A psychophysically calibrated controller for navigating through large environments in a limited free-walking space

David Engel* Cristóbal Curio*
MPI for Biological Cybernetics
Tübingen

Lili Tcheang†
Institute of Cognitive Neuroscience
University College London

Betty Mohler* Heinrich H. Bülthoff*
MPI for Biological Cybernetics
Tübingen

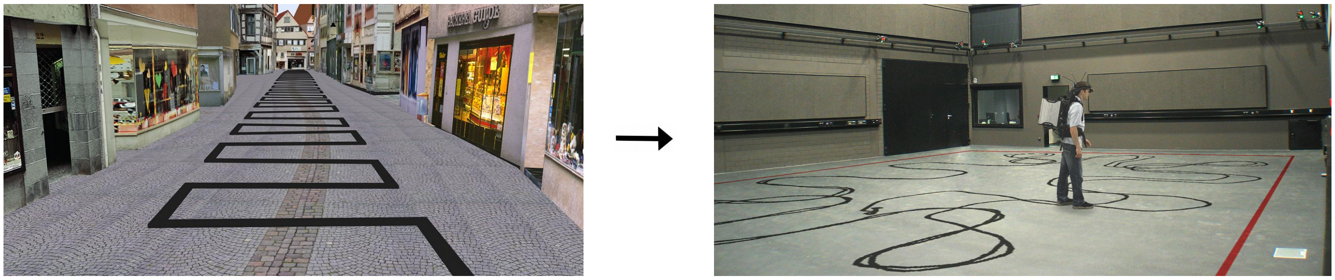

Figure 1: Left: *A subject's view along a path through a virtual city.* **Right:** *The subject's path bent into a tracked space using our controller.*

Abstract

Experience indicates that the sense of presence in a virtual environment is enhanced when the participants are able to actively move through it. When exploring a virtual world by walking, the size of the model is usually limited by the size of the available tracking space. A promising way to overcome these limitations are motion compression techniques, which decouple the position in the real and virtual world by introducing imperceptible visual-proprioceptive conflicts. Such techniques usually precalculate the redirection factors, greatly reducing their robustness. We propose a novel way to determine the instantaneous rotational gains using a controller based on an optimization problem. We present a psychophysical study that measures the sensitivity of visual-proprioceptive conflicts during walking and use this to calibrate a real-time controller. We show the validity of our approach by allowing users to walk through virtual environments vastly larger than the tracking space.

CR Categories: I.3.7 [Computer graphics]: Three-Dimensional Graphics and Realism—Virtual reality;

Keywords: motion-compression, rotational gains, virtual reality

1 Introduction

One goal of virtual reality applications is to convey a strong sense of immersion. Immersion is the feeling of being present in the virtual reality and experiencing it as real. Several interesting areas such as architectural design, interactive entertainment and situational awareness training would benefit from an increased sense of immersion. Studies have shown that walking is the most natural way to explore a virtual environment and provides the best level of immersion [Slater et al. 1995; Usoh et al. 1999; Ruddle and Lessels 2006]. Over the last few years tracking technology has become much better and cheaper, thus opening up new and interesting possibilities for virtual reality applications. Nevertheless, even though large tracking spaces are possible, the available physical tracking space is a limiting factor with reference to the size of the explorable virtual world. This is a severe hindrance and several techniques have been proposed to allow exploration of large virtual worlds in a natural way.

One approach uses technical solutions such as the CyberSphere [Fernandes et al. 2003], moving robotic floor tiles [Iwata et al. 2005] and motion carpets [Schwaiger et al. 2007]. However, such technical solutions are costly and can only support one user. Because of these limitations they will probably remain in the prototype stage for the foreseeable future. Other approaches such as [Razzaque et al. 2002; Interrante et al. 2007; Williams et al. 2007] compress the space by introducing either translational or rotational gains. These methods are based on the observation that in the presence of conflicts between proprioception and vision, the brain is heavily biased towards visual information, preferring a stable visual interpretation of the conflict. This strategy might be explained by the proposition that the proprioceptive system is not calibrated as suggested by e.g. [Bernier et al. 2005]. Using this inability of the human brain to detect such conflicts allows us to decouple positions in the virtual and real world. VR setups allow us to introduce such conflicts in a clever way, allowing us to manipulate the path in the real world associated with a virtual path a user follows.

Several approaches based on such conflicts have been proposed but their main drawback is that they compute the reorientation factors statically before trials begin. This way the system is not able to cope with deviations from the path. To counteract this problem we propose a novel way to dynamically determine reorientation factors (in our case rotational gains), in order to steer a user away from the physical boundaries of the tracking lab, whilst keeping the introduced manipulations as small as possible. If the virtual environment enforces enough turns our technique allows users to theoretically explore infinitely large virtual worlds in a limited tracking space.

*e-mail: firstname.lastname@tuebingen.mpg.de
†e-mail:ltcheang@gmail.com

For paths that contain only few corners our approach would need to be extended to incorporate schemes that allow the controller to steer the user during straight stretches. A major advantage of this method over other approaches is that it is based on a dynamic minimization of cost function, making this system robust against certain deviations from the predefined path. Our system operates on rotational gains, which are multiplicative factors that are applied to head rotations during the exploration of a virtual world. We conducted a psychophysical study that investigates the noticeability of a wide range of rotational gains and show how to translate such findings into a cost term for the optimization function of our controller.

To summarize, this paper is divided into four parts. Section 2 introduces related work and different approaches to similar problems. Section 3 details the theory underlying our method, including the selection of a suitable virtual environment and the on-line calculation of the gains. It also details the psychophysical study conducted to measure the noticeability of rotational gains. Section 4 presents results obtained by our framework and shows data from real-world trials. Finally, Section 5 provides an outlook to further experiments and extensions to this project.

2 Related Work

Perceiving a stable world during active rotations and translations is a multisensory process (e.g. [Sun et al. 2004; Tcheang et al. 2005]). The brain integrates information from vision, proprioception and audition to determine the contributions of self-motion and object motion to the change in the visual world. Conflicts between these channels of information can often occur in virtual environments and have to be dealt with by the perceptual system ([Wallach 1987]). Most of the methods described in this section are based on the observation that such conflicts are often decided in favor of perceiving a stable world.

Several approaches exploiting such conflicts to fit a virtual path into a limited tracking space have been proposed, but none of them optimize the reorientation factors during the exploration. Our approach differs from them since we adapt the reorientation factors dynamically during runtime, making it robust against perturbations. [Interrante et al. 2007] used translational gains applied to the forward motion to augment locomotion through virtual environments. [Razzaque et al. 2002] proposed a technique called *redirected walking* which continuously shifts the target while the subject is walking along a straight virtual path. The immersed user walks along an infinite straight path in a virtual world while walking in a circle in the real world. This method yields impressive results but needs a very large free walking space. [Field and Vamplew 2004] tested different *redirected walking* algorithms and their influence on the necessary size of a tracking space. [Nitzsche et al. 2004; Groenda et al. 2005; Steinicke et al. 2008] proposed ways to employ *redirected walking* and applied gain factors during rotation to create collision-free real world paths corresponding to given virtual paths. [Kohli et al. 2005] demonstrated a way to deform the path in such a way that real-world objects line up with their virtual counterparts, effectively combining *redirected walking* with passive haptics. [Williams et al. 2007] proposed a method exploiting rotational gains for 'resetting' individuals in virtual environments, overcoming the physical limitations of the tracking space. Recently, [Peck et al. 2008] showed that introducing distractors during a reorientation phase enhanced the naturalness and the feeling of presence in the virtual environment.

Teleoperated robotics is a related field where it is also desirable to overcome the physical limitations of the available local space. To overcome these limitations [Nitzsche et al. 2004] and [Su and Luo 2005] employed a technique called *motion compression* which, similarly to *redirected walking*, continuously displaces the target to steer the user inside a tracking space while exploring a much larger virtual world. [Groenda et al. 2005] applied motion compression techniques to a video game enhancing the feeling of presence in the virtual world.

Recent results by [Jaekl et al. 2005] and [Steinicke et al. 2008] investigate the thresholds for translational and rotational gains that are perceived as stable and what the smallest circles are, a user can be led along without noticing that they are being redirected. In this paper we too investigate the noticeability of rotational gains but do not aim at finding the absolute thresholds. Since we use it as part of the optimization scheme, we have tried to derive a function relating a rotational gain to a scalar, that represents the probability that it will be noticed by the user. It is generally accepted that these thresholds become much smaller if the attention of the user is directed towards the conflict. Since this has to be the case during the study the reported thresholds should be seen as lower bounds. However, the size of the circles which are perceived as straight during "redirected walking" makes an approach based solely on this technique undesirable, especially for more complex paths.

A further characteristic of immersive virtual environments is that humans underestimate distance presented through a head-mounted display (e.g. [Loomis and Knapp 2003; Thompson et al. 2004]). Furthermore, a study by [Riecke and Wiener 2007] shows that subjects easily lose their orientation in virtual worlds. This difference in perception of real and virtual environments may be one of the sources why rotational gains are not noticed. If that is the case it may be necessary to continually monitor detection thresholds as virtual reality systems improve.

3 Methods

This section describes the building blocks of our system and the rationale for using them. We first introduce the concept of optimizing rotational gains dynamically. Then, we show the perceptual study we conducted to calibrate our algorithm. Finally, we detail the complete controller that calculates the optimal rotational gains during the navigation.

3.1 Basic concept: Dynamic optimization of gains

Based on the observation that humans often put more trust in their visual system during the presence of visual-proprioceptive conflicts we plan to introduce such conflicts in a way that decouples the positions in the real and virtual world. Our approach is based solely on rotational gains and does not use redirected walking techniques. Our goal is to fit a given path through a virtual world into a tracking space in the real-world by applying different rotational gains at appropriate times. We use the standard definition of rotational gains as being the quotient between the rotation in the virtual world and the rotation in the real world:

$$gain = \frac{rotation_{\text{virtual}}}{rotation_{\text{real}}} \quad (1)$$

The gains are applied only to the subjects' yaw axis. Thus, a head-rotation in the real world by α degrees in the real world will result in a head rotation of *gain* x α degrees in the virtual world. Therefore, the virtual world rotates with the same speed as the real world when the gain is 1. Gains larger than 1 mean that the virtual world rotates faster than the real world and gains smaller than 1 imply the opposite. Given the known size of the tracking lab and a predetermined path, one can precompute the rotational gains for all turns in the virtual world yielding an optimal set of gains. However, we have run preliminary experiments which show that a static precalculation approach is only feasible for very short paths. To determine the gains, such an approach has to be passed the exact sequence of turns the user will make. Since more experienced users adopt a faster and sloppier walking style, that deviates from the predefined path (for example by cutting corners), the precalculation approach no longer applies. This is because the errors no longer cancel out, leading to an accumulation of errors similar to "dead reckoning" in robotics.

To solve this problem we propose a framework that dynamically determines the gains required to keep the user from colliding with a wall. This is a significant advantage over other methods that precompute the gains and are consequently unable to compensate for deviations from the path. Our dynamic approach is not able to find a global optimum of gains but offers robustness against natural walking in a limited space. The approach must achieve multiple goals: First, it must ensure that the user does not collide with the walls of the tracking lab while walking along the virtual path. Second, it should minimize the overall discomfort introduced by the rotational gains. By discomfort we mean anything the user might experience that distracts him/her from the task or reduces the feeling of immersion such as situations where the user perceives visual-proprioceptive conflicts or cybersickness. Lastly, the dynamic approach must be computationally efficient in order to determine the gains whilst the user walks through the virtual world.

The main contribution of this paper is a way to determine rotational gains dynamically using a controller based on an optimization problem. Optimization problems are well understood in the artificial intelligence community and efficient solvers for different kinds of problems exist. The challenge is to translate the problem into a cost function which is minimal for the optimal combination of gains. Our cost function is comprised of a part that maximizes the distance to the walls and a second one that minimizes the discomfort caused by the introduced gains. In the following section we demonstrate how to determine psychophysically the part of the cost function that relates the discomfort during the exploration to a scalar cost. The complete cost function will be discussed in Section 3.3.

3.2 Perceptual study: Determining the cost-of-rotational-gains function

This section details how we generated a cost term that relates a rotational gain to its associated level of discomfort. We aim at explicitly maximizing the comfort during the exploration. Since it is hard to measure perceived level of comfort directly we try to minimize the overall noticeability of the rotational gains. The two measures are expected to be correlated but the latter can be directly determined via a psychophysical study. The main goal of this term is not to force the gains as close to the neutral gain of 1 as possible, since this does not automatically mean the overall level of comfort will be maximal. Studies report that gains between 0.8 and 1.4 are not noticeable [Nitzsche et al. 2004] (although they only investigated head rotation).

The experiment was performed in a tracked space of dimensions 7.55 m x 6.15 m. We track the head of the participants using a state-of-the-art optical tracking system composed of 12 Vicon cameras that capture the position and orientation of the tracking helmet at a frame rate of 120 Hz. The virtual environment was displayed via a head mounted display (HMD) that is connected to a laptop mounted on a backpack the subject wears. The laptop communicates wirelessly with the tracking system and renders the virtual world, allowing the user to walk unhindered by cables. During the experiment the tracking hall was darkened and a curtain attached to the tracking helmet blocked all remaining light sources in the real world. Participants wore earplugs that cancelled out any possible real-world sound sources. The full setup can be seen in Figure 2.

Figure 2: *Participant in the Tracking Lab, equipped with tracking helmet, HMD, and notebook computer mounted on a backpack.*

The experiment took place in small 4 m x 4 m virtual rooms (see Figure 3). In a reset phase, subjects had to walk to one of four marked starting positions in the room with the neutral gain 1 applied to their rotations. During the test phase, they walked around a block that appeared either to the left or the right of them. Furthermore, a rotational gain was applied to their head movements. After the turn they were asked to report whether they turned more or less then 90 degrees in the real world by pressing one of two buttons on a joypad. During a familiarization phase before the actual experiment the participants were introduced to the virtual environment. They were instructed to walk around the corner according to the visual information and at the same time pay attention to conflicts with their real world turning speed.

There was no interaction with the investigator during the test phase of the experiment to ensure that participants could only use their visual and proprioceptive information to determine the gains. The order of left and right turns and the sequence of rotational gains was randomized to minimize possible adaptation effects.

The participants were ten right handed students (four female, six male) from the University of Tuebingen with normal or corrected to normal eyesight. Each completed a total of 90 trials with 9 different gains. The gains were selected manually to cover important rotational gains. The extreme values (0.5 and 2.0), which are always detectable, were determined by a pilot study. Sampling density increased around the detection thresholds given by the literature. The aim was not only to determine the detection thresholds but to arrive at a curve relating rotational gain to noticeability. Furthermore, the division in the gain factor calculation leads to an asymmetry for gains larger and smaller then 1 which makes uniform distribution of the gains not the best way to sample the space. Accordingly, we mirrored the gains we tested at 1. The final set of rotational gains

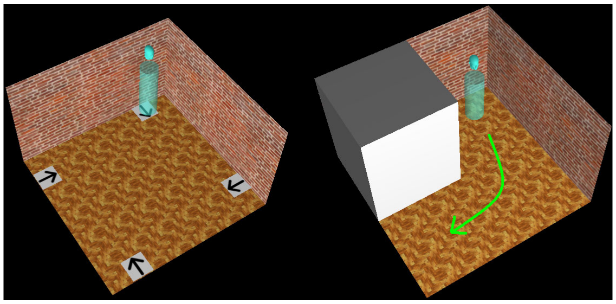

Figure 3: Left: *Virtual room displayed during reset phase.* **Right:** *Virtual room for testing phase. The green arrow marks the trajectory during a right turn.*

we tested were: $1.75, 1.5, 1.3, 1.15, 1, \frac{1}{1.15}, \frac{1}{1.3}, \frac{1}{1.5}, \frac{1}{1.75}$.

Post-experiment questionnaires indicated that participants did not know where the were in the tracking lab. Subjects also unanimously reported that rotational gains smaller than 1 were perceived as being less comfortable.

We used a two-alternatives-forced-choice (2AFC) answer scheme and obtained the response probabilities to which we fitted the psychometric function shown in Figure 4. A psychometric function is a commonly employed fit, that is used to measure how accurately and consistently subjects are able to make a discrimination. We set the answers to be 0 and 1 for the gains 0.5 and 2.0 respectively, since those were determined to be always noticeable during our pilot experiments. We fitted a psychometric function to the *relative gains* which we define as the relative deviations from 1. The transformation of absolute gains to relative gains is given as:

$$\text{gain}_{relative} = \begin{cases} \text{gain}_{abs} - 1 & : \quad \text{gain}_{abs} >= 1 \\ 1 - \frac{1}{\text{gain}_{abs}} & : \quad \text{gain}_{abs} < 1 \end{cases} \qquad (2)$$

We used relative gains instead of the absolute values applied to head rotations since gains are distributed along a multiplicative scale. The psychometric function appears to fit the relative gains well, but more data would be needed to verify this.

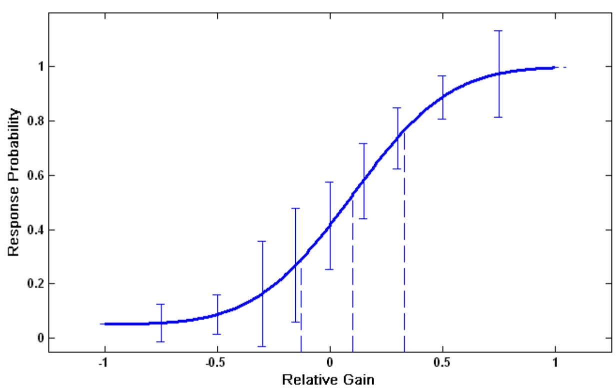

Figure 4: *Psychometric function fitted to the responses of the subjects. The x-axis shows the relative gains described earlier. The y-axis shows the probability of the subjects answering that they experienced a gain larger than 1. The error bars denote the standard deviation with 10 subjects.*

The absolute detection thresholds we measured for active walking were approximately 0.85 and 1.35 with the point of subjective equality being slightly above 1 which is close to the values reported by [Nitzsche et al. 2004] for head rotations. It is critical to note that we do not aim to determine the detection thresholds for rotational gains but want to obtain a continuous function that we can use as a cost term in our control problem. To obtain such a function we calculate the z-values for a densely sampled set of relative gains. Z-values denote the distance between the means of two normal distributions and are used here to determine how the response statistic of the subjects differed from chance level, thus indicating how easily such a gain is detectable. We normalized the z-values to the range between 0 and 1, and incremented the cost for negative relative gains since these were perceived as less comfortable according to the questionnaire. The resulting cost function is shown in Figure 5.

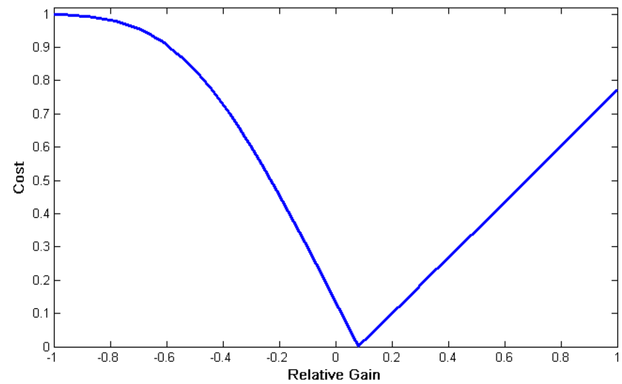

Figure 5: *z-values of the psychometric function in Figure 4 translated into a cost function that relates a relative gain to a single cost value which is correlated with the noticeability.*

This psychometric function and the derived cost term defines the lower bound for the detectability of rotational gains. In this experiment the subjects were informed of conflicting visual and proprioceptive information during the introduction and were explicitly instructed to pay full attention to the occurring visual-proprioceptive conflicts. During the real application the users would be naïve with regards to rotational gains and would perform tasks such as navigation or exploration in a visually rich environment, both of which could function as a distracters. Therefore, the detection rates reported here represent the maximal performance of participants. Consequently, greater rotational gains could be applied during real applications without reducing the level of immersion. Although it would be more useful to measure the detection rates during such ”real application” situations, it is not possible to probe noticeability without directing attention to the concerned factor. Hence, we view these curves as being the worst-case scenario for non-naïve subjects.

An observation of the individual data reveals strong individual differences. It is possible that each subject has different sensitivities to rotational gains. Should this be the case it might be useful to tailor the described cost function explicitly to the specific user before running them in further experiments. This could help to reduce the feeling of cybersickness. Three of the participants experienced dizziness. This could either be due to a general discomfort in virtual environments or to the visual-proprioceptive conflicts caused by the rotational gains. In the case of the latter, avoiding certain rotational gains might solve the problem. Lastly, the reported values only show which rotational gains are likely to be noticed. In the questionnaire subjects mostly reported that they could not detect any gain in about 30% of the cases. However, even if a subject does not consciously notice a conflict it does not automatically mean that it has no influence on human perception. It has been shown that subjects can adapt to rotational gains which are below the reported detection thresholds. This suggests that even non-noticeable rotational gains can lead to a disorientation in the real world.

3.3 Final controller: Dynamic spatial compression

Given the cost-of-gain function described in Section 3.2 we can now compile the complete cost function. It is designed to simultaneously achieve the conflicting goals of minimizing the noticeability of the invoked rotational gains whilst keeping the user safely away from the walls. Furthermore, since we want to determine the gain continuously during the exploration we need to be able to evaluate the function efficiently. A last desirable property of a cost function is differentiability. This enables gradient descent techniques for minimization which can find local minimums very efficiently by exploring the cost function in the direction of the negative gradient.

The output of our optimization process is two gains, one for head rotations to the left and one for head rotations to the right. We allow different gains for left and right turns, to open up more possibilities for manipulation. In particular, when the user is looking around during exploration different gains for the two directions will have an effect since they will generate a drift of the virtual environment with respect to the real world. These drifts are similar to the unnoticeable shifts introduced by the "redirected walking" technique of [Razzaque et al. 2002].

The input into the cost function is the current position in the real world, the current position in the virtual world and the path in the virtual world. The user is expected to walk along the path. Using a dynamic optimization method allows the system to handle certain deviations from the path, but in order to be able to keep the user from hitting the walls, the algorithm needs to be able to predict where the subject intends to go. Equation 3 shows the complete cost function. It is composed of four linearly weighted terms.

$$\mathcal{C}(\mathbf{g},\mathbf{p}) = \alpha_1\varphi(\mathbf{g},\mathbf{p}) + \alpha_2\psi(\mathbf{g},\mathbf{p}) + \alpha_3\mu(\mathbf{g}) + \alpha_4\lambda(\mathbf{g}), \quad (3)$$

where $\mathbf{g} = [g_{\text{left}}, g_{\text{right}}]^T$ is a vector of gains which are applied to turning to the left and right respectively. The vector $\mathbf{p}$ denotes the current state of the user, specifically the path in the virtual world and the position of the user in the real and virtual world. The terms are designed to achieve different goals, $\varphi(\mathbf{g},\mathbf{p})$ controls the distance to the physical wall, $\psi(\mathbf{g},\mathbf{p})$ tries to keep the user parallel to the walls, $\mu(\mathbf{g})$ ensures that the gains for left and right turns do not differ greatly and finally $\lambda(\mathbf{g})$ penalizes gains according to their noticeablity. These terms will be explained in more detail later in this section.

As can be seen in Equation 3 only two terms depend on the current state $\mathbf{p}$ of the user. In our case, the path is parameterized as a set of straight lines connected by turns. This allows an efficient and easy way to introduce rotational gains and translate the path from virtual to the real world. The α's are linear weights that control the importance of the terms in relation to each other. We are seeking the $\mathbf{g}$ that minimizes the cost function 3 such that:

$$\mathbf{g} = \underset{\mathbf{g}}{\textit{argmin}}\,\mathcal{C}(\mathbf{g},\mathbf{p}) \quad (4)$$

The first term $\varphi(\mathbf{g},\mathbf{p})$ of $\mathcal{C}(\mathbf{g},\mathbf{p})$ represents the distance of the subject to the next wall of the tracking lab, if the gains $\mathbf{g}$ were introduced at the next turn. To this end we project the path that the user would walk in the real world if she/he would follow the virtual path and the gains would be introduced at the next turn. We superimpose this path on the tracking space and calculate the distance to the intersection with the next wall. The final cost term is

$$\varphi(\mathbf{g},\mathbf{p}) = \frac{1}{\textit{distance to wall}}. \quad (5)$$

Consequently, we can maximize the distance to the walls by minimizing the cost term. Furthermore, we have a non-linear cost function that penalizes short distances to the walls more severely than larger distances. At first glance, it might look promising to also introduce gains to consecutive corners, but this would increase the computational cost exponentially since we would have to optimize the gains for all subsequent corners independently. We could introduce the same set of gains $\mathbf{g}$ to all consecutive corners, but for an unknown path there is no reason to believe that this would be a good set of gains. Also, since we continuously recompute the gains, we will optimize the gains for the next turn after the current one.

The first term of the cost function encourages gains that lead the main direction of the paths through the center of the tracking space, since this is the longest possible path. As it turns out, this is not the best way to proceed, as it leads the user into the corners of the tracking hall where more severe gains are needed to avoid collisions. In order to keep the overall manipulation closer to the natural gain of 1 we introduced the second cost term $\psi(\mathbf{g},\mathbf{p})$. It is designed to keep the general direction of the path parallel to the nearest physical wall. Figure 6 shows the advantage of this part of the function. By penalizing paths traversing the center center of the tracking hall, we can avoid the corners and need only apply smaller gains. This term also encourages the resulting circular paths that can be seen in Figures 9 and 10. We defined the cost as the scalar product between the direction of the closest wall and the overall direction of the path. The overall direction of the path can be obtained either by projecting a point further along the virtual path into the tracking lab or if the path is simple (as it was in our case) it can just be set by hand. In both cases the gains are applied to the next corner before projection into the real world.

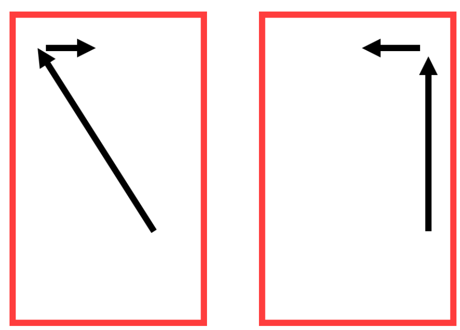

Figure 6: *Bold arrows show the main walking direction of the user.* **Left***: The first term of the cost function would lead the participant through the center into the opposite corner where more severe gains are necessary.* **Right***: Keeping the subject parallel to the closest wall counteracts this.*

The third term is the cost function $\mu(\mathbf{g})$ for the gains which has already been described and grounded in perception in Section 3.2.

We explicitly allow different gains for turning to the left and to the right and check all possible combinations independently. This can lead to situations where the two gains differ greatly from each other even though it might not be necessary. Such situations might be uncomfortable for the user and should be avoided if possible. To address this, we integrated the final term $\lambda(\mathbf{g})$ into the cost function. It penalizes situations where the gains for turning left and right differ greatly. It is computed as the difference between the two gains normalized to 1 as follows:

$$\lambda(\mathbf{g}) = \frac{\| g_{left} - g_{right} \|}{1.5} \quad (6)$$

Simulations and tests have led to a rather small α_4 that still ensures a good trade-off between stability of gains over time and keeping gains equal for both sides.

We manually optimized the α's during simulations and tests in the tracking hall. There seems to be a wide range of alphas which yield a cost function that keeps the subject inside the tracking space during the experiment. The optimal combination of alphas should be determined by a psychophysical experiment, but determining a four dimensional minimum in a very shallow space is difficult. Consequently, the alphas were set to the manually optimized values of $\alpha_1 = 3$, $\alpha_2 = 0.5$, $\alpha_3 = 2$ and $\alpha_4 = 0.1$.

Thus, we arrive at a complete cost function for which a minimum has to be determined in each time step. The algorithm operates on the assumption that the user performs the least amount of turning necessary to follow the path. Since the algorithm manipulates the world only during rotations, this is the worst possible situation. It is important to optimize the turns for this worst case scenario. This way, the algorithm cannot invoke the lowest possible rotational gains but it can ensure that the user does not collide with the walls, which is the more important goal. It is important to notice that even though we optimize the manipulation at the next turn the gains are constantly applied to head rotations. Consequently, if the user does not walk along an exactly straight line (e.g. compensating for overshooting after a turn) or looks around, the algorithm manipulates the relation between real and virtual world. If the user stands still and just looks around, the applied rotational gains would finally align the virtual path ahead in a way that the algorithm determines as optimal. Since humans never walk perfectly straight this helps to reduce the gains needed at the corners.

The complete cost function can be computed very efficiently. Our C++ implementation can determine the rotational gains at 60 Hz on the backpack-mounted laptop without producing any lags. Thus, our framework allows us to compute and introduce rotational gains without hindering the exploration of the virtual world. Since we can compute the cost function fast enough, we can do an exhaustive search through a discreet set of gains. However, all parts of the cost function either are or could be reformulated to be differentiable. This would allow us to use more sophisticated gradient descent optimization techniques, which could be given the previous solution as a starting point, for rapidly finding a minimum. This might become necessary if further cost terms are integrated to achieve more complex or subtle manipulations of the subject in the virtual environment. A certain amount of stability of the rotational gains is guaranteed by the continuity of the parts of the cost function. Adjacent poses (position and orientation) in the real and virtual world will have very similar cost function terms and will therefore yield similar minimums. Still, another regularization term penalizing the derivative of the gains could easily be introduced to keep the gains of adjacent time steps close to each other.

3.4 Suitable virtual environments

Our algorithm needs to be given at least part of the path the user will walk along. A dynamic optimization approach allows certain deviations from the path but the algorithm needs to be able to predict the effect of introducing gains. The system is based on manipulating the turns during navigation through a virtual environment. Consequently, the algorithm would fail for straight paths longer than the size of the tracking lab, thus a combination of our technique with a "redirected walking" scheme might be more appropriate. For our demonstration we present a virtual world which includes enough turns to ensure that the user is safely kept away from the walls, whilst allowing the algorithm to produce rather small rotational gains. We chose a simple meandering path with 2.5 meters of straight stretches joined together by 90° turns (for an example see Figure 7). Note that our algorithm does not only work with meandering paths but with all paths that incorporate enough turns. We chose this rather unconventional path for our demonstrations since it is simple and enforces lots of turns. Furthermore, it allows an easy prediction of the overall direction which is needed for the term that keeps the general direction of the path parallel to the walls in Equation 3.

A second requirement to the virtual environment is that it should allow easy visual orientation. The technique is based on the observation that during conflicts between the visual and proprioceptive inputs, the human brain often favors the visual channel. To further encourage this, we present a visually rich environment that allows easy visual orientation. Large city models like Virtual Tübingen (http://virtual.tuebingen.mpg.de) are especially useful since they do not only provide a rich environment but also encourage the user to look around. As mentioned earlier the algorithm also introduces rotational gains while the user is looking around and can consequently bend the space even when the subject is stationary.

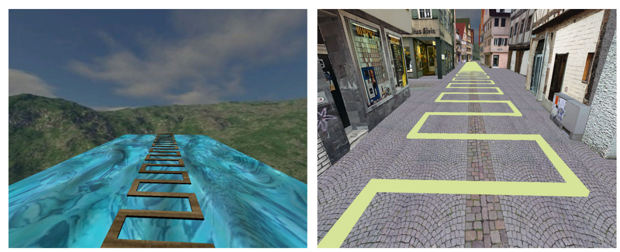

Figure 7: *Virtual environments used for our experiments. Both feature meandering paths through rich open spaces.* **Left***: Path over a lake inside a skybox.* **Right***: Path through Virtual Tübingen.*

4 Results

To test the capabilities of our algorithm we implemented a framework that simulates the movement of a person along the meandering paths described in Section 3.4. It uses the described cost function to determine the gains it applies at the corners. It only simulates straight segments joined by 90° turns, which represents the worst case scenario, since all other ways to follow the path inevitably result in a larger accumulated turn angle. During the simulations the algorithm always kept the user within the tracking space boundary.

It is obvious that the tracking lab dimensions will directly influence the gains required to keep the user inside the tracked space. The simulation allows us to investigate the relation between the size of the tracking lab and the average gains. Figure 8 shows this relationship for meandering paths from randomly drawn starting points in a quadratic tracking lab.

To show the applicability of our system we also implemented a demonstrator that allows users to walk through the virtual worlds in our tracking hall. These tests were done in a larger tracking lab which provides a 9 m x 12 m tracking space. We track the head of the subject with a state-of-the-art Vicon optical tracking system

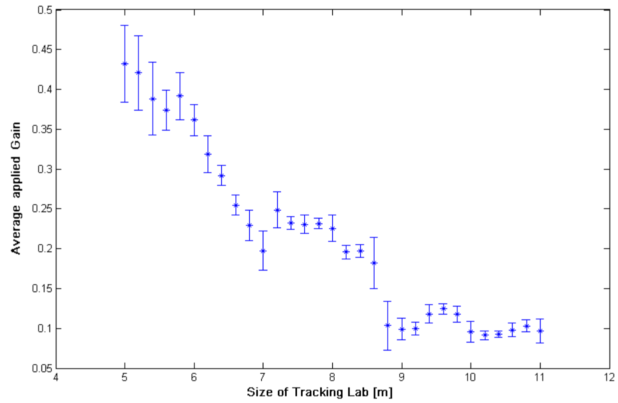

Figure 8: *The relationship between the size of the tracking lab and the average gain that is applied by our algorithm. The size of the simulated tracking lab is shown on the x-axis. The error bars denote the standard deviation over 100 simulated runs with random start positions and orientations for each tracking lab size.*

while projecting the virtual reality with the same head mounted display (HMD) as described in Section 3.2.

Naïve users reported that they did not notice the dynamically changing gains and could explore an infinitely long meandering path through the virtual world, whilst the algorithm always enforced a safe distance to the walls. Figure 9 shows the first 20 meters of a recorded path of a naïve test subject walking through a virtual city (see also the supplementary video). It is interesting to note that even though some turns involved detectable gains outside the range of 0.8 to 1.4, the users did not notice them during the exploration. We believe that this is due to the fact that the noticeable thresholds are further apart if the attention of the user is not explicitly directed to the visual-proprioceptive conflict. If they adopt a natural walking style and do not pay special attention to their proprioception even non-naïve users do not realize the conflicts.

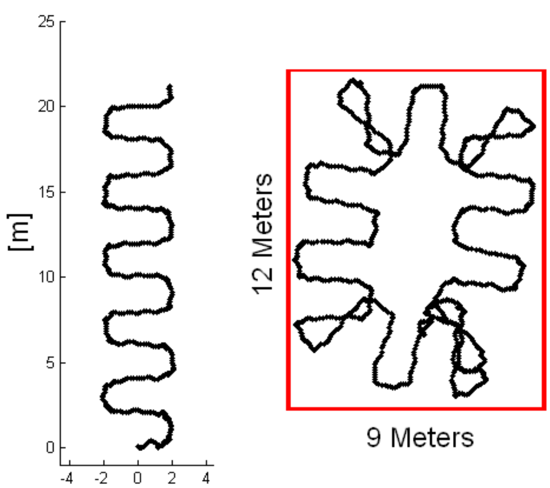

Figure 9: *Recorded data from a naïve subject walking through Virtual Tübingen.* **Left***: Path of the subject in the virtual environment.* **Right***: Path in the tracking space (trackable area shown in red).*

As mentioned before the dynamic optimization method makes the approach stable against perturbations. This allows users to adopt a sloppy walking style. Figure 10 shows the recorded trajectories of an experienced user that does not adhere to the presented path but just walks along an approximate meandering path inside the convex hull of the original path. Previously proposed methods that determine the reorientation factors a priori, would not be able to accommodate such a walking style. Consequently, other methods would run into problems if used by more experienced users, or would have to introduce reset conditions should the accumulated errors become too large and the reorientation factors have to be recomputed.

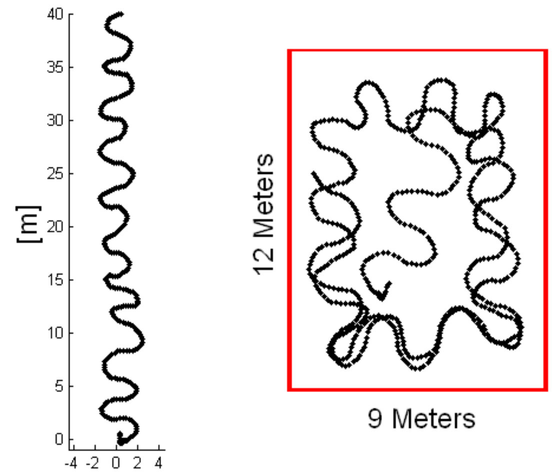

Figure 10: *Recorded data from non-naïve subject walking fast and sloppily along a meandering path through a virtual world.* **Left***: Path of the subject in the virtual environment.* **Right***: Path in the tracking space (trackable area depicted in red).*

The user trials showed that the gains become more evident when real world light or sound sources are available, such as people talking in a corner of the tracking lab or the lights not being completely turned off. The users employ such real world landmarks subconsciously for orientation in the tracking space. It seems that under such conditions, human perception relies less on the visually presented virtual environment and therefore tolerates only smaller gains before noticing the conflict. This observation supports the notion that spatial orientation is a multi-sensory integration process. Consequently, only in the absence of all other cues, can the visual pathway dominate over proprioception. Thus, when designing such free walking systems, one also has to pay attention to the other senses. Audition for example, should be suppressed by either using earplugs such as in our experiments or headphones capable of rendering 3D sounds which should further increase the range of tolerable gains. Furthermore, we believe that view angle of the HMD is an important factor in the amount of immersion that can be achieved, correlating with the users' trust in the virtual scene.

A nice feature of our control framework is that the currently involved terms need to know only the next few meters of the path through the virtual world to determine the optimal gains at the current location. Nevertheless, knowing some part of the future path is a crucial necessity for all such optimization algorithms based on rotational gains or redirected walking. Otherwise the controller cannot predict the outcome of the introduced manipulations. Consequently, it is theoretically impossible to allow a user to explore a virtual world in an unconstrained way.

5 Outlook

A primary extension to our framework will be to move away from the static meandering paths, as they constitute an unnatural way to walk through a world. At the moment these paths are required to ensure that the subject takes enough turns during the exploration of the virtual environment. After parameterization of the path the user will be able to select a path interactively. A dynamic path routing will be integrated into a global optimization framework.

Secondly, we plan to further investigate the cost-of-gain function which gives a level of discomfort for each allowed gain using psychophysical paradigms. This function can be different for each subject and may also be situation dependent (e.g. depending on the current walking speed, the acceleration, the current articulated body-state or the presence of an additional task). Measuring it by psychophysical experiments should allow the introduction of larger gains in a non-perceptible fashion.

Lastly, at the moment our technique requires corners and we consider combining it with the *redirected walking* method from [Razzaque et al. 2002]. If we can integrate the advantages of both techniques by bending the space dynamically during straight walking as well as in corners, we will be able to use overall lower gains, making the virtual environment more immersive and the experience even more natural.

Acknowledgements

The authors would like to thank Michael Weyel and Stephan Streuber for their help. This work was supported by EU-Project BACS FP6-IST-027140.

References

Bernier, P.-M., Chua, R., and Franks, I. M. 2005. Is proprioception calibrated during visually guided movements? *Exp Brain Res 167*, 2, 292–296.

Fernandes, K. J., Raja, V., and Eyre, J. 2003. Cybersphere: the fully immersive spherical projection system. *Commun. ACM 46*, 9, 141–146.

Field, T., and Vamplew, P. 2004. Generalised algorithms for redirected walking in virtual environments. In *International Conference on Artificial Intelligence in Science and Technology*.

Groenda, H., Nowak, F., Rössler, P., and Hanebeck, U. D. 2005. Telepresence techniques for controlling avatar motion in first person games. In *INTETAIN*, 44–53.

Interrante, V., Ries, B., and Anderson, L. 2007. Seven league boots: A new metaphor for augmented locomotion through moderately large scale immersive virtual environments. In *IEEE Symposium on 3D User Interfaces*, IEEE, 167–170.

Iwata, H., Yano, H., Fukushima, H., and Noma, H. 2005. Circulafloor. *IEEE Computer Graphics and Applications 25*, 1, 64–67.

Jaekl, P. M., Jenkin, M. R., and Harris, L. R. 2005. Perceiving a stable world during active rotational and translational head movements. *Exp Brain Res 163*, 3, 388–399.

Kohli, L., Burns, E., Miller, D., and Fuchs, H. 2005. Combining passive haptics with redirected walking. In *International conference on Augmented tele-existence*, ACM, 253–254.

Loomis, J. M., and Knapp, J. 2003. Visual perception of egocentric distance in real and virtual environments. In *Virtual and Adaptive Environments*, Erlbaum, Mahwah, NJ, 21–46.

Nitzsche, N., Hanebeck, U. D., and Schmidt, G. 2004. Motion compression for telepresent walking in large target environments. *Presence: Teleoper. Virtual Environ. 13*, 1, 44–60.

Peck, T., Whitton, M., and Fuchs, H. 2008. Evaluation of reorientation techniques for walking in large virtual environments. In *IEEE Virtual Reality Conference*, 121–127.

Razzaque, S., Swapp, D., Slater, M., Whitton, M. C., and Steed, A. 2002. Redirected walking in place. In *Workshop on Virtual environments*, Eurographics Association, 123–130.

Riecke, B., and Wiener, J. 2007. Can people not tell left from right in vr? point-to-origin studies revealed qualitative errors in visual path integration. In *IEEE Virtual Reality*, 3–10.

Ruddle, R., and Lessels, S. 2006. For efficient navigational search, humans require full physical movement, but not a rich visual scene. In *Psychological Science*, vol. 17, 460–465.

Schwaiger, M. C., Thmmel, T., and Ulbrich, H. 2007. Cyberwalk: Implementation of a ball bearing platform for humans. In *HCI (2)*, Springer, J. A. Jacko, Ed., vol. 4551 of *Lecture Notes in Computer Science*, 926–935.

Slater, M., Usoh, M., and Steed, A. 1995. Taking steps: the influence of a walking technique on presence in virtual reality. *Transactions on Computer-Human Interaction 2*, 3, 201–219.

Steinicke, F., Bruder, G., Ropinski, T., and Hinrichs, K. H. 2008. Moving towards generally applicable redirected walking. In *Virtual Reality International Conference*, 15–24.

Su, J., and Luo, Z. 2005. Incremental motion compression for telepresent walking subject to spatial constraints. In *ICRA*, 69–74.

Sun, H.-J., Campos, J., and Chan, G. 2004. Multisensory integration in the estimation of relative path length. *Exp Brain Res 154*, 246–254.

Tcheang, L., Gilson, S. J., and Glennerster, A. 2005. Systematic distortions of perceptual stability investigated using immersive virtual reality. *Vision Research 45*, 16, 2177–2189.

Thompson, W. B., Willemsen, P., Gooch, A. A., Creem-Regehr, S. H., Loomis, J. M., and Beall, A. C. 2004. Does the quality of the computer graphics matter when judging distances in visually immersive environments? *Presence: Teleoperators and Virtual Environments 13*, 5, 560–571.

Usoh, M., Arthur, K., Whitton, M. C., Bastos, R., Steed, A., Slater, M., and Frederick P. Brooks, J. 1999. Walking > walking-in-place > flying, in virtual environments. In *SIGGRAPH*, ACM, New York, NY, USA, 359–364.

Wallach, H. 1987. Perceiving a stable environment when one moves. *Annual Review of Psychology 38*, 1–29.

Williams, B., Narasimham, G., Rump, B., McNamara, T. P., Carr, T. H., Rieser, J., and Bodenheimer, B. 2007. Exploring large virtual environments with an hmd when physical space is limited. In *APGV*, ACM, New York, NY, USA, 41–48.