, K. Y. Tang, Y.-Y. Fu, C.-W. Fan, S.-Z. Xu, et al. Overcoming spatial constraints in vr: A survey of redirected walking techniques. *Journal of Computer Science and Technology*, 39(4):841–870, 2024. 2

[28] J.-L. Lugrin, F. Kern, C. Kleinbeck, D. Roth, C. Daxer, T. Feigl, C. Mutschler, and M. E. Latoschik. A framework for location-based vr ap-plications. In *Virtuelle und Erweiterte Realitat: 16. Workshop der GI-Fachgruppe VR/AR (Berichte aus der Informatik)*, pp. 148–159. 2

[29] M. Lutfallah, M. Ketzel, and A. Kunz. Redirected walking for exploration of unknown environments. *Frontiers in Virtual Reality*, 4:1259816, 2024. 3

[30] J. Messinger, E. Hodgson, and E. R. Bachmann. Effects of tracking area shape and size on artificial potential field redirected walking. In *2019 IEEE conference on virtual reality and 3D user interfaces (VR)*, pp. 72–80. IEEE, 2019. 3

[31] 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. 2

[32] 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. 2

[33] M. Qi, Y. Liu, and J. Cui. A mapping-based redirected walking algorithm for large-scale vr. *Virtual Reality*, 27(3):2745–2756, 2023. 2

[34] S. Razzaque. *Redirected walking*. The University of North Carolina at Chapel Hill, 2005. 5

[35] S. Razzaque, D. Swapp, M. Slater, M. C. Whitton, and A. Steed. Redirected walking in place. In *Egve*, vol. 2, pp. 123–130, 2002. 2

[36] S. Ri, K. Matsumoto, T. Narumi, and H. Kuzuoka. Redirected walking method considering the interaction between users. *Frontiers in Virtual Reality*, 6:1304780, 2025. 3

[37] 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. 3

[38] Q. Sun, A. Patney, L.-Y. Wei, O. Shapira, J. Lu, P. Asente, S. Zhu, M. McGuire, D. Luebke, and A. Kaufman. Towards virtual reality infinite walking: dynamic saccadic redirection. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018. 2

[39] J. Thomas, C. Hutton Pospick, and E. Suma Rosenberg. Towards physically interactive virtual environments: Reactive alignment with redirected walking. In *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–10, 2020. 3

[40] J. Thomas and E. S. Rosenberg. A general reactive algorithm for redirected walking using artificial potential functions. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 56–62. IEEE, 2019. 3

[41] T. Van Gemert, N. C. Nilsson, T. Hirzle, and J. Bergström. Sicknificant steps: a systematic review and meta-analysis of vr sickness in walking-based locomotion for virtual reality. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pp. 1–36, 2024. 2

[42] X. Wang, Y. Li, and H.-N. Liang. Magicmap: Enhancing indoor navigation experience in vr museums. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 881–891. IEEE, 2024. 1

[43] N. L. Williams, A. Bera, and D. Manocha. Arc: Alignment-based redirection controller for redirected walking in complex environments. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2535–2544, 2021. 3

[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. 2, 3

[45] N. L. Williams, L. C. Stevens, A. Bera, and D. Manocha. Sensitivity to redirected walking considering gaze, posture, and luminance. *IEEE Transactions on Visualization and Computer Graphics*, 31(5):3223–3234, 2025. 2

[46] S.-Z. Xu, K. Huang, C.-W. Fan, and S.-H. Zhang. Saferdw: Keep vr users safe when jumping using redirected walking. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 365–375. IEEE, 2024. 2

[47] M. Zmuda, E. Bachmann, E. Hodgson, and J. Wonser. Improved resetting in virtual environments. In *IEEE Virtual Reality*, pp. 91–92, 2013. 2