CHAPTER 10

Locomotion and Navigation

10.1 INTRODUCTION

When we use traditional interactive systems, like computers and games, our physical bodies are typically not moving. We sit on a couch or at a desk, watching a screen and manipulating interactable objects to effect change in the system we are observing. When we use computers for much of the day, first to work and then to play, we end up sitting for hours at a time, which is not healthy. Some computer and game systems, like the Wii, PlayStation Move, and Microsoft Kinect, have been developed to encourage users to move more when using computer systems, mostly by converting body movements into control events. Even if the avatar of ourselves in the game we are playing is instructed to move around the simulated world we are observing, we ourselves are not moving (or are only moving a proxy amount), and there is a separation between our commands to move our avatar and the motion we see in the simulation. In some situations, we may choose to walk while we are playing a hand-held game or using our phone, but this is generally considered dangerous and could increase the chance of tripping or bumping in to things. Augmented reality software—systems that add a layer of information to the physical world through a device like a smartphone—can make a connection between our motion in the physical world and what is shown to us through our screens, although the location and direction of the phone or AR device is the primary way in which the user's motion is understood by such system.

As may be familiar to the reader by now, virtual reality (VR) is different. In much the same way that reaching out to pick up an object feels different than pressing a button to cause an avatar to pick up an object, walking around a room feels fundamentally different than moving a joystick to move an avatar around the room. Our sensorimotor contingencies reinforce our perception of the virtual world as if it were real, and a deeply natural, intuitive mode of interaction becomes available to us—we can walk over to something interesting and interact with it.

The challenge, of course, is that the physical spaces we are in when we use VR are usually not the same shape and size as the virtual world we are in, so walking about very quickly causes us to break presence by bumping into walls or tables, or by allowing us to walk through a virtual wall that should stop us. Because of

DOI: 10.1201/9781003261230-10

these mismatches, developing methods to move about in virtual worlds has become a significant source of innovation and frustration.

The ability to perform **locomotion**, or to move around an environment, is necessary not just to accomplish tasks that may be out of reach but also to increase the extent to which environments can be perceived—we can look around corners and walk into rooms to see what we couldn't see before we moved. Well-implemented movement in a virtual environment, when combined with good practices for immersing sensory units, can enhance the user's sensation of presence within the environment. Conversely, poorly implemented or unsuitable movement can be detrimental to the user's overall experience. As such, it is important for VR developers to understand several aspects of locomotion, including the following:

- The physical and psychological mechanism by which locomotion occurs
- The purpose(s) of locomotion in a VR scenario
- The different categories of locomotion strategies that can be used
- Design principles and considerations for selecting and implementing locomotion strategies
- Techniques for analysing the effectiveness of locomotion strategies across different types of virtual environments
- The benefits, drawbacks, and overall implications of locomotion strategies

The psychological component of locomotion, which is comprised of identifying a path between one's origin and a destination, as well as orienting themselves within a physical space, is known as **wayfinding**. In order to effectively wayfind, a user must be able to spatially analyse an environment, understand their current location relative to their goal, and plan how to move from here to there. The process of moving towards a destination along a desired route is the physical component of locomotion, which is known as **travel**. Travel can also be described as the observable output of a locomotion or navigation system. Often, wayfinding and travel happen together, or in cycles (check the map, walk to the next waypoint).

Historically, in the physical world, wayfinding has been a challenging task requiring accurate timekeeping, direction finding, surveying, and map making. Improvements in each of these tasks have led to technologies (the chronograph and the compass) that have improved these and other aspects of our lives, but the development and launch of satellite-based location systems like GPS and GLONASS, and the subsequent removal of dithering, has meant that it is essentially impossible for anyone on the planet who owns a smartphone to be lost. GPS signals do not function well indoors, but technology such as Bluetooth beacons mean that indoor automated wayfinding is also available.

In the physical world, locomotion is usually intuitive and subconscious. Once a destination and route have been identified, ambulatory individuals[1] can simply walk along this route if the destination is nearby. If the destination is sufficiently distant that walking would be impractical, then a bicycle or motor vehicle might be used instead. These vehicles, despite eliminating much of the physical exertion in walking (or walking-adjacent motion), still contain hardware that translates a user's physical movement into travel (via mechanical or control assistance). Additionally, travel in the real world is only ever restricted by external factors, which can either be temporary (e.g. passers-by and construction zones) or permanent (e.g. buildings, bodies of water, and cliffs).

In contrast, the extent to which walking can be performed in a virtual environment is constrained by the size of the user's physical play area. As a result, unless the virtual environment's size is also limited, walking on its own is usually insufficient for exploring the bounds of the virtual environment and can rapidly lead to the user interacting with the boundary of their playspace. Such a discrepancy can negatively affect the user's experience, especially if an environment's inaccessible areas contain stimuli with which the user expects to be able to interact. As such, the vast majority of VR experiences contain alternative strategies for moving about in the virtual world. These strategies generally provide greater locomotive range than walking, but often at the cost of performance in other areas, including ease of use, appeasement of the vestibular system, reduction of plausibility illusion, and overall effect on the sensation of presence. Thus, it is necessary to understand the benefits, limitations, design considerations, and overall ramifications of each particular locomotion strategy.

The Suitability of the Term "Walking" in VR

We often make a general assumption that the majority of VR users are ambulatory, while also acknowledging that this is not universally the case, and that designers may need to take this into account when planning locomotive methods. Non-ambulatory users may be able to move themselves by way of surrogate devices, such as wheelchairs or electric scooters, but many "walking" locomotive experiences in VR assume that users can walk and therefore may be exclusionary if the user is unable to do so. Throughout this chapter, "walking" will be shorthand for one-to-one movement in reality and VR. The accessibility implications of locomotion strategies that require such one-to-one movement are discussed in Section 8.5

This chapter begins with a discussion of the concept of travel, including its most common tasks, its overall purposes, and its ability to be separated into subtasks and subsystems. Next, the most prevalent locomotion strategies in modern VR experiences are categorized. The chapter then contains a discussion on various techniques

[1]In much the same way that vision is typically prioritized in VR hardware development, walking is often prioritized in VR locomotion systems. Not everyone can walk, and developers should keep human diversity and difference in mind when designing locomotion systems.

for quantifying the efficacy of locomotion strategies, with proposed metrics including outright performance, utility, effect on presence, and synthesis with other strategies. Following this section, design principles for maximizing the ease of use and contextual justification of locomotion strategies are explored, as is the concept of movement being a mechanic in and of itself. Additionally, a framework for the environment-influenced suitability of locomotion strategies is proposed, which includes environments designed specifically to accommodate particular strategies. Finally, prevalent adverse effects of locomotion strategies are discussed from a holistic perspective, including vection and VR sickness, ergonomic concerns, and accessibility issues.

10.1.1 Locomotion terminology

Because locomotion is such a fundamental aspect of VR experiences and remains difficult to address, there are many approaches to any given locomotion problem, and a few terms have emerged in common practice that may be confusing for the new developer. Further, some of these terms, although seemingly generic, have been used to describe specific implementations of specific techniques.

Continuous motion is sometimes used to mean an artificial motion that is not discrete, for example when a user uses a thumstick to indicate a direction and the view drifts in that direction. This type of motion is called “continuous” to separate it from discrete motion such as teleportation or angular rotation. We will use the term “continuous” to refer to any motion or rotation that is not discrete, whether caused by the thumbstick, some other control, or involuntarily. “Voluntary continuous motion” (VCM) is the term we will use to describe manipulating a control, like a thumbstick, to drift through the environment.

Natural locomotion is sometimes used to refer to a specific implementation of a technique of swinging your arms or walking in place to move. The rate and direction in which you swing your arms translates to the direction in which the view drifts. We refer to this technique as “arm cycling” or “walking in place.”

10.2 PURPOSES OF LOCOMOTION

At its core, the goal of locomotion is to get the user from one place in a virtual world to another. The scope of this general task is broad, ranging from walking across the room to pick up an object to travelling around the world or across space to visit another location. The core aspect in all of these tasks is the movement of a user’s agency from one context to another, that is allowing the user to affect change in another virtual location. Some tasks that appear to require locomotion (like reaching an object on the far side of a room) may be implemented instead by *action at a distance* as discussed in Section 11.4. The developer should consider whether locomotion itself is required or whether moving objects or contexts may also be possible.

10.2.1 Travel Tasks

There isn't one correct way to implement locomotion in VR. Many different VR locomotion strategies exist because suitability is dependent on the task for which they are used. As such, it is important to understand both the different types of travel tasks, the reasons for which each task is performed, and the implementation options available.

Besides simply seeking a destination, travel tasks can also be classified as *exploration*, *search*, or *manoeuvring*, as shown in Figure 10.1. Exploration means you are learning about the environment, searching means you are looking for something specific in the environment, and manoeuvring means you are accomplishing some specific task within the environment.

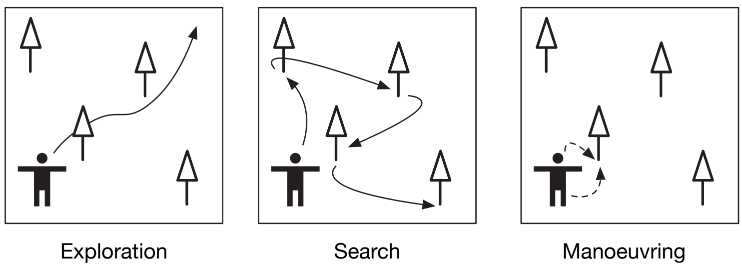

Figure 10.1: Common locomotion tasks.

Exploration tasks are characterized by the user lacking an explicit goal for their locomotion. Instead, the primary purpose of the user's movement is to obtain information about the environment such that any explicit goals may be more easily accomplished. Exploration tasks are primarily performed when the user first encounters an environment; however, they may be performed again in latter stages of an experience if new goals or stimuli appear, or if the environment changes.

Sometimes, exploration is initiated by unexpected stimuli—a user may abruptly deviate from their current path if they encounter something interesting in the environment. This control should enable users to alter their destination and orientation at any point throughout the travel process and carry minimal cognitive load. Exploration tasks are necessary for VR experiences with large open-world environments and a non-linear progression of gameplay, but for experiences where the user must perform tasks in a smaller or well-known environment, searching may be a preferred method.

Search tasks, in contrast to exploration tasks, are goal directed; they involve travel to a specific location or locations (possibly unknown) within the environment. Prior to initiating travel, the user consciously chooses a desired destination which may or may not contain the goal state; thus, search tasks involve both travel and wayfinding. There are two subcategories of search tasks:

- **Primed search tasks** are tasks in which the user knows the location of their target destination, as well as a path by which to reach it.

- **Naive search tasks** are tasks in which the user lacks knowledge of either the destination's precise location or a path by which to reach it.

Naive search tasks are different from exploration tasks because the user does, at the very least, have a destination in mind for their travel; thus, the path of locomotion observed in a naive search task will usually be linear. Primed search tasks may also contain elements of exploration if the user wants to explore an alternate route, but the final destination is fully explicit, and the user still possesses knowledge of at least one method of access. In contrast, exploration tasks are entirely unstructured, and the path of travel may be linear or cyclical.

Despite these differences, exploration and search share several common aspects. Both types of tasks involve user-directed movement for the purpose of obtaining or corroborating environmental information; the user's findings are reinforced by sensorimotor information at each step in the locomotion process. As such, locomotion strategies with wayfinding techniques and feedback systems may be better suited for search tasks than those with purely general mechanics.

Finally, **manoeuvring tasks** are performed whenever the user must precisely position their avatar in order to perform a specific action in the virtual environment. For example, if the user is playing a VR basketball game and wants to shoot a three-point shot from the corner, they must first navigate to the approximate location, which is a primed search task. However, the process of carefully positioning their avatar such that it is behind the three-point line, and turning to face the net in order to maximize their chances of making the shot, is a manoeuvring task. Manoeuvering tasks are perhaps the most difficult of the three travel task types to implement effectively, as many locomotion strategies enable a wide range of motion in exchange for low fidelity and poor precision, while others are advantaged with respect to precision but suffer in the areas of speed and range. In order to accommodate manoeuvering tasks, locomotion strategies should feature a balance between speed and precision.

10.2.2 Movement Subtasks

Beyond simply identifying a destination and establishing a route to it, there are numerous characteristics of travel tasks and the overall process of locomotion. These characteristics may influence the duration of wayfinding and exploration or be directly tied to the exact locomotion strategy that the user invokes if several locomotion modalities are present in the same experience. Because most physical playspaces are small relative to the virtual worlds they represent, many locomotion strategies are invoked artificially (i.e. by some mechanism other than physical motion). When multiple locomotion options are available, the user must both plan a path to a destination and choose which technique to use to follow that path. Some examples of these deductions are listed below.

10.2.2.1 Determining an ideal range of locomotion

The suitability of a locomotion strategy for travel in a VR environment is influenced not only by accessibility and usability but also by the properties of the environment itself. If the virtual environment's area is comparable to the user's playspace, then direct locomotion may be sufficient. In environments with larger (but still limited) dimensions, a strategy that either is invoked artificially or enforces a difference between real-world and virtual-world displacement is likely necessary. In open environments with primarily long-ranged travel tasks, locomotion strategies with velocity control may be more efficient.

Note that for manoeuvering tasks specifically, the size of the virtual environment doesn't matter since these tasks may always be accomplished by short-ranged physical motion. However, if a manoeuvring task is associated with other travel tasks, it may be more appropriate to align or combine manoeuvring with another strategy or remove the requirement for manoeuvring by having the local alignment of final destination be "right" with respect to whatever task may be needed at that location. If a user is running from place to place and taking cover at each location, the system could automatically put the user in a crouched position at each new location.

10.2.2.2 Explicit selection of a destination

An important factor in the transition between wayfinding and travel in virtual environments is whether the destination can be explicitly selected or not. In locomotion strategies such as teleportation, wherein the user points their controller at a location in the environment and is discretely transported to it, explicit destination selection could be argued to alleviate the user's cognitive load, as it makes the process of travel passive. In contrast, most locomotion strategies with continuous transitions force the user to actively facilitate each stage of their travel. The range of teleportation here is a critical design strategy, since short-range teleportation may still require wayfinding, while long-range teleportation may allow the user to select a destination and involve a cutscene rather than motion through the environment.

The primary caveat of locomotion strategies with explicit destination selection and passive travel is that the discrete nature of the motion may cause errors in movement precision. Teleportation that requires direct selection of a destination has the same challenges as object selection techniques (see Section 11.4.1) and can therefore have challenges with accuracy and control, when compared to continuous motion. If locomotive accuracy is an integral feature of a VR experience, then it is necessary to consider whether the physical and cognitive advantages of explicit destination selection outweigh its higher propensity for errors.

10.2.2.3 Reorientation and recalibration

The ability to view a virtual environment at different orientations is just as important for wayfinding, exploration, and manoeuvring as the ability to view at different positions. In the physical world, both repositioning and reorientation are accomplished via physical motion. The former two metaphors are likely still suitable in environments with comparable size to the user's designated play space, but in larger environments, long-distanced travel is generally invoked by artificial means, while walking is best suited for manoeuvring tasks.

In an open virtual environment which affords primarily artificial locomotion, the question of how reorientation should be performed gains several layers of complexity. On its own, physical rotation of the head-mounted display is advantaged with respect to simplicity and cognitive load. People are used to performing head rotations in the real world; a one-to-one head rotation requires no user translation of a physical action to a virtual interface, and the ocular system's sensorimotor feedback matches feedback received by the vestibular system. However, if physical head rotation is combined with an artificial locomotion strategy for translation, then cognitive load may actually increase due to the presence of multiple interaction metaphors—especially if the artificial translation occurs in a non-continuous manner, such as in teleportation.

If you push a thumbstick forward to begin to drift, the direction of drift may be specified by the direction of the thumbstick, or the direction you are looking. It may be more natural to walk in one direction and look in another, but controlling drift direction with a thumbstick may require an alignment to a world coordinate system that the user might not be able to perform.

In contrast, artificial rotation's primary weakness is its tendency to cause VR sickness if its transitions are continuous, as the changes observed by the ocular system do not stimulate the vestibular system. If the rotation occurs in discrete segments, then this problem is largely mitigated, as the task of rotation becomes passive. It is worth noting that, in a similar vein to explicit destination selection, some teleportation systems (see Section 10.3.3) additionally enable the user to preemptively choose an orientation for their avatar post-teleport. In addition to relegating reorientation to a passive task, this strategy enables amalgamation of search tasks with manoeuvring tasks, which saves time and reduces cognitive load. However, any significant change to orientation will prevent the user from immediately analysing the accuracy of the change in position, and vice versa. Consequently, error recovery increases in difficulty.

Other advantages of artificial rotation include timeliness, as the user need not consciously navigate between multiple interaction metaphors, and ergonomics, as the lack of physical motion mitigates prolonged strain on the user's cervical spine.

10.2.2.4 Travel as a subtask

There exists a subset of search tasks in which the user is not consciously aware of the fact that they are selecting a destination and travelling to it. In these scenarios, travel is a necessary subtask of some overarching supertask—in other words, it is a means to an end. For example, in a VR experience designed to teach the principles of recycling, the user may intuitively develop an awareness of the recycle bin's location, as well as the simplest path by which to access it. However, the user's objective is not to simply arrive at the recycle bin—what they really want to do is dispose of the plastic water bottle in their hand. In situations where travel is a subtask, it is important that the selected locomotion strategy does not distract from or interfere with the user's focus on the supertask. For example, if the controls for holding the bottle and for activating locomotion are on the same hand, the user may not be able to do both at the same time.

10.2.3 Wayfinding Aids

In its simplest form, wayfinding is a combination of *perception* (identifying clues along a route and recalling their significance), *attention* (interpreting new stimuli and building a mental model), and *memory* (comparing current stimuli to recalled information from previous exposure to the environment).

Just as wayfinding in the physical world is difficult without the aid of compass, chronograph, GPS, or signposts, well-constructed virtual environments may also be too complex for the user to independently identify what sources of information will best guide their travel. Furthermore, complex virtual environments are prone to causing disorientation and inaccurate perceptions of depth, especially if the available locomotion strategies cause a mismatch between the ocular and vestibular systems. As a result, virtual environments benefit from the inclusion of *wayfinding aids*, which assist users in forming mental models and maintaining a sense of bearing. There are two primary types of wayfinding aids:

Environmental wayfinding aids are sources of navigational information that are diegetically (see Section 8.1.1) present in the virtual environment, and whose existence is justified in the context of the overall experience. Examples include road signs, buildings, trees, terrain changes, and other unique signifiers. Environmental wayfinding aids are often visual, but other wayfinding aids exist, such as auditory (the sound of traffic indicating the presence of a nearby road), tactile (bumps on a sidewalk indicating a curb), or olfactory (following the smell of a bakery).

Personal wayfinding aids are sources of navigational information carried with the user, like a map or a compass. These may be diegetically present in the virtual environment, or they may be part of an overlay or heads-up display (HUD), like a mini-map, goal indicator, or route marker.

Mini-maps and game mechanics

A mini-map is a small representation of the game world that the user can see. Traditional gaming scenarios usually situate the mini-map on the user's HUD or makes the map diegetically available on a tool (like a handheld screen or GPS) available to the avatar, and subsequently seen by the user. Mini-maps are usually top-down and usually two dimensional, appearing much like a traditional map or GPS display in the physical world. Topographical features may be indicated by lines or shading, but the location of the avatar and objects or enemies in the world is represented in two dimensions. As such, players may not have complete information, such as which floor of a building an enemy is on. VR provides the opportunity for three-dimensional mini-maps, which can be used not just for information display but also for interaction. The "world-within-world" design pattern offers the user the chance to interact with the mini-map and in some cases even interact with the game world directly through the mini-map. "A fisherman's tale" (Innerspace VR/Vertigo Games, 2019) provides a compelling example of this design pattern. A fisherman is in a shack, and on the table is a model of the shack he has been working on. When he lifts the roof of the model, the roof of the shack he is in lifts, and he gazes up to see an enormous version of himself holding the roof. Likewise, down in the model, he sees a tiny version of himself interacting with an even smaller model of the same shack. The game allows clever gameplay events where one version of the fisherman can pass objects to a smaller or larger version, similarly acquiring a smaller or larger version of the item.

If an environment lacks wayfinding aids, or the structural model of its wayfinding aids conflicts with the user's interpretation of previous environments, then it may be necessary to alleviate the process of mental mapping and simply display the information directly. In such circumstances, personal wayfinding aids are ideal. A map system on a HUD that updates in real time and displays the user's position, for example, ensures that the user can easily identify their location regardless of the degree to which their mental model of the environment has developed. A compass is another alternative, which provides a sense of direction without being tied to environmental stimuli. Due to the relative ease of operating a compass, especially while in motion, a compass can be implemented as a physical object within the environment. Depending on the context, this may be a more favourable design choice than a non-diegetic HUD element with respect to inducing the sensation of presence.

10.2.3.1 Markers

A *marker* is an object or interface element which helps a user remember or navigate to a location. Markers may be added by the user themselves, or by an agent or task within the experience. Markers straddle the line between environmental and personal wayfinding aids, because they can appear on local maps or out in the world, or both. An object becomes a marker when it is used to navigate—thus a building can be a marker if the navigation task includes "turn right at the bakery." Markers can exist non-diegetically in the form of pushpins or coloured shapes on a HUD map, or diegetically in the form of flags and territory markings. Markers can also take on the form of trails, which suggest a path to take rather than a location to find. If trails are marked by directional cues, such as footprints or arrows, then they can provide even more information.

Markers enhance the user's ability to develop a mental model of the virtual environment. Providing users with numerous markup options with respect to form, colour, and size provides the opportunity for a "legend" to be established, which can improve spatial comprehension. As well, providing a variety of marker types enables users to use different types of markers for different purposes, either defined by the application or by the user themselves. For instance, users may decide to use yellow markers to indicate areas that have already been visited, while cyan markers indicate areas that look interesting and need to be visited later. Experiences which afford such a degree of interactive analysis can accelerate wayfinding and ensure that minimal time is lost due to travel-related confusion.

10.3 ARTIFICIAL LOCOMOTION STRATEGIES

In the physical world, locomotion is simple. We walk, run, roll, or ride to get from one place to another. Sometimes, we are in control of where we are going, and sometimes, someone (or something) else makes those decisions for us. In the virtual world, we are not constrained to these methods, and therefore there are a large variety of possibilities when it comes to transitioning our agency from one place to another.

One of the primary factors for choosing a locomotion technique is its tendency to cause VR sickness. Any kind of situation where the user's eyes see one thing, and their vestibular system feels another, can lead to feelings of nausea and dizziness, as described in more detail in Section 10.6. For this reason, many developers choose locomotion techniques that are less prone to causing VR sickness. Because some users have developed "VR legs" and are less susceptible to VR sickness, developers often add several locomotion techniques for the user to choose from, which can also lead to confusion and differential accessibility (see Section 8.5).

We begin by considering the locomotion methods that are inherited from traditional flat-screen games and experiences and are very different from walking in the real world. We refer to these techniques as **artificial locomotion**, where the user themselves remains largely stationary, in the physical world, and travel in the virtual world is invoked by some type of gesture or control interaction. These are in contrast to **physical locomotion (described later in this chapter**, where the user must move in some way in the physical world, for their avatar to move in the virtual world).

The advantages of artificial locomotion are that the user can stay in one place in the physical world, meaning the technique is more useful for small playspaces, and more generalizable for users with different playspaces or ambulatory abilities. Users can travel across large virtual spaces, and users do not experience the fatigue associated with physical locomotion like walking and running. Although artificial locomotion strategies tend to be easy to use, they are often unfamiliar at the start, which means training may be necessary, and interaction reminders may need to be available. Artificial locomotion tends to have lower plausibility than physical techniques (since artificial motion techniques in general do not exist in the physical world), and occasionally artificial locomotion can lead to coupling of tasks, since a controller or gesture may be used for both moving and interacting, or a movement may be tied to an interaction or vice versa.

Artificial strategies can be broken down into **continuous motion** methods, which allow the view to change smoothly as the user drifts through the virtual world, and **discrete motion** methods which make the user's view jump from location to location or direction to direction. Additionally, these methods include options for whether the user is in control of the motion or not and how the motion is explained in the narrative.

Inverse Locomotion

While the majority of artificial locomotion strategies are framed as moving the user's viewpoint within the virtual space, there are sometimes when it can make conceptual sense to imagine the space itself being moved, rather than the user's point of view. This is called **inverse locomotion** and can take many forms. In some cases, inverse locomotion can be perceptually identical to regular user-centric locomotion, and although the user is controlling the movement of the land rather than the movement of the avatar, they can't tell the difference (See Section 4.1.8). If the user can control more than just the position of the viewpoint, inverse locomotion becomes more useful as a concept. With the use of two hands, the position, orientation, and scale of the world can be modified. In some implementations, the environment can be continuously moved by alternating grab-and-pull actions with each hand, much like flick scrolling through a document on a touchscreen device. This can be effective for experiences involving data manipulation and visualization, or rapidly exploring areas of interest from arbitrary viewpoints, for example when issuing commands to units on a battlefield in a strategy game. Developers must be cautious, though, because the unique mechanics required for inverse locomotion can carry a higher learning curve.

10.3.1 Voluntary Continuous Motion (VCM)

The user's viewpoint moves smoothly through the virtual world, and the user is in direct control of that movement. VCM can be suitable for VR experiences where exploration of a vast environment is necessary, such as an open-world platformer, or for experiences where the user's movements must be precise, such as a first-person combat game. Advantages include ease of use and familiarity, especially when the input is a joystick, as this button mapping is likely similar to that of countless flat-screen video games the user may be familiar with. Disadvantages include moderate risk of inducing VR sickness, although the increased level control delegated to the user can mitigate this, and reduced biomechanical symmetry depending on the method of implementation. This may force certain tasks to be coupled by the same action, which decreases intuitiveness to the user. Some examples of VCM implementations are listed here:

Gaze-directed VCM. The direction of the user's gaze specifies the forward vector for motion. In other words, pushing the joystick forward will move the user in the direction in which they are currently looking. This should be combined with the ability to *strafe* or move along the left, right, and backward vectors, based on the position of the joystick. There is usually no need for the input device to control rotation or vertical translation, as this can be handled by any HMD with six degrees of freedom. Indeed, if rotation is controlled both with gaze direction and with the joystick, this can add to user confusion. Gaze-directed

VCM is a familiar metaphor for experienced gamers, especially if motion is constrained to the two horizontal axes, and it does not require any peripheral hardware. Its primary weakness is the fact that the user is unable to look in one direction and travel forward in another direction. As such, gaze-directed steering may not be suitable for tasks requiring movement relative to objects in the virtual environment.

Hand-directed VCM. The orientation of the user's hand or controller specifies the forward vector for motion. This can be more natural than gaze-directed VCM because the user can look around the virtual world while travelling in a different direction. The user's sense of proprioception can inform them on the direction in which their hand is pointing—even if the HMD occludes explicit vision of it. Hand-directed VCM mitigates the coupling issue that was described for gaze-directed VCM and is generally more flexible as well. Hand-directed VCM has the potential to induce strain on the user's hand, wrist, and shoulder due to prolonged stretching. The other challenge appears when a user must use an object while they are travelling: if the user attempts to throw an object or place it in a precise location, their travel direction may be affected, even if using the other hand. In an extreme example, the user may need to aim a tool or weapon with one hand and dictate the direction of travel with the other hand. Since both actions involve precisely extending and orienting an arm, the user's cognitive load may increase in such experiences.

Torso-directed VCM. The orientation of the user's torso specifies the forward vector for motion. Typically, a tracking device is attached to the user's waist, and the forward orientation of this device is relayed to the VR system. This decouples gaze direction and travel direction, while also leaving the user's hands free to manipulate objects, but requires an additional controller to detect the direction of the hips or torso. An alternative (called "Lean-directed VCM," see also "Human Joystick" in Section 10.4.3) allows the system to measure the direction in which the user is leaning, and use that to indicate the direction of travel. This can lead to instability in the user's stance and can make it difficult for the user to reorient their direction of view more than a few degrees while travelling.

Hardware-directed VCM. A steering prop is integrated into the virtual environment to specify the forward vector for motion. This typically occurs in flight simulators and other experiences that simulate vehicles controlled by hardware, but can also be used in other experiences with standard VR controller hardware. The direction of the user's motion is influenced by the manner in which the user operates the external device. Cockpit-based steering props can simulate driving a car, being at the helm of a ship, or flying an aircraft. Delegating the control of movement to a dedicated hardware device can add realism (since sensorimotor contingencies are more accurately supported) as well providing haptic and restive feedback to the user, depending on the hardware. Such hardware interfaces can be expensive, and the physical interface is well aligned with the virtual interface only in scenarios that are supported: driving a car with a hands-on throttle and stick (HOTAS) is likely to be non-intuitive, especially if the virtual avatar is holding on to a steering wheel.

It is worth noting that many VR developers have attempted to recreate such steering props in a purely virtual manner. These virtual steering devices may assist users in understanding the interface of a vehicle, but are generally much more difficult to control due to their lack of proprioceptive force feedback. Nonetheless, the inclusion of a realistic interface for a vehicle in virtual space is useful in preventing VR sickness, as the interface is a constant frame of reference for the user.

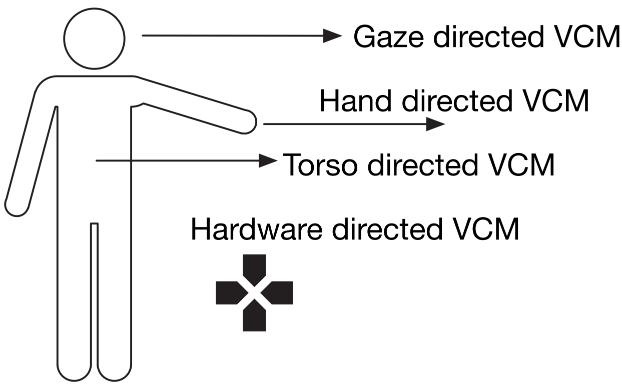

Figure 10.2: Options for controlling voluntary continuous motion.

10.3.2 Involuntary Continuous Motion (ICM)

Rather than the user controlling their movement, the user is passively moved through the virtual environment without the intervention of input or agency. Since there is no input from the user, this method is easy to implement and highly accessible. Although ICM can be thematically appropriate for cinematic experiences or experiences where the user is (for example) a passenger in a vehicle they are not controlling, this method has a high risk of inducing VR sickness, since not only is there conflict between visual and vestibular input, but the user cannot anticipate the motion, only react to it. Consistent linear ICM can be fairly comfortable or be adapted to quickly, but rapid changes in direction and high rates of speed can lead to the onset of VR sickness. Common mitigation techniques (discussed below) can help, including reduced field of vision and visual context cues.

VR Roller Coasters

When a VR enthusiast wants to convince one of their friends that VR is amazing, they will usually select a short, easy experience to show off the technology. Often, this experience is a virtual roller coaster. It seems ideal—the user does not have to do anything or be taught any controls; all they have to do is sit and experience the wonder. A track extends in front of them, and they ride around a fantastical coaster that snakes its way through alien landscapes or a dense jungle of skyscrapers. The VR novice sits down and puts on the headset and is amazed by the visuals, but as soon as the ride starts, they are very quickly nauseous. The VR enthusiast, of course, has developed a resistance to VR sickness over many hours of gameplay, but as a first experience, a roller coaster is almost guaranteed to induce VR sickness if the user is prone to it at all. It is somewhat ironic that what seems like a great first experience may in fact be discouraging people from experimenting further with VR.

10.3.3 Discrete Motion

While continuous motion is more conceptually straightforward, the fact that it often can lead to VR sickness has caused many developers to consider discrete techniques, where the user's viewpoint jumps from place to place. The most common implementation of this technique is **teleportation**, where the user specifies a destination within the virtual environment, via some input device, and is instantly or near-instantly transported to the new destination. Although initially developed to avoid continuous motion that leads to VR sickness, teleportation can be an effective locomotion technique in itself, especially when considering the challenges of traversing vast virtual environments.

There are two design decisions to be made when implementing teleportation: how to choose a destination and how to transition to the destination. Selecting a destination can be done using any pointing technique (see Section 11.4.1.2), and teleportation targets can either be individual points on the landscape or the entire surface in front of the user. Targets may also be range limited, if it is important for the situation that users not be able to travel too far. An alternative could be to select the destination by choosing a direction and a distance or even to select a direction and then to jump a standardized distance in that direction. Some options for destination selection are shown in Figure 10.3.

Usually, the orientation of the user's view is maintained between the origin and the destination of the transition; however in some cases it may be useful to allow the user to reorient as they transition. Jumping behind an enemy is not as useful if you then have to turn around to face them, but if you can turn around as you are jumping, you can be both in the right place and facing the right direction. Figure 10.4 shows the difference between teleportation with and without a rotation.

Once a destination has been chosen, the user's viewpoint is then shifted to the new location. This shift in view can itself be discrete or continuous. The simplest and quickest transition is to immediately change the view with no indication or delay, often called a "snap change." This works well for rotating in place (a snap turn), but can be very disorienting for teleportation. A "blink" is a slightly delayed transition, where the user's view is momentarily occluded, often by fading out and back in again. Allowing the user to maintain their view during the transition can significantly reduce

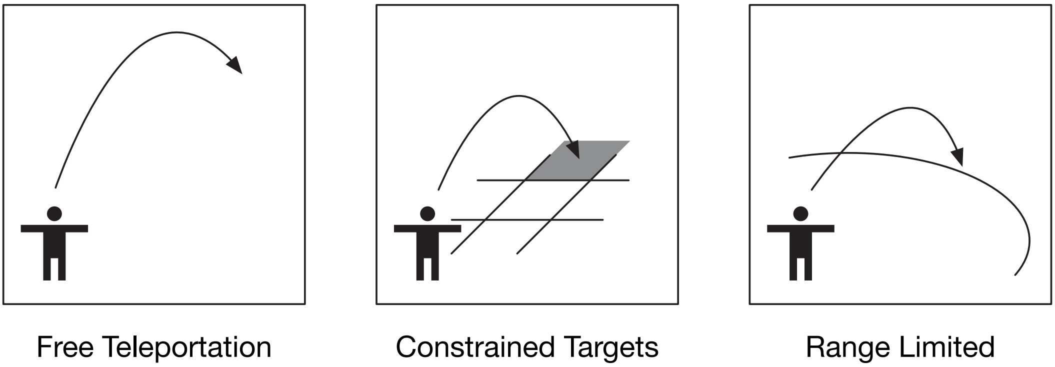

Figure 10.3: Options for selecting a destination for teleportation.

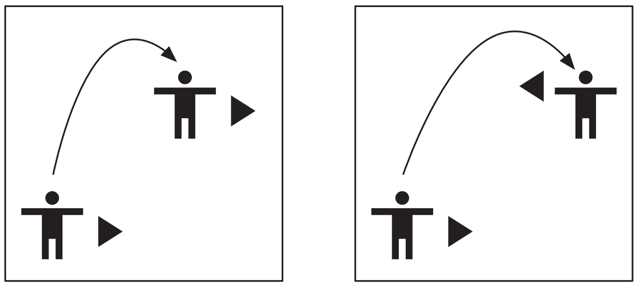

Figure 10.4: Teleportation without rotation (left) is restrictive, but teleportation with rotation (right) requires additional control and can be disorienting.

disorientation, because the user can see where they were and where they are going; however, this method becomes very similar to ICM and can lead to VR sickness. A faster transition, with smooth changes, can mitigate these challenges. Another clever solution to the problem of transitioning the user's viewpoint is **portalling**, where the user can see the view as it will be from the new location as they prepare to teleport. Portalling sometimes requires a surface to transition through, with a portal on the surface revealing the new location behind it; alternatively, a spherical portal can be placed showing the viewpoint of the new location from any direction around the portal. Portalling removes the requirement for transitioning from the old location to the new location, but can add complexity to both implementation and interaction.

While teleportation can be a versatile and comfortable alternative to continuous motion, it comes with a few drawbacks. It can be difficult for a user to accurately control the destination of a teleportation, especially if the target is far away (See Fitt's law, Figure 11.24), and it can be disorientating, especially if the method of transitioning the viewpoint is not carefully considered. Teleportation can also lead to reduced sense of presence, since teleportation is unlike any method of locomotion used in the real world. It can be seen by the user as implausible and takes away from the feeling of being in a realistic world. Many developers make a special effort to justify the existence of a teleportation system in-world, in order to allow users to suspend their disbelief. This can be achieved by giving the user a teleportation tool, like a handheld beaming device or a glove; however, such a justification can add additional design criteria which may conflict with the usability and comfort criteria described above.

10.4 PHYSICAL LOCOMOTION STRATEGIES

Physical locomotion encompasses methods of travel in which the movement of the avatar within the virtual world is directly tied to movement of the player in the physical world. In some cases, this can be as simple as allowing the user to walk around and have that change in location reflected in the virtual space; this can also be as complicated as full body tracking to calculate the speed and direction the avatar should move based on how the user is moving in the physical world.

Physical locomotion strategies are both highly intuitive and highly plausible, since the user is employing very familiar actions already conceptually aligned with getting from place to place. Additionally, physical locomotion strategies are (usually) much less prone to VR sickness because stimulation of the vestibular system is more closely aligned with what the user sees as they travel. These methods are constrained, however, by the physicality of the user, and therefore can be less versatile and accessible and contribute to fatigue or balance issues. Additionally, if walking in VR is tied to walking in the real world, the range of travel is bounded by the size of the playspace. Some physical locomotion strategies require expensive or bulky hardware that may be impractical from a consumer perspective, and these methods tend to exclude to non-ambulatory users.

Physical locomotion strategies can be categorized as **gait-direct** techniques, in which the user walks around a physical space and that motion is translated into motion in the virtual space; **gait negation** techniques, in which additional hardware like treadmills or low-friction surfaces allow the user to walk somewhat naturally in-place; and **indirect** techniques, where some other aspect of the user's motion in the physical world is translated into movement in the virtual world.

These strategies can also be considered based on their ability to mitigate VR sickness. In general, any technique that involves physical motion will lead to a reduction in the conflict between the vestibular and ocular systems, but the more closely aligned the motion of the user, the lower the likelihood of experiencing VR sickness. Gait-direct techniques tend to mitigate VR sickness most completely, because the motion of the user's physical body is most closely aligned with the changes in the user's vision as a result of motion in the virtual world. Gait negation techniques can also work well, but are highly dependent on the responsiveness and implementation of the negation technique itself. Indirect methods are highly variable, depending on the techniques as described below. All physical motion techniques are effective in mitigating VR sickness to the extent that the technique fulfils these criteria:

- The user is in direct control of the direction and speed of travel.
- The physical motion matches the virtual motion.
- The locomotion is perceived as natural and plausible.

Plausibly is important here because the user will be more comfortable with a particular technique to the extent that they can use it without thinking about it, and therefore their motions will become more natural, and hence more closely aligned with their expectations in the virtual space.

Each technique has benefits and drawbacks, and developers may benefit from considering modifications to their planned experiences to enable the use of more natural techniques. For example, if a virtual space were the same shape and size as the physical space a user is in, gait-direct techniques may be applicable in all circumstances, and the user's experience would greatly benefit. Not all users share the same size playspace, however, and thus it is usually appropriate to implement a variety or combination of techniques.

10.4.1 Gait-Direct Implementations

The most natural and comfortable methods of locomotion are the methods that are most familiar. Most people walk to get places, and walking creates motion in the user's headset that is easily translated to well-aligned changes in the visual field of the virtual environment. Walking also requires no training, no controller, no decision-making, and no conscious effort. We must remember that not all users are ambulatory, and some may have challenges walking for long periods of time, but in general, walking is close to an optimal locomotion technique for VR experiences, assuming the user has the space for it.

Room-Scale Walking. The user walks through the virtual environment. Real-world displacement is directly mapped to virtual displacement, which unfortunately means that the locomotive range of room-scale walking is limited to the user's play area. As well, room-scale walking is less accessible to users for whom walking may be a problem, than most artificial strategies. Room-scale walking is one of the few strategies that appropriately stimulates the vestibular system, meaning there is little to no risk of VR sickness. The walking metaphor is also advantaged with respect to intuitiveness, accuracy, sense of bearing, and presence. On its own, room-scale walking is best suited for small virtual environments; however, if combined with a reorientation strategy, it may be a viable choice for larger environments as well.

Modified Walking. The user walks around their playspace as in room-scale walking, but their movement in virtual space is not directly aligned with their movement in the physical world. Small changes to the user's perception of virtual space encourage them to stay within the bounds of their physical space. There are a wide variety of redirected walking techniques, some examples of which are listed here:

Scaled walking. The user's motion in virtual space is different from their motion in physical space. Usually, small steps in the physical world are translated to larger steps in the virtual world, which can increase the range of travel but reduce plausibility

Redirected walking. The user's direction in virtual space is different from their direction in physical space. A user may walk in a straight line in virtual space, but the virtual view is subtly changed to encourage the user to change the direction they are walking in the physical world, orbiting around the centre of their physical playspace. Users quickly lose directional reference while in VR and can be convincingly redirected by the virtual space without losing plausibility.

Figure 10.5 shows an example of redirected walking. A user is presented with a virtual world, and they begin to walk forward. The virtual world shifts its orientation very slowly and subtly, requiring the user to turn their head in the other direction in order to continue walking in the direction they perceive as forward. The result is that the user is redirected in a curving path around their room, while in the virtual space, they feel as though they are walking in a straight line beyond what their physical space would allow. Because the changes to orientation are gradual and subtle, the directional sensing of the vestibular system has low fidelity compared to the visual system, and thus the body can be tricked into changing direction.

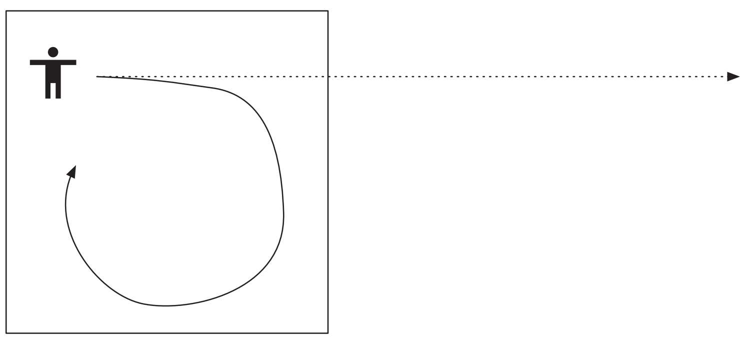

Figure 10.5: Redirected walking.

Contextual Redirection. The user is prevented from walking beyond the physical bounds of their playspace by events or objects in the virtual world. These events can then distract the user while the virtual world shifts subtly around them, allowing them to continue to walk after the event, allowing the virtual world to seem larger than the physical playspace.

Portalling. The user walks to a doorway which opens into another room, but the doorway encourages them to turn around, thus expanding the size of the virtual space while continuing to restrict motion within the physical space.

Walking in Place. The user moves their feet to simulate walking while remaining in the same location in their playspace. There can be some disagreement as to whether this is a gait-direct technique or an indirect technique, since the user isn't actually walking, and depending on the implementation it may even be considered a gait negation technique, since the user is walking but not moving. Regardless of the classification, the technique is popular because it is almost as natural as room-scale walking, without requiring a restriction or modification of the virtual space. Walking in place can be difficult to implement, however, since most VR equipment does not track the feet. If a user is considering adding foot tracking to implement walking in place, a more comprehensive gait negation technique could also be used without too much extra effort. Without additional hardware, walking in place can be identified by the up-and-down movement of the headset, and a more aggressive walking technique (like marching) can be encouraged to improve the reliability of this tracking. The direction of motion must be manually determined, since walking in place does not result in any actual motion. Similarly to VCM, the direction of motion can be dictated by the user's gaze, torso orientation, pointing direction, or controller input.

10.4.2 Gait Negation Implementations

When a user walks, the natural human gait results in forward motion. Walking in place restricts that motion but changes the user's gait. Gait negation[2] is any technique which attempts to remove (or *negate*) the forward motion of a user's walking without unnecessarily modifying the user's gait. Gait negation usually requires specialized hardware that allows a user's feet to travel past the ground rather than pushing against the ground. The most common and familiar hardware that accomplishes gait negation is the **treadmill**.

Simple Treadmills. The simplest way to negate forward motion of the human stride is to use a treadmill. Basic treadmill locomotion in VR is straightforward to implement, but has significant limitations. First, a treadmill allows forward motion but does not provide any mechanism for direction, orientation, strafing, or other non-forward motion, and so the developer would need to apply techniques from continuous motion to provide direction. Second, the treadmill controls the speed of walking, rather than the user, and so this basic treadmill locomotion suffers from poor plausibility. The user must manually adjust the speed of the treadmill to walk faster or slower. Treadmills are more likely to be found in users' homes, however, and walking on a treadmill while using a VR experience can improve cardiovascular health.

Alternative Treadmills. Treadmills are most commonly found in exercise gyms or in home workout areas, but other machines that simulate motion can also be found there, including rowing machines, stationary bicycles, and stair climbers. Any of these devices could be connected to a virtual experience to improve the experience of exercise or to improve the experience of travel within a VR scenario. Additional features can be added to such treadmills (and to simple

[2] It is not really the gait that is being negated, rather the motion resulting from the gait. The term is somewhat awkward, but it is in popular use and has a nice verbal symmetry to it, so we will continue to use it.

treadmills to further improve plausibility), such as a fan to simulate the feeling of wind as you cycle through France.

Omnidirectional Treadmills. Simple treadmills only provide forward motion, but in a VR scenario, it is reasonable to expect a user to be able to walk forward, backward, left, and right. An omnidirectional treadmill allows motion along both horizontal axes. Such treadmills may be passive, relying on the user's mass and momentum to move, or active, being directly controlled by the VR system which moves the surface of the treadmill whenever a change in the user's direction is observed. Passive omnidirectional treadmills are typically constructed with a low-friction surface and concave structure and may require the user to don special low-friction footwear, as well as motion trackers on their feet. These treadmills are often called "socks-in-a-bowl" and also require a frame and gantry to support the user as they lean in whatever direction they are virtually travelling.

Active omnidirectional treadmills look more like simple treadmills, but have a surface that can move left and right as well as forward and backward. These treadmills attempt to anticipate the motion of the user and move the surface accordingly, but often suffer from latency and balance issues, as adjustment of the treadmill's surface is a mechanical process. A harness or safety railing is often present to prevent users from losing their balance.

In general, omnidirectional treadmills demonstrate promise in the area of enabling unrestricted "walking" in virtual environments. Such devices enable a walking cycle that is very similar to our natural gait and, depending on the durability and construction, can support running, crouching, and jumping as well. These devices are typically large and expensive, require a dedicated space, and may be restrictive on who can reasonably make use of them.

10.4.3 Indirect Implementations

While gait-relative locomotion techniques attempt to replicate human walk cycles in some form or other, indirect locomotion techniques make use of other motions in the physical world to create travel in the virtual world. Almost any trackable physical motion can be translated to virtual locomotion, but some common techniques are presented here.

Human Joystick. As the name implies, this technique reinterprets the users' upright body as a joystick, and the user can lean or shift in one direction or another to invoke travel in that direction. The system keeps track of the relative position of the user's head, using this information to control horizontal direction. The more the user leans or shifts, the faster the invoked travel. This technique is compatible with small play spaces and can work even when the user is seated, increasing accessibility. While the natural mapping of direction of lean to direction of travel can increase presence, this method is prone to VR sickness, since the user sees motion but does not feel it. User agency reduces

(a) A passive omnidirectional treadmill supports the user's centre of mass. *Maurizio Pesce, CCBY2.0*

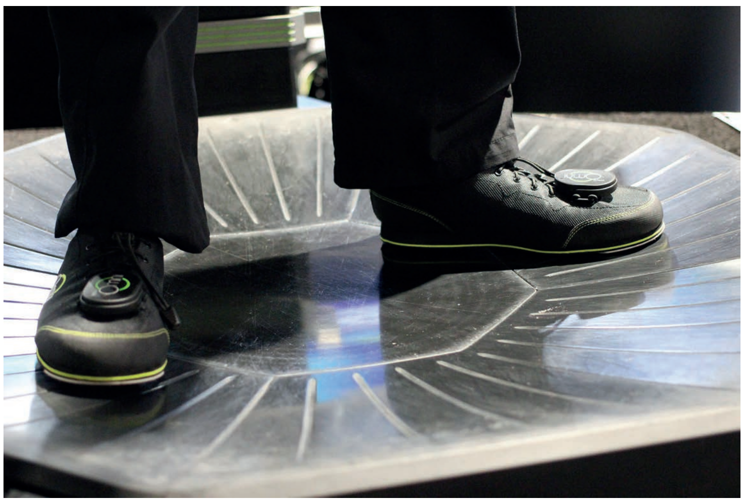

(b) Users wear low-friction footwear with motion trackers. *Maurizio Pesce, CCBY2.0*

Figure 10.6: Passive omnidirectional treadmill.

VR sickness somewhat, but other mitigation techniques as described below may be necessary. Additionally, if the user takes a step rather than just leaning, the VR sickness effect can be reduced. A challenge with this method is that for the user to stop, they must return to the location the system considers the "centre" of the "joystick." Hardware joysticks use springs to automatically centre when released, and if a hardware joystick malfunctions and doesn't return to the centre, the system experiences drift. With a human joystick, there is no automatic centring, so the user must seek the centre manually; to make matters worse, people usually perform such local navigation tasks visually, and the whole point of this system is to move the user's viewpoint to another area. Finding the centre again can be very difficult, so a large dead zone may be necessary; otherwise users may experience involuntary motion as a result.

Arm Cycling. Related to walking in place, this technique replicates the upper body motion of walking without requiring the feet to move. Hands are commonly tracked in modern VR systems, while feet are not, and so arm cycling can be accomplished with standard hardware rather than requiring additional third-party equipment. Forward motion is applied to the user's viewpoint when the user swings their arms, moving the VR system's tracked controllers (or employing the system's hand tracking). Turning can be accomplished using any of the techniques described for VCM. Alternatively, the direction of travel can be indicated by the vertical displacement between the two controllers at the apex of motion, or the difference in amplitude or speed between the motion of the two controllers.

Arm cycling is quick to learn and easy to use and is surprisingly effective at mitigating VR sickness, since the user is not only in control of their motion but also experiences some vestibular displacement as a result of the arm motion—users naturally tend to bounce up and down a bit when they swing their arms. As with other physical methods, the increased physical effort can improve cardiovascular health but can also lead to fatigue and be a barrier for persons with physical disabilities, but this technique can also work well when seated, increasing accessibility. One additional problem with arm cycling is that it can restrict the utility of other interactions a user might need to do while they are travelling. It's hard to aim a tool or a weapon or activate a switch or look at a wrist display, while you are swinging your arms about. One-arm swinging may need to be temporarily invoked, which may change the utility of different steering mechanisms

Non-Walking Locomotion

Although most locomotion techniques discussed here replicate the process of moving horizontally across a landscape, many experiences have been developed that use artificial or physical locomotion strategies to move in different ways. Rock climbing simulators allow a user to reach up with one hand, grab the side of the wall, and pull themselves up to the next handhold. Kayaking and rowing simulators make use of paddles or oars to propel a boat along a waterway. Other types of unusual locomotion involve control of a vehicle like a sailboat or a parachute, which require detailed hand motions that serve as controls to a motion technique, one step removed from actual locomotion. A person may swing through the jungle on vines, or through the city on spiderwebs, or skateboard along a street or snowboard down a mountain. In each of these cases, specific actions by the user serve as direct or indirect control to the user's viewpoint moving through the virtual space, and whenever the viewpoint moves through the space, consideration must be made for usability and to mitigate VR sickness.

10.5 QUANTIFYING LOCOMOTIVE EFFICACY

With a diverse collection of locomotion techniques available, it is worthwhile to consider how a developer might choose between them. In the descriptions above, we have noted some advantages and disadvantages of each, but what follows is a collection of metrics that can be used to evaluate these techniques and others based on the constraints of the specific VR experience being considered for a locomotion mechanic.

10.5.1 Locomotive Range and Scope

In the previous section, we explained that the major drawback of room-scale walking and walking-adjacent travel in VR experiences is that its range is limited to the size of the user's play space, which results in two common problems. First, users sometimes overstep the boundaries of their play area when walking through a virtual environment, which causes presence breaks at minimum and potential injury or property damage at worst. Secondly, long-distanced travel is impossible without the inclusion of a reorientation strategy or another locomotion method. This limits the scope of virtual environments that only afford room-scale walking, as the user's overall experience may suffer if there is a disconnect between what is visible and what is accessible.

Visible range is the area of a virtual environment that the user can see (or otherwise sense), while **locomotive range** is the area of a virtual environment that the user can access via a particular locomotion strategy. Locomotive range varies depending on the size of the environment, but it is not dependent on the speed or control of the locomotion method in question, only by whether or not you can get there. Indoor virtual environments have small ranges (both visual and locomotive), while exterior experiences have much larger ranges. In an ideal implementation, you can go anywhere you can see, and visible range would equal locomotive range. Locomotive range might be restricted by the locomotion technique being considered; however, there are also conceptual or narrative reasons that may limit the range of where a player can travel in a world. It may not be optimal to allow a player to travel to the top of a mountain or across the sea, and indeed it may not be practical to implement accessibility to all regions that can be seen. On the other hand, **open-world** games are specifically designed to allow a player to go almost anywhere they can see, and limitations are usually narrative or diegetic, rather than mechanical—if a character has no climbing equipment, they cannot climb the large wall that is in front of them, and instead they must find another way up or around the mountain. Additionally, users may expect that every object they see be interactable, and this may be a similar challenge to the developer—either every object the user sees can be manipulated or manipulable objects need to be signified in some way. In the same way, either all locations the user can see should be reachable or areas the user can't reach should be signified in some way.

The **locomotive scope** of a strategy is the ratio of its locomotive range to the environment's visual range. This is a measure of the difference between what a user can see and where a user can go. A smaller locomotive scope means the user can see lots of areas they can't go to, which can be frustrating unless explained diegetically. When a person looks up at the sky, they generally don't assume they can visit the stars; however, if they are playing a space simulator, it might be more reasonable to assume that they can visit the stars they can see. Similarly, if the background environment of a simulation shows rolling hills, but it is made clear to the user that they can only visit locations that are within a set boundary, this is not a reduction in locomotive scope.

Locomotive Scope Example

Imagine a VR experience where the user is inside a large room, with interesting things on each wall. The experience supports room-scale walking, scaled walking, and teleportation. If the virtual room were, say, 4 metres by 5 metres, the visible range of the experience would be $4 \times 5 = 20$ square metres. If the user's physical playspace for VR was, say, 2 metres on each side, room-scale walking would give them a locomotive range of $2 \times 2 = 4$ square metres. The locomotive scope of room-scale walking would therefore be $\frac{4}{20} = \frac{1}{5}$ or 20%, meaning that a user could only experience a fifth of the room that they could see if they were walking; depending on where they started, they may not be able to access any of the interesting things on the walls at all. For scaled walking, if every step the user took in the physical world was stretched two-to-one in the virtual world, the user would be able to reach $4 \times 4 = 16$ square metres, resulting in a locomotive scope of $\frac{16}{20} = \frac{4}{5}$ or 80%. The user could get almost anywhere in the room, but there would still be a fifth of the room they couldn't get to. Finally, teleportation would provide direct access to the whole room, for a locomotive scope of 100%. The trade-offs with these methods must be carefully considered, and it may be most appropriate to implement a hybrid approach, where the user could teleport to a corner of the room and then walk around interacting with things and then teleport back to the far side of the room when necessary. Figure 10.7 shows the different locomotive scopes for each option.

Thus, based on the principles of locomotive range and scope, gait-direct techniques have a disadvantage over other techniques, and artificial techniques tend to be favoured when range is an issue, as they alleviate the presence-breaking disconnect that can occur between visible range and locomotive range. The trade-off, however, is that locomotion methods that are not one-to-one have the potential to induce VR sickness. Nonetheless, some gait-relative techniques like redirected walking have the potential to allow increased scope while still maintaining the high degree of presence that walking allows.

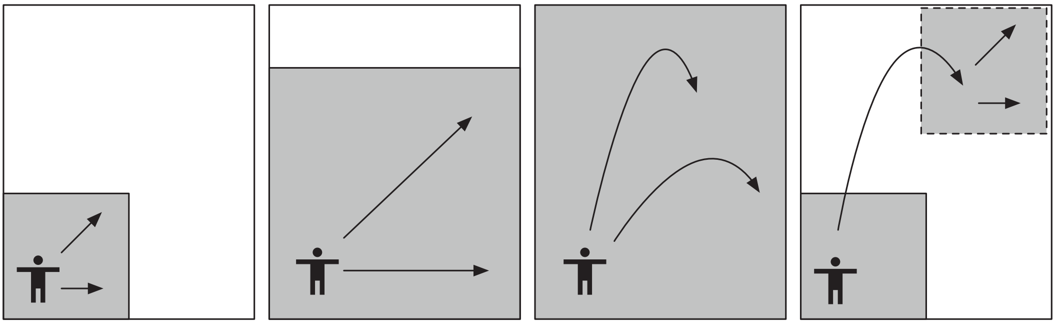

Figure 10.7: An example showing locomotive scope for different travel techniques.

10.5.2 Locomotive Utility

Locomotive utility refers to the overall suitability of a locomotion strategy for any particular VR experience. It can be represented as a function of performance-based characteristics that are indicative of the strategy's effectiveness and usefulness, but not directly tied to usability considerations like ergonomics or the prevention of VR sickness. The key characteristics of locomotive utility are as follows:

Range, described above, measures the extent of the virtual environment that can be accessed by a locomotion strategy. The locomotive range of gait-related and indirect physical strategies is directly dependent on the size of the play area, while boundaries usually need to be established on the backend of the experience for the other subcategories. In limited environments, all strategies are likely to have full locomotive range. In environments whose size exceeds that of the play space, only room-scale walking will have insufficient range, as scaling and redirection methods are likely able to alleviate the space disparity.

Velocity relates to how quickly travel is completed after it is invoked through a certain locomotion strategy expressed as a function of virtual displacement over time. Artificial strategies generally possess higher velocities than physical strategies, with discrete strategies being the fastest, as motion occurs near-instantaneously, and indirect physical strategies usually being the slowest, as their required motions are less intuitive than those of gait-related strategies.

Fidelity expresses a locomotion strategy's ability to perform small increments of travel. It is a contributor to accuracy, since small displacements enable precise error correction without affecting the user's sense of bearing. Room-scale walking has the highest fidelity, as displacement and rotation are mapped from the real world to the virtual environment one-to-one. Continuous artificial strategies generally possess high fidelity as well, while modified walking strategies may perform slightly worse depending on the sensitivity of the scaling or redirection algorithm. Non-continuous artificial strategies such as teleportation have the lowest fidelity and, thus, the least reliable accuracy.

Adjustability describes the extent to which the various components of locomotion can be tuned by the user, especially while they are in use. Factors influencing a locomotion strategy's adjustability include the ability to change one's position while rotating, the ability to change one's orientation while translating, the possible speeds of translation and rotation, and whether termination of travel occurs abruptly or not. In contrast to the previous three factors of utility, it is difficult to establish general principles of adjustability for each locomotive subclass. Instead, not only is adjustability strategy dependent, in some cases, it may even be implementation dependent. For example, all of the techniques described within VCM enable a variety of travel speeds; gaze-directed VCM couples travel direction and the direction the user is looking, while torso-directed VCM is more cumbersome to terminate than the hardware-controlled implementations. Thus, gaze- and torso-directed VCM have lower adjustability than the other types of VCM. Similarly, due to the latency issues present in active omnidirectional treadmills, their adjustability is lower than their passive counterparts. Furthermore, implementations of teleportation that enable explicit destination and orientation selection have higher adjustability than those with only explicit destination selection.

The Velocity-Fidelity Continuum

A locomotion strategy can possess both high range and high velocity, or high range and high fidelity; however, high velocity and high fidelity are almost always mutually exclusive (see Fitt's law, Figure 11.24). Think of velocity and fidelity as components of a microscope. Velocity enables swift travel to the approximate area of one's destination, much like the coarse focus of a microscope is used to quickly bring an image into near focus. Meanwhile, fidelity facilitates precise correction of the subtle inaccuracies of high-velocity locomotion, in the same way that the fine focus accomplishes the necessary adjustments to bring an image into full focus. In order to ensure that both high-velocity and high-fidelity travel are possible, especially in vast environments, it may be necessary to include multiple locomotion strategies in the same experience.

Sometimes, certain utility factors are more desirable than others. For example, if exploration and searching tasks are necessary for an experience, then high range and velocity are essential. Similarly, experiences with reaction-based actions, like dodging attacks or defeating waves of enemies, are most enjoyable if their locomotion strategies feature high fidelity and adjustability. Although trade-offs between the characteristics are usually necessary to some degree, it is still good practice, as much as possible, to design for a balance of range, velocity, fidelity, and adjustability, as such versatility usually results in a higher degree of control for the user. This, in turn, can reinforce the plausibility of a virtual environment and also improve the usability characteristics of the locomotion strategy.

10.5.3 Travel and Presence

Presence is maximized when the user interprets the virtual environment as a real location (place illusion), and interactions within the experience as real events happening to them (plausibility illusion). A major component of interpreting an environment as a real place is the ability to move through it in a realistic manner. In the context of sensorimotor contingencies, realistic movement may involve any of the following actions while the user is fixed in the same location in an environment:

Turning one's head to look at a different part of the environment

Strafing one's head to cause parallax to distinguish foreground stimuli from background stimuli

Crouching and craning to look around and behind objects

In addition, the user may choose to move their body in order to take up a new position to get a better view of an object or stimuli in the environment. When the goal of the travel is to collect more information, this aspect of travel is also a sensorimotor contingency.

All of these sensorimotor contingencies have three things in common: they are invoked by physical movements; changes to one's bearing are observed continuously; and the user has full control over the speed and magnitude at which displacement occurs. Unsurprisingly, these aspects pertain to each of the sensory illusions that contribute to presence. Continuous observation of positional and rotational updates based on the user's physical movements represents validation of sensorimotor contingencies, which strengthens place illusion.

Room-scale walking, as a locomotion strategy, will maximize presence, since it corresponds directly to the user's experience of movement in the real world. Redirected walking strategies can also benefit presence if the redirection or scaling is subtle by design; for example, larger arcs are better at masking the fact that redirection is occurring than smaller arcs. Even more extreme methods like non-Euclidean geometries (Section 6.4) may also align with the user's experience of real-world walking, depending on the implementation, since although the virtual environment may not behave realistically, the act of moving through the environment does. Gait negating strategies still perform well with respect to plausibility, since although the user's vestibular system is not stimulated in a way that aligns directly with their motion in the virtual space, the motions are similar, and users can quickly get used to the differences.

Artificial strategies are widely considered to be the least realistic VR locomotion strategies, as they are not invoked by natural motions, there is no synchronization between real and virtual displacement, and vestibular senses are not stimulated at all. Nonetheless, if all other factors point to an artificial strategy being suitable, then continuous methods lead to increased presence compared to discrete methods, even though they can be more prone to causing VR sickness. Low realism is teleportation's greatest weakness as a locomotion strategy, as it performs well in the areas of accessibility, speed, range, adjustability, and VR sickness mitigation. The plausibility of teleportation as a *mechanic* can be circumvented by providing some form of diegesis or in-world justification for it, but as a *locomotion method*, it is often still difficult for users to grasp as reality.

10.5.4 Combining Locomotion Strategies

As mentioned previously, implementing different locomotion strategies and allowing the user to combine them can increase the range and adjustability of the experience's overarching locomotion, although it is worthwhile to consider establishing a unified interaction metaphor for invoking each strategy. If the goal is to enhance locomotive range without sacrificing VR sickness mitigation or presence, for example, then principles of room-scale walking, scaled walking, and gaze-based VCM could be combined. In this hypothetical system, the user's real displacement would only be scaled in virtual space if a particular button is held, which has potential to alleviate VR sickness because it provides extra control. Furthermore, presence can be maintained if the scaling mechanic is explained in the game's narrative—perhaps the user is travelling along a conveyor belt that they can turn on or off whenever they please, for example.

Teleportation alone usually results in the user's location changing, while all other aspects (rotation, gaze/view, status) are maintained. If a user needs to be facing a different direction when they arrive, it may be appropriate to combine teleportation with some form of rotation. Preserving the user's original orientation may assist in maintaining a sense of bearing, but requiring the user to look around before choosing their next teleport can extend the process of wayfinding. To implement this combined method, the user needs to be able to control two variables simultaneously: the location of the destination of the teleport and the rotation of the final view as compared to the current view. One possible mapping might be to have the user press the thumbstick to invoke the teleport targeting system, aim the destination by pointing with the controller, and set the destination rotation using the thumbstick. The teleport itself could then be activated by releasing the thumbstick. Such a system would eliminate the synchronization issue and also allow the user to rely on artificial actions for both translation and rotation. In teleportation systems lacking rotation, the user must break a single action into two pieces: teleport and then rotate (either physically or with a controller), which can increase cognitive load.

Combining locomotion strategies may also serve to simplify the interaction and goal forming of the user, although the specific controls may become more complex as a result. There are very few real-world scenarios in which the user will be rapidly shifting from physical motion to passive travel. As such, forcing the user to maintain and switch between multiple locomotion strategies can cause confusion and frustration.

10.5.5 Multi-Modal Locomotion Systems

Different users may have different experiences of locomotion. Some users are particularly affected by the physiological effects of motion in VR, while others may appreciate a more rapid and diegetic locomotion technique. For this reason, **multi-modal locomotion systems**—systems in which two or more locomotion strategies are implemented and can be evoked simultaneously—enable greater creativity, customization, and expression within the user's movement. There are two main types of multi-modal locomotion systems:

Complementary locomotion systems are systems whose locomotion strategies are intended to resolve each other's limitations or deficiencies and are used in different situations. For example, a complementary locomotion system might include teleportation, which has high range but low fidelity, and room-scale walking, which has high fidelity but low range. This system is designed such that users will teleport for long-ranged exploration and search tasks, but walk for short-ranged manoeuvring tasks. The idea is that the deficiencies of one method are complemented by the benefits of the other method. Another example might be implementation of two different teleportation strategies, one to "nudge" the user locally and the other to allow long-distance jumps.

In complementary systems, the locomotion strategies usually map to mutually exclusive use cases, and therefore the intuitiveness of the combination of techniques is not as important. Nonetheless, design recommendations for complementary locomotion systems include simplifying the transitions between the different strategies, establishing a common interaction mechanism whenever possible, and diegetically justifying the presence of any seemingly non-realistic strategy—perhaps even going as far as explaining the connection of two dissimilar strategies in the experience's narrative.

Supplementary locomotion systems are systems where two or more locomotion strategies are implemented which serve largely the same purpose. In these cases, the choice of which technique to use may depend on user preference or other factors. One strategy may be preferred over the other(s) for certain types of tasks and manoeuvres. One common supplementary locomotion system is the combination of some form of VCM with teleportation. Both of these strategies have good range and reasonable adjustability, which makes them best suited for exploration and search tasks, but VCM has higher fidelity and affects presence more positively, while teleportation has higher velocity and better VR sickness mitigation. As such, the user may benefit from alternating between these two strategies if the experience suggests different types of travel tasks. Consider a VR dungeon crawler, for example, in which the user travels down labyrinths, solves puzzles, and fights enemies. For tasks that solely involve moving through the environment, the user may prefer teleportation because it is faster than VCM. However, teleportation may be undesirable for combat, since it can cause the user to lose their sense of bearing with respect to the enemy. In these situations, VCM would likely be more suitable.

An important factor to consider when implementing supplementary systems is that the user may initially struggle to understand when a certain strategy is preferred. If this occurs, then a diegetic indication of the preferred strategy may be of assistance. Of course, simply forcing the user to use a particular strategy is an option as well, but we consider this to be poor practice, as it defeats the purpose of multi-modal locomotion systems in the first place. As well, it is better to allow the user to seamlessly switch between any of the system's strategies, as forcing them to pause their session and toggle settings in a menu are both cumbersome and presence breaking. More guidelines on designing for multi-modal locomotion, including specific considerations for when the user should be using a certain strategy, can be found in Section REF.

A final consideration in multi-modal systems is that of equity in competitive gameplay. If two supplementary options are offered to allow a user to choose fidelity over comfort, for example, it may be the case that one locomotion strategy has strategic benefit over the other. In combat scenarios, both teleportation and VCM have benefits and drawbacks, and it may be difficult to balance the advantages given to one over the other. Teleportation, especially combined with rotation, will allow users to sneak behind their opponent, and so a slight time delay may be appropriate for balance. Alternatively, different leagues or leaderboards could be implemented based on which locomotion strategy is used. When multiplayer games become competitive, there is an incentive to gain any edge, and a responsibility for balance on the part of the developer.

10.6 PHYSIOLOGICAL EFFECTS OF MOVING IN VR

As has been mentioned several times throughout this chapter, one of the key challenges with locomotion in VR is that sometimes, when people move in VR, they feel sick. This section describes the likely causes of this effect as well as techniques for mitigation.

10.6.1 VR Sickness As a Result of Locomotion

In Section 4.1.8, we introduced the concept of *vection*, where a user feels the sensation of motion strictly based on visual input. Although vection can lead to uncomfortable symptoms of VR sickness (described below), vection can be beneficial for the user's experience in a VR environment, since making a user feel like they are moving can enhance presence if implemented well. A good implementation of vection typically involves synthesis of ocular, proprioceptive, and vestibular stimuli—in other words, the user sees their perspective change, infers new positions of their body parts, and perceives some type of motion via the inner ear all at the same time.

With gait-direct locomotion techniques, this sensory synthesis is essentially automatic since the viewpoint the user sees is changed relative to the position of the head as they move around. Most modern VR hardware tracks the head position sufficiently that this happens without input from the developer, although it should be noted that frame rate and refresh rate are critical to maintaining this illusion. If, as a developer, you create an experience so packed with detail that the headset fails to keep up and the frame rate starts to drop, this in itself can lead to problems with vection.

In all other circumstances, however, there will be a disparity between movement in the physical world and movement in the virtual environment. This disparity can be particularly problematic for artificial strategies, since the vestibular system senses no real-world movement that might corroborate the virtual travel that these strategies invoke. Even in gait negation and indirect physical strategies, although there is real-world motion, that motion is not aligned with the motion that would directly correspond to what the user is seeing.

For example, suppose that the user is playing through a VR experience with gaze-based VCM. If they invoke the relevant input for forward motion, then the environment is likely to provide realistic optic flow stimuli for the eyes. Consequently, the ocular system perceives that motion is occurring. However, since the user is not moving in the real world, the inner ear does not perceive motion. This sensory mismatch often leads to **VR sickness**, which has similar symptoms to motion sickness or seasickness, but actually represents the opposite physiological effect. In motion sickness, the vestibular system detects motion that the eyes do not see, usually as a result of visual obstruction—people who get carsick feel it more intensely when they are reading a book in the car, and it can be alleviated by looking out the window. Similarly, seasickness is experienced most acutely when belowdecks, and the symptoms may fade if the person comes up on deck to look at the horizon.

The physiological symptoms that VR sickness shares with motion sickness include headache, nausea, disorientation, vomiting, and perspiration. Some of its psychological symptoms that are not commonly observed in traditional motion sickness include dejection (feeling depressed), confusion, apathy, and a desire for fresh air. These symptoms are detrimental to a VR user's experience and usually result in the user quitting the experience (and sometimes quitting VR entirely); therefore, preventative measures for VR sickness are necessary when designing a locomotion system. At the same time, users who have been using VR for a long time and are familiar with locomotion can, over time, develop "VR legs" (analogous to "sea legs"), rarely experiencing VR sickness and even feeling disoriented when coming out of VR. Such users may prefer more active methods of locomotion when given the choice and may find mitigation techniques described below to be distracting. Because of this, sickness reduction methods should be provided as an option which can be turned off, and different techniques of locomotion should be offered when the narrative permits. VR games are sometimes rated on a "comfort" scale, from comfortable to intense, based on the amount and type of motion involved in the experience.

Developer familiarity and VR sickness

VR sickness is a situation particularly prone to the challenges of developer familiarity. When a developer builds a system, they become highly familiar as a result of the experience of building it. It is often difficult to remember or imagine what it might be like for a user to experience your system for the first time. This is problematic in any usability scenario, but is significantly worse in the context of VR sickness because developers are often familiar not just with their own system but with many other VR systems and forms of locomotion. Someone who has used VR for many hours may be said to have developed "VR legs"—analogous to sea legs, when a sailor is used to living on a ship and therefore does not get seasick. In the same way, a developer with VR legs may not be susceptible to VR sickness in the same way that a novice might, and because different locomotion methods can cause VR sickness in different ways, an experienced developer may not even notice that the locomotion method they are using may make people sick. For this reason, it is very important for developers to test their system with many users regularly, to identify problems the developer may never have considered.

10.6.2 VR Sickness Mitigation Techniques

Although the exact physiological mechanisms of VR sickness are still being investigated, the most prevalent theory is that a mismatch between sensory inputs is interpreted by the brain as a problem, possibly interpreting this as something you ate, and the response is to cause nausea to remove the problem. This is a prevailing theory around motion sickness and seems a reasonable explanation of the cause of VR sickness; additionally, mitigation techniques based on this assumption have been successful, lending evidence to this cause.

The primary sensory mismatch in VR sickness is the mismatch between visual and vestibular stimuli. Simulating motion with vision is straightforward because vection is so effective an illusion, but there is currently no substitute method of stimulating the vestibular system, nor is there any way to prevent the vestibular system from perceiving motion—or lack thereof—in the real world while other sensory units are immersed in VR. Because the inner ear is always perceiving the user's physical motion, and comparing it to motion perceived by vision, the strategies for reducing this mismatch are focused on either mitigating the effects of vection or finding some way for the user to move in the physical world that is "close enough" for the vestibular system to not raise the alarm. Three visual mitigation techniques outlined in this section are **field-of-view (FOV) reduction**, **depth-of-field blurring**, and **reference points**, all of which present different advantages, disadvantages, and potential use cases. Motion-based mitigation techniques are also considered, but these primarily use physical motion to drive virtual motion, assuming that the physical motion will be sufficient to reduce the experience of VR sickness.

10.6.2.1 FOV Reduction

In FOV reduction, the user's FOV is artificially reduced during motion. The reasoning behind this technique is that visual motion is primarily perceived by the rods, and since there are more rods than cones in the peripheral field (see Section 4.1.3), restricting motion in the periphery of vision will reduce vection. Common implementations of FOV reduction include *tunnelling*, in which the space outside the reduced FOV is fully occluded with a dark restrictor, and *vignetting*, in which the brightness and saturation of peripheral objects are reduced. Tunnelling and vignetting can be implemented to always be present or to be invoked only when the user is moving.

This repeated appearance and disappearance of the tunnel or vignette can be distracting and is typically not diegetic, and as a result presence can be reduced with this technique. Additionally, information otherwise available to the user may be occluded during travel, and the feeling of motion due to vection will be reduced, meaning the user's experience may be less dynamic. This method is very effective, however, and has become a standard in VR experiences which require continuous motion of some kind.

FOV reduction strategies work best when combined with eye-tracking technology, as the visible centre of the FOV restrictor can be adjusted based on the user's eye gaze instead of their head gaze. This enables a wider area for which visual scanning is possible.

10.6.2.2 Depth-of-Field Blurring

In depth-of-field blurring, objects beyond a predefined distance from the user's point of view are blurred to simulate reduced fidelity of faraway objects. The technique was developed in order to combat the **accommodation-convergence conflict**, which arises due to spatial inconsistencies between virtual displays and the real world. Three-dimensional rendering presents objects at different virtual distances from the user, and when the user focuses on a specific object, the angle of their eyes change so that the object aligns with the fovea of each eye. In the physical world the eyes would also focus at that distance, by adjusting the shape of the lens, but when viewed through a VR headset, the distance from the screen to the eye is always the same, and so the two factors (eye angle and eye focus) are in conflict, which can cause VR sickness.

Depth-of-field blurring effects can lessen the strain that is caused by the accommodation-convergence conflict, reducing visual fatigue and potentially improving the quality of the VR experience. Eye movements must be tracked, and detecting the depth-level attention of the user is particularly challenging. Moreover, the user must be able to look at an object in the distance and have the focus field realign almost instantaneously, in order for the experience to not be distracting. Similar to how active treadmills must predict a user's motion before they make it, depth-of-field blurring must predict what the user will be looking at before they look. Additionally, if the transition to the new depth of field is not immediate, the failed sensorimotor contingency can lead to increased VR sickness if the effect is too excessive. The strategy works best when the blur implementation itself closely resembles what humans see in the real world.

10.6.2.3 Reference Points

A reference point is a fixed object or display item within the virtual environment that doesn't move relative to the user's position in the virtual space. These reference points may be spatial UI elements (e.g. outline of a helmet or glasses and firearm reticle) or relate to larger objects or contexts in the scene (e.g. cockpit of an aircraft and steering wheel of a car). The advantage of reference points is that they assist in dictating the direction of the user's eye gaze, as well as providing justification for

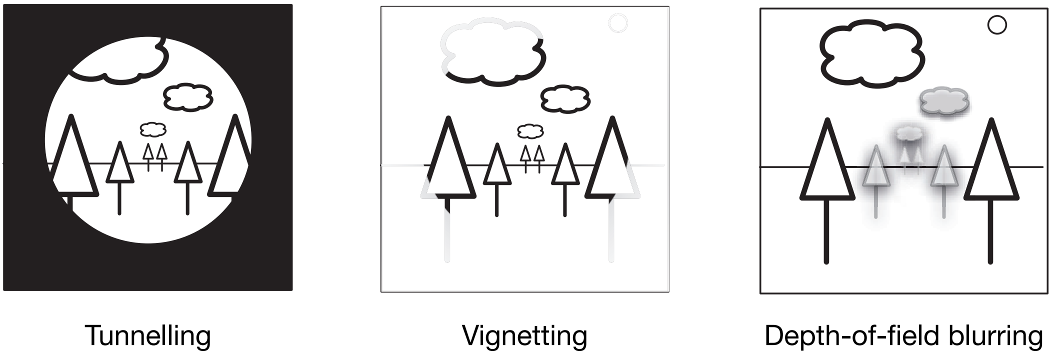

Figure 10.8: Methods to reduce field of view while travelling.

reduced vestibule sensation, which may reduce the intensity of sensory mismatches. Indeed, adding a virtual “nose” to the scene in the location where a user’s nose would be, as unusual as it sounds, can be sufficient to reduce VR sickness.

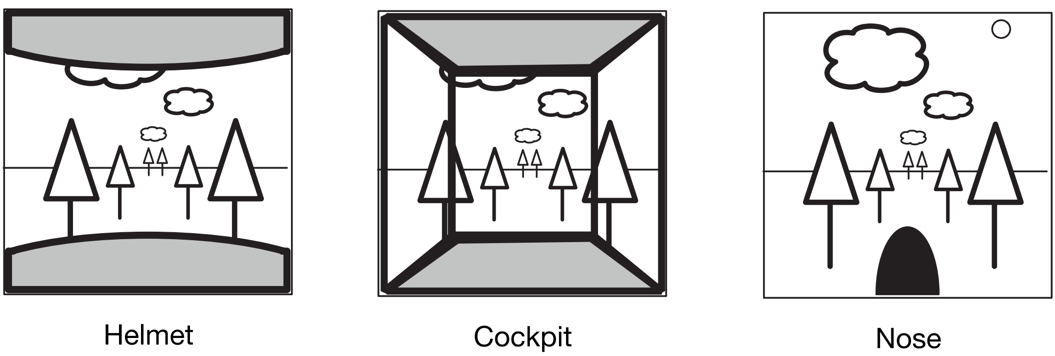

Figure 10.9: Reference view objects.

A smaller reference object, like a helmet or reticle, is usually less effective in preventing VR sickness than a larger one like a vehicle or aircraft cockpit; however, if the reference object is too large, it may impede the visibility of other objects in the environment and sometimes detract from presence as well. Reference objects may be locationally linked to the user’s head (like the nose example above) or the user’s immediate environment (like the vehicle cockpit), but in either case, the key is that when the user initiates motion in the environment, the reference object stays fixed. If you fly a spaceship through an asteroid field, the motion instigated by the user is the motion of the spaceship, not the motion of the user’s avatar. As such, significant pitch, roll, and yaw rotations can take place without VR sickness, if the windows of the spaceship’s cockpit can be seen. If the user is floating in space with no reference, these motions are very likely to cause VR sickness. The location, size, persistence, and occlusion of a reference object are highly context dependent and can increase or reduce presence based on diegesis of the object.

Reference points and field-of-view reduction can be combined, in the example of a multi-window vehicle cabin, where only the windows you are looking out of are clear. The other windows are frosted over or blocked out, reducing vection and mitigating VR sickness.

10.6.3 Ergonomic Concerns

Since the motion of the body itself is an input source in VR, developers need to take into account the physical comfort of users in a VR experience. One significant ergonomic consideration for VR locomotion is whether prolonged use of a strategy can cause strain, discomfort, or fatigue. Although many developers have created VR-based fitness and workout experiences, where the goal is to exercise and an argument can be made that fatigue is a good thing, most VR experiences that involve significant motion or locomotion use unfamiliar or constrained movements, which can lead to strain or discomfort faster than expected. Many locomotion strategies require such constrained movements, and many VR experiences encourage long-term play experiences, both of which can result in ergonomic issues:

- Keeping the arm outstretched for long periods of time causes lactic acid to accumulate in the muscle tissues of the arm, which causes discomfort and fatigue. This is an important factor to consider with respect to both **hand-directed VCM** and **teleportation**. In the former, the user needs to keep their arm at a very specific orientation if they intend to travel in the same direction for a long period of time, and in the latter, the user may feel inclined to repeatedly crane their arm above their head (or contort their wrist) in order to maximize distance, increase accuracy, or teleport in rapid succession.

- **Lean-directed VCM**, as its name implies, requires the user to physically lean in the desired direction of travel. Spending long periods of time with one's back in a non-neutral position, whether sitting or standing, puts pressure on the spine and may cause back pain as a result.

- The unusual stride required for certain gait negation devices, particularly **passive omnidirectional treadmills**, can be uncomfortable for new users. In some cases, the stride resulting from the use of these devices has been found to be closer to skating than walking. As well, due to the fact that many passive omnidirectional treadmills have concave surfaces, the user's feet often land on an incline when a step is taken. This is also unnatural and may cause foot and ankle pain during a prolonged play session.

In situations where highly precise and versatile movement is required, or users of varying physical ability are expected to play through an experience, it is important to evaluate the ergonomic implications of locomotion. The user's sense of enjoyment may deteriorate if the required actions of the experience result in physical discomfort, whether in the form of VR sickness, pain, fatigue, or muscle strain. Even something as simple as sitting or standing may pose problems, as standing for long periods of time can cause fatigue, while sitting may be uncomfortable if the user is required to continuously look around the virtual environment. If the user must hold a tool in a particular way for a long time, their arms may get tired.

Nonetheless, a locomotion strategy's ergonomics is most likely to be acceptable if the action required to execute it is similar to a corresponding action in the physical world. The walking and steering metaphors are common, and individuals with experience using either metaphor may be able to translate this experience to virtual space without a second thought. When selecting a locomotion strategy or system for their experience, developers should be ergonomically conscious: to evaluate the resemblance of a strategy's invocation method to some action in the physical world and to consider the expected physical abilities and prior experience of their target demographics.

10.7 SUMMARY

Being able to move through a virtual world is probably the greatest advantage of VR. Seeing and hearing the world is one thing, but being able to plan a route and navigate means you are a participant in the space, rather than just an observer. Locomotion in VR is core to the experience of many scenarios, but is also highly variable in terms of the way people relate to the experience. Users should be given options to choose which style of locomotion is most appropriate for them, and consideration should be given to how users with one mode of locomotion may interact with users with a different mode of locomotion. In the next chapter, we see what we can do once we get where we are going.