Predictive multiuser redirected walking using artificial potential fields

Christian Hirt[1]*, Noah Isaak[1], Christian Holz[2] and Andreas Kunz[1]

[1]Innovation Center Virtual Reality, Department of Mechanical and Process Engineering, ETH Zurich, Zurich, Switzerland, [2]Sensing, Interaction & Perception Lab, Department of Computer Science, ETH Zurich, Zurich, Switzerland

Real walking is considered as the best locomotion metaphor to explore virtual environments in terms of user experience. In addition to being intuitive for the user, walking captures the true feelings of motion since the visual and proprioceptive sensations are harmonized well. The major disadvantage of choosing walking over other locomotion metaphors involves the physical constraints of the available space, which is usually considerably smaller than the virtual environment. To address this issue, redirected walking (RDW) introduces slight mismatches between a user's visually perceived path and their actual walking pattern, compelling them to subconsciously compensate for the inconsistency by adjusting their walking trajectory. As a result, users are steered to a certain degree, and expansive virtual environments are effectively compressed into smaller physical spaces. Among others, particularly predictive RDW offers immense potential for growth since it unites various algorithmic systems, whereas many approaches from literature depend on drastic restrictions like single-user constraints or architectural limitations to ensure real-time performance. This work presents two novel predictive RDW systems that allow multiple physically colocated users to explore independent and unconstrained virtual environments. The systems rely on two new implementations of prediction systems based on clothoid trajectory generation combined with a cost-based planning concept built on non-harmonic artificial potential fields (APFs), which inherently allow non-convex and dynamic physical environments. Using the APFs, three additional RDW conditions popular in the literature are implemented for comparison purposes. The five RDW concepts are then validated in an extensive user study with 150 participants conducted in 75 pairs. The results indicate that the novel predictive RDW systems outperform the three systems from literature, except for particular sections of the virtual environment with specific architectural traits.

KEYWORDS

virtual reality, human locomotion, redirected walking, redirection techniques, user study

1 Introduction

Choosing an ideal locomotion metaphor to navigate in virtual reality (VR) while ensuring the highest user presence possible has been a challenge for years. Multiple metaphors have been investigated as substitutes for walking, such as simple teleportation or flying using a handheld controller (Usoh et al., 1999; Bolte et al., 2011) to more elaborate systems like motion platforms (Bouguila et al., 2002; Bouguila and Sato,

FIGURE 1
Predictive multiuser redirected walking—utilizing the benefit of incorporating a user's trajectory prediction into the determination of effective and efficient redirections based on non-harmonic artificial potential fields in colocated user settings.

2002) or omnidirectional treadmills (Souman et al., 2008). However, none of these metaphors have managed to fully capture the naturalness and intuitiveness of real walking to date (Usoh et al., 1999). Furthermore, they often invoke simulator sickness (Kennedy et al., 1993) due to the inherent inconsistencies between the visually perceived motion and a user's vestibular sensations. More comprehensive overviews and relevant discussions on the advantages and disadvantages of the various locomotion metaphors have been reported in multiple reviews (Langbehn, 2019; Boletsis, 2017; Nilsson et al., 2018).

Although real walking is the most natural motion and is frequently preferred over other metaphors, it poses major challenges in VR applications. Allowing the users to explore virtual environments while really walking induces technical constraints since the users must never collide with real objects, other users, or boundaries of the available physical space. Most importantly, the physical space requirements do not adhere to the extensive virtual environments and *vice versa* and have therefore taken relevance over other challenges. A potential solution to this issue was proposed more than 20 years ago in the form of redirected walking (RDW) (Razzaque et al., 2001). Even though this was only a crude execution, RDW held strong promise and was rapidly developed into a full research field, enhancing the existing field of locomotion in VR. RDW exploits the fact that humans tend to rely more on their vision than on other senses and aims to deviate the virtual from the real pathway by inducing intentional, small mismatches between motion and vision (so-called RDW gains; Williams et al., 2006; Razzaque, 2005) to compel the users to subconsciously compensate while walking. As long as these induced mismatches remain adequately low, the users are mostly unaffected by simulator sickness, while their locomotion can be decoupled between the environments to a certain extent (Steinicke et al., 2008). Applying these so-called RDW gains in succession and complying with a steering paradigm essentially provides a decision basis on where to actively steer the users during their exploration. These steering paradigms can be distinguished into three base categories as reactive, predictive, and scripted RDW. These categories differ from each other by how they incorporate knowledge about a user's future path. Reactive RDW only evaluates a user's current state (i.e., position and heading) in a given environment and applies the gains accordingly. Scripted RDW follows a narrative approach and mostly prevents the users from deviating

from the intended path, thus fixing the trajectory a user can actually follow; in this approach, the RDW gains are calculated and optimized before actual exposure to users, ensuring the best redirection possible. Naturally, this paradigm fails if the users deviate considerably from the intended paths (e.g., by meandering or veering). Predictive RDW is arguably the least investigated category owing to its complexity and error-proneness originating from user spontaneity; however, it still offers immense potential for growth. Since predictive approaches mostly comprise two base algorithms (prediction and steering algorithms), each individual and the combinations thereof can be solved in several ways.

In this work, we present a novel approach to predictive RDW based on clothoidic trajectories (see Figure 1) that allows multiple users to be steered simultaneously in a common artificial potential field (APF). Specifically, we highlight the following contributions:

- two different technical implementation strategies of the prediction concept;
- a novel combination of prediction, steering, and non-harmonic APF-based costs that is fully functional in virtual open spaces with multiple colocated users;
- a user study with 150 participants highlighting the merits of the new algorithms.

The remainder of this manuscript is organized as follows. The related works are presented in section 2, which focus on RDW mainly in the predictive landscape. Then, subsection 3.1 introduces the APFs, subsection 3.2 discusses prediction and its implementation, and subsection 3.3 summarizes and combines the APFs with prediction and redirection implementation. Section 4 presents a user study, and section 5 provides an overview of the gathered data and its analysis. Finally, the findings are summarized in section 6, and the potential applications and consequences of RDW for future works are discussed in section 7.

2 Related works

In this section, we describe the work related to RDW, mainly focused on the predictive perspective.

2.1 Redirected walking

Over the years, it has been shown that the walking trajectories of users can be manipulated by applying RDW gains, such as linear, rotational, or curvature gains in the virtual environment (Razzaque et al., 2001; Razzaque, 2005; Nitzsche et al., 2004; Williams et al., 2006, 2007; Langbehn et al., 2017), which effectively compress the expansive virtual environment into a small physical space. Furthermore, when these RDW gains are only applied within certain thresholds, the users tend to not notice these redirections, while mostly avoiding simulator sickness (Steinicke et al., 2008); (Kennedy et al., 1993). However, a crucial example of applying redirection outside these detection thresholds are the so-called resets (Williams et al., 2007); resets are introduced as emergency strategies to prevent collisions shortly before impact and resemble external interventions, often breaking the sense of presence of a user. Such resets can be designed in multiple ways, and the most popular type is rotational reset, which demands that the user spin in place through an interface in the head-mounted display (HMD). Although this virtual reorientation often visually corresponds to a full 360° turn to ensure easier continuation of the experience, a rotational gain is applied to divert the real rotation to a different target orientation (i.e., toward an open space). Although this rotational reset has a rather simple implementation, it can be incorporated well in many redirection concepts. Furthermore, since resets tend to invoke at least partial breaks in presence, they are well-established as quality criteria for evaluating the performances of redirection strategies.

Arguably, the most important concern to date is the choice of when a technique must be optimally applied and at what intensity, as this strongly depends on the virtual architecture (Hodgson et al., 2014)). Generally, there are three core concepts in redirection techniques: reactive (Razzaque, 2005; Thomas and Rosenberg, 2019; Bachmann et al., 2019), predictive (Peck et al., 2012; Zmuda et al., 2013; Nescher et al., 2014,) and scripted (or semi-predictive) (Azmandian, 2018; Yu et al., 2018).

Since the focus of this work is mainly predictive approaches, we primarily discuss related examples. Peck et al. (2012) presented RFED as the first algorithm that resembled the predictive nature; they derived a skeleton graph consisting of nodes and segments that described the walkable paths within the available area in the virtual corridors to create a prediction to the next node based on the user's head orientation. Similarly, Zmuda et al. (2013) created a skeleton graph with probabilities assigned to the crossings, which were used to prioritize the final score of a predicted and redirected path; this final score was calculated by evaluating the distance a user could walk straight ahead at the end of a redirected prediction multiplied by the path probability. Finally, Nescher et al. (2014) introduced MPCRed as a method that recursively evaluates a cost function consisting of actions, such as applying redirection, while also following predefined bidirectional skeleton graphs. Recently, approaches utilizing reinforcement learning have been reported (Lee et al., 2019, 2020) that have drastically reduced real-time calculations; however, these algorithms demand sophisticated setups with neural networks that must be trained based on previously recorded user data, which makes them less applicable to more general use cases.

2.2 Artificial potential fields

APFs were first introduced to RDW in 2019 via two slightly different implementations (Thomas and Rosenberg, 2019; Bachmann et al., 2019) and follow the simple principle of assigning high potentials to potentially dangerous physical environments and low potentials to safe equivalents. Following the derivatives of such potential fields always guarantees flow from higher to lower potentials, which can be described by a force vector "pushing" and "pulling" the user. Accordingly, aligning the redirection with such force vectors ensures that the users are always steered away from dangerous regions while being pulled toward safety. These two first implementations follow a classic reactive approach and straightforwardly align a redirection vector with the repulsive force vectors generated by obstacles. They only differ in the manner in which the force vectors are generated but essentially apply the same principles. Later, Dong et al. (2020) enhanced this concept of repulsion by adding attractive forces to influence the redirection vector toward a real open space. APFs inherently offer the advantage of superimposing multiple layers of potentials, both static and dynamic, which enables incorporation of multiple users roaming the same physical space. All of these APF RDW applications have only considered the current states of the users and are thus categorized as reactive RDW approaches, while any predictive attempts were only suggested conceptually (Hirt et al., 2019a).

3 Methodology

This section presents different algorithms and describes how the elements of these algorithms are implemented. First, an overview on the APF is shown, followed by introduction of the predictive algorithm and its details, and finally the description of the complete process with all concatenated elements.

3.1 Artificial potential field

The physical environment in which a user is tracked while walking and exploring the virtual environment has been hard-coded in early RDW approaches and is always assumed to be static and convex. Naturally, such static tracking spaces provide stable and uncomplicated boundary conditions but restrict dynamic adjustment of the tracking space size while mutually excluding multiple users exploring virtual environments in colocated spaces. There have been some attempts to study how such physical tracking space boundaries could be made more dynamic and accessible (Yang et al., 2019; Cheng et al., 2019; Hirt et al., 2018), but an entirely explorative RDW (e.g., RDW in the wild) is still subject to research (Lutfallah, 2023). However, given fixed boundary conditions around a tracking space, APFs are an elegant option to describe the space dynamically and were initially presented in robotics (Khatib, 1986; Krogh, 1984); they are indispensable in many modern robotic path-planning approaches. APFs can incorporate non-convex tracking boundaries (Messinger et al., 2019) and allow multiple users in colocated tracking spaces

TABLE 1 Comparison of calculation times for different types of APFs.

APF	Time in A	Time in $4 \times A$
Non-harmonic	3.0 ms	11.0 ms
Harmonic	3.5 ms	12.0 ms
Wave expansion	16.5 ms	105.0 ms

(Thomas and Rosenberg, 2019; Bachmann et al., 2019) by modeling the respective second users as dynamic obstacles influencing each user's own potential field. Since there exist many options to model APFs, we briefly investigated three approaches, namely wavefront expansion (Barraquand et al., 1992), harmonic (Kim and Khosla, 1992), and non-harmonic (Thomas and Rosenberg, 2019; Bachmann et al., 2019) APFs, with some basic limiters added to avoid local minima. We compared the calculation time in a static environment encompassing 10^4 grid cells, with three obstacles inside (see Table 1). Although the harmonic and non-harmonic algorithms perform sufficiently similarly with slight advantages for the non-harmonic APF, the wavefront approach lags considerably by a factor of almost 6×. Quadrupling the environment size shows linear increases in the calculation times for the harmonic and non-harmonic approaches, with the wavefront approach again being more resource-intensive by a factor 10× compared to the others. Notably, none of the APFs were specifically optimized and were only implemented with their basic functionalities. From these qualitative comparisons, it was decided to pursue only non-harmonic APFs in the following implementations.

3.2 Prediction

Unlike the prediction algorithms employed in existing predictive RDW approaches, our system works in unrestricted virtual open spaces. Accordingly, we refrain from using the popular skeleton-graph-based predictions and implement two different algorithms derived from a recently suggested and evaluated concept (Hirt et al., 2019b, 2022). Accordingly, we employ clothoid trajectory shapes, which are bound by a lemniscate shape in the first case and by the inverse of their curvature in the second case. Both approaches were designed for a short-term prediction (Nescher and Kunz, 2012) to favor higher calculation frequency over long-term predictions.

3.2.1 Smoothing

To ensure more accurate and precise predictions, both algorithms relied on incorporating some historic data points recorded from the user's previously walked trajectory. These historic data points are stored in the HMD data buffer with fixed sizes, and new recordings are captured while removing the oldest data points. If the HMD data buffer remains unprocessed, then the data acquired through the HMD contains unstable and oscillating gait artifacts from the user's natural head veering (Hirasaki et al., 1999; Murray et al., 1964) and noise. Utilizing such perturbed data when extrapolating from the HMD data buffer can quickly and sustainably impair the results. Accordingly, a double-exponential smoothing approach intended for similar purposes (Nescher and Kunz, 2013) was extended by calibrating a damping factor

(Hyndman and Athanasopoulos, 2018; Taylor, 2003) to capture the trend of a given path without compromising the smoothing result.

3.2.2 Lemniscate Path Prediction (LPP)

In LPP, the prediction origin is always defined by a number of steps T backwards in time, which can naturally be varied on the basis of how much the HMD data buffer is integrated in the prediction. The prediction itself creates a bundle of clothoids contained within a lemniscate shape, which are then evaluated based on a path similarity measure to identify the best trajectory.

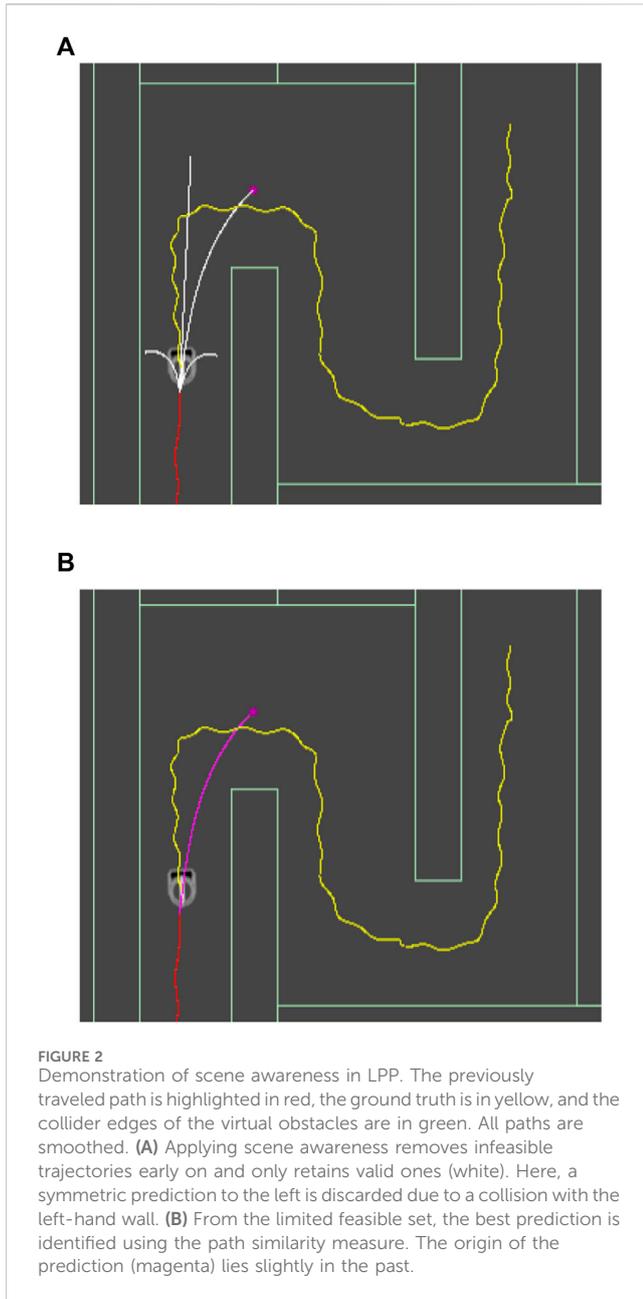The lemniscate is given by Eq. (1) and is created anew with each prediction iteration.

$$
\begin{aligned}
x(t) &= A_{Lem} \cdot \frac{sin(t)}{1 + cos^2(t)} \\
y(t) &= A_{Lem} \cdot \frac{sin(t)\ cos(t)}{1 + cos^2(t)},
\end{aligned}
\tag{1}
$$

where A_{Lem} denotes a constant defining the size of the lemniscate, and the running variable $t \in (0, \pi)$ incorporated only the positive half of the lemniscate. In theory, A_{Lem} can be used to directly relate the size of the lemniscate with the current walking speed of the user; instead, we designated a rigid horizon of 3 m to reduce the number of variables and hence improve comparability among the approaches. To create the trajectories within the lemniscate, a set of endpoints is identified along the contour of the shape and generated by uniform discretization, notably with an odd number of endpoints to ensure that a straight prediction is always available in the symmetry. This discretized lemniscate is now transformed to the chosen user state at $t - T$ and overlaid on the virtual environment, where a collision check is performed using Unity's stacked overlap boxes. Collision events can thus be found, and the endpoints colliding with the virtual obstacles are labeled as infeasible trajectories for consequent removal. To create the clothoid trajectories, an external C# library called *Curves*[1] was used and modified for seamless integration with Unity. *Curves* contains a convenient function that allows creation of 2D clothoids given a starting point (i.e., user position at $t - T$), direction (i.e., user heading at $t - T$), and final point (i.e., discretized endpoints). From this generated set of trajectories, a single best prediction is isolated using a simple mean-squared error enhanced by a discount factor. This ensures fast computation while providing stronger weights to more recent buffer states. The LPP with the functionalities of scene awareness and path isolation is shown in Figure 2.

3.2.3 Forward Path Prediction (FPP)

In FPP, a more analytical approach is chosen that inherently results in a single predicted trajectory. FPP revolves around solving the non-holonomic dynamic system of equations (Arechavaleta et al., 2006) presented in Eq. (2):

---

1 https://github.com/SEilers/Curves, accessed: 01.08.2022

FIGURE 2
Demonstration of scene awareness in LPP. The previously traveled path is highlighted in red, the ground truth is in yellow, and the collider edges of the virtual obstacles are in green. All paths are smoothed. **(A)** Applying scene awareness removes infeasible trajectories early on and only retains valid ones (white). Here, a symmetric prediction to the left is discarded due to a collision with the left-hand wall. **(B)** From the limited feasible set, the best prediction is identified using the path similarity measure. The origin of the prediction (magenta) lies slightly in the past.

The general parametric definition of a curvature κ given a curve $\gamma(t) = (x(t), y(t))$ is

$$\kappa = \frac{\dot{x}\ddot{y} + \ddot{x}\dot{y}}{\left(\dot{x}^2 + \dot{y}^2\right)^{\frac{3}{2}}} \tag{3}$$

and its derivative $\dot{\kappa} = \frac{d\kappa}{dt}$ is given by

$$\dot{\kappa} = \frac{2\left(\dot{x}^2 + \dot{y}^2\right)\left(x^{(3)}\dot{y} + 2\ddot{x}\ddot{y} + y^{(3)}\dot{x}\right) - 6\left(\ddot{x}\dot{y} + \dot{x}\ddot{y}\right)\left(\dot{x}\ddot{x} + \dot{y}\ddot{y}\right)}{2\left(\dot{x}^2 + \dot{y}^2\right)^{\frac{5}{2}}}, \tag{4}$$

where the primes in both equations refer to derivatives with respect to time; the prime$^{(3)}$ corresponds to the third derivative.

Finally, in accordance with Wilde (2009), the clothoid constant A_{Cl} is related to u_2 as follows:

$$A_{Cl} = \frac{1}{u_2} = \frac{1}{\dot{\kappa}}. \tag{5}$$

To solve Eq. (2) using Eqs. (3–5), the derivatives of x and y need to be computed. These first- to third-order derivatives can naturally be found by various methods, such as the finite difference approximation of derivatives (Wilmott et al., 1995; Olver, 2014). Following the basic method of central finite differences (CFD) (Fornberg, 1988), the derivatives of the differential equations are numerically approximated with the coefficients $C_{CFD,x}$ (with $x = d1$ for the first derivative, and analogs for second and third) given by

$$
\begin{aligned}
C_{CFD,d1} &= \left(\frac{1}{280} \quad -\frac{4}{105} \quad \frac{1}{5} \quad -\frac{4}{5} \quad 0 \quad \frac{4}{5} \quad -\frac{1}{5} \quad \frac{4}{105} \quad -\frac{1}{280} \right) \\
C_{CFD,d2} &= \left(-\frac{1}{560} \quad \frac{8}{315} \quad -\frac{1}{5} \quad \frac{8}{5} \quad -\frac{205}{72} \quad \frac{8}{5} \quad -\frac{1}{5} \quad \frac{8}{315} \quad -\frac{1}{560} \right) \\
C_{CFD,d3} &= \left(-\frac{7}{240} \quad \frac{3}{10} \quad -\frac{169}{120} \quad \frac{61}{30} \quad 0 \quad -\frac{61}{30} \quad \frac{169}{120} \quad -\frac{3}{10} \quad \frac{7}{240} \right).
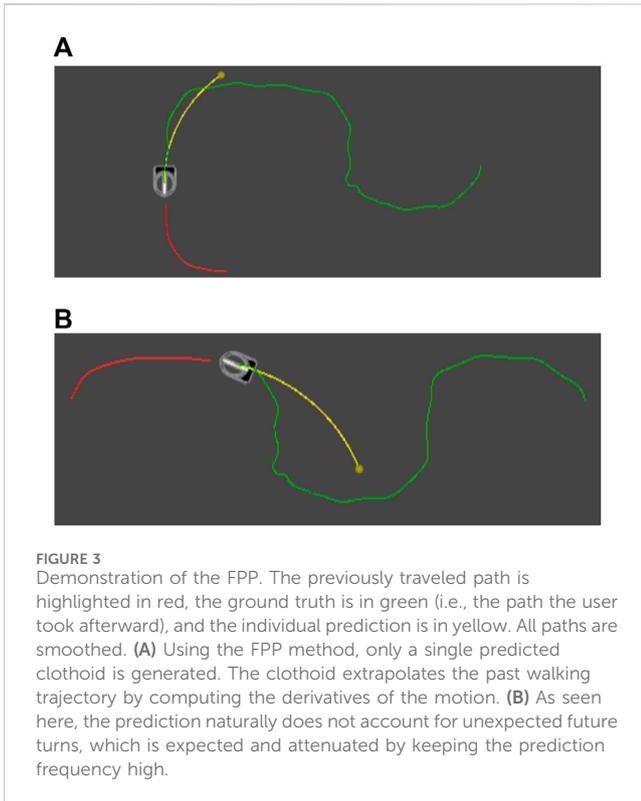\end{aligned} \tag{6}
$$

The CFD method allows computing the derivatives within the given time steps from t to $t - 8$, with the resulting derivatives corresponding to the time step $t - 4$ (c.f. "central"). Accordingly, the position vectors $\vec{x}$ and $\vec{y}$ are read from the HMD data buffer as

$$\vec{x} = \begin{pmatrix} x_{t-8} \\ x_{t-7} \\ \vdots \\ x_{t-1} \\ x_t \end{pmatrix} \quad \text{and} \quad \vec{y} = \begin{pmatrix} y_{t-8} \\ y_{t-7} \\ \vdots \\ y_{t-1} \\ y_t \end{pmatrix}. \tag{7}$$

Once the coefficients $C_{CFD,x}$ are determined and positional vectors $\vec{x}$ and $\vec{y}$ are available, the derivatives are computed as

$$
\begin{aligned}
\dot{x}_{t-4} &= C_{CFD,d1} \cdot \vec{x} & \dot{y}_{t-4} &= C_{CFD,d1} \cdot \vec{y} \\
\ddot{x}_{t-4} &= C_{CFD,d2} \cdot \vec{x} \quad \text{and} & \ddot{y}_{t-4} &= C_{CFD,d2} \cdot \vec{y} \\
x^{(3)}_{t-4} &= C_{CFD,d3} \cdot \vec{x} & y^{(3)}_{t-4} &= C_{CFD,d3} \cdot \vec{y}.
\end{aligned} \tag{8}
$$

These derivatives in Eq. (8) often manifest evident oscillations or other tracking noise originating from amplifications through the CFD method. Hence, the tracking inputs are smoothed more rigorously at risk of losing relevant data or the derivatives are alternatively treated again with the extended double-exponential smoothing filter. To avoid the risk of data loss, the latter option is preferred and applied, which eliminates the occurrence of high-frequency content even while introducing a slight but tolerable latency.

$$\begin{pmatrix} \dot{x}_T \\ \dot{y}_T \\ \dot{\varphi}_T \\ \dot{\kappa}_T \end{pmatrix} = \begin{pmatrix} \cos\ \varphi_T \\ \sin\ \varphi_T \\ \kappa_T \\ 0 \end{pmatrix} u_1 + \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} u_2, \tag{2}$$

where x, y, φ, κ denote the user's x and y positions as well as heading and curvature, respectively. The control input $u_1 = v_T$ is the user's linear velocity, and $u_2 = \dot{\kappa}_T$ is the curvature change rate. If the variables $\varphi, \kappa, u_1, u_2$ are known, a clothoid is fully determined and can be constructed with the corresponding parameters. Here, the user velocity and heading are directly acquired from the smoothed HMD data buffer, while the current curvature κ and control input u_2 are determined from the general definition of a curvature (Eq. (3)) and its derivative (Eq. (4)).

FIGURE 3
Demonstration of the FPP. The previously traveled path is highlighted in red, the ground truth is in green (i.e., the path the user took afterward), and the individual prediction is in yellow. All paths are smoothed. **(A)** Using the FPP method, only a single predicted clothoid is generated. The clothoid extrapolates the past walking trajectory by computing the derivatives of the motion. **(B)** As seen here, the prediction naturally does not account for unexpected future turns, which is expected and attenuated by keeping the prediction frequency high.

By substituting the smoothed derivatives from Eq. (8) into Eqs. (3, 4), the clothoid curvature κ and control input u_2 are fully determined. Furthermore, using Eq. (5), the clothoid constant A_{Cl} is computed. Utilizing *Curves*, the predicted clothoid is constructed through a different method that requires the starting position (x, y), heading φ, curvature κ, clothoid constant A_{Cl}, and clothoid length L. By definition, the clothoid length is inversely linear to the curvature; thus, the length L is simply clamped between a maximum length $L_{max} = 3$ m (for consistency with the LPP) and inverse of κ as

$$L = min\left(L_{max}, \frac{1}{\kappa}\right). \qquad (9)$$

The final FPP with its single prediction is demonstrated in Figure 3. Notably, the FPP approach neglects the virtual architecture in its prediction but is faster than the LPP despite the calculation of multiple derivatives as effectively fewer computational steps are required in the determination of a single prediction instead of many.

3.3 Predictive redirection

Conceptually, the predictive RDW entails a simple approach:

- it predicts a single path T_{pred};
- T_{pred} is redirected based on an action set U consisting of multiple redirection techniques and different gains, resulting in T_{red};
- a cost-based analysis of T_{red} is used to identify the best redirection $\pi_{optimal} \in U$;
- $\pi_{optimal}$ is applied to the user.

TABLE 2 Various redirection gains contained in the action set U.

Gain	Lower bound	Upper bound
Translation	0.86	1.26
Rotation	0.8	1.49
Curvature	−7.5	7.5
Null redirection	1	1

In this work, we present two predictive approaches, *PredRed LPP* and *PredRed FPP*, which differ in their predictive components; thus, T_{pred} originates from the respective LPP and FPP algorithms.

3.3.1 State update

The action set U of redirection techniques utilized in our approach consists of rotation gains, translation gains, curvature gains, and a combination of translation and curvature gains (Grechkin et al., 2016). Each gain $\pi \in U$ is only applied within its respective noticeability thresholds (Steinicke et al., 2009), as shown in Table 2. Naturally, resets are applicable above the respective rotational thresholds depending on the requirements determined for each reset. The specific process of injecting the particular redirection gains is included in the Redirected Walking Toolkit (Azmandian et al., 2016) and employed identically in our case. It is important to note that the size of the state space in which redirection is applied is manually limited by allowing only a single redirection action along the complete predicted path. Although the predicted path could potentially be divided into shorter segments, with each being subject to a separate evaluation, the computational demand grows exponentially with the size of the action set U, which would in turn linearly propagate with the number of segments that a path is divided into. Therefore, a finer segmentation of the path was rejected, and the state space was restrained in favor of faster iterations. The action set finally contains a null redirection gain and six different gain values per gain, i.e., the upper and lower noticeability threshold values (Steinicke et al., 2009) with a uniform distribution in between.

3.3.2 Redirection vector

In accordance with the performance comparison results of APFs in subsection 3.1, non-harmonic APFs are utilized (Bachmann et al., 2019; Thomas and Rosenberg, 2019), and a redirection vector $\vec{F}_{red}$ is defined. This redirection vector $\vec{F}_{red}$ is a total force vector and is the sum of all repulsive forces exhibited by the obstacles i in the physical space O, as shown in Eq. (10):

$$\vec{F}_{red} = \sum_{i \in O} \vec{F}_{rep,i}, \qquad (10)$$

where the magnitude $\|\vec{F}_{red}\|$ denotes the total force acting on a specific place in the physical environment and is considered as its APF value.

The contribution of a single repulsive obstacle i is formulated as in Eq. (11):

$$\vec{F}_{rep,i} = \begin{cases} a_o \cdot exp\left(-b_o \cdot d_o^2\right) \cdot \vec{e}_o, & \text{if } d_o \leq d_d, \\ 0, & \text{else} \end{cases}, \quad (11)$$

where a_o defines the maximum APF value generated by a single obstacle, and the constant b_o denotes the distribution width with respect to how quickly the APF value dissipates and approaches zero with increasing distance from the obstacle i; d_o describes the Euclidean distance between a specific obstacle's position p_i and the evaluated position p_{eval}; $d_o = \|p_i - p_{eval}\|$ and d_d is the threshold range within the repulsive force that is considered to still contribute to the total force. Here, $\vec{e}_o$ is the normalized direction vector from the evaluation position p_{eval} to the obstacle position p_i and is given as

$$\vec{e}_o = \frac{p_i - p_{eval}}{\|p_i - p_{eval}\|}. \quad (12)$$

3.3.3 Cost function

The cost function is a crucial element in determining the best redirection among the redirected trajectories T_{red}. Generally, minimizing the total cost J_{total} over all redirection actions π within the action set U results in the best redirection $\pi_{optimal}$, as shown in Eq. (13):

$$\pi_{optimal} = \underset{\forall \pi \in U}{\operatorname{argmin}} \ J_{total}(\pi) \quad (13)$$

The overall cost function J_{total} thus consists of the sum of all individual costs J_i over the complete redirected path T_{red}, which is evaluated for each data point from $t_1, \ldots t_N$, with N being the final point of the trajectory. To incorporate the increasing uncertainty along the prediction into the cost function, a discount factor $\alpha = 0.8$ is implemented similar to that in literature (Nescher et al., 2014), as shown in Eq. (14):

$$J_{total} = \sum_{i=1}^{N} \alpha^i \cdot J_i. \quad (14)$$

The cost for each point is designed by separately considering the APF, the user's current heading penalization owing to deviation toward the redirection vector, and the resets and redirection gains. The cost for the individual data point is thus given by Eq. (15):

$$J_i = J_{APF,i} + J_{Heading,i} + J_{Reset,i} + J_{Gain,i}, \quad (15)$$

where $J_{APF,i}$ denotes the APF value and is given by the magnitude of the redirection vector $\vec{F}_{red}$:

$$J_{APF,i} = \|\vec{F}_{red}\|. \quad (16)$$

Accordingly, the magnitude of the redirection vector directly translates the user's proximity to the obstacles and thereby penalizes the positions in danger of collisions. Furthermore, $J_{Heading,i}$ is given by the dot product between $\vec{F}_{red}$ and current heading $\vec{\theta}_i$:

$$J_{Heading,i} = h_0 \cdot \frac{1}{2} \cdot \left(1 - \frac{\vec{F}_{red}}{\|\vec{F}_{red}\|} \cdot \frac{\vec{\theta}_i}{\|\vec{\theta}_i\|}\right). \quad (17)$$

With this definition, the heading cost is limited between 0 and 1 as well as multiplied by h_0, which is considered as the design parameter for heading penalization. This later offers the option to vary the weighting balance between the heading and APFs, with low

heading costs to steer users toward open spaces rather than favoring positions closer, yet parallel to the boundaries with low APF costs. For the current implementation, h_0 was chosen to be 1.

The resets are penalized through a binary decision of whether the evaluated data point lies within the tracking room boundaries or not, as shown in Eq. (18). Through this approach, collisions that unavoidably invoke a reset upon first impact are delayed as the collisions accumulate costs linearly with the amount of path lying outside the walkable space. Therefore, a late reset induces a lower cost in contrast to an early reset.

$$J_{Reset,i} = \begin{cases} 1000, & \text{if outside tracking room boundaries} \\ 0, & \text{otherwise} \end{cases}. \quad (18)$$

Finally, the cost results from the different gains are formulated in $J_{Gain,i}$. Accordingly, $J_{Gain,i}$ and $J_{Heading,i}$ were set to 0 for simplicity, but it would be reasonable to provide containers for these in the initial formulation of the PredRed implementation for later investigations. We aim to include psychometric functions (Neth et al., 2012; Steinicke et al., 2008) or similar means to account for the noticeability in deciding the redirection strength.

With all components elaborated, the two new predictive RDW algorithms PredRed LPP and PredRed FPP are complete. Along with these, three more conditions were implemented, namely NULL as well as the two reactive variations of the APF RDW steer-to-center (S2C) (Razzaque, 2005) and steer-to-gradient (S2G) (Thomas and Rosenberg, 2019; Bachmann et al., 2019). Although S2C has been established over the years as a simple yet powerful algorithm that is popular for its adequate performance comparisons in literature, S2G appears to be an obvious choice owing to its APF-based algorithmic proximity and inherently powerful premise. In both reactive approaches, the algorithms compare the current real user state to the redirection vector $\vec{F}_{red}$ and apply the redirection that best closes the deviation between the two. In the case of S2C, $\vec{F}_{red}$ is implemented by placing the coordinate system's origin at the center of the tracking space and inverting the user's position vector, which inherently forces $\vec{F}_{red}$ to always point toward the center of the space. In the case of S2G, $\vec{F}_{red}$ is simply calculated as shown in Eq. (10). NULL naturally refrains from applying any redirection while walking and only utilizes resets to prevent collisions. For all conditions, the reset algorithm was maintained consistently as reset-to-gradient (R2G) (Thomas and Rosenberg, 2019; Bachmann et al., 2019), which recalibrates the user by aligning their heading with $\vec{F}_{red}$ via the reflex angle of the initial deviation. In our case, the reflex angle is specifically chosen over the acute and obtuse angles since it allows weaker rotational gains, consequently placing less strain on the user perceptions. Notably, applying R2G to a NULL redirector seems inappropriate at first since it also incorporates the potential influence of a second user, but we intended to maintain the conditions in addition to the redirectors consistent rather than implementing a simplistic alternative like the original 2:1-turn reset (Williams et al., 2006).

4 User study

To validate the newly presented PredRed redirectors, a user study was conducted to acquire relevant data. The participants

used a wireless HTC Vive Pro driven by two desktop towers, each with an i9-10900K and 64 GB of RAM as well as an NVIDIA GeForce RTX 3090, and handheld controllers were not provided. This setup ensured a steady frame rate of 90 fps and full resolution for both eyes. The system implemented SteamVR but the safety chaperones were deactivated to increase the presence of the users, with the only safety measure being a reset prompt that asked users to stop and perform a reset (see Figure 4A). Naturally, the reset direction is governed by the redirection vector and can either indicate clockwise or counterclockwise reorientation. The physically available tracking space was 5.5 m × 8.0 m, with a 0.2 m reset margin.

4.1 Study design

The participants were invited in pairs and explored separate yet identical virtual environments colocated in the same physical tracking space. The separation of the virtual environments was particularly chosen to avoid any bias from the users interacting with each other in VR (e.g., avoiding virtual collisions). Each pair of participants experienced only a single RDW condition to avoid learning or accommodation effects. All participants followed the same study protocols and filled a pre-questionnaire, were instructed before being exposed to the VR, and answered a post-questionnaire. The pre-questionnaire consisted of general demographic information, a self-assessment of their susceptibility to motion sickness, their experiences with VR systems, and a simulator sickness (SS) questionnaire (Kennedy et al., 1993). In the instructional part, the participants were introduced to the VR system, their task, and their expected behavior of collecting stars (see Figure 4B). They were further instructed to be especially responsive to the reset prompts, take short steps (since only their head was tracked), and avoid passing through virtual walls. The participants were aware that they shared a common tracking space, that they cannot see each other in the virtual environment, and that only their physical positions were exchanged between the systems to prevent collisions. They were not informed about the purpose of the study regarding RDW, so that the advance questions were unanswered. After receiving instructions, pairs of participants started their virtual experiences in opposite corners of the tracking space facing the center and aligned with the initial virtual corridor. Following the virtual task, they filled the post-questionnaire consisting of another round of the SS questionnaire and the standardized questionnaires on the simulation task load index (SIM TLX) (Harris et al., 2020) and cognitive absorption (CA) (Agarwal and Karahanna, 2000).
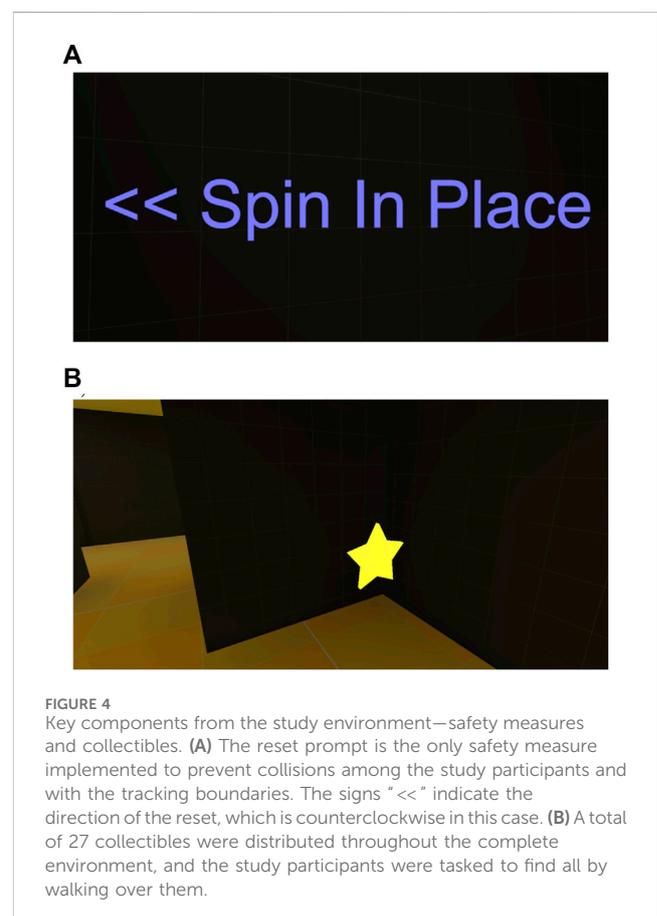
4.2 Study environment

The study environment consisted of seven rooms, each with distinctly different architectures varying from corridors and maze-like environments to local obstacles like pillars and mostly or even fully open spaces, as shown in Figure 5 (in order: Hallway, Yellow01, Yellow02, Red01, Red02, Blue01,

and OpenSpace Forest). Although this delivered some variety to the participants to make the exploration more interesting, the rooms were also evaluated separately to highlight the performances in each of the architectural layouts. In addition to the aesthetics, the floor colors helped with orientation and provided some feedback regarding study progression to the investigators. The participants were asked to gather collectibles distributed throughout the environment by simply walking over them. Overall, the 27 stars that were to be collected were carefully distributed in the five rooms connected to the starting point (i.e., Yellow01, Yellow02, Red01, Red02, and Blue01), ensuring that the users explores the full environment without missing any corners or corridors. All the stars in a room were visible simultaneously, while the various rooms were connected by sliding doors that opened automatically after all the stars in a room were collected. The participants were always informed through an interface in the HMD regarding the number of stars they had already collected and that remaining to be discovered in a given room. The final area of the study was populated with trees and stones, suggesting independent exploration without a clear goal.

5 Results and discussion

In this section, we present and discuss the analytical results of the data acquired from the user study entailing the study population

FIGURE 4
Key components from the study environment—safety measures and collectibles. **(A)** The reset prompt is the only safety measure implemented to prevent collisions among the study participants and with the tracking boundaries. The signs "<<" indicate the direction of the reset, which is counterclockwise in this case. **(B)** A total of 27 collectibles were distributed throughout the complete environment, and the study participants were tasked to find all by walking over them.
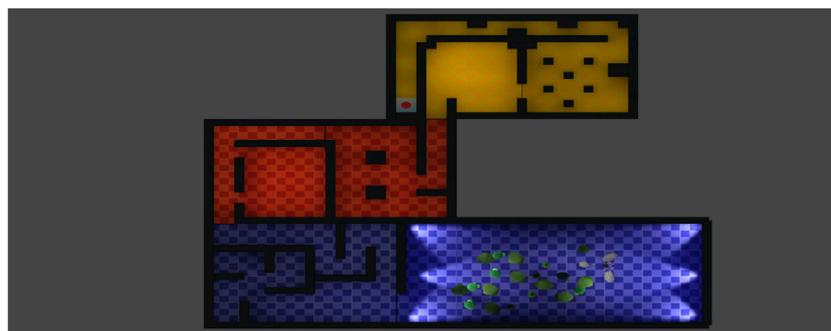
FIGURE 5
Study environment, in which a participant starts at the red dot in the yellow area and subsequently makes their way through the various rooms (Yellow01, Yellow02, Red01, Red02, Blue01, and blue OpenSpace Forest). Each room varies in obstacle density, while the rooms are separated by sliding doors that automatically open after all the stars in the present room are collected.

followed by the subjective and objective metrics. The video (Supplementary Video S1) displays a simulated top-down view of the tracking space using the recorded positions of two simultaneous study participants.

5.1 Study population

In total, 150 subjects participated in the user study (25.5 ± 5.5 years, 45 females). Since the participants enrolled in pairs, each condition was encountered 15 times, with an overall number of 30 recordings per condition. Among the participants, 31 wore contacts, 37 used glasses, and none of them experienced any remarkable issues with the VR system. The participants were recruited through university channels, which was reflected in the 138/150 participants who were currently studying. The complete study procedure roughly required an hour, for which the participants were compensated fairly. Even though the participants' experiences with VR varied considerably (15 without VR experience, 103 below 20 h, 35 above), apart from the distinctly different navigation speeds, all participants collected all 27 stars each. Furthermore, although some of the participants considered themselves to be exceptionally susceptible to motion sickness, none of the study procedures were interrupted or aborted because of sickness.

5.2 Subjective metrics

The SS, SIM TLX, and CA questionnaires were filled out by the participants during the different study phases, and the results are clustered by redirection condition and summarized in Table 3. None of the participants aborted the study, which is confirmed by the fact that none of the collected delta values ΔSS between the pre- and post-exposure conditions indicate significant occurrence of SS. Furthermore, the SIM TLX scores (0–100) and CA (0–30) were all clustered around similar values without apparent advantages for any of the conditions, indicating that the new PredRed LPP and FPP methods can be integrated well with existing and proven algorithms. With the exception of presence and CA, lower values

are deemed better for all conditions. Based on more qualitative feedbacks, the study participants notably complained that the NULL condition required them to spin a lot, which was hardly unexpected considering that some of the participants registered 100 or even more resets during their exposure time. On the other hand, the study participants at the lower end of reset counts, especially in PredRed LPP conditions, were often surprised with the size of the explored environment compared to the tracking space since they lost track of their real surroundings quickly. Furthermore, some participants realized and noted that they could feel the redirection, especially in the S2C and S2G conditions, but not in an overly disturbing sense.

5.3 Objective metrics

To assess the performances of different RDW concepts and steering algorithms, a popular measure for comparison is the mean distance between resets d_{Reset}. Unlike the total number of resets, which was commonly used in early performance evaluations, d_{Reset} normalizes the number of resets over the total walking distance for each user and thereby better accounts for varying walking speeds, including stand still or other personal human walking traits like veering. On average, each participant covered 381 m during their exploration, but d_{Reset} was calculated for each participant separately based on their individually covered distance. Figure 6 shows the overall distribution of d_{Reset} against the different RDW conditions, and Table 4 presents the respective numerical mean values. Notably, in rare occasions, the participants who reset very quickly would overspin their initially indicated reset, instantaneously provoking a second successive reset. In such cases, these back-to-back resets were summarized into a single reset and applied to all conditions consistently.

Evidently, the newly introduced predictive RDW algorithm PredRed LPP managed to outperform the remaining algorithms overall. This is further emphasized by the one-tailed t-test on the normally distributed data, evaluating d_{Reset} in the presented order (PredRed LPP–PredRed FPP–S2G–S2C–NULL). The numerical results of the p-values are shown in Table 5, and a significance of $\alpha = 0.05$ was chosen from RDW literature. With this, we confirm that PredRed LPP showed statistically significant best performance ($p < \alpha$),

TABLE 3 Results of the subjective questionnaires clustered by RDW condition: simulator sickness (SS), SIM task load index (0−100), and cognitive absorption (0−30). Aside from presence and CA, lower values are considered to be better.

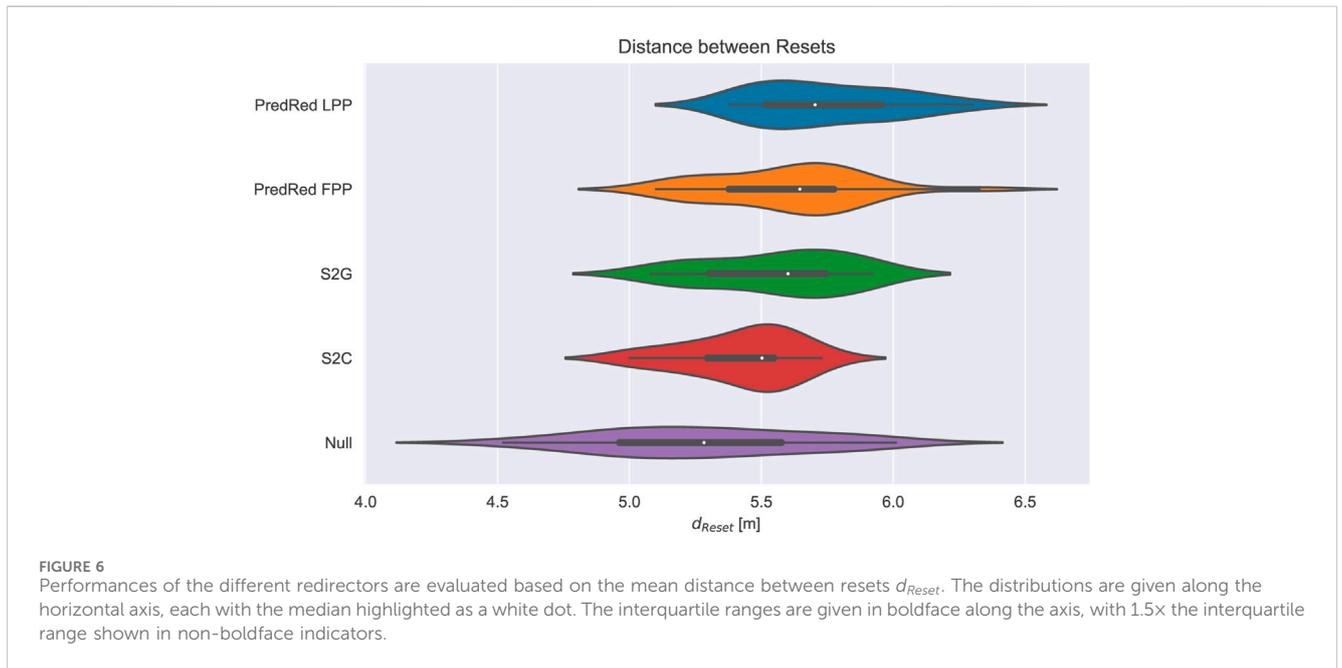| RDW Condition	Δ SS	Mental Demand	Physical Demand	Task Compl.	Stress	Distr.	Navig.	Presence	CA
PredRed LPP	15.7	25.3	13.3	9.3	12.0	18.0	21.7	72.0	22.4
PredRed FPP	9.2	25.0	13.7	8.0	13.0	17.3	18.0	71.3	23.6
S2G	12.3	22.3	16.3	6.0	8.3	15.7	21.0	76.3	23.4
S2C	8.6	32.3	16.3	7.0	12.0	18.3	28.0	70.7	22.2
NULL	9.6	24.0	20.0	8.7	10.0	20.7	21.7	77.7	22.2

FIGURE 6
Performances of the different redirectors are evaluated based on the mean distance between resets d_{Reset}. The distributions are given along the horizontal axis, each with the median highlighted as a white dot. The interquartile ranges are given in boldface along the axis, with 1.5× the interquartile range shown in non-boldface indicators.

TABLE 4 Results of the objective measure for each RDW condition over all rooms: mean distance between resets d_{Reset} including standard deviation.

RDW Condition	d_{Reset}
PredRed LPP	5.75 ± 0.26 m
PredRed FPP	5.60 ± 0.27 m
S2G	5.56 ± 0.26 m
S2C	5.42 ± 0.20 m
NULL	5.28 ± 0.39 m

TABLE 5 Results of a one-tailed t-test investigating the statistical significances of the respective algorithms based on d_{Reset}. A significance level of $\alpha = 0.05$ is considered. The table is read from left to right, e.g., PredRed LPP vs PredRed FPP results in $p = 0.031$. Green indicates $p < \alpha$, and yellow is $p > \alpha$.

p-value	PredRed FPP	S2G	S2C	NULL
PredRed LPP	0.031	0.012	6.4E-05	2.1E-06
PredRed FPP		0.312	0.013	4.7E-04
S2G			0.047	0.002
S2C				0.056

while PredRed FPP and S2G performed similarly ($p > \alpha$) even as they significantly outperformed S2C and NULL redirection ($p < \alpha$ in both cases). Noticeably, S2C showed no statistical superiority over NULL redirection. Given that only resets were employed in NULL redirection, this may indicate that continuous application of redirection is less powerful or less effective than usually agreed upon in literature, and we speculate that the resets may possibly have a much larger significance than commonly expected.

Furthermore, d_{Reset} was evaluated by considering the different architectural spaces separately (see Figure 7). On the left, the overall values are shown as references, and the rooms are ordered from left to right in consecutive order of exploration during the user study. Notably, especially in the initial hallway, PredRed LPP's strength of integrating scene awareness was apparent. Even though the performance was not the best for each room, PredRed LPP often showed considerable advantage over the other algorithms; although it suggested an apparent weakness in the final open space, the

FIGURE 7
Individual performances in the form of d_{Reset} for the different RDW conditions broken down per room. On the left, d_{Reset} over all rooms is given as a reference. From left to right, the different rooms are in the order presented to the study participants. The colors indicate the performance distribution and are highlighted from yellow (best) to blue (worst).

algorithm performed well for the other more open spaces in Yellow02 and Red02. PredRed FPP, on the other hand, while having only a slightly higher d_{Reset} over S2G, appeared to struggle most with bent trajectories in guided environments (e.g., in Blue01 and Yellow01). Interestingly, S2C showed a particularly high d_{Reset} in the straight hallway at the beginning but deteriorated drastically in the two yellow rooms.

6 Conclusion

In this paper, we introduced two new RDW concepts addressing challenges encountered in predictive RDW. These mainly consist of implementing simplifications like virtual constraints (e.g., mazes) to limit the choices of the predictive components, thus lowering the computational demand and ensuring real-time application, as well as applying prediction in colocated multiuser environments. We present the detailed implementations of the core elements and discuss how these are connected to the overall concepts forming the redirectors PredRed LPP and PredRed FPP. Both approaches combined cost-based evaluations correlating with a non-harmonic APF representation of the physical tracking space. The main difference between PredRed LPP and FPP was in the technical implementations of their respective path prediction algorithms. While LPP relied on multiple clothoid trajectories contained in a lemniscate shape, the application of a simple scene awareness to sort out infeasible trajectories, and a path similarity measure, FPP numerically approximated a single clothoid trajectory forecasting the user path based on their recent explorative behaviors.

Aside from the technical procedures, we designed and conducted a user study involving 150 participants. In addition to the two new concepts, we recreated three known RDW concepts from literature using the presented APFs: NULL, S2C (Razzaque, 2005), and S2G (Thomas and Rosenberg, 2019; Bachmann et al., 2019). We matched all five conditions with a consistent R2G reset strategy. For these five distinct conditions, we invited pairs of participants to obtain 15 pairs, with 30 individual recordings per condition. None of the participants aborted the user study due to sickness, and all participants completed the given search and collect task fully. Through recordings of the explorations, we showed that PredRed LPP outperformed the other conditions statistically significantly with regard to the mean distance between resets. Although PredRed FPP showed superior results over S2G, S2C, and NULL, it only showed statistical significant results over S2C and NULL. Furthermore, both PredRed approaches could be implemented well in locations with proportionally high amounts of open spaces without apparent spikes in their computational demands, other technical issues (e.g., dropping frame rate), or obvious deteriorations in the RDW performances.

7 Outlook

Both PredRed LPP and FPP algorithms introduce a new take on predictive RDW, enabling investigation of further opportunities. For example, in the current implementation, the second user was modeled as a simple, dynamic, and local obstacle whose influence on the redirection vector was equally weighted

to the boundaries. A simple yet potentially powerful consideration would entail a priority system, which in its simplest implementation would contain a rigid order, in which some of the users receive precedence over others. Alternatively, since each user has their own velocity and prediction, incorporating a stronger weight along the expected motion axis or even full consideration of the predicted trajectory in the APF propagation would potentially improve the overall redirection strategy. This implies that the APF can be locally propagated to better reflect the effective future states of the environment in the planning iteration and sharing the propagated APFs with each other. Furthermore, as explained in subsection 3.3, instead of firmly locking the state space in favor of computational demand, segmenting the predicted path and allowing the full action set to be considered for each individual segment along the prediction would create a more elaborate redirection. Naturally, this requires careful integration into the cost function as changing the applied redirection gains in rapid succession can potentially become more noticeable and break a user's presence quickly. On the prediction side, in the case of FPP, an additional integration of a system allowing scene awareness could create benefits with respect to accuracy and precision of prediction. This clearly imposes an algorithmic overhead and may most likely impair the computational demand slightly; however, since FPP is already less resource-intense, this should not pose a major issue. The complete procedure also entailed elaboration of many placeholder containers, such as A_{Lem} in Eq. (1) to incorporate walking speed in the lemniscate prediction or heading cost $J_{Heading,i}$ from Eq. (17) to address user alignment with the boundaries. Each of these placeholder containers, if calibrated properly, may again improve the overall performances of the new predictive systems, but further isolated investigations are necessary. As such, during calibrations of the placeholder containers, an ablation study could enrich understanding of the impacts of individual cost terms on the global cost function J_{total}. Regarding the resets, actively integrating the so-called predictive resets may potentially improve the redirection strategy considerably by preventing users from navigating into difficult situations, such as a corner of the tracking space. Otherwise, predictive resets could also take a user priority system into consideration and could be employed to actively balance the number of resets between the users or simply be integrated in a segmented evaluation connected to an unrestricted state space.

The algorithms presented herein have only been tested in virtually separated environments. Accordingly, including two users in the same virtual environment could potentially increase the users' presence further, but we currently believe this to occur regardless of the applied RDW condition and therefore not affect the objective RDW performance measures, such as the mean distance between resets. Another interesting topic could be the incorporation of certain aspects of RDW alignment (Williams et al., 2021; Thomas and Rosenberg, 2020), which would eventually allow the users to physically interact with the real environment or even with each other.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Ethics statement

Ethical approval was not required for the studies involving humans because there was no personal information collected that could potentially be followed back to the individual participants. The studies were conducted in accordance with the local legislation and institutional requirements. All participants provided written informed consent to participate in this study.

Author contributions

CHi: conceptualization, data curation, formal analysis, funding acquisition, investigation, methodology, project administration, resources, software, supervision, validation, visualization, writing–original draft, and writing–review and editing. NI: conceptualization, data curation, formal analysis, methodology, software, validation, visualization, and writing–review and editing. CHo: project administration, resources, supervision, and writing–review and editing. AK: funding acquisition, project administration, resources, supervision, and writing–review and editing.

Funding

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors, and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/frvir.2024.1259429/full#supplementary-material

References

Agarwal, R., and Karahanna, E. (2000). Time flies when you're having fun: cognitive absorption and beliefs about information technology usage. *MIS Q.* 24, 665–694. doi:10.2307/3250951

Arechavaleta, G., Laumond, J.-P., Hicheur, H., and Berthoz, A. (2006). "The nonholonomic nature of human locomotion: a modeling study," in *The first IEEE/RAS-EMBS international conference on biomedical robotics and biomechatronics, 2006* (IEEE), 158–163.

Azmandian, M. (2018). *Design and evaluation of adaptive redirected walking systems*. Ph.D. thesis.

Azmandian, M., Grechkin, T., Bolas, M., and Suma, E. (2016). "The redirected walking toolkit: a unified development platform for exploring large virtual environments," in *Everyday virtual reality (WEVR), 2016 IEEE 2nd workshop on* (IEEE), 9–14. doi:10.1109/wevr.2016.7859537

Bachmann, E. R., Hodgson, E., Hoffbauer, C., and Messinger, J. (2019). Multi-user redirected walking and resetting using artificial potential fields. *IEEE Trans. Vis. Comput. Graph.* 25, 2022–2031. doi:10.1109/tvcg.2019.28987642898764

Barraquand, J., Langlois, B., and Latombe, J.-C. (1992). Numerical potential field techniques for robot path planning. *IEEE Trans. Syst. man, Cybern.* 22, 224–241. doi:10.1109/21.148426

Boletsis, C. (2017). The new era of virtual reality locomotion: a systematic literature review of techniques and a proposed typology. *Multimodal Technol. Interact.* 1, 24. doi:10.3390/mti1040024

Bolte, B., Steinicke, F., and Bruder, G. (2011). The jumper metaphor: an effective navigation technique for immersive display setups. *Proc. Virtual Real. Int. Conf.* 1, 2–1.

Bouguila, L., and Sato, M. (2002). "Virtual locomotion system for large-scale virtual environment," in *Proceedings IEEE virtual reality 2002* (IEEE), 291–292.

Bouguila, L., Sato, M., Hasegawa, S., Naoki, H., Matsumoto, N., Toyama, A., et al. (2002). "A new step-in-place locomotion interface for virtual environment with large display system," in *ACM SIGGRAPH 2002 conference abstracts and applications*, 63. doi:10.1145/1242073.1242098

Cheng, L.-P., Ofek, E., Holz, C., and Wilson, A. D. (2019). "Vroamer: generating on-the-fly vr experiences while walking inside large, unknown real-world building environments," in *2019 IEEE conference on virtual reality and 3D user interfaces (VR)* (IEEE), 359–366. doi:10.1109/vr.2019.8798074

Dong, T., Chen, X., Song, Y., Ying, W., and Fan, J. (2020). "Dynamic artificial potential fields for multi-user redirected walking," in *2020 IEEE conference on virtual reality and 3D user interfaces (VR)* (IEEE), 146–154. doi:10.1109/vr46266.2020.00033

Fornberg, B. (1988). Generation of finite difference formulas on arbitrarily spaced grids. *Math. Comput.* 51, 699–706. doi:10.1090/s0025-5718-1988-0935077-0

Grechkin, T., Thomas, J., Azmandian, M., Bolas, M., and Suma, E. (2016). "Revisiting detection thresholds for redirected walking: combining translation and curvature gains," in *Proceedings of the ACM symposium on applied perception* (ACM, NY, USA: ACM), 113–120. doi:10.1145/2931002.2931018

Harris, D., Wilson, M., and Vine, S. (2020). Development and validation of a simulation workload measure: the simulation task load index (sim-tlx). *Virtual Real.* 24, 557–566. doi:10.1007/s10055-019-00422-9

Hirasaki, E., Moore, S. T., Raphan, T., and Cohen, B. (1999). Effects of walking velocity on vertical head and body movements during locomotion. *Exp. brain Res.* 127, 117–130. doi:10.1007/s002210050781

Hirt, C., Ketzel, M., Graf, P., Holz, C., and Kunz, A. (2022). "Heuristic short-term path prediction for spontaneous human locomotion in virtual open spaces," in *2022 IEEE conference on virtual reality and 3D user interfaces abstracts and workshops (VRW)* (IEEE), 636–637.

Hirt, C., Zank, M., and Kunz, A. (2018). "Geometry extraction for *ad hoc* redirected walking using a slam device," in *International conference on augmented reality, virtual reality and computer graphics* (Springer), 35–53. doi:10.1007/978-3-319-95270-3_3

Hirt, C., Zank, M., and Kunz, A. (2019a). "Prewap: predictive redirected walking using artificial potential fields," in *2019 IEEE conference on virtual reality and 3D user interfaces (VR)* (IEEE), 976–977. doi:10.1109/vr.2019.8798118

Hirt, C., Zank, M., and Kunz, A. (2019b). "Short-term path prediction for virtual open spaces," in *2019 IEEE conference on virtual reality and 3D user interfaces (VR)* (IEEE), 978–979. doi:10.1109/vr.2019.8797709

Hodgson, E., Bachmann, E., and Thrash, T. (2014). Performance of redirected walking algorithms in a constrained virtual world. *IEEE Trans. Vis. Comput. Graph.* 20, 579–587. doi:10.1109/tvcg.2014.34

Hyndman, R. J., and Athanasopoulos, G. (2018). *Forecasting: principles and practice*. Melbourne, Australia: OTexts.

Kennedy, R. S., Lane, N. E., Berbaum, K. S., and Lilienthal, M. G. (1993). Simulator sickness questionnaire: an enhanced method for quantifying simulator sickness. *Int. J. Aviat. Psychol.* 3, 203–220. doi:10.1207/s15327108ijap0303_3s15327108ijap0303_3

Khatib, O. (1986). "The potential field approach and operational space formulation in robot control," in *Adaptive and learning systems* (Springer), 367–377. doi:10.1007/978-1-4757-1895-9_26

Kim, J.-O., and Khosla, P. (1992). Real-time obstacle avoidance using harmonic potential functions. *IEEE Transctions Robotics Automation* 8, 338–349. doi:10.1109/70.143352

Krogh, B. (1984). "A generalized potential field approach to obstacle avoidance control," in *Proc. SME conf. On robotics research: the next five years and beyond* (Bethlehem, PA), 11–22.

Langbehn, E. (2019). *Walking in virtual reality: perceptually-inspired interaction techniques for locomotion in immersive environments*. Ph.D. thesis.

Langbehn, E., Lubos, P., Bruder, G., and Steinicke, F. (2017). Bending the curve: sensitivity to bending of curved paths and application in room-scale vr. *IEEE Trans. Vis. Comput. Graph.* 23, 1389–1398. doi:10.1109/tvcg.2017.2657220

Lee, D.-Y., Cho, Y.-H., and Lee, I.-K. (2019). "Real-time optimal planning for redirected walking using deep q-learning," in *2019 IEEE conference on virtual reality and 3D user interfaces (VR)* (IEEE), 63–71. doi:10.1109/vr.2019.8798121

Lee, D.-Y., Cho, Y.-H., Min, D.-H., and Lee, I.-K. (2020). "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)* (IEEE), 155–163. doi:10.1109/vr46266.2020.00034

Lutfallah, M. (2023). "[dc] maximizing natural walking in virtual environments," in *2023 IEEE conference on virtual reality and 3D user interfaces abstracts and workshops (VRW)* (IEEE), 989–990.

Messinger, J., Hodgson, E., and Bachmann, E. R. (2019). "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)* (IEEE), 72–80. doi:10.1109/vr.2019.8797818

Murray, M. P., Drought, A. B., and Kory, R. C. (1964). Walking patterns of normal men. *JBJS* 46, 335–360. doi:10.2106/00004623-196446020-00009

Nescher, T., Huang, Y.-Y., and Kunz, A. (2014). "Planning redirection techniques for optimal free walking experience using model predictive control," in *3D user interfaces (3DUI), 2014 IEEE symposium on* (IEEE), 111–118. doi:10.1109/3dui.2014.6798851

Nescher, T., and Kunz, A. (2012). "Analysis of short term path prediction of human locomotion for augmented and virtual reality applications," in *2012 international conference on cyberworlds* (IEEE), 15–22. doi:10.1109/cw.2012.10

Nescher, T., and Kunz, A. (2013). "Using head tracking data for robust short term path prediction of human locomotion," in *Transactions on computational science XVIII* (Springer), 172–191. doi:10.1007/978-3-642-38803-3_10

Neth, C. T., Souman, J. L., Engel, D., Kloos, U., Bulthoff, H. H., and Mohler, B. J. (2012). Velocity-dependent dynamic curvature gain for redirected walking. *IEEE Trans. Vis. Comput. Graph.* 18, 1041–1052. doi:10.1109/tvcg.2011.275

Nilsson, N. C., Serafin, S., Steinicke, F., and Nordahl, R. (2018). Natural walking in virtual reality: a review. *Comput. Entertain. (CIE)* 16, 1–22. doi:10.1145/3180658

Nitzsche, N., Hanebeck, U. D., and Schmidt, G. (2004). Motion compression for telepresent walking in large target environments. *Presence Teleoperators Virtual Environ.* 13, 44–60. doi:10.1162/105474604774048225

Olver, P. J. (2014). *Introduction to partial differential equations*. Springer.

Peck, T. C., Fuchs, H., and Whitton, M. C. (2012). The design and evaluation of a large-scale real-walking locomotion interface. *IEEE Trans. Vis. Comput. Graph.* 18, 1053–1067. doi:10.1109/tvcg.2011.289

Razzaque, S. (2005). *Redirected walking*. University of North Carolina at Chapel Hill. Ph.D. thesis.

Razzaque, S., Kohn, Z., and Whitton, M. C. (2001). "Redirected walking,", 9. Manchester, UK: Proceedings of EUROGRAPHICS, 105–106.

Souman, J. L., Giordano, P. R., Schwaiger, M., Frissen, I., Thümmel, T., Ulbrich, H., et al. (2008). Cyberwalk: enabling unconstrained omnidirectional walking through virtual environments. *ACM Trans. Appl. Percept. (TAP)* 8, 1–22. doi:10.1145/2043603.2043607

Steinicke, F., Bruder, G., Jerald, J., Frenz, H., and Lappe, M. (2008). "Analyses of human sensitivity to redirected walking," in *Proceedings of the 2008 ACM symposium on Virtual reality software and technology* (New York: ACM), 149–156. doi:10.1145/1450579.1450611

Steinicke, F., Bruder, G., Jerald, J., Frenz, H., and Lappe, M. (2009). Estimation of detection thresholds for redirected walking techniques. *IEEE Trans. Vis. Comput. Graph.* 16, 17–27. doi:10.1109/tvcg.2009.62

Taylor, J. W. (2003). Exponential smoothing with a damped multiplicative trend. *Int. J. Forecast.* 19, 715–725. doi:10.1016/s0169-2070(03)00003-7

Thomas, J., and Rosenberg, E. S. (2019). "A general reactive algorithm for redirected walking using artificial potential functions," in 2019 IEEE Conference on virtual Reality and 3D user interfaces (VR) *(osaka, Japan: ieee)*, 56–62. doi:10.1109/VR.2019.8797983

Thomas, J., and Rosenberg, E. S. (2020). *Reactive alignment of virtual and physical environments using redirected walking*. Atlanta, GA, USA: IEEE, 317–323. doi:10.1109/VRW50115.2020.00071

Usoh, M., Arthur, K., Whitton, M. C., Bastos, R., Steed, A., Slater, M., et al. (1999). "Walking> walking-in-place> flying, in virtual environments," in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (ACM Press/Addison-Wesley Publishing Co.), 359–364. doi:10.1145/311535.311589

Wilde, D. K. (2009). "Computing clothoid segments for trajectory generation," in *2009 IEEE/RSJ international conference on intelligent robots and systems* (IEEE), 2440–2445.

Williams, B., Narasimham, G., McNamara, T. P., Carr, T. H., Rieser, J. J., and Bodenheimer, B. (2006). "Updating orientation in large virtual environments using scaled translational gain," in *Proceedings of the 3rd Symposium on applied Perception in Graphics and visualization (ACM)*, 21–28. doi:10.1145/1140491.1140495

Williams, B., Narasimham, G., Rump, B., McNamara, T. P., Carr, T. H., Rieser, J., et al. (2007). "Exploring large virtual environments with an hmd when physical space is limited," in *Proceedings of the 4th Symposium on applied Perception in Graphics and visualization*. Cambridge, UK: ACM, 41–48. doi:10.1145/1272582.1272590

Williams, N. L., Bera, A., and Manocha, D. (2021). Arc: alignment-based redirection controller for redirected walking in complex environments. *arXiv Prepr. arXiv: 2101.04912* 27, 2535–2544. doi:10.1109/tvcg.2021.3067781

Wilmott, P., Howison, S., and Dewynne, J. (1995). *The mathematics of financial derivatives: a student introduction*. Cambridge, UK: Cambridge University Press.

Yang, J., Holz, C., Ofek, E., and Wilson, A. D. (2019). "Dreamwalker: substituting real-world walking experiences with a virtual reality," in *Proceedings of the 32nd annual ACM symposium on user interface software and technology*, 1093–1107. doi:10.1145/3332165.3347875

Yu, R., Duer, Z., Ogle, T., Bowman, D. A., Tucker, T., Hicks, D., et al. (2018). "Experiencing an invisible world war i battlefield through narrative-driven redirected walking in virtual reality," in *2018 IEEE conference on virtual reality and 3D user interfaces (VR)* (IEEE), 313–319. doi:10.1109/vr.2018.8448288

Zmuda, M. A., Wonser, J. L., Bachmann, E. R., and Hodgson, E. (2013). Optimizing constrained-environment redirected walking instructions using search techniques. *IEEE Trans. Vis. Comput. Graph.* 19, 1872–1884. doi:10.1109/tvcg.2013.88