

Roy A. Ruddle

School of Computing
University of Leeds
Leeds, U.K., LS2 9JT

Justin C. D. Savage

Dylan M. Jones

School of Psychology
Cardiff University
Cardiff, U.K., CF10 3YG

Evaluating Rules of Interaction for Object Manipulation in Cluttered Virtual Environments

Abstract

A set of rules is presented for the design of interfaces that allow virtual objects to be manipulated in 3D virtual environments (VEs). The rules differ from other interaction techniques because they focus on the problems of manipulating objects in cluttered spaces rather than open spaces. Two experiments are described that were used to evaluate the effect of different interaction rules on participants' performance when they performed a task known as "the piano mover's problem." This task involved participants in moving a virtual human through parts of a virtual building while simultaneously manipulating a large virtual object that was held in the virtual human's hands, resembling the simulation of manual materials handling in a VE for ergonomic design. Throughout, participants viewed the VE on a large monitor, using an "over-the-shoulder" perspective. In the most cluttered VEs, the time that participants took to complete the task varied by up to 76% with different combinations of rules, thus indicating the need for flexible forms of interaction in such environments.

I Introduction

One important component of interaction in virtual environments (VEs) is object manipulation, for which many different types of interface device and forms of interaction have been developed. Examples include arm extension techniques such as the "go-go" (Poupyrev, Billingham, Weghorst, & Ichikawa, 1996), ray casting techniques such as HOMER (Bowman & Hodges, 1997), and image plane interaction (Pierce et al., 1997).

A notable limitation of these algorithms is that they have been developed and evaluated from the point of view of manipulating objects in open spaces that are largely free of obstacles. In contrast to this, many types of VE application necessitate the use of environments that are cluttered, and particular examples are the simulation of manual materials handling and the design of industrial plant so that it is easy to assemble or maintain.

A fundamental premise of VEs used for these types of application is that interaction should have real-world pragmatics. At the object level, this means that collisions must be detected and objects prevented from penetrating each other. At a user's level, it is important that they should be embodied in the VE because the space they occupy sometimes plays a critical role in determining the manipulations of an object that can be made. In the present study, this was achieved by letting participants "be" a virtual human (a 3D mannequin) and

travel through the environments while directly manipulating a virtual object that they effectively held in their hands. A tethered (over-the-shoulder) view perspective was used to help participants see their surroundings in the VEs while still interacting from an egocentric perspective. The following sections of this paper outline applications that involve object manipulation in cluttered VEs, describe rules of interaction that promote the manipulation of objects in such VEs, and present data from two experiments that evaluated participants' performance when they used the rules.

2 Object Manipulation in Cluttered VEs

The two types of application that we will outline are the simulation of manual materials handling (MMH) and the use of VEs for studies of design-for-assembly (or maintenance). In both cases, VEs have great potential for allowing designers to experience performing the tasks themselves, thereby gaining insights into the details of—and the problems caused by—a design that would otherwise not become apparent until a physical prototype was constructed.

MMH involves one or more people in carrying an object through an environment. Examples include the movement of automotive subassemblies from storage racks to a manufacturing cell, within a production line, and the evacuation of a casualty on a stretcher from an offshore gas platform (Hubbold & Keates, 2000). Although the environments themselves are very different, the MMH tasks shared certain common characteristics. First, the objects being manipulated are generally bulky, and the impoverished field of view (FOV) provided by most VE systems means that it would not be possible to see the whole of the object at one time if a human's eye perspective was used. Instead, an over-the-shoulder perspective can be adopted, allowing the object to be seen together with a person's immediate surroundings in a VE, and this has been shown to significantly aid user interaction (Ruddle, Savage & Jones, in press). Second, interaction involves both manipulation of a virtual object and travel of a virtual human through the environment, and the space occupied by the virtual human af-

fects how the object can be manipulated. At least nine degrees of freedom (DOFs) are involved: six for the object (3D position and orientation) and three for the human (position on a plane and orientation in that plane).

In design for assembly and design for maintenance, designers have to assess the installation or removal of equipment from a restricted space, such as a maintenance engineer reaching into an aircraft engine to remove some pipework (McNeely, Puterbaugh, & Troy, 1999). The engineer remains in one position, simplifying the requirements for a VE user interface, but the space occupied by his or her hands, arms, and tools is of paramount importance.

These are some of the requirements of applications that involve the manipulation of objects in cluttered virtual spaces. Although the tasks performed in different applications introduce variations in the requirements for interaction, there is currently a distinct lack of studies into user interaction in any type of cluttered space. The following experiments begin to address this void, using a task based on what is known as "the piano mover's problem" (Lengyel, Reichert, Donald, & Greenberg, 1990) to study object manipulation in cluttered VEs. The task, well known within the field of path planning, was one in which a person had to move a large, tightly fitting object through parts of a virtual building. Clearly, the piano mover's problem is just one of a range of tasks that could have been chosen, but it does capture many of the elements that are present in any form of MMH.

3 Rules of Interaction

The primary differences between manipulating objects in open and cluttered spaces is the frequency with which collisions occur, and the precision with which movements must be made in cluttered spaces to avoid collisions. This paper investigates rules of interaction that affect the manipulation of objects in VEs in which collisions are likely. The rules with which the paper is concerned are collision feedback, collision response, person-object rotation, physical compatibility, clutching, and inertia. Table 1 summarizes the options

Table 1. Summary of the Rules of Interaction for Object Manipulation in Cluttered VEs

Rule	Brief description of the primary options
Collision feedback	Haptic, visual, or auditory
Collision response	If a collision occurs, all objects are prevented from moving (stop-as-a-whole) or only movement of the colliding objects is prevented (stop-by-parts). Alternatively, objects can be automatically guided around (or along) each other.
Orientation constancy	When users turn around in a VE, the orientation of the object they are carrying remains constant in either the user's (local) or the environment's (global) reference frame.
Physical compatibility	Whether or not there is 1:1 correspondence between a user's physical and virtual movements
Clutching	The interface contains a clutch that allows users to change their physical posture without manipulating anything in the VE.
Inertia	Objects have virtual inertia, which limits the rate at which they can be manipulated.
Human locomotion	The mechanism by which users change their position and orientation in a VE, as opposed to manipulating the object

for each of these, and the following subsections describe the likely effects of each rule on a user's ability to carry a large virtual object through a cluttered VE.

3.1 Collision Feedback

Collision feedback provides information to a user that a collision has occurred and, therefore, helps explain why an object cannot be moved in a particular direction or why it appears to be "stuck." In VEs, the three primary types of collision feedback that can be provided are haptic, visual, and auditory. Haptic feedback is the "natural" option and the focus of much research. However, the technological limitations of current devices create difficulties in implementing haptic feedback for a task such as the piano mover's problem because of the scale of movement that is required. The immediate surroundings in which users interact in the piano mover's problem is approximately a $2 \times 2 \times 2$ m cube, whereas the six-DOF version of the best known type of haptic device, the PHANToM, has a working volume of only $0.2 \times 0.27 \times 0.38$ m and does not per-

mit full 360 deg. rotation. One solution is to scale movements of the physical interface (such as translations of the PHANToM) so that they produce correspondingly larger movements of a virtual object, but this introduces haptic instabilities. A second solution is to provide a clutch that allows the user to reposition and reorient the input device relative to the virtual object. Evidence from informal user tests indicates that this second solution has considerable promise (McNeely et al., 1999; personal communication, W. A. McNeely, 27 November 2001).

Visual feedback requires no technological advancement and does not suffer from the physical constraints of an input device. Graphical highlighting is the most common technique for providing visual feedback, and it can be used to indicate either the general region in which a collision has occurred or the exact parts of objects that are in collision. The most common form of auditory feedback is a simple tone that indicates a collision. Spatial sound could also be used to indicate the approximate direction in which the collision has occurred, but, although localization is relatively good in

frontal azimuth, accuracy is unacceptably low in elevation and rear azimuth (Blauert, 1997). Auditory feedback has the advantage of providing feedback for collisions that are outside a user's field of view. Visual feedback provides feedback that is more precise but adds yet more information and, potentially, distracting clutter to a person's visual channel.

3.2 Collision Response

If a collision has occurred, there are, generally speaking, two types of system response. The first type comprises rules that aid interaction by facilitating movement in noncolliding directions. Examples are allowing one object to slide along another using *constraint-based modeling* (Thompson, Maxfield, & Dew, 1998) and using *slip* or *force field* algorithms to automatically guide a user around obstacles (Jacobson & Lewis, 1997; Xiao & Hubbold, 1998).

The other type simply prevents objects from moving into a colliding position, but within this there are two levels of resolution: *stop-as-a-whole* and *stop-by-parts*. With the former, all movement in a graphics frame is prevented if any collisions take place, no matter where they are. (The whole world freezes). With the latter, only the colliding objects are prevented from moving. From a developer's perspective, stop-as-a-whole is substantially easier to implement. However, stop-by-parts is a much more flexible rule of interaction. For example, it allows users to reposition themselves in a VE even if the virtual object they are carrying collides with part of the environment and this, in turn, can help the user move the object to a noncolliding position.

Clearly, the applicability of facilitating and prevention rules varies from setting to setting. If interaction is to be made as easy as possible, then a facilitating rule should be chosen. On the other hand, if the purpose of a particular application is to simulate how difficult it would be to maneuver an object into a particular position in the real world, then a facilitating rule is not appropriate because it would make the virtual version of the task artificially easy. Instead, a prevention rule such as stop-by-parts should be chosen.




Figure 1. Example of a situation (carrying an object round a 90 deg. bend in a corridor) in which either local or global constancy is likely to allow the easiest interaction, depending on the type of collision response rule (*stop-as-whole* or *stop-by-parts*) that is being used. With global constancy, the user explicitly has to turn the virtual human and the object, but with local constancy the object turns with the virtual human. This means that, when local constancy is used in combination with *stop-by-parts* collision response, the time taken to perform the turn at the bend is reduced; but when *stop-as-whole* response is used, the collisions of the object that occur when the human turns will inhibit all movement.

3.3 Orientation Constancy

This rule defines how an object moves when the user carrying it turns round in a VE. The object's orientation can remain constant either in the user's (local) frame of reference, or within a global reference frame. (For a video, see appendix A.) The distinction is particularly important if a nonimmersive display such as a monitor (desktop) or flat-screen projector is being used because, with these types of display, a user always faces the same physical direction. This makes it more difficult for the user to control simultaneously their direction of view, direction of travel, and the orientation of an object than in an immersive VE, wherein all of the degrees of freedom are controlled using natural body movements.

In a cluttered VE, either local or global orientation constancy could be optimal, depending on the type of collision response rule that is being used. Consider the example shown in figure 1, where a user wants to carry

an object around a 90 deg. bend in a narrow corridor. With global constancy, the user turns around in the VE without changing the global orientation of an object. This means that separate inputs to the interface are required to change the orientation of the object so that it can be moved around the 90 deg. bend, but there is no difference in the inputs required for stop-as-a-whole and stop-by-parts collision response. With local constancy, the object rotates in conjunction with the user. If stop-as-a-whole is being used, each rotation of the user that causes the object to collide with the walls of the bend will cause all movement in the VE to cease. In turn, this means that a ratcheting motion will have to be adopted to move both the user and the object around the bend, and this is likely to be a substantial impediment to interaction in the VE system. However, if stop-by-parts is being used, movement around the bend will be only a two-stage operation. First, the person turns through 90 deg., with the object turning as far as it can until it collides. Then the object is manipulated through the remainder of the 90 deg. Of the four combinations of rule, interaction with local orientation constancy and stop-by-parts collision response is likely to be quickest because this will minimize the amount of user input required to manipulate the object. Interaction will be slowest with the local-whole combination because of the ratcheting that is required.

3.4 Physical Compatibility

Physical compatibility refers to the difference between a user's physical and virtual hand positions. If physical compatibility is preserved, the two positions are identical, but in a cluttered VE this can lead to situations in which the user moves his or her hand to a position that causes the held virtual object to penetrate part of the environment (like a wall), which is something that cannot be allowed to happen if the laws of physics are to be maintained. The solution is to allow physical compatibility unless this would cause an object to be moved to an "impossible" position. (For illustrative videos, see appendix A.) When this occurs, the object remains in its last valid position and visual feedback (the wireframe image in figure 2) is turned on to indicate the




Figure 2. Example showing physical compatibility feedback. The object (shaded) is in its last noncolliding position, and the wireframe image shows its physically compatible position. The wall that is in collision with the object is highlighted (wireframe triangle strips). The environment is the C-shaped VE used in the experiments and the walls between the viewpoint and the virtual human are transparent.

physically compatible position of the virtual object. Provision of the visual feedback is a signal to the user that a collision has occurred. A further problem occurs when the physically compatible position becomes valid once more because, potentially, this will cause the virtual object to “jump” suddenly in position. The solution here is to limit the rate at which the object is allowed to move using a rapid controlled movement algorithm (Mackinlay, Card, & Robertson, 1990) so that the adjustment between the old and new valid positions occurs gracefully but quickly. An added advantage of using controlled movement is that it prevents an object from leaping from one side of a solid barrier such as a wall to another if collisions are checked using only the object's position at the end of each frame. The alterna-

tive would be to perform collision detection using the volume that each object travels through, which is technically difficult and computationally time consuming.

If physical compatibility is not preserved, then, conceptually, there is an offset between a user's physical and virtual hand positions. This simplifies the software that controls interaction because, when a collision occurs, the offset is redefined so that the user's last hand position is always veridical and controlled movement is not required.

It is not possible to predict whether physical compatibility will aid or impede interaction. On one hand, physical compatibility forces a user to maintain a natural posture, which is likely to make interaction more realistic. On the other hand, a lack of physical compatibility would allow virtual objects to be manipulated from comfortable physical positions even though the same operation in the real world would be substantially more awkward (such as adjusting something near the limit of a user's reach). However, this is clearly not desirable in many ergonomic applications that seek to simulate the movements that people will have to make in the real world (such as to remove a component from an aircraft engine).

3.5 Clutching

Clutching is the process by which a user changes their physical posture without modifying anything in a VE. It is most commonly used to allow an interface prop to be reoriented by temporarily disconnecting the control of the virtual object that is under its governance (Hand, 1997). Clutching is particularly useful for avoiding awkward postures and allowing buttons on the prop to remain in convenient (fingertip) locations irrespective of a virtual object's orientation. If physical compatibility is not being preserved (see previous), then clutching can also be used to modify the offset between a user's physical and virtual hand positions, again allowing interaction to be less awkward. Although clutching is generally accepted to be beneficial in an interface, its use has seldom been studied. In one study, however, participants used the clutch for approximately 10% of the total trial time when they had to rotate and translate an object from

one place to another while using a gloved-based interface (Zhai, Milgram, & Buxton, 1996).

3.6 Inertia

Virtual objects are "weightless," which means that they can be manipulated much more rapidly than equivalent, physical objects. In addition, small changes of orientation produce large changes in the position of parts of the virtual object a long way from the axis of rotation. One side effect of this is a substantial decrease in the precision with which large virtual objects can be manipulated, and this is of particular importance in cluttered VEs.

A solution is to limit the rate at which virtual objects are allowed to be manipulated, a form of *virtual inertia*. This can be implemented using a controlled movement algorithm (Mackinlay et al., 1990), with visual feedback indicating to a user the amount of movement that remains to be made. (For a video, see appendix A.) For example, if a user attempted to rapidly turn a virtual object through 90 deg., the new orientation could be indicated using a wireframe version of the object (feedback) while the object gradually rotated to that orientation, at the rate prescribed by the virtual inertia. Once the object reached the new orientation, the feedback would be switched off.

The benefits of virtual inertia are likely to become more evident as the space in which an object is manipulated becomes more restricted, because small amounts of noncolliding movement can take place even if the final position that the user was trying to move the object to is in collision with another part of the VE. (See figure 3.) Without inertia, the object will frequently appear to be "stuck" unless the user's manipulations are precise and collision free.

3.7 Human Locomotion

All of the preceding rules of interaction govern the way in which an object is manipulated in a VE, but another important component of interaction in a task such as the piano mover's problem is the way in which a user travels through a VE. Interfaces for human locomotion




Figure 3. A situation in which virtual inertia is likely to aid object manipulation. (a) The object is being rotated clockwise from the orientation shown by the solid line to the orientation shown by the dashed line. However, at the end of the first graphics frame, the object has the orientation shown by the dotted line, which is in collision with the left-hand wall (the object is not positioned centrally between the two walls), so no movement is allowed and the object remains stuck in its initial position (solid line). (b) Virtual inertia allows only a small amount of movement in each frame. Movement in the first (dotted line) and second (dashed line) frames occurs because these orientations of the object do not collide with the wall, meaning that inertia allows the user to make some of the movement they intended. In both (a) and (b), the arrow shows the intended direction of rotation and collision detection is performed at the end of each graphics frame.

can be divided into those that allow movement to take place only in a user's direction of view, those that decouple a user's view and travel direction, or those that allow a user to move directly in any direction irrespective of the orientation of their virtual body (such as side-

step around obstacles). The latter include "walking" interfaces. (For a review, see Templeman, Denbrook, and Sibert (1999).)

In cluttered VEs, view-direction (gaze-directed) movement has been shown to improve performance compared with body-direction movement when participants had to explore room-sized spaces and find target objects (Ruddle & Jones, 2001). Walking interfaces theoretically enhance maneuverability but require very specialized laboratory facilities. Investigation of different rules for human locomotion is outside the scope of the study reported in the present paper. Therefore, for performance and practical reasons, view-direction movement was used throughout.

3.8 Summary of Rules

The preceding subsections describe rules of interaction that are likely to affect the ease with which objects can be manipulated in cluttered VEs. Some of the rules state principles that may seem obvious; for example, collision feedback will almost certainly help a user understand why a virtual object cannot be moved to a particular location. Others of the rules promote flexibility in interaction by increasing the likelihood that users can make progress in their manipulations of the object at any moment in time. Stop-by-parts collision response does this by allowing noncolliding objects to move if at all possible. Local constancy does this by reducing the amount of user input required to move the virtual human and object around a corner. Virtual inertia does this by breaking manipulations down into increments.

Flexible rules of interaction also have disadvantages. As flexibility (and, inevitably, complexity) is added to an interface, it becomes more difficult to explain to a novice, and thereafter the interface is more difficult to learn. Complex interfaces are also substantially more difficult to implement in a robust manner and to test thoroughly.

Even if flexible rules improve performance, the magnitude of those improvements cannot usually be predicted without running human participants in behavioral experiments. Given that most aspects of rule flexibility are not required if an environment contains

Table 2. Rules of Interaction Investigated by Each Experiment

Experiment	Rules inherited	Rule(s) investigated
1	-	Orientation constancy (local versus global) Collision response (stop-as-a-whole versus stop-by-parts) Clutching
2	Constant local orientation Stop-by-parts Clutching	Physical compatibility (on versus off) Inertia

few obstacles to movement (if collisions are unlikely, then flexible forms of response are not required), experimental studies are needed to determine the amount of clutter for which performance benefits to the deployment of such rules begin to occur. Failure to do this will lead to interaction interfaces containing functionality that is not required and, therefore, development resources being wasted.

4 Experimental Methodology

The remainder of this paper describes two experiments that used the piano mover's problem to investigate the effects of some of the preceding rules of interaction. This section describes the general methodology, and the subsequent sections describe the detail of each experiment in turn.

In each experiment, participants performed a series of trials in which they moved objects through two shapes of VE (*offset* and *C-shaped*; for videos, see appendix A). One of the objects was a cube, and the others were both an abstract shape of a type that is known as a Shepard-Metzler (SM) object (Shepard & Metzler, 1971). Each experiment investigated certain combinations of the rules of interaction. (See table 2.) In experiment 1, one of these combinations was stop-as-a-whole collision response and global orientation constancy. This provided a baseline for participants' performance because it is the most straightforward to implement and, therefore, typical of the rules of interaction that are implemented in current VE applications.

In all the experiments, the primary dependent variable was the time taken to complete each trial, but real-time recording of participants' interactions with the VE system also allowed analysis to be performed of participants' behavior. For each participant, the experimenter demonstrated how to move each object through the VEs. This meant that participants only had to re-create that pattern of movement, using the "expert's" tips. A much more difficult task would have been one in which participants had to determine whether it was possible for particular objects to be moved through a given VE, but that is left for later studies.

The VE software was a C++ Performer application that was designed and programmed by the authors and ran on a SGI Maximum IMPACT workstation. The position and orientation of the interface prop was measured using a Polhemus FASTRAK sensor and the MR Toolkit (Green, 1995). The application update rate was 30 Hz. The next subsections describe the detail of the environments, the display, and the general user interface.

4.1.1 Environments. Interior views of the offset and C-shaped environments are shown in figures 2 and 4, and plan views of their layouts are shown in figure 5. The cube object measured $0.5 \times 0.5 \times 0.5$ m, and the dimensions of the two SM objects are shown in figure 6. The ordinary SM object was used in both experiments, but the large SM object was used in only experiment 2. These two objects were the same general shape but different sizes.


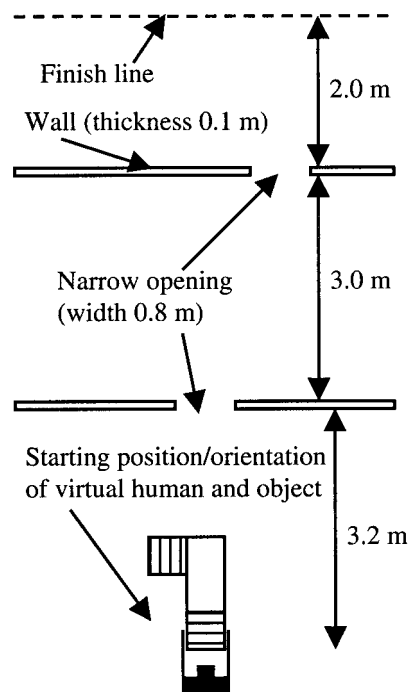


Figure 4. A view inside the offset VE. The virtual human is standing in the first opening and traveling toward the second.


4.1.2 Display and View Perspective. In the task used in the present study, users had to change their direction of view only occasionally. Unlike studies of navigation in virtual buildings, there was little requirement (or reason) for frequent “head” movements such as glancing to one side. For this reason, a large-screen monitor (86 cm/34 in.) was chosen for the display, rather than a head-mounted display (HMD). If the study had been performed with an HMD, the length of time involved for the experimental trials would have been expected to cause severe symptoms of VE sickness in a number of the participants, even if frequent rest periods had been incorporated into the procedure.

Each participant performed the experiment in an erect posture, facing a large-screen monitor that was positioned 1.6 m away on a table. The participant’s position and orientation within the VE was portrayed by a 3D model of a virtual human that held the object being manipulated, as shown in figures 2 and 4. In other words, the virtual human was the participant’s embodiment within the VE.

The participant’s viewpoint was positioned 3 m behind the human, connected by an egocentric tether (an over-the-shoulder view). This meant that the participant’s direction of view was always the same as that of



(a) Offset VE



(b) C-shaped VE

Figure 5. Plan views of the offset (a) and C-shaped VEs (b). In both cases, the ceiling was at a height of 2.4 m, and the narrow openings were 2.0 m high.

their virtual counterpart, but the participant was able to see the virtual human’s immediate surroundings in the VE despite the impoverished field of view (48×36




Figure 6. Dimensions of the ordinary (left) and large (right) Shepard-Metzler objects used in the experiments. In both objects, the stubs were at 90 deg. to each other. The large object was used in only experiment 2.

deg.). The use of a tethered view meant that the participant's viewpoint was sometimes on the opposite side of a wall to the virtual human. When this occurred, the walls in question were rendered as semitransparent using an alpha value of 0.2 (0.0 and 1.0 were fully transparent and opaque, respectively).

4.1.3 User Interface. Participants held a small box (100 × 75 × 40 mm, an interface prop) in their hands. The box contained a FASTRAK sensor and four buttons. If a participant held down one button, they accelerated forwards (that is, in their direction of view) at 0.5 ms^{-2} , to a maximum speed of 0.5 ms^{-1} . If they held the second button, they accelerated backwards at the same rate. The third button acted as a clutch. When it was held down, the participant could reposition and reorient the prop in relation to their hands without changing the position of the object. The fourth button was used to change the mode of the FASTRAK sensor. When this button was held down, changes of the prop's orientation caused the participant's direction of view to be rotated. At all other times, the prop directly controlled the orientation of the object.

When collisions did not occur or the clutch was not being used, there was a 1:1 correspondence between changes in the participant's physical hand position and orientation, and that of the object, even though the ob-

ject was being manipulated remotely. (It was 3 m away.) This is known as *hand-centered manipulation* (Bowman & Hodges, 1997).

Collisions of the cube and SM objects with the environment or the virtual human were detected using the RAPID software library (Gottschalk, Lin, & Manocha, 1996). Two types of response, corresponding to stop-as-a-whole and stop-by-parts collision response, were implemented. For stop-as-a-whole, all movement was prevented until the object was moved in a noncolliding direction, whereas, for stop-by-parts, the virtual human could still be moved if the object was in collision with the environment. Each time a collision occurred, and for both types of response, yellow graphical highlighting indicated which geometric (such as tri-strip) primitives were in collision, and the offset between a participant's physical and virtual hand position was redefined.

Collisions of the virtual human with the environment were detected by software written by the authors. When this type of collision occurred, a slip response algorithm allowed the virtual human to continue moving in a direction that was tangential to the colliding surface. The focus of these studies was on the difficulty users have in manipulating an object through a cluttered VE, whereas, in real life, it is trivial for people themselves to avoid walking into walls while carrying something. In summary, the slip rule ensured that the virtual human could move freely at all times, and the versions of the stop rule that were implemented for object collisions ensured that the users experienced the full difficulty of maneuvering the object in each VE.

5 Experiment 1

Experiment 1 investigated local and global orientation constancy, and stop-as-a-whole and stop-by-parts collision response. Each participant was randomly allocated to one of four groups, each of which used one combination of these rules (i.e., *local-whole*, *local-parts*, *global-whole*, and *global-parts*). Clutching was implemented for all of the combinations. Each participant moved cube and the ordinary SM object through the

offset and C-shaped VEs (for illustrative videos, see Appendix A).

5.1 Method

5.1.1 Participants. A total of 22 participants (four men and eighteen women) took part in the experiment. Their mean age was 23.8 ($SD = 7.1$). All the participants volunteered for the experiment and were paid an honorarium for their participation. The first sixteen participants were randomly allocated to the four conditions, and the remaining participants replaced those who timed out in a test trial. (See subsection 5.2.)

5.1.2 Materials and Procedure. The materials were as described in section 4. Participants were run individually by the same experimenter and performed the experiment over two separate days. On the first day, they performed the piano mover trials, and this took approximately two hours. On the second day, a participant performed a spatial ability test (Smith & Whetton, 1988). The data for this test are reported in section 8.

At the start of the first day, the experimenter demonstrated how to perform the piano mover's task, using physical scale models of the SM object, and the offset and C-shaped environments. For the piano mover's trials, there were four combinations of object and VE, and each participant performed trials in these combinations in increasing order of difficulty (cube-offset, cube-C, SM-offset, and then SM-C). In each trial, a participant carried the object from the starting position until the virtual human had crossed the finishing line. (See figure 5.)

The first set of trials (cube-offset) started with the experimenter demonstrating the interface and explaining how to move the object through the VE. The participant then performed three practice trials in which they were given advice and instruction by the experimenter as the trial progressed. This combination was used solely for training, so no test trials were performed.

Next, the experimenter demonstrated how to carry the cube through the C-shaped VE, and the participant then performed three practice trials and four test trials. The participant was given advice and instruction during

the practice trials (as for the cube-offset combination) but given no help at all in the test trials. If the participant had not completed a test trial after 300 sec., the trial was terminated and the participant progressed to the next trial.

The participant then performed the SM-offset and SM-C trials. For each combination, the experimenter demonstrated how to carry the object through the VE, and then the participant performed practice and test trials as just described. As before, each test trial was terminated after 300 sec. if it had not been completed.

5.2 Results

With each combination of the rules of interaction, four participants completed every test trial within the 300 sec. time limit. Of the other participants, one exceeded the time limit in an SM-offset trial, and the other five exceeded the time limit in one or more SM-C trials. The data for the participants who never exceeded the time limit and those who exceeded it in at least one trial are reported in subsections 5.2.1 and 5.2.2, respectively.

5.2.1 Time Limit Never Exceeded. The data for the participants who never exceeded the time limit was initially analyzed using mixed factorial analyses of variance (ANOVAs) that treated the collision response rule (whole versus parts) and the orientation constancy rule (local versus global) as between-participants factors, and the test trial number as a repeated measure. In the SM-offset condition, participants completed the test trials significantly faster as the trials progressed ($M = 90, 81, 87, \text{ and } 73$ sec. for trials one to four, respectively), $F(3, 12) = 3.28, p = .03$, indicating an effect of training. However, there was no effect of trial number for any of the other combinations of object and VE, or significant interactions of the trial number with either of the two rules. For the data reported later, participants' mean performance in the four test trials was treated as a single dependent variable and analyzed using two-factor (collision response \times orientation) ANOVAs. Interactions between the two rules are reported only if they were significant.




Figure 7. Mean time taken to complete the test trials in experiment 1. C = C-shaped VE; offset = offset VE; error bars indicate the standard error (SE).

Figure 7 shows the amount of time that participants took in the test trials for the three combinations of object and environment. In the cube-C trials, the time taken did not differ significantly between the whole and parts collision rules ($F(1, 12) = 1.05, p = .32$) or the local and global orientation rules ($F(1, 12) = 0.52, p = .48$). Similarly, in the SM-offset trials, the time did not differ significantly between whole and parts collision ($F(1, 12) = 0.29, p = .60$) or between local and global orientation ($F(1, 12) = 0.55, p = .47$). However, participants performed the SM-C trials significantly more quickly with the parts rule than the whole rule: $F(1, 12) = 13.59, p < .01$. In these trials, there was no effect of orientation ($F(1, 12) = 0.03, p = .86$), but there was a significant interaction between the orientation and response rules ($F(1, 12) = 11.5, p < .01$).

To further investigate these effects, the time data for the SM-C trials were divided into periods when participants were holding down the forward button on the interface prop, the backward button, or neither button. (See figure 8.) Included within the forward time were periods when participants were attempting to move through the VE but were prevented from doing so because the object was in collision. With the parts rule, the neither button time included periods during which participants were adjusting the position of the virtual human relative to the object by moving their hands while the object was in collision. Overall, participants spent less than 1% of their time holding down the backward button. For the time that participants spent holding

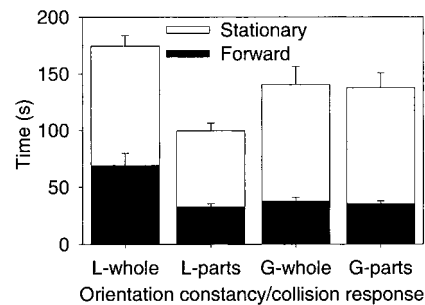


Figure 8. Mean time spent pressing the forward button, or neither button ("stationary") in the SM-C test trials of experiment 1. L = local; G = global; error bars indicate the SE.

down the forward button, there was a significant interaction between the collision and orientation rules: $F(1, 12) = 8.15, p = .01$. Overall, participants spent less time holding down the button with the parts rule than with the whole rule ($F(1, 12) = 10.38, p < .01$) and with global orientation than local orientation ($F(1, 12) = 5.89, p = .03$) but further analysis showed that these main effects were solely caused by the time spent by the local-whole group. Analysis of the amount of time in which participants did not hold down either button showed there was no significant difference between whole and parts collision ($F(1, 12) = 2.77, p = .12$) or between local and global orientation ($F(1, 12) = 1.89, p = .19$).

Two other types of data were analyzed for the participants who never exceeded the time limit. The first of these was the percentage of the total trial time that participants spent using the clutch. (See table 3.) In none of the three combinations of object and environment was there a significant effect of response rule or orientation, but there was a significant interaction between the two rules in the SM-offset trials: $F(1, 12) = 6.10, p = .03$. This is likely to have been caused by the increased difficulty participants had in completing the task in the global-whole condition.

The second type of data was the size of the position discrepancy between participants' physical and virtual hand positions in the SM-C trials. This was a measure of the extent to which participants adopted a physically incompatible position. In every trial, this was deter-

Table 3. Mean (Standard Deviation) Percentage of Time Spent Clutching During the Test Trials in Experiment 1

	Combination of rules			
	Local-whole	Local-parts	Global-whole	Global-parts
Cube-C	5.4 (5.4)	8.9 (1.8)	7.1 (3.6)	5.5 (2.4)
SM-offset	15.8 (4.9)	10.0 (1.3)	9.2 (3.3)	11.2 (1.8)
SM-C	10.4 (3.1)	9.4 (3.1)	9.9 (2.1)	12.4 (3.3)

Table 4. Root Mean Square (Standard Deviation) Distance between Participants' Physical and Virtual Hand Position for Each Combination of Rules in the SM-C Test Trials (Units are mm) in Experiment 1

	Local	Global
Whole	442 (31)	315 (66)
Parts	320 (39)	342 (55)

mined by calculating the root mean square (RMS) value of the position discrepancy for each graphics frame. An ANOVA showed there were not significant differences between the response rules ($F(1, 12) = 3.75, p = .08$) and orientations ($F(1, 12) = 4.48, p = .06$) but there was a significant interaction between the two ($F(1, 12) = 9.03, p = .01$). (See table 4.) This interaction was probably caused by the increased difficulty participants had in completing the task in the global-whole condition. Across all four conditions, further investigation showed that 80% of the discrepancy occurred in the lateral direction, relative to participants' (and the virtual human's) body.

5.2.2 Time Limit Exceeded. Of the six participants who exceeded the trial time limit, five did so in the SM-C combination of object and environment encompassing, between them, all four rule conditions. In total, these five participants exceeded the limit in nine of their twenty test trials. To provide information on where these participants experienced difficulties, each trial was divided into five stages:

- (i) traveling towards the narrow opening,
- (ii) maneuvering the far stub through the opening,
- (iii) rotating the object in the opening (one stub on each side),
- (iv) maneuvering the near stub through the opening, and
- (v) traveling towards the finish line.

The percentage of time that these participants spent in each stage for trials in which they succeeded and timed out is shown in table 5. These data indicate that participants encountered most of their difficulties maneuvering the far stub through the opening and rotating the object in the opening.

5.3 Discussion

The most important finding from experiment 1 was that local orientation constancy proved superior to global constancy, but only when a flexible collision response algorithm (stop-by-parts) was also used. This effect was anticipated and can be explained by considering the sequence of movements through which the virtual human and virtual object had to be moved to rotate the latter through a large angle (such as 90 deg.) in a cluttered VE. With the local-parts combination of rules, both the human and the object could be rotated simultaneously. By contrast, with the local-whole combination, a participant had to make incremental rotations in a ratcheting fashion while in between moving the object away from a colliding position. With global constancy, there was little difference between stop-by-parts and stop-as-a-whole. In both these cases, movement was a

Table 5. Mean Percentage of Time that the SM Object was in Each Part of the C-shaped VE for Trials that Participants either Completed Successfully or Timed Out. The Mean Trial Times were 229 sec. (Completed) and 300 sec. (Timed Out). All of these Data are for Participants who Timed Out in at Least One SM-C Trial in Experiment 1

Stage of Trial	Percentage of Time	
	Trials that were completed	Trials in which participants timed out
(i) Traveling towards opening	21.2	20.2
(ii) Maneuvering far stub	23.8	33.3
(iii) Rotating object	24.9	39.9
(iv) Maneuvering near stub	9.9	2.4
(v) Traveling towards finish line	20.2	4.2
TOTAL	100	100

two-stage process. The virtual humans had to be rotated, followed by rotation of the object.

There were significant effects for the time data in trials in which the SM object was moved through the C-shaped VE, but not for the cube object in that VE or the SM object in the offset VE. In other words, significant differences between the rules of interaction occurred for only the most cluttered combination of the object and VE that was studied, and this provides an indication of the amount of clutter that must exist within a VE before flexible rules of interaction become advantageous.

All of the participants used the clutching facility of the interface. Even in the simplest of the piano mover's tasks in which test trials were performed (cube-C), participants clutched for 7% of the time. Although trials without clutching were not run, these data provide a clear indication of the benefits that clutching provides in any form of virtual object manipulation.

Of concern was the proportion (27%) of participants

who failed to complete trials even though they had undergone a substantial amount of training with the rules of interaction they had been allocated. Clearly, this is most unlikely to have happened if the task had been performed in the real world, indicating that there is substantial room for improvement in the design of interfaces for the manipulation of objects in cluttered VEs. Experiment 2 focused on the rules of physical compatibility and inertia. In experiment 1, participants held the virtual object in a position that was physically incompatible by approximately 350 mm (the mean distance between participants' physical and virtual hand position, relative to their body), and this may have contributed to the difficulties that participants experienced. Inertia reduces the rate at which the object changed position and orientation, and was predicted to help participants to make small, incremental manipulations in the most difficult parts of the tasks.

6 Experiment 2

Experiment 2 used the same tasks as experiment 1, but with different rules of interaction. Participants were allocated to one of four groups, which all used the stop-by-parts and local orientation constancy rules (the combination that had proven most effective in experiment 1). For one of the groups (feedback-PC), the virtual object was maintained in a physically compatible position whenever it was not colliding with any other object in the VE. When a collision did occur, a wireframe version of the object was used to indicate the physically compatible position (see figure 2), and, after a collision, a rapid controlled movement algorithm ensured that the object moved smoothly to its physically compatible position. A second group (inertia) also used physical compatibility, but the rate at which the object could be manipulated was constrained. For illustrative videos of the feedback-PC and inertia conditions, see appendix A. Physical compatibility was not implemented for the other two groups. One of these (no-PC) was identical to the local-parts group in experiment 1. With the other (feedback-no-PC), a wireframe version of the object was displayed each time a collision occurred and indicated

where a participant was trying to move the object. With this condition, unlike the feedback-PC condition, there was only ever a small difference between the position of the object and its wireframe version because, without physical compatibility, the offset between a participant's physical and virtual hand position was redefined each time a collision took place. During the experimental program, the inertia group was run after the other three groups (it was a supplementary manipulation of experiment 2, investigated as a consequence of this experiment's results) but, for the purposes of reporting, all four groups are treated as a single experiment.

6.1 Method

6.1.1 Participants. A total of 35 participants (13 men and 22 women) took part in the experiment. Their mean age was 20.9 ($SD = 3.3$). All the participants volunteered for the experiment and were paid an honorarium for their participation. None had taken part in experiment 1. Thirty-two of the participants were randomly allocated to the four conditions, and the remaining participants replaced those who timed out in a SM-offset or SM-C test trial.

6.1.2 Materials. The experiment used the same hardware and software as experiment 1. The implementation of the rules of interaction for the no-PC, feedback-no-PC, and feedback-PC condition is as described in section 3. The inertia rule of the fourth group restricted the rate at which the virtual object moved to an angular (rotational) speed of 22.5 deg./sec. and a linear speed of 0.621 m/sec., and gave the impression that the object was being manipulated in a clear but viscous fluid. The angular speed was the slowest at which acceptable manipulation performance could be achieved, and the linear speed was calculated as $1.5 \times \tan(22.5^\circ)$; the SM object was 1.5 m long. If a participant tried to move the object quicker than this, a wireframe version of the object was displayed in its physically compatible position (see figure 2) and the object continued moving to this position at its maximum rate. At all times, the object moved toward the most recent physically compatible position. If a collision occurred, all residual rota-

tional movement was canceled by modifying the rotational offset between a participant's physical and virtual hand orientations (see subsection 4.1.3). However, due to the physical compatibility rule, any residual translational movement was not canceled.

6.1.3 Procedure. As in experiment 1, participants were run individually and performed the experiment over two separate days. On the first day, they performed the piano mover trials, and this took approximately 2.5 hours. On the second day, a participant performed the spatial ability test.

The difference between the two experiments was that, after completing the SM-C trials, a participant did a further set of practice and test trials in which they maneuvered a large version of the Shepard-Metzler object through the C-shaped VE. This is referred to as the large-C combination, and the dimensions of the object are shown in figure 6.

6.2 Results

Two participants in the no-PC group and one in the feedback-PC group timed out in the SM-C test trials. In total, these participants timed out in seven of the trials, and one also timed out in an SM-offset trial. These participants were excluded from the analyses reported here. Seven other participants successfully completed all the SM-offset and SM-C test trials, but each timed out in one of the large-C trials. These participants are included in the following analyses, with their trial time set to 300 sec. for the trials in which they timed out. Thus, the same time-out criterion was used as in experiment 1, in which participants never performed trials with the large-C combination. Initial analyses of the time data showed no effect of trial number, so the data reported here were analyzed using single-factor ANOVAs that treated the experimental group as the independent variable.

Figure 9 shows the amount of time that participants took in the SM-offset, SM-C, and large-C test trials. In the SM-offset trials, there was a main effect of group: $F(3, 28) = 3.31, p = .03$. Further analysis (Tukey; $p < .05$) showed a significant difference between the

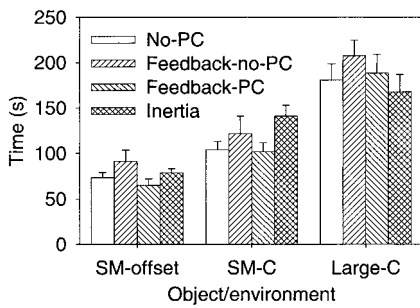


Figure 9. Mean time taken to complete the test trials with the normal (SM) and large Shepard-Metzler objects in experiment 2. Error bars indicate the SE.

feedback-PC and feedback-no-PC groups, but none of the other pairwise comparisons were significant. There was a marginal effect of group in the SM-C trials ($F(3, 28) = 2.81, p = .06$), but no effect in the large-C trials ($F(3, 28) = 0.85, p = .48$).

6.3 Discussion

Overall, there was little difference between the four combinations of rules that were investigated in experiment 2. The largest difference occurred for the simplest of the tasks involving the SM object, and it should be noted that this difference remained similar when participants who had timed out in the SM-offset or SM-C trials were included in the analysis. The most likely cause of the relatively poor performance of the feedback-no-PC group was that the brief appearances of the wireframe object (every time a participant moved farther into a collision) added visual clutter to the display, which was difficult to comprehend. By contrast, in the feedback-PC condition, the feedback remained displayed for much longer periods.

Although the implementation of virtual inertia had no significant effect on participants' performance, there was a noticeable increase in the time participants took to complete the SM-C trials, relative to the other conditions, and a slight decrease for the most difficult task (large-C). This may be explained as follows. By limiting the rate at which the virtual objects could be manipulated, inertia caused an increase in the minimum

amount of time in which it was theoretically possible for participants to perform each task. Therefore, inertia produced a performance penalty if participants found a task straightforward. However, if a task were difficult (such as the large-C condition), then inertia was beneficial. As explained previously (see also figure 3), inertia increased the likelihood that some movement of an object took place in each graphics frame because, when a participant tried to make a large movement, the VE software broke that down into a series of incremental movements that took place in a sequence of frames. Collision detection was performed on a frame-by-frame basis, meaning that the object moved by all increments of movement up to the first collision. The net result of this is that inertia slowed down participants who were skilled at manipulating the virtual objects, but was probably beneficial to participants who experienced difficulty in performing the experimental tasks.

The conclusions that can be drawn from experiment 2 are as follows. First, physical compatibility has little effect, either detrimental or advantageous, on object manipulation in the type of cluttered VEs that were under investigation. Second, substantial problems remain with the rules of interaction that were investigated, as shown by the participants who continued to experience great difficulty in performing the trials, even in the easiest two combinations of object and environment.

7. Spatial Ability Test

The primary focus of the present study was to evaluate the effect that different rules of interaction had on participants' performance when they carried bulky objects through cluttered VEs. In this, individual differences in ability represent a nuisance factor because they add noise to the experimental data. However, there is also considerable interest in the magnitude and cause of these individual differences themselves. Of particular interest are psychometric tests that could be used to predict individuals' likely performance in various VE tasks, and which could be used to identify those people who are either best suited to using VEs for particular design tasks, or are likely to require additional training.

Within the field of VEs, individual differences have most often been measured for tasks that required participants to navigate large-scale spaces, and the most comprehensive of these is a study by Waller (2000). However, large-scale navigation is a fundamentally different type of task to the piano mover's problem because the latter tests participants' ability to visualize the movement of a 3D object and their dexterous skill to execute those movements.

To provide information about the individuals who participated in the present study, all were requested to return to the laboratory on a second day and perform a spatial ability test (Smith & Whetton, 1988). This test contained 22 questions, the first two of which were practice questions. In each question, a participant had to imagine what a flat pattern would look like if it were cut out and folded into a 3D object. Each question consisted of drawings of a flat pattern and four similar 3D objects. The participant had to indicate whether each of the four 3D objects could be constructed by folding the flat pattern. Both the questions and the answer sheet were presented on paper, and participants were permitted 20 min. in which to answer the twenty test questions.

In total, 41 of the 57 participants returned to the laboratory to complete the spatial ability test, including six of the eight who had timed out during the SM-C task. Participants' scores in the test ranged from 47 to 80 ($M = 71.4$; $SD = 7.4$), and the maximum possible score was 80. The scores were compared with the mean time that participants took to complete their SM-C test trials, with a time of 300 sec. used for each trial in which a participant timed out. Analysis showed that there was a significant correlation between the two sets of data: $r(41) = -.32$, $p = .04$.

A second correlation was performed using time data that were adjusted to take account of the differences among the rules of interaction. For this, each participant's mean time in the SM-C trials was recalculated as $t' = t * n/m$, where t was the actual mean time (used in the first correlation), n was the mean time of all participants who took part in the quickest condition (physical compatibility, experiment 2), and m was the mean time of all participants who took part in the present participants' condition. The adjustment factors (n/m) ranged

from 0.670 to 1.0. Analysis showed these adjusted times were not correlated with participants' score in the spatial ability test: $r(41) = -.25$, $p = .12$.

This spatial ability data provide some evidence for a link between individuals' abilities, as assessed through the use of a spatial ability psychometric test, and the speed at which they were able to perform a spatial-motor task in a VE. However, the results should be treated with caution until more comprehensive studies are performed that use both a battery of psychometric tests and a substantially larger number of participants.

8 General Discussion

The present study adopted a new paradigm—the piano mover's problem—for evaluating rules of interaction for object manipulation in cluttered VEs. The basic problem was for a user to move a bulky virtual object through a restricted space. The user had to separately control the movement of a virtual human through a VE and manipulation of a virtual object that the human held in their hands. Thus, the experimental task made similar demands on users, in terms of the type of user input that was required, to those made by some VE applications, such as ones that seek to simulate the manual handling of materials on an engineering production line.

The more flexible the rules governing a user's interactions the quicker they are likely to be able to move an object from one place to another. However, increases in flexibility come hand in hand with increases with complexity. Flexible interfaces are substantially more difficult to implement in a VE system and test in a thorough manner, and they can also be more difficult for a user to learn. Added to this, the potential benefits of flexibility will depend on the difficulty of the task being performed and, within the piano mover's problem, this is closely related to the clearance between the object and the environment.

The results of experiment 1 provide information on the amount of clutter that is required before simple rules (such as the stop-as-a-whole collision response) significantly inhibit interaction. In restricted spaces (the SM-C trials), these simple rules increased the amount of

time required to manipulate an object by up to 75%, but, in less restricted spaces (such as SM-offset), the differences between the rules were negligible.

Physical compatibility had little effect (positive or negative) and neither did the introduction of virtual inertia. However, three potential uses of inertia should still be noted. First, it increases the realism of interactions that occur because large objects can no longer be manipulated as if they have no mass, and in this there is substantial scope for implementing an algorithm that models inertia in a mathematically more correct manner to take into account an object's physical inertia and a user's strength. Second, it aids interaction for unskilled users because it increases the likelihood that some movement takes place in each graphics frame when the environment is cluttered. Third, a notable difference between inertia and constraint-based modeling is that the latter assumes that no collisions occur between the constrained degrees of freedom (for example, components are free to slide along one particular axis), but with the inertia rule no such assumptions are made and collisions are checked each frame. With the inertia approach, it is not necessary to know a priori whether an object can be physically manipulated through a particular space. It follows that inertia has great potential as a rule of interaction in situations in which a person such as a designer needs to determine whether it is possible for tightly fitting components to be assembled together.

On this latter note, all participants are likely to have had greater difficulty completing the trials if they had had to determine whether it was possible to maneuver particular virtual objects through a VE (that is, to solve the piano mover's problem) rather than re-create a sequence of movement (a known solution) that they were shown by the experimenter. Further studies are planned that address this issue, as are studies in which two users collaborate to solve the piano mover's problem.

Acknowledgments

This work was supported by grant GR/L95496 from the Engineering and Physical Sciences Research Council, and per-

formed while Roy Ruddle was employed in the School of Psychology at Cardiff University.

References

- Blauert, J. (1997). *Spatial hearing: The psychophysics of human sound localization*. Cambridge, MA: The MIT Press.
- Bowman, D. A., & Hodges, L. F. (1997). An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. *Proceedings of the 1997 Symposium on Interactive 3D Graphics*, 35–38.
- Gottschalk, S., Lin, M., & Manocha, D. (1996). OBB-Tree: A hierarchical structure for rapid interference detection. *Proceedings of the 1996 ACM Conference on Graphics (SIGGRAPH '96)*, 171–180.
- Green, M. (1995). *The MR Toolkit Version 1.4* [Computer program]. Department of Computing Science, University of Alberta, Canada.
- Hand, C. (1997). A survey of 3D interaction techniques. *Computer Graphics Forum*, 16, 269–281.
- Hubbold, R., & Keates, M. (2000). Real-time simulation of a stretcher evacuation in a large-scale virtual environment. *Computer Graphics Forum*, 19, 123–134.
- Jacobson, J., & Lewis, M. (1997). An experimental comparison of three methods for collision handling in virtual environments. *Proceedings of the Human Factors and Ergonomics Society's 41st Annual Meeting*, 1273–1277.
- Lengyel, J., Reichert, M., Donald, B. R., & Greenberg, D. P. (1990). Real-time robot motion planning using rasterizing computer graphics hardware. *Computer Graphics*, 24, 327–335.
- Mackinlay, J., Card, S. K., & Robertson, G. G. (1990). Rapid controlled movement through a virtual 3D workspace. *Computer Graphics*, 24, 171–176.
- McNeely, W. A., Puterbaugh, K. D., & Troy, J. J. (1999). Six degree-of-freedom haptic rendering using voxel sampling. *Proceedings of the 1999 ACM Conference on Graphics (SIGGRAPH '99)*, 401–408.
- Pierce, J. S., Forsberg, A., Conway, M. J., Hong, S., Zeleznik, R., & Mine, M. R. (1997). Image plane interaction techniques in 3D immersive environments. *Proceedings of the 1997 Symposium on Interactive 3D Graphics*, 39–44.
- Poupyrev, I., Billinghamurst, M., Weghorst, S., & Ichikawa, T. (1996). The go-go interaction technique: Non-linear mapping for direct manipulation in VR. *Proceedings of ACM User Interface Software & Technology (UIST '96)*, 79–80.

- Ruddle, R. A., & Jones, D. M. (2001). Movement in cluttered virtual environments. *Presence: Teleoperators and Virtual Environments*, 10, 511–524.
- Ruddle, R. A., Savage, J. C. D., & Jones, D. M. (in press). *Levels of control during a collaborative carrying task*. Presence: Teleoperators & Virtual Environments.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701–703.
- Smith, P., & Whetton, C. (1988). *Spatial booklet from the "General Ability Tests"*. Windsor: ASE, A Division of NFER-NELSON (National Foundation for Educational Research).
- Templeman, J. N., Denbrook, P. S., & Sibert, L. E. (1999). Virtual locomotion: Walking in place through virtual environments. *Presence: Teleoperators and Virtual Environments*, 8, 598–617.
- Thompson, M. R., Maxfield, J. H., & Dew, P. M. (1998). Interactive virtual prototyping. *Proceedings of Eurographics UK '98*, 107–120.
- Waller, D. (2000). Individual differences in spatial learning from computer-simulated environments. *Journal of Experimental Psychology: Applied*, 6, 307–321.
- Xiao, D., & Hubbold, R. (1998). Navigation guided by artificial force fields. *Proceedings of the Computer Human Interfaces Conference (CHI '98)*, 179–186.
- Zhai, S., Milgram, P., & Buxton, W. (1996). The influence of muscle groups on performance of multiple degree-of-freedom input. *Proceedings of the Computer Human Interfaces Conference (CHI '96)*, 308–315.

Appendix A

MPEG videos illustrating some of the rules of interaction, and trials in the offset and C-shaped VEs, can be accessed from www.comp.leeds.ac.uk/royr/video/. None of the videos contain sound.