


THIRD ANNUAL REPORT FOR:
TRAINING SPATIAL KNOWLEDGE
ACQUISITION
USING VIRTUAL ENVIRONMENTS

(1 FEBRUARY 1998 TO 31 JANUARY 1999)



NATHANIEL I. DURLACH, ET AL.,

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
RESEARCH LABORATORY OF ELECTRONICS
77 MASSACHUSETTS AVENUE, ROOM 36-767
CAMBRIDGE, MASSACHUSETTS 02139
617-253-2534
DURLACH@CBGRLE.MIT.EDU

PREPARED FOR:
DR. T. ALLARD
COGNITIVE & NEURAL SCIENCES
OFFICE OF NAVAL RESEARCH (CODE 342)
800 N. QUINCY STREET
ARLINGTON, VA 22217-5660
703 696 4502 VOICE
703 696 8343 FAX

20000301 063

1. EXECUTIVE SUMMARY

This report summarizes the work done by MIT in Year 3 (1 February 1998 through 31 January 1999) of the ONR Grant N00014-96-1-0937 entitled "Training Spatial Knowledge Acquisition Using Virtual Environments." It has been prepared by Dr. Thomas E. v. Wiegand, Nathaniel Durlach (PI), Sam Madden, and Rebecca Lee Garnett.

The main effort during this period has been directed towards the development of a hardware/software system (the "VE Construction System") to facilitate the creation of complex virtual building environments. Virtual environment modeling of complex spaces is presently a manually intensive task in which required labor increases markedly with the level of three-dimensional detail that must be represented. As most real-world spaces worth modeling are of high complexity, the cost of this process can have a significant negative impact on the scope and timeliness of such projects. The tools we have developed during this year of our program greatly reduce the amount of time required to build the complex photorealistic VE models necessary for performing our envisioned VSPAN experiments. Also, of course, these tools should prove useful for other programs conducted at other laboratories.

Additional efforts have been devoted to further analysis of experimental data obtained in Year 2 and the preparation of a report on these data (Koh et al, 1999).

In addition, work on a "white paper" was initiated that is aimed at assisting ONR in designing the overall research program in the area of VE-assisted training of spatial behavior (this paper will be completed within the coming year).

The body of this report is devoted entirely to a description of the VE construction system. The paper by Koh et al. (1999), which has been submitted to the journal *Presence: Teleoperators and Virtual Environments*, is included in Appendix 1. A summary of all work accomplished during this grant, as well as plans for future work over the next three years, can be found in our renewal proposal.

2. THE VE CONSTRUCTION SYSTEM

2.1 Overview

The most time-consuming task in the creation of any VE simulation is the construction of the 3D models and the acquisition of photo-realistic textures for these models. The VE construction system, which we have nearly completed this past funding period, simplifies both of these tasks by allowing users to import two-dimensional DXF floorplan files (commonly available for many installations, including all of the buildings at MIT) and then to automatically attach textures to walls, doors, windows, and other objects within the building.

TOADS (Three-dimensional Open-ended Architectural Database System) is the software component of this system. TOADS can import plan files of various types (due to the Open-ended design of the program). The indoor worlds of the current MIT VSPAN work are based on DXF files available from the MIT space accounting office. Using the graphical interface tools provided within TOADS, we can group these raw floorplans, which initially consist of uncoordinated lines and polygons, into logical objects that are classified by object type (e.g., room, door, wall.) A graphical interface allows floorplans to be further augmented with paths indicating the areas through which VE subjects will normally walk and with movable polygonal objects that represent non-permanent structures (such as furniture) within the building.

The hardware component of the system supports the texture acquisition function. It consists of a rack-mounted scanner camera which runs on a linear path. A software controller (the "Grabber" program) steps the camera along a section of wall, capturing image frames at regular intervals, and pastes them together to form long strips of texture for each wall section. These texture files are automatically managed by the software, and are associated with specific polygons within the floorplan.

Once textures have been associated with the floorplan, the datafiles can be exported as a list of 3D polygons and associated textures in any format required by the real-time VE rendering software. Currently, we are using VRML-format output from the system for use with the Cosmo Player VRML browser (non immersive viewing) and DIVE (for immersive HMD-based viewing).

The following subsections describe the components of this system in finer detail.

2.2 Scanning Rack

In Y3, we have completed our work on the rack-based texture scanning system begun in Y2. During this work, we have solved a number of problems which previously degraded the captured images, and we have refined the controlling interface to be more closely integrated with the TOADS software (discussed in detail below). The present scanning system utilizes a Macintosh G3 laptop computer with a video capture card for image capture and general interfacing and control functions, while the control of the stepper motor and position sensing along the rack are controlled by a "BASIC Stamp" microcontroller via a serial port.

Our work on the scanning rack has resulted in two slightly different versions each with particular characteristics. Rack 1, our main rack, utilizes a 5.5 foot toothed rack and a mating pinion gear on the stepper motor. The camera and motor are mounted on the moving carriage (see Figs. 1 and 2). Rack 2 is a 4 foot unit, belt driven, in which the motor is stationary; only the camera (and an optional microphone) is mounted on the moving carriage. Both racks utilize identical electronics to control the motor with only a slight change in the firmware to accommodate the different rack lengths. Both racks include similar provisions for making initial adjustments to the tilt and rotation of the camera, as well as a simple crossed-beam laser alignment array.




Figure 1. Photograph of belt-driven scanning rack. Insert on left depicts one set of positioning lasers. Insert on right depicts head-on view of camera mounted on carriage.




Figure 2. Close-up photograph of camera/carryage of gear driven rack. One set of positioning lasers appears to the lower right.

Having examples of two different drive systems is enabling us to compare the relative advantages of each. In general, although the geared rack is very precise, the belt driven rack may be more practical because of the lower electrical and mechanical noise achievable by the remote motor mounting. Currently, we are using Rack 2 for our scanning and we are making adjustments to Rack 1.

One departure we have made from the earlier scanning systems is to include a video-capture card and video camera as a replacement for the QuickCam. Although the QuickCam was usable for some purposes, it was not suitable for experimenting with different (pre-distorting) lenses. In addition, the resolution was about half that of even an NTSC camera, and the compression required to send the QuickCam output over the serial port was seen as an unnecessary constraint. The current video-capture system takes advantage of the special hardware characteristics of the G3 powerMac, allowing fast and flexible video processing, with our choice of video camera. Rack 2 utilizes a fairly standard 1CCD color NTSC camera fitted with a fisheye lens, Rack 1 contains a more elaborate Y/C camera which contains a dataport for controlling camera parameters (white balance, exposure, etc).

The purpose of using the fisheye lens is to achieve better utilization of the available pixels within the field of view. In this way, a larger percentage of pixels represent the

world at and near eye level, and fewer pixels are dedicated to the periphery (i.e., floor and ceiling areas). This compression at the periphery is compensated for within the capture program, so that objects maintain their natural proportions regardless of where they appear in the field of view. Only the image resolution changes (decreases) as the object moves to the periphery.

In addition to the inverse distortion processing, the capture algorithm performs some simple filtering to minimize thermal noise (seen at very low light levels) and the effects of vibration of the camera (which is responsible for occasional bumps in the captured image).

Lastly, a provision for generating synchronized acoustic impulses has been included in the form of an instrumented cap pistol. Although the report of the pistol is not as distinctively impulsive as an electronically generated sound (e.g. a spark or piezo generator), the output is very high, and the average shape of the waveform is predictable enough to determine the room impulse response if the known waveform of the pistol is taken into account.

2.3 DXF Parser

The parsing engine provided is designed primarily to work with floorplans provided by the MIT space accounting office (<http://web.mit.edu/afs/athena.mit.edu/org/o/ofms-space/www/>). The parser is limited in the types of DXF objects it supports; in particular, it lacks:

- Support for rotated objects
- Support for hatched and patterned objects
- Support for blocks (object groups).

The parser has not been tested with files other than MIT blueprints and may or may not behave properly when non-MIT files are used.

2.4 Opening, Viewing, and Navigating DXF Files

Selecting “Open...” from the File menu will bring up a standard file dialog, from which a text file can be selected. If an invalid or uninterpretable DXF file is selected, nothing will happen. Otherwise, a window like the one show in Figure 3 will open. A tool window, as show in Figure 4, will also appear. The Layers menu will be initialized with any layers in the DXF document.




Figure 3. A sample DXF file, as it appears in the TOADS application. The rectangular box frames the drawing according to the scaling ratio between the width and height of the drawing,




Figure 4. The Tool window, showing the tools available for editing.

Table 1 gives a brief description of each tool.

Once a file has been opened, it can be manipulated via any of the tools shown in Table 1. Navigating the file is accomplished via the zoom tool and the scroll bars. The zoom tool doubles the size of drawn objects and tries to center the picture at the clicked point. If the option key is held down, objects are made half as large (the document is zoomed out).

Objects can be selected via the Select Tool or the Select Box Tool. The Select Tool selects any objects beneath the clicked point on the screen. The Select Box Tool selects any objects completely enclosed in the dragged rectangle. Selected objects can be removed by hitting the delete key.

#	Tool Name	Description
1	Zoom Tool	Zooms into or out of (by holding down the option key) the selected point in the document.
2	Room Tool	Defines a room as a closed polygon of points. Brings up the Room Creation dialog.
3	Select Tool	Selects objects under the selected point in the document/
4	Grabber Tool	Creates the 5'x5' texture grabber box, which can be dragged along lines in the document and used to associate the texture grabber dialog with a specific piece of an object.
5	Junk Tool	Creates polygonal objects in the document which represent movable, non-floorplan specified items (such as furniture, boxes, and trash) in the environment.
6	Path Tool	Creates parallelogram-segment based paths. Paths have two collision detecting walls, and two non-collision detecting walls. They define habit-trails where people can move freely through junk.
7	Box Select Tool	Allows objects within a specified rectangle to be selected.

TABLE 1. Description of tools.

2.5 The DXFObject Menu

Several menu items become available when an object is selected. The DXFObject menu (see Fig. 5) controls settings specific to individual objects.




Figure 5. The DXFObject menu.

The "Set Object Length.." item allows a correspondence to be established between DXF coordinates (as arbitrarily defined in the DXF file) and real world distances (as measured in the real building). This setting is only important in determining the size of the grabber rectangle (see subsection X on "Acquiring Textures" below.) Initially, 1 DXF unit is assumed to be equal to 1 meter in length. Fig. 6 shows the dialog through which this default scaling factor can be modified.




Figure 6. The Object Length Dialog allows a mapping to be established between DXF and real-world coordinate spaces. By selecting an object of known length in the plan, and entering the value of this length in the box, the scaling factor is determined.

Selecting "Get Object Info..." brings up the Object Info Dialog for the first-selected item, as shown in Figure 7. The Object Info Dialog allows a PICT file-based texture to be

associated with an object, allows an object to be given a specific height, and allows collision-detection to be turned on or off for the object.




Figure 7. The Object Info Dialog allows modification of the name, texture, height, and collision detection flag for a selected object.

The “Collision Detect” item allows the collision detection flag to be turned on or off for any number of selected items. A check mark will appear next to this item if the first selected object has collision detection turned on.

The “Set Object Type” item allows a group of object to be set to a specific object type. Object type affects the default texture associated with an object, as set through the preferences dialog (see the next subsection). Currently supported object types include walls, doors, windows, elevators, stairs, junk, and paths.

2.6 The Preferences Dialog

In addition to object specific settings, there are a few document-level settings and defaults that are accessible through the File menu under “DXF Preferences...” Selecting this item will bring up the dialog shown in Figure 8. Default textures can be set for each of the object types specified above. A default ceiling height for objects which specify no height can also be set. Finally, the mapping between real-world and DXF-space coordinates can be set, as in the Object Length dialog described above.




Figure 8. Document specific preferences dialog allows the user to set default textures and heights.

2.7 Adding Junk and Rooms to DXF Files

The Junk and Room tools have similar interfaces. They define closed polygons on the floorplan which correspond to the outline of junk or the boundaries of a room. The tools work much like the polygon tools from common drawing programs: each mouse click on the screen defines a point in the polygon. A double-click closes the polygon.

When the room tool is in use, the Room Settings dialog, as shown in Figure 9, will appear after the polygon has been closed. This dialog can be used to specify a name, color, ceiling height, and texture for a room. The “Rooms” menu, described in Section VIII below, can be used to add objects to rooms or change the settings for a room. As described above, rooms are a logical unit that contain objects – objects inherit height and texture information from their parent rooms.




Figure 9. The Room Settings Dialog allows the user to specify defaults for objects in rooms.

Once a Room or Junk polygon has been defined, it can be dragged freely around the drawing. Holding down the option key and dragging a polygon will cause it to rotate about its center. This is particularly useful for junk, which may be moved in the physical building and need to be updated in the floorplan.

2.8 The Rooms Menu

The rooms menu, as shown in Figure 10, allows objects to be added to rooms, rooms to be selected, and the settings for rooms to be changed. Selecting a room from the “Add Item(s) to Room” submenu will cause the currently selected objects to become members of the specified room. Notice that objects can belong to several rooms, because there are sometimes walls which are on the boundary between two logical room areas, although objects only inherit height and texture information from the first room to which they are added (see the section on property inheritance in “Exporting to 3D Formats” below). There is currently no way to remove an object from a room once it has been added.

The “Select Room” submenu selects objects which belong to the chosen room. Similarly, the “Room Settings...” item shows the Room Settings dialog for the selected item.




Figure 10. The Rooms Menu.

2.9 Adding Paths to DXF Files

The Paths tool works somewhat differently from the other tools in the TOADS application. The process for defining a path is as follows:

- (1) Select the path tool;
- (2) Click and hold the mouse button at the first point in the path;
- (3) Drag to define the first non-collision-detecting edge;

- (4) Release the mouse button and position the cursor so that the two collision detecting edges are located appropriately;
- (5) Click the mouse button and the second non-collision-detecting edge will appear;
- (6) The next path segment is now being defined, as in step 4;
- (7) Press the escape key when no more path segments are desired.

Paths are created as a number of “perimeter lines” (rather than polygons), so that individual lines defining the paths can be selected, deleted, and have toggle collision detection on or off.

2.10 Acquiring Textures: The Grabber

Textures are added via the Grabber tool. When the Grabber tool is initially used, it places a 5'x5' rectangular box on the drawing at the mouse location. This box represents the texture acquisition frame and can be dragged around the drawing. It will automatically snap to the nearest line object in the drawing, or to the center of the nearest junk polygon. It rotates to match the slope of the line it is currently snapped onto. Pressing the left and right arrow keys with the Grabber tool selected will cause the Grabber rectangle to move parallel to the line it is currently snapped to, in increments of the rack length.

Once the Grabber tool has been snapped to an object, the “Grab Texture...” item in the DXFObjects menu becomes enabled. Select this item to bring up the Texture Acquisition dialog, as shown in Figure 11.




Figure 11. The Texture Acquisition Dialog

The left scrolling list show QuickTime movies which have been scanned using the Grabber Program (see subsection 2.11, below) or added from an external source. Click the "Add" button to bring up a standard file dialog which can be used to select any QuickTime movie. The "Scan" button launches the Grabber program (if it isn't already running) and commands it to acquire a texture-movie. In either case, a new item is added to the left scrolling list.

The right scrolling list contains PICT files which represent swaths of texture. They are obtained by tiling frames of captured texture movies into a single file. Click on a movie in the left list and press the "Copy >>" button to launch the QuickTime Viewer application (See section XI below) and open the selected movie. QuickTime Viewer processes the movie to create texture swaths, which are sent back to the dialog and added to the "Processed Textures" list. Use the "Up" and "Down" buttons to arrange pictures so that textures corresponding to adjacent wall segments are next to each other in the list. Once textures are properly arranged, click the "Make" button to concatenate them into a single texture file.

2.11 Using The Grabber

The Grabber program is designed to capture QuickTime movies via QuickTime compliant FrameGrabber devices such as the Connectix QuickCam. When the Grabber program first launches, it attempts to open the first available frame grabber device. If it finds no such device, it exits. Otherwise, a small window with the input of the current device appears, as shown in Figure 12.



Figure 12. The Grabber Main Window.

To change settings for the current frame grabber, or switch to a new frame grabber, select "Camera Settings..." from the Video menu. This will bring up a QuickTime Input Settings dialog box, which includes several panels. The "Source" panel, as shown in Figure 13, can be used to select the current input device.




Figure 13. The Source panel of the QuickTime Input Settings dialog box.

The Grabber application is designed to control a simple stepper motor via the Mac serial port. It drives a Basic Stamp based controller via a simple serial interface. Selecting "Robot Settings..." from the Robot menu brings up the dialog shown in Figure 14, which

can be used to select the serial port for the motor controller and the size of the steps which the motor makes.




Figure 14. The Robot Settings dialog box.

The Grabber program can also capture sound impulses as it records video. The "Impulse Settings..." item in the Sound menu shows the dialog in Figure 15 which allows specification of the duration of sound impulses. Impulses are associated with video frames as capture occurs. Impulses apply to all video frames from the time they are captured until the time the next impulse is captured.




Figure 15. The Impulse Settings dialog.

Once the settings have been properly configured, movie capture can begin. Select "Start Recording..." from the File menu. A standard file dialog will appear requesting the location of the recorded movie. Then, the QuickTime Input Settings dialog will appear. Go to the "Compression" panel and select the desired compression method – settings other than "None" may not work. Exit the dialog by clicking "OK" to begin recording. For maximum recording efficiency, events are not processed and the main video window

is not updated. To add an audio impulse during recording, type ⌘-I. The sound input device in the sound control panel will record for the specified period of time (no visual cue will be given). To stop recording, type ⌘-E. The movie will be written to disk and the program will return to normal operation.

2.12 Using QuickTime Viewer Program

The QuickTime Viewer application allows examination of all data stored in the augmented QuickTime movies generated by the Grabber program. It takes these movies and concatenates rectangular sections of frames together to form panorama-style PICT images.

When a Grabber-generated movie is opened in QuickTime Viewer, a list of frames in the movie is presented, as shown in Figure 16. Double-clicking on a frame displays the frame and some basic information about (such as its number, time into the movie, and whether it has an impulse associated with it.) Figure 17 shows this Frame Info window.




Figure 16. The frame list window.




Figure 17. The frame info window.

To generate a Panorama, highlight the frames which are to be included and select “Build Panorama...” from the File Menu. A new window like the one shown in Figure 18 will appear. Generally, it is not desirable to include the entire frame in a generated texture file, so frames can be clipped to a rectangle using the “Select Frame Size...” option from the Panorama menu. Once the clipping is set up properly, select “Save To PICT” from the Panorama menu to save the texture to disk.




Figure 18. The panorama window.

The “Compute Frame Size...” option in the Panorama menu shows the dialog displayed in Figure 19. By specifying the wall-to-camera distance, camera field-of-view, and camera movement per frame, the program will automatically compute the appropriate frame size to exactly tile the movie.




Figure 19. The “Compute Frame Size...” dialog.

Finally, the Panorama menu includes the items “Rotate Clockwise” and “Rotate Counterclockwise” – these options rotate each frame 90° in the specified direction, to compensate for a horizontally oriented camera (which allows more vertical resolution.)

2.13 Exporting to 3D Formats

The TOADS program does not directly support any existing 3D file format, but instead generates an intermediate file format which contains information about the 3D polygons and textures which should be included in any 3D file. The file format and the C++ interface to it are available in the on-line TOADS documentation. The details of this process are transparent to the user once the appropriate output module is written and installed. The current module we are using generates VRML files.

To generate an intermediate 3D file, select “Write 3D Data...” from the File menu. The program will request a destination file via a standard file dialog. One polygon for each line or segment of a polygon is written out. All other DXF objects (circles, arcs, solids, text, etc.) are ignored. The height and texture to specify for each polygon is determined from the inheritance chart shown in Figures 20 and 21 below.




Figure 20. Height inheritance for polygons in 3D format.




Figure 21. Texture inheritance for polygons in 3D format. Rooms only specify textures for wall objects, or when objects have no default type.

2.14 Using DIVE

The current output module from TOAD generates VRML files that we are using in conjunction with DIVE real-time immersive VE rendering software. The DIVE package runs both on the SGI platform as well as on Windows machines. However, we are primarily using the DIVE/WindowsNT/Intergraph platform that has recently been assembled in our lab.

We have moved toward DIVE and away from Coryphaeus software for a number of reasons. DIVE is the product of academic research and is not part of a closed-

environment for-profit enterprise. We can get source code to modify the software to our needs, and there is a community of users to help with implementation issues. DIVE is available for both Unix and Windows/NT platforms, and indeed, we have it running on our SGI machine as well as the Intergraph.

DIVE also comes pre-configured to support collaborative VEs with multicast sound and video. The ability to create "invisible participants" and the ability to record user responses and activity (motions, viewpoints, actions) are important features for use in an experimental environment. Furthermore, many specialized features have been added (and are available at no extra cost) by the users, such as spatialized audio support.

World descriptions in DIVE are VRML based, further enabling crossplatform sharing and development. VRML browsers are available for all platforms, and has been accepted as the current standard for VE design and distribution. Creating our environments in VRML means that we will be able to extend them using environments others have built and distribute them without the inclusion of bulky, platform-dependent, executables. DIVE also includes an API whereby many of the limitations of VRML can be overcome via C and tk/tk programming. This API enables many of the multi-user and multicast features of DIVE and also allows highly interactive, dynamic environments, which are not well supported by raw VRML, to be implemented.

3. APPENDIX

Koh, G., von Wiegand, T.E., Shinn-Cunningham, B., Garnett, R.L., and Durlach, N.I. (1998). Use of Virtual Environments for Acquiring Configurational Knowledge about Specific Real-World Spaces: I. Preliminary Experiment. Submitted to *Presence: Teleoperators and Virtual Environments*.

**USE OF VIRTUAL ENVIRONMENTS FOR ACQUIRING CONFIGURATIONAL
KNOWLEDGE ABOUT SPECIFIC REAL-WORLD SPACES:**

I. PRELIMINARY EXPERIMENT

Glenn Koh*

Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge MA 02139

Thomas E. von Wiegand

Massachusetts Institute of Technology

Barbara Shinn-Cunningham

Boston University

Rebecca Lee Garnett

Nathaniel I. Durlach

Massachusetts Institute of Technology

Correspondence to:

Nat Durlach
Research Laboratory of Electronics
Massachusetts Institute of Technology
77 Massachusetts Avenue 36-709
Cambridge MA 02139
(617) 253-2534
durlach@cbgrle.mit.edu

*currently at Morgan Stanley, New York.

USE OF VIRTUAL ENVIRONMENTS FOR ACQUIRING CONFIGURATIONAL KNOWLEDGE ABOUT SPECIFIC REAL-WORLD SPACES:

I. PRELIMINARY EXPERIMENT

ABSTRACT

A relatively simple architectural space was modeled and used to compare the effects of spatial training in simulations versus training in the real world. Thirty-five subjects were trained in one of the following conditions: real world (RW), virtual environment (VE), nonimmersive virtual environment (NVE) and Model (Mod). The VE condition made use of a head-mounted display to view the simulated environment, while the NVE condition used a desktop monitor. In the Mod condition, the subject viewed and could manipulate a 3D model of the space, viewed from a desktop display. The training-transfer tasks, performed after brief unstructured exposure to the actual space or to one of the simulations, consisted of estimating the bearing and range to various targets in the real space from various spatially-distributed stations, each such pair of estimates constituting a subtask of the overall transfer task. Results obtained from each of the four training conditions proved to be roughly the same. Training in any one of the simulations was comparable to training in the real world. Independent of training condition, there was a strong tendency among subjects to underestimate range. Variability in range errors was dominated by differences among subjects, whereas variability in bearing errors was dominated by differences among subtasks. These results are discussed in the context of plans for future work.

1 Introduction

As evidenced both in this special issue of PRESENCE and in a recent past special issue of PRESENCE [issue 7(2)], considerable attention is being given to the use of virtual environment technology for training spatial behavior.

In our research program, we are primarily concerned with the use of VE technology to train individuals' spatial behavior in the real world. We are interested in spatial behavior in virtual worlds only to the extent that it relates to behavior in the real world. More specifically, our focus is on *transfer* from experiences in the virtual world to behavior in the real world. Furthermore, within this general domain, we are concerned with two distinct but related sets of objectives. The first set focuses on the use of VE technology for familiarizing individuals with *specific* spaces. The second set focuses on the use of VE technology to improve an individual's spatial behavior *in general*. In this paper, attention is confined entirely to the first set.

Within the domain of specific-space familiarization, we intend not only to explore how training that exploits VE technology compares to training conducted in the real world, but also how it compares to training based on simplified and less costly technology. It should also be noted that we do *not* assume that the best VE trainer is the one that provides the highest fidelity or the greatest degree of realism. Rather, we believe that, as in many other forms of training, simplification, emphasis, and other types of "distortions" may enhance training.

This paper is divided into four sections: (2) Background, (3) Experiment, (4) Results, and (5) Discussion. In Sec. 2, we have attempted not only to help the reader put our own work in context, but also to provide a brief overview of the wide variety of types of VE/spatial work now being conducted. In Secs. 3 and 4, we report on the results of a preliminary experiment that was performed as a first step in our developing program to explore the use of VEs to train spatial behavior in real specific spaces. In Sec. 5, we summarize our main results and comment on implications and plans for future work.

2 Background

Spatial behavior has been studied extensively for many years in essentially all phyla of the animal kingdom and with respect to all kinds of sensory systems. Even when restricted to humans, the literature on spatial perception, cognition, and behavior, which we shall refer to simply as "spatial behavior," is enormous.

In this section, we provide an overview of some recent experimental studies that have been concerned both with spatial behavior and with virtual environments. In several of these studies, the focus is the same as in this paper: the use of virtual environments to train spatial behavior in the real world (training transfer). In others, the focus is on improving people's spatial behavior in virtual environments. In still others, virtual environments are regarded primarily as an experimental tool that can be used to improve basic understanding of spatial behavior. Studies in the first category are considered in some detail below. Studies in the second and third categories are considered briefly at the end of this section.

Training and assessment of spatial behavior using virtual environments (primarily concerned with route following) has been addressed by Witmer, Bailey, and Knerr (1995) and Witmer, Bailey, Knerr, and Parsons (1996). Subjects were trained to follow a specified complex route through a real building and were divided into three groups: a building rehearsal group, a VE rehearsal group, and a symbolic rehearsal group. Half the subjects in each group studied a map before the rehearsal. In the individual assessment phase, subjects were asked about their sense of direction and were given the Building Memory Test (Eckstrom, French, Harmen, and Dermen, 1990). After rehearsal, the VE group filled out questionnaires on simulator sickness and the subjective sense of presence. In the study phase, the designated route was studied (route directions, photos of landmarks, maps for the map subgroups) for 15 minutes. In the rehearsal phase, the building group rehearsed in the real building, the VE group rehearsed in a high-quality VE representation of the building (SGI reality engine, stereoscopic display with a large field of view using a Fakespace Boom2C), and the symbolic group verbally rehearsed the route directions out loud while viewing the landmark photos. In the transfer-test phase, the subjects traversed the route in the real building, stopping to identify destinations along the way. Measures of performance included number of wrong turns, route traversal times, misidentification of destinations, and distance traveled. Subjects were also required to demonstrate knowledge of the building configuration (even though they had not been trained for this) by estimating distance and direction of various goal sites from various reference locations.

All groups evidenced some route learning during rehearsal as a function of rehearsal trial number. However, the performance of the VE group was much worse than that of the symbolic group, which in turn was worse than that of the real building group. In the transfer tests, the building group performed best, the VE group second best, and the symbolic group the worst (differences were larger for the wrong-turns measure than for the traversal-time measure). Configuration knowledge was correlated to some degree with gender, building memory test score, and landmark recall; however, it was not affected by the training method employed. Overall, very little configuration knowledge was gained from the route learning experience in the VE. [The fact that extensive route knowledge does not necessarily lead to a high level of configurational knowledge has been demonstrated previously in studies of hospital nurses and taxi drivers (Chase, 1983; Moeser, 1988).]

The use of VEs for training firefighters to navigate and operate in specific spaces has been studied by Bliss, Tidwell, and Guest (1997) and by Tate, Sibert and King (1997). In the study by Bliss et al., VE training for route navigation in an unfamiliar building was compared to training by the use of blueprints and to no training. The test task consisted of rescuing a mock baby (a life-sized doll) in the building. Whereas the VE and blueprint groups were instructed to follow the route that had been trained, the no-training group was told to proceed as they would in a real situation when they had no previous experience with the building. All subjects wore goggles sprayed with paint to simulate the degraded visibility that occurs in real fires. The VE system used an SGI RE2 engine, an HMD with a 30 degree field of view and 100% stereo overlap, a 6 DOF Ascension head tracker, and a mouse to control motion through the VE. The route to be learned was 389 feet long, and included 16 decision points (with 13 changes of direction). During training, both the VE group and the blueprint group traversed the prescribed route twice under the guidance of verbal instructions from the experimenter. Test performance measures included time to execute the rescue and number of wrong turns. The results of these experiments showed that VE training and blueprint training were roughly equal, and that both these types of training were much superior to no training (e.g., the average navigation times for

blueprint, VE, and no training were 100, 115, and 177 secs, respectively).

In the study by Tate et al., the environment was the Navy test ship ex-U.S.S. Shadwell, and the subjects were previously trained naval firefighting personnel. In part 1 of the study, the task was purely a navigation task; in part 2, a real firefighting task was included. In both parts, comparisons were made between (1) traditional mission review and rehearsal and (2) traditional mission review and rehearsal plus VE rehearsal. The VE system made use of an SGI RE2, a VR4 HMD, a Polhemus tracking device, joystick motion control, and a "glove avatar" which, together with the joystick, provided both a "fly where you point" metaphor for movement control and a means for opening and closing doors. The VE rehearsal involved a "magic carpet ride" through the space, a user-controlled trip through the space, and a user-controlled trip in which visibility was reduced by the introduction of simulated fire and smoke. The results of the study showed that supplementing the traditional review and rehearsal procedures by the VE rehearsal improved performance. In part 1, the VE group took 1:54 minutes to complete the task, whereas the non-VE group took 2:38 minutes; in part 2, the VE group took 9:26 minutes to extinguish the fire, whereas the non-VE group took 11:43 minutes. In both parts, the VE group took many fewer wrong turns than the non-VE group.

Darken and Banker (1998) explored the use of VEs for terrain familiarization in natural environments, using experienced male map readers with varying amounts of experience in the sport of orienteering. The test task was to locate nine control points in a specified order, all of which were located in a 1200 x 700 m² portion of the former Fort Ord in California. Three training methods were employed: map only, VE plus map, and real world plus map. The map was an orienteering map that showed the environment in great detail. The VE had high visual fidelity, used a desktop PC with a mouse and keyboard, and a very low frame rate (3 frames/sec). The views included a top-down map view with a you-are-here indicator, as well as an egocentric view. All subjects had one hour to study the terrain and prepare for the test. At the start of the test, they had access to a map and compass for one minute to allow for initial orientation, and they were permitted at any time to ask for the map or compass for a 30-second interval. A

differential geographic positioning system (15 meter accuracy) was used to collect positional data on the subjects every two seconds. Although interpreting the data with confidence was relatively difficult, the authors concluded that (1) ability level was more important than training method in determining performance, (2) intermediate ability-level users benefited most from the VE training, and (3) the ability to compress travel time in VEs, and thus to experience the traversal of greater areas in a given time interval, may provide an important advantage to VE training over real world training.

In a related study, Goerger, Darken, Boyd, Gagnon, Liles, Sullivan, and Lawson (1998) considered an architectural rather than natural environment, as well as effects of exposure duration. Whereas subjects in the map group studied floor plans of a seven-story building for 30 minutes, subjects in the map-plus-VE group were simultaneously exposed to a VE representation of the building using a 3-screen, rear-projection system with a 145 degree FOV and a movement control system that made use of verbal commands. Inasmuch as the study period for the latter group was the same as for the former, i.e., 30 minutes, time spent on the VE subtracted from the time spent on the floor plans for this group. Thus, although the map-plus-VE subjects had the additional potential benefit of the VE, they had to pay for this potential benefit by decreasing the time used for studying the plans.

Transfer tests in the real seven-story building involved navigating from an entry point through a sequence of four target locations. In addition, subjects were required to point to various target locations (egocentric estimates), and to construct and navigate a path back to the entry point from the last target location. Finally, after completing the above-described tasks, subjects were required to estimate relative distances and bearings among the targets by placing numbered magnets on a metal whiteboard (exocentric estimates).

The results of this study showed that the map group made fewer route-following errors than the map-plus-VE group, that errors in estimation of direction were roughly the same for the two groups (errors on the egocentric task were roughly double that on the exocentric task for both groups), and that the map group was significantly more accurate than the other group in the

distance-estimation tasks. The extent to which the performance of the map-plus-VE group relative to that of the map group would have improved if the study time had been much longer than 30 minutes (i.e., long enough for the benefits of spending additional time on the floor plan to saturate) was not determined.

In a further related study, Goerger (1998) studied the differential effects of three training conditions on the acquisition of spatial knowledge and navigation performance in the natural terrain orienteering course at the former Fort Ord. The map group was trained by studying the detailed orienteering map; the real-world group was trained by both studying this map and exploring the actual terrain; and the VE group was trained by both studying this map and making use of a VE representation of the terrain. The map used by all groups was specially developed using a digital aerial photo and ground reconnaissance, and was unusually detailed, making use of orienteering symbols and a highly magnified scale (1:5000 compared to traditional military maps of scale 1:25,000 and to competition orienteering maps of scale 1:15,000). The VE system was high-end, making use of an SGI RE-2 workstation, a three-screen rear projection display system with an FOV of roughly 100°, and a 6 DOF joystick for control of movement and orientation. Included in the viewing options was a top-down view, as well as teleportation to a variety of locations within the virtual terrain. Fifteen male subjects (all but one in military service) were given various tests of spatial abilities and then distributed evenly among the three training groups according to results obtained with the Guilford-Zimmerman Spatial Orientation Aptitude Survey (GZ).

In the training phase of the experiment, the subjects were given one hour to study a variety of materials (including the map, the list of tasks to be performed, a course clue sheet, etc.) and to plan a route to navigate through 9 specified control points in a specified order. Whereas the real-world group was given the opportunity to explore the real terrain, the VE group was given the opportunity to explore the virtual terrain. However, as in the study cited above, the total amount of training time was the same for all groups.

In the testing phase of the experiment, the subject had to perform 9 planned tasks

(locating the 9 control points in the proper order) as well as 3 unplanned tasks of the same type. Evaluation of route-following performance included measurement of the number of deviations and the distance deviated from the planned routes. Subject behavior was recorded by use of a differential global positioning system and a helmet-mounted camera, as well as by direct observation by human observers. Determination of survey knowledge was based on directional pointing tasks, unplanned route selection tasks, unplanned navigation tasks, and exocentric spatial tasks using the magnet-whiteboard technique used in the Goerger et al. (1998) study.

Overall, the results failed to show any substantial benefits of using a portion of the training time to explore the real world or the virtual environment; participants in the map group tended to do at least as well as participants in either the real-world group or the VE group. Whereas the real-world group tended to outperform the VE group with respect to route knowledge, the VE group evidenced slightly better performance than the real-world group with respect to survey knowledge. It was also found that the results of the tests of spatial abilities applied were positively correlated with the results obtained in the natural terrain spatial tasks considered.

In a study by Waller, Hunt, and Knapp (1998), the effects of different training conditions on spatial knowledge of a 14 x 18 foot maze was explored. Test tasks in this study included navigating the maze while blindfolded and taking a true-false test in which subjects were required to determine whether a given map correctly represented a portion of the maze. Six training conditions were considered: (1) blind (no training), (2) real world training, (3) map training, (4) VE desktop training, (5) VE immersive training (2 minute exposure), and (6) VE long immersive training (15 minute exposure). The results of this study for the blindfolded navigation task (measured in terms of time to completion of task) suggested that all the forms of training tested were useful, and that the real world and VE-long training were somewhat better than the map, desktop, or VE-short training. The results for the true-false test showed best performance for the map group and worst for the VE-short group. Although there was a slight gender effect in the true-false test (men 69% correct, women 64% correct), and a substantial

gender effect in the maze performance for the VE training conditions, there was no gender effect in the maze performance for the real-world training condition. It is also worth noting that the Guilford Zimmerman Test of Spatial Orientation was not predictive of performance for most of the subjects.

In addition to the above studies, there have been many studies of spatial behavior that involved the use of VEs but were not explicitly concerned with training transfer to the real world. Brief comments on most of these studies are provided in the following paragraphs.

Henry (1992), concerned with spatial perception in architectural environments, found judgments of room size to be smaller, and judgments of angle to be more variable, in a simulated environment than in the real environment (although systematic biases in angular judgments were found to occur in the real environment as well as the simulated one).

Peruch, Vercher, and Gauthier (1995), in a study comparing spatial knowledge in a VE acquired through active free exploration with that acquired through passive viewing, found (consistent with much classical work) that actual exploration was superior.

May, Peruch, and Savoyant (1995) showed that map misalignment (deviations from the natural correspondence between "up" and "forward") reduced speed and accuracy of navigation in a VE.

Satalich (1995) compared a number of methods for exploring a VE (self-exploration, active guided, passive guided) to a control condition in which subjects only had access to a map of the space. Results on orientation tasks, distance estimation tasks, and wayfaring tasks in the VE generally showed that performance with the control condition was better than or equal to performance under any of the conditions involving experience in the VE.

Aginsky, Harris, Rensink, and Beusmans (1996) studied route learning in a driving simulator. They found the environmental information in the vicinity of choice points is more likely to be retained than information at other points and that two strategies emerged: a "visual strategy" which relies primarily on visual recognition of actual intersections along the route and a "spatial strategy" which relies on a mental map that incorporates the environment's spatial

structure.

Darken and Sibert (1993, 1996a, 1996b) performed a variety of studies directed towards improved navigation in VEs, particularly large-scale ones. They developed a navigation tool set and examined the use of various environmental cues and navigational aids in a complex nautical search task, using performance measures that included search time, percentage area searched, and errors made in map production. In general, they found that principles previously developed in connection with navigation in real-world spaces could be usefully applied to the problem of navigating in VEs.

Stoakley, Conway, and Pausch (1995) implemented and explored the use of a virtual hand-held miniature model of the VE (referred to as a "world in miniature," or WIM) that provides the user with an exocentric view of the VE that can be changed at will by simply rotating or zooming in on the model. In many respects a WIM can be regarded as an extension and elaboration of a map that effectively exploits modern VE technology. Among the many ways in which a WIM can be exploited in a VE (a large selection of which are discussed by the authors), some are clearly useful for improving spatial behavior in the VE.

Tlauka and Wilson (1996), interested in differences in spatial knowledge achieved by means of maps and by means of navigation through the space mapped, found that bearing judgment errors by the map-group subjects were smaller than those by the navigation-group subjects (in contrast the results of Thorndyke and Hayes-Roth, 1982, who found directional judgments to be less accurate for a map group than a navigation group in a real-world environment), but that there was no significant effect of alignment in either group.

Pausch, Proffitt, and Williams (1997) examined the relative effectiveness of immersive HMD displays and desktop displays when searching for targets in a heavily camouflaged room. It was found that task completion time (time to find target) was roughly the same for the two types of display on trials in which the target was present, but that it was smaller for the HMD display on trials when no target was present (time to search the entire room and conclude that no target was present).

Colle and Reid (1998), concerned with the acquisition of configuration knowledge from exploration and navigation in a VE, found that recall of direction information was more accurate when the situation to be recalled was one in which the object to be pointed at by the subject was in the same virtual room as the subject rather than in a separate room (the so-called "room effect"). It was also found that pointing responses based on the analysis of maps that the subjects had drawn following their experiences in the VE were more accurate than those based simply on recall of the VE experience itself.

Witmer and Kline (1998), in a study concerned with the perception of distance in virtual environments, examined distance estimates achieved from static visual cues and from cues obtained by traversing the distance to be estimated. In the static case, estimates of the virtual target distance were only 47% of the "correct" virtual distance compared to a figure of 72% in comparable real-world tests. However, the scatter of the data in the VE tests was smaller than in the real-world tests. In terms of resolving power (as opposed to response bias), the VE results appear to have been at least as good as the real-world results. In the distance-traversed experiments (with traversal achieved by means of a joystick, a treadmill, or teleportation), the underestimation of distance was less severe; however, surprisingly, transversal by means of the treadmill did not produce better results than transversal by joystick or by teleportation. It was also found, as one would expect, that an auditory cue which reported distance traveled (one beep every 10 feet of movement), drastically reduced both the bias and the scatter evident in the data.

Ruddle, Payne, and Jones (1997, 1998, 1999), concerned with navigation in large complex virtual buildings, examined a variety of effects, including the effects of navigational experience and orientation aids. In addition, they compared the performance obtained when using an HMD to that obtained when using a desktop monitor.

In the 1997 study, an experiment was performed using a desktop VE system that paralleled in considerable detail the experiment conducted previously in the real world by Thorndyke and Hayes-Roth (1992). Among other things, Ruddle et al. found that route-finding ability improved substantially with navigation practice in the VE, that such experience produced

more accurate estimates of VE route distance than VE Euclidean distance, that estimates of direction and distance based on navigation experience were slightly less accurate than the equivalent results in the Thorndyke and Hayes-Roth experiment, that subjects with map experience rather than navigation experience made more accurate judgments of Euclidean distance in the VE, and that the accuracy of the map group estimates in the VE were roughly equivalent for both distance and direction to those of the map group in the Thorndyke and Hayes-Roth study. The navigation group in the VE study, although it showed great variability in the distance estimates, showed no overall tendency to underestimate or overestimate distance; however, these subjects were provided with specific scale information (certain distances were specified in feet or meters). Two further experiments in the 1997 study focused on the benefits provided by different types of landmarks. Based on all three experiments, the authors concluded that alternative navigational aids, beyond presentation of landmark cues, need to be developed for successfully navigating in VEs. Included as candidates for such a purpose were virtual suns, compasses, VE maps, and WIMS (Stoakley et al., 1995).

In the first experiment in the Ruddle et al. 1998 study, a within-subjects design was used to study errors in orientation estimates for simple paths containing 1, 2, or 3 turns. Results showed an increase in mean error from about 20 degrees to 33 degrees as the number of turns increased, independent of whether the VE desktop system used employed a 45 degree or 90 degree field of view (compared to a mean error of 29 degrees in the 1997 study). In a second experiment, performed in a large-scale VE using the 90 degree FOV, it was found that route-finding ability and survey knowledge (VE orientation and VE Euclidean distance) improved with experience in a given VE, as well as transferring to some extent from a first to a second (distinct but somewhat similar) VE. However, no effect of providing a compass to the subjects was found on the various performance measurements, although the availability of the compass did appear to influence certain aspects of behavior during the tests and also the overall comfort of the subjects. In discussing these results, the authors suggest that the results of Presson and Montello, 1994, and Rieser, 1989, concerning differences in the maintenance of directional information after real

or imagined rotations, might imply perhaps that global orientation can be better maintained in immersive VEs than in desktop VEs. They point out, however, that in one of their own previous studies (Ruddle, Randall, Payne, and Jones, 1996), route-finding performance and accuracy of direction estimates were roughly the same for the two kinds of systems.

In the 1999 study, performance in navigating complex building VEs with an HMD was compared to performance using a desktop monitor. With the HMD (VR4, 247 x 230 resolution, Polhemus head tracker, 50 degree FOV), direction and speed of movement was controlled by means of a handheld button box. With the desktop display (21", 1280 x 1024 resolution, 40 degree FOV), viewing direction was controlled by means of a mouse and movement direction and speed by use of a keyboard. With both displays, viewing direction and movement direction were decoupled. No stereographics were employed in either display. Each of twelve subjects first navigated around one virtual building four times using one display (HMD or desktop) and then navigated around a second virtual building four times using the other display (with appropriate counterbalancing of display order, building order, and display/building combination). The task required of the subject in each of the four runs assumed greater familiarity with the building as the run number increased from 1 to 4, and included estimation of directions and straight-line distances during run 4. The results of this study showed that the HMD led to slightly more efficient route-following behavior in terms of travel time, but not in terms of travel distance. Increased travel time with the desktop display appeared to be related to an increased tendency to stop before altering direction. With respect to acquisition of survey knowledge, the HMD was found to produce more accurate estimates of relative straight-line distances (there was no consistent tendency to underestimate or overestimate these distances with either display); however, the mean orientation errors observed (which, in general, were quite large, ranging from about 40 degrees to 65 degrees) were insignificantly smaller for the HMD case.

Chance, Gaunet, Beall, and Loomis (1998) examined how directional estimates in virtual mazes are affected by the mode of locomotion used to move through the maze. In the walk mode, subjects walked through the maze wearing an HMD with tracking information being used

to update the visual imagery. In the visual-turn mode, subjects moved through the virtual maze using a joystick to control turns. In the real-turn mode, regarded as intermediate between the first two modes, subjects physically turned in place to steer while translating through the maze. Although the errors in the absolute directional estimates (from terminal position to various target positions) tended to be very large, often exceeding 60 degrees in tests for which totally random behavior would have led to a mean absolute error of 90 degrees, there was some indication that proprioceptive information was useful in making the directional estimates, e.g., the errors for the walk mode tended to be somewhat smaller than the errors for the visual-turn mode. The walk mode was also less likely to cause simulator sickness in the tests.

Loomis, Golledge, and Klatzky (1998), in a study directed towards the development of a navigation aid for visually-impaired persons, examined guidance performance as a function of four display modes, one using spatialized sound from a virtual acoustic display and three involving verbal commands issued by a synthetic speech display. Using time-to-complete-course and distance-traveled as performance measures, they found that performance was best with the virtual acoustic display, and that providing verbal guidance information without appropriate heading information caused a significant degradation in performance.

In general, the number of studies concerned with spatial behavior that involve the use of VEs is increasing very rapidly. However, as indicated above, most of these studies focus on improving navigation in VEs and/or in using VEs as a research tool for studying spatial behavior. Relatively few are focused on the use of VEs to train spatial behavior in the real world (either in a specific space or in general) or, equivalently, on issues related to training transfer to the real world. The extent to which the development of methods or aids that improve spatial behavior in VEs (but are not available to the subject in the real world) will lead to improved spatial behavior in the real world, is, of course, an open question and is likely to depend on the specific method or aid in question. In addition, most of the previous work on transfer to the real world is focused on route knowledge, not configuration knowledge. And little attention has been given to the use of 3D maps or WIMs in the VE training.

3 The Experiment

3.1 The Real Space

The space chosen for our preliminary study consisted of a portion of the 7th floor of Building 36 at MIT. This space, the layout of which is shown in Fig. 1, was chosen for convenience; it is easily accessible and close to the Virtual Environment facilities in our laboratory. Although the space is relatively simple (compared, for example, to the complex spaces employed by Ruddle and associates), we believed it was adequate for our preliminary work.

3.2 The Test Task

The test task employed was chosen to probe subjects' acquisition of configurational knowledge. In general, we regard configurational knowledge as more important than route knowledge for two reasons. First, there exist tasks for which performance depends more on configurational knowledge than on route knowledge (e.g., designing a ventilation system for a building, planning to demolish a building, etc.). Second, whereas configurational knowledge may imply route knowledge (including knowledge of secondary routes when primary routes are blocked), route knowledge does not generally imply configurational knowledge.¹

For all training conditions considered (see Sec. 3.3, below), the subject, located at a given reference position ("station") within the real space, was required to estimate the location of a given landmark within the real space by reporting the bearing (θ) and the range (R) of the landmark relative to the location of the station. Thus, the type of knowledge being probed in these tests was Euclidean-metric configurational, rather than, for example, topological or city-block-metric configurational.

The stations and landmarks used are shown in Fig. 1. Of the 20 possible pairs of stations

¹One could argue that configurational knowledge does not necessarily imply familiarity with all landmarks relevant to route knowledge, or that route knowledge, if sufficiently complete (and sufficiently metric) does imply extensive configurational knowledge. However, we believe that, at least to a first-order approximation, our statement is correct.

and landmarks (4 reference locations x 5 landmarks), only the 13 shown in Table 1 were used in the experiment. In all these cases, the landmarks were obscured from view by walls.² The various (station, landmark) pairs will be referred to as "subtasks" and identified by number, consistent with the numerical assignments shown in Table 1.

TABLE 1: Subtasks (about here)

During testing, subjects were brought to the stations in the order in which they are labeled (A, B, C, D), and were transported from one station to the next in a wheelchair while blindfolded. The blindfold was then removed during a subject's attempt to estimate the location of the landmarks. Directional estimation was accomplished by use of a pointing rod and protractor mounted on a tripod. Feedback was not provided to subjects about errors made in estimates of bearing or range. The total time consumed by each subject to perform the 14 subtasks was less than one hour. The order in which the subtasks were performed is that shown in Table 1.

3.3 The Training Conditions

Four training conditions were employed, each with a different set of subjects:

- | | | |
|-----|-----|------------------|
| (1) | RW | Real World |
| (2) | VE | Immersive VE |
| (3) | NVE | Non-Immersive VE |
| (4) | Mod | Model |

For all these training conditions, the subjects were informed about the task they would be asked to perform after the training was completed, although they were not informed about the specific locations to be used as references or the specific objects to be used as landmarks. For all conditions, training included 10 minutes of free exploration under the specified training

²Responses were also collected for the case (A, Water Fountain) as a benchmark, the only case in which the landmark was visible from the station. These data are discussed at the end of Sec. 4, below.

condition. For the RW Condition, this was the only training provided. In the other three cases, an opportunity to become familiar with the technology was provided prior to the training period in which the synthetic version of the 7th floor space was explored. Subjects in these groups were familiarized with the equipment by permitting them to explore a vastly simplified VE unrelated to the test space.

In the VE condition, subjects used a headmounted display (HMD) for visualization and a joystick for navigation. The NVE condition used the same equipment and procedures, except that a 21-inch monitor was used to view the simulation, rather than the HMD. For both conditions, the joystick was configured so that a forward push moved the subject forward in the scene, a backward push moved the subject backward, and left and right pushes rotated the subject in the corresponding direction. In the Mod condition, subjects were provided with a virtual miniature 3D model of the space that could be manipulated using a mouse and monitor (similar in some ways to the WIM of Stoakley et al.). Subjects were able to rotate and view the model from any vantage point as well as zoom in closely to examine specific details. The graphical model of the space used for this condition was essentially identical to that used in the two virtual "walkthrough" conditions VE and NVE mentioned above. The only differences were the point-of-view change and removal of the ceiling so that the "internal space" could be seen from the outside. The viewpoint for the Mod condition was controlled as follows: to translate the model up (down), click the right mouse button with the cursor at the top (bottom); to rotate the model left (right), click the middle mouse button with the cursor at the left (right); and to zoom in (out) on the model, click the middle mouse button with the cursor at the top (bottom).

3.4 Technical Details

The virtual environments in this study (VE, NV, and Mod) were run on modified Easyscene software from Coryphaeus running on an SGI Onyx RE/2. The Onyx was equipped with 128 MB of RAM and two R4400 processors operating at 150 MHz. The Easyscene software was modified to allow for first-person navigation through the developed model, with collision detection and joystick support added. Modifications were also made to enable support for the

HMD and adaptation of multiple viewpoints and control methods for the various experimental training conditions.

The HMD used in the VE condition was a Virtual Research VR4 (horizontal resolution 350 lines, vertical resolution 230 lines, 60 degree field of view, weight 33 ounces). The HMD was not stereo-enabled for this study.³ A Polhemus Fastrak provided orientation and position information on the HMD. The Polhemus positional information was discarded and the image presented by the HMD was controlled by the Polhemus orientation variables and the joystick variables. In the NVE condition, the 21-inch monitor was operated in a 640 x 480 graphics mode and the image was controlled solely by the joystick. In the Mod condition, the same monitor was used; however, the joystick control was replaced by the mouse control.

An architectural model of the 7th floor of building 36 was constructed from blueprints using Coryphaeus' Designers Workbench software. Accessible areas of the virtual environment were modeled with open doorways. Any closed doors in the virtual environment signified inaccessible areas irrelevant to the experiment. Motion of the subject was confined to the X and Y plane. Each room in the model was populated with objects created from within Designers Workbench. All objects were fully texture mapped using Designer's Workbench. Actual textures were obtained by using a Nikon N6006 camera with film scanned at 1280 x 1024 or by using a Kodak DC20 digital camera. Texture maps were edited using Kodak PhotoEasy software and Adobe Photoshop, and imported into Designer's Workbench for application to the 7th floor model.

Figures 2 and 3 show illustrative views of the space obtained in the Mod training condition. Figure 4 shows a view of the space by the elevator doors obtained in the VE (or, equivalently, the NVE) condition. This same region of space can be seen in the lower left corner of Fig. 2.

3.5 Subjects

³Because of all the other depth cues available, such as those provided by perspective, it seemed unlikely that stereoscopy would have had much of an effect on performance in this experiment. See also the comments in Brooks, 1992.

Thirty-six MIT students were recruited for in this study. Data from one subject was excluded due to apparent disorientation in all subtasks. Of the remaining 35, 10 were trained using the RW condition, 10 using the VE condition, 7 using the NVE condition, and 8 using the Mod condition. (The unequal numbers were not part of the experimental plan; they resulted from time constraints imposed on the duration of the experiment).

4 Results

In this section, we present the main results of our experiments. Tests of statistical significance are presented in Appendix A.⁴

Fig. 5 shows plots in the X,Y plane of the estimated and actual locations of the landmarks relative to the stations for each of the 13 subtasks specified in Fig. 1 and Table 1. Each point represents the response of one subject; the different symbols differentiate the training groups. Figs. 6 and 7 present histograms of the bearing and range errors segregated according to training method but pooled over subtasks and subjects to give roughly 100 responses per histogram.⁵

Because errors in range estimation tend to increase with range in a more or less consistent

⁴ We have relegated the results of these tests to an appendix because we do not regard them as crucial to the discussion. Although we certainly want to avoid emphasizing results that are not statistically significant, we do not believe that results which are statistically significant are necessarily important. For example, a difference in the mean bearing error for the conditions VE and NVE might be statistically significant but fail to be of interest because the statistically significant difference in means may still be small relative to the standard deviations of the two groups. The fact that for the given number of trials the difference between the means is large enough relative to these standard deviations to make the difference statistically significant does not imply that it is large enough to produce a value of the sensitivity index d' that reaches a usefully-defined difference threshold. Stated differently, the fact that the difference in the means is statistically significant does not imply that knowing to which population an item belongs produces a significantly positive value of information transfer. Also, of course, a statistically significant difference in the means between the VE and NVE groups might fail to be of interest because it is small relative to other types of differences between VE and NVE (e.g., related to comfort or cost).

⁵ Two points should be noted about these histograms. First, in order to take account of the fact that the total number of trials for each training method was not the same, we have plotted percentage of responses in a bin rather than the number of responses in the bin. Second, slightly more data are included in this histograms than in the X,Y plots (less than 2% more) because the X,Y plots only show data when estimates of both range and bearing were available and only when the point fell within the X,Y region displayed.

fashion (i.e., the ratio of the range error to the actual range is relatively independent of the actual range, consistent with Weber's law), range data were examined in terms of the logarithm of the ratio of the estimated range to the actual range, so that a positive error corresponds to an overestimation of the range. For each of the two estimate components, bearing and range, we examined both the signed errors and the absolute values of these errors. Whereas in all cases the range errors are unitless (because they are expressed in terms of a ratio), the bearing errors are throughout the paper always expressed in degrees.

Figs. 8 and 9 show contour plots of the signed errors as a function of both training-condition/subject and subtask.

Table 2 presents means and standard deviations. The entries in square brackets in each subsection of the table give the means and standard deviations corresponding to the pooled histograms shown in Figs. 6 and 7. A detailed description of all the entries in this two-page table is provided in the legend for this table.

TABLE 2: Means and Standard Deviations (see legend) (about here)

The first result that is immediately evident from these data is that there are few, if any, major differences across the different training conditions (see the histograms in Figs. 6 and 7, and the summary of the grand means and standard deviations corresponding to these histograms shown in Table 3). There appears to be a slight overall bias in bearing estimate of roughly -2 degrees, and in log range estimate of roughly -0.06 (corresponding to a range estimate that is only 85% of the true range). Results from the RW condition are worse than results from all the other conditions with respect to both mean errors and standard deviations, for both bearing errors and range errors, and for both signed and absolute errors. The differences that appear most pronounced are in the log range data: the mean error in the signed data is relatively large for both RW and Mod, the mean error in the absolute data is relatively large for RW, and the standard deviation of the RW data is relatively large for both the signed and absolute errors. It appears

from the data shown in Figs 5 and 6 and Table 3 that the synthetic environments used for training are no worse than the real environment, and that there are no dramatic differences in the training effectiveness of the different synthetic environments. With respect to the latter of these statements it should be noted that all the values of the sensitivity index d' corresponding to differences among training groups (difference in means divided by the average of the standard deviations) are less than 0.3 for all the bearing data, less than 0.4 for the signed log range data, and less than 0.6 for the absolute log range data (the largest value being obtained for RW vs. NVE in the absolute log range data).

TABLE 3: Grand Means and Standard Deviations (about here)

A second important result to note is that the variance in the overall pooled results cannot be interpreted simply as repeated-measures variability. In other words, it is not the case that the scatter in the data can be explained by assuming a single underlying random variable with a probability density that is blind to differences across subtasks and subjects.

The most obvious structural element in the data that contradicts such a simple model is the remarkably strong dependence of signed bearing error on subtask. This dependence is evident in the X,Y plots shown in Fig. 5, in the contour plots shown in Figs. 8 and 9, and in the results listed in Table 2. For example, whereas there is a large negative bearing error on subtask 3, there is a large positive error on subtask 5. According to the results listed in table 2. the mean signed bearing errors on subtask 3 for the four training groups are -16.1, -21.1, -19.8, and -12.7, whereas the comparable results for subtask 5 are 14.0, 16.8, 18.6, and 7.5. On the average (across training groups), the means and standard deviations for these two tasks are: $m(\text{task 3}) = -17.3$ degrees, $m(\text{task 5}) = 14.2$ degrees, $\sigma(\text{task 3}) = 7.9$ degrees, and $\sigma(\text{task 5}) = 11.2$ degrees. Using these averages to make a crude estimate of d' , one obtains a value of $d'(\text{task 3, task 5})$ greater than 3. Obviously, there is a strong effect of subtask in the signed bearing errors.

Further results concerning the role played by subtask are presented in Tables 4 and 5.

Table 4 shows the correlation coefficients between pairs of training groups for the dependence of mean errors (averaged across subjects in the group) on subtask, for each of the four error types (signed bearing error, absolute bearing error, signed log range error, and absolute log range error). the values of the correlation coefficient ρ for the signed bearing errors ($0.82 \leq \rho \leq 0.96$; Ave $\rho = 0.91$) are remarkably high. Apparently, something about the structure of the subtasks played a very powerful role in determining the values of the signed bearing errors. The absolute bearing errors and the signed log range errors also show some substantial correlations, but they are not nearly as pronounced.

TABLE 4: Correlation ρ between pairs of Training Groups for the Dependence of Mean Errors (across subjects) on subtasks (about here)

Table 5 compares the overall standard deviations (shown inside the square brackets in Table 2) to the standard deviations arising from variations across subjects (the mean value of column (c) in Table 2) and the standard deviations arising from variations across subtasks (the mean value of column (e) in Table 2). As can be seen in the results for signed bearing errors, most of the variation in the pooled results (characterized by the grand σ) is the result of variation across subtasks. For example, in the RW condition, for which the grand σ is 15.0, the σ associated with variation across subtasks is 14.0 whereas the σ associated with variation across subjects is only 9.9. This is also true for the absolute bearing error, but to a much smaller extent. As was evident in the above correlation analysis (Table 4), it is also clear from this table that subtask plays a smaller role in range estimates than in bearing estimates. A way of describing the strong dependence of signed bearing error on subtask is outlined briefly in Sec. 5.

It is also clear that the results depend strongly on the subject. Because the tests with the different training methods used different subjects, one cannot examine the correlation between methods for the dependence of errors (averaged over subtasks) on subject. However, as would be

expected on the basis of previous research in this area, and as can be seen by examining the means and standard deviations in Table 2, intersubject differences constitute a substantial source of overall variance.

The fact that variation among subjects plays an important role, particularly in the range estimates, is also evident in Table 5. In contrast to the manner in which the intersubtask standard deviation (σ_{subt}) is the main contributor to the grand standard deviation (Grand σ) for the signed bearing errors, the intersubject standard deviation (σ_{subj}) is the main contributor to the grand variance for the signed log range errors.

TABLE 5: Comparisons of Grand Standard Deviations to Standard Deviations Across Subjects and Across Subtasks (about here)

The manner in which the bearing and range data relate to the variations across subtasks and subjects is further displayed in Table 6.

TABLE 6: Values of $\sigma(m) / m(\sigma)$ (about here)

In the upper portion of this table we have tabulated the values of $\sigma(m) / m(\sigma)$, where $\sigma(m)$ denotes the standard deviation of the means shown in column (b) of Table 2, and $m(\sigma)$ denotes the mean of the standard deviations shown in column (c) of Table 2. Large values of $\sigma(m) / m(\sigma)$ correspond to the case in which the differences among different subtasks are large relative to the variation across subjects. In the lower portion of the table, we have tabulated the values of $\sigma(m) / m(\sigma)$, where $\sigma(m)$ denotes the standard deviation of the means shown in column (d) of Table 2, and $m(\sigma)$ denotes the mean of the standard deviations shown in column (e) of Table 2. Here, large values of $\sigma(m) / m(\sigma)$ correspond to the case in which the differences among different subjects are large relative to the variations across subtasks. Loosely speaking, the values of $\sigma(m) / m(\sigma)$ in the top portion of the table can be interpreted as average values of

d' between pairs of subtasks, and in the bottom portion, as average values of d' between pairs of subjects. As can be seen from Table 6, as well as from Table 5, intersubtask differences are more pronounced for bearing errors, whereas intersubject differences are more pronounced for range errors. For example, confining attention to signed errors, one finds that the mean d' for bearing errors is roughly 4 times greater between subtasks than between subjects (1.21 compared to 0.34), whereas for log range errors it is roughly 4 times smaller between subtasks than between subjects (0.30 compared to 1.21). The same phenomenon is evident graphically in Figs. 8 and 9: the striations in the bearing errors are primarily horizontal, whereas the striations in the range errors are primarily vertical.

In addition to the above observations, it is important to note that for a given subject, the correlation between bearing errors and range errors is relatively weak. For the case of signed errors, the correlation ρ lies in the range $-0.56 \leq \rho \leq 0.80$ for the 35 subjects ($Ave = 0.22$, $\sigma = 0.32$). For the case of absolute errors, the range is $-0.60 \leq \rho \leq 0.65$ ($Ave = -0.01$, $\sigma = 0.35$). Although for any one subject the correlation may be significant, across subjects there is no systematic trend in the relationship between either signed bearing errors and signed log range errors or between absolute bearing errors and absolute log range errors.

As indicated in Sec. 3, we did not in this preliminary experiment pay attention to the manner in which subjects were assigned to groups, nor did we administer any spatial-abilities tests to the subjects selected. Thus, we are not in a position either to interpret intersubject differences or to guarantee that our results are not biased by the composition of the various training groups.

A cursory examination of the effects of gender in the RW and Mod groups (the first of which contained 6 males and 4 females, the second of which contained 4 males and 4 females) showed no pronounced gender effect in either the means or the variances, except possibly for the mean results in the signed log range data which suggested a possible tendency for females to underestimate distance to a slightly greater degree than males. This result is consistent with the difference between these two groups and the VE and NVE groups (each of which contained only

1 female) with respect to the mean signed log range errors shown in Table 3. Whereas the first two groups had a mean error of -0.09, the last two had mean errors of -0.02 and -0.03. (In both the RW and Mod groups, the mean log range error for the males was -0.06 and the mean log range error for the females was -0.12.) Obviously, serious study of individual differences related to gender or any other subject descriptor will require extensive pretesting of spatial abilities and systematic subject selection, as well as an increased number of subjects.⁶

Finally, in considering the above results it should be noted that before these data were collected, each subject was required to perform the task of estimating the bearing and range of a landmark that was visible from the station (the Water Fountain from station A). As one might expect, the bias in signed bearing estimate for this visible-landmark task tended to be very small: the mean estimated bearing for each of the training groups differed from the actual bearing (23 degrees) by less than 1 degree. In contrast, however, the bias in the range estimate was surprisingly large: the mean signed log range error for the groups was -0.20 (RW), -0.11 (VE), -0.08 (NVE), and -0.21 (Mod). Not only do these errors seem large on an absolute basis (e.g., -0.21 corresponds to a range estimate that is only about 60% of the actual range of 25 feet), but they are large relative to the mean errors encountered in the tests analyzed above in which the landmarks were always hidden from view and the estimates depended on the subject's memory. [The corresponding figures for these tests according to the results shown in Table 3 were -0.09 (RW), -0.02 (VE), -0.03 (NVE), and -0.09 (Mod)]. This suggests, perhaps, that (1) the subjects had an inflated picture of the length corresponding to 1 foot, and (2) with respect to this inflated unit of length, they tended to *overestimate* the range of the landmarks that were not visible. Alternatively, one might simply conclude that landmarks tend to be judged as nearer when they are visible than when they are hidden from view.

⁶We have included these comments on possible gender effects only because we know that if we had omitted them, many readers would have wondered about the gender composition of the groups and how it related to the results. The study was not designed, and the results cannot be used, to evaluate gender effects in a statistically meaningful manner.

In order to determine the extent to which the estimates in this visible-target task (which we refer to as subtask 0) were correlated on a subject-by-subject basis with the mean value of the estimates for the invisible-target tasks 1-13, correlation coefficients were computed for both bearing and log range, for both signed and absolute errors, and for both the overall pool of subjects and the individual training groups RW, VE, NVE, and Mod. The results of these computations (see Table 7) show that the correlations are relatively high for the RW log range data (both signed and absolute) and, to a slightly lesser extent, for the Mod signed bearing data and the log range data (both signed and absolute). The case that exhibits the lowest correlation is the absolute bearing errors.

TABLE 7: Correlations with Subtask 0 (Target Visible) (about here)

In order to determine the extent to which the results would have changed had we normalized them by subtracting out the errors made on subtask 0 (on a subject by subject basis), we recomputed the means and standard deviations for each of the cases RW, VE, NVE, and Mod, for both the signed and absolute versions of these errors. The results of these further computations are shown in Table 8.

Overall, the changes due to normalization do not seem large. The change from negative to positive results for the means of the signed log range data is consistent with the fact, mentioned previously, that the underestimation of distance was exceptionally large for subtask 0. Also, the reduction in the standard deviation for RW in the signed log range case appears consistent with the large positive correlation for this case shown in Table 7. It is interesting to note, however, that large correlations in Table 6 for the RW absolute log range case and for both the signed and absolute versions of the Mod log range case are not accompanied by significant changes in σ when the data are normalized. It should also be noted that normalization increases σ in a number of cases (e.g., the VE absolute log range case). In general, none of the main results previously discussed appear to be substantially altered by the normalization process.

TABLE 8: Results of Normalizing Data According to Estimates Made in Subtask 0 (about here)

5 Discussion

In general, it is clear from the results summarized in Sec. 4 that differences among the training methods considered are trivial in comparison to the variability of the data resulting from intersubtask and intersubject differences. Training with the use of a virtual environment appears to be as effective as training with the real environment, and there appear to be no dramatic differences among the results obtained with the different virtual environment systems.

Although we do not find these general results surprising in the light of the previous work cited in Sec. 2, it is difficult to make meaningful detailed comparisons with these previous studies because of the many important differences between these studies and our own. Not only were relatively few of the previous studies concerned with training transfer from the virtual environment to the real world, but even within this small group of studies, the training test conditions were radically different from our own. For example, in the studies by Witmer, Bailey, and Knerr (1995), Witmer, Bailey, Knerr, and Parsons (1996), Bliss, Tidwell, and Guest (1997), and Tate, Sibert, and King (1997), the focus was on the acquisition of route knowledge rather than configurational knowledge.

If we had been forced to rate the training conditions in advance based on our intuition, we would have guessed that RW would have been best, NVE worst, and VE and Mod intermediate. We would have guessed that VE would be superior to NVE because of the often cited advantages of immersion; we would have guessed that Mod would do relatively well because of its similarity to a 3D map (and the work by Stoakley, Conway, and Pausch, 1995); and we would have guessed that RW would have been best, given the technical limitations of current simulation systems. Eventually, we expect that a combination of a good VE system and a good Mod system (and an appropriate method for transitioning between the two) will prove substantially superior to RW, and that this superiority will be evident for a much wider variety of

test tasks than those considered in our preliminary experiment.

Additional results concern the nature of the response errors. Obviously, it would be worthwhile developing methods to eliminate them. Furthermore, to the extent that they cannot be eliminated, it would be worthwhile building models to predict them. In general, such efforts will serve both to improve spatial behavior and to improve our ability to distinguish between different training methods by reducing variability.

As discussed in Sec. 4, bearing errors were dominated by the subtask, more or less independent of subject and training method. Furthermore (see Fig. 5), the bias in bearing estimation can be summarized by stating that subjects tended to "square the rectangle" (subjects thought that the rectangular space was wider than it actually was). Similarly consistent biases in bearing estimates have been noted previously, e.g., see Henry, 1992.

Also, as shown in Sec. 4, range errors usually consisted of underestimates, and, unlike the bearing errors, the variability in these estimates was dominated by intersubject differences, not by intersubtask differences. Although underestimation of range was to be expected on the basis of previous work (e.g., see Henry, 1992; Witmer and Kline, 1998), we were surprised that the RW group showed the most pronounced bias.

These biases in the bearing and range estimates are clearly of interest from the general viewpoint of the psychology of spatial perception. However, we suspect that they will prove to be relatively unimportant from the viewpoint of VE-assisted spatial training, because we believe that it will be possible to train subjects to eliminate these biases. In particular, we see no reason why a subject who consistently underestimates range cannot learn to achieve better accuracy by correct-answer feedback training. Moreover, we believe that effecting this correction in one spatial environment would transfer to most other environments. To the extent that this notion is correct, and to the extent that attention is focused on response bias, it matters relatively little whether the distance response chosen for the experiment is phrased in terms of "estimated feet to the target" or in terms of something that is more ecologically valid. If the bias is easy to eliminate, the source of the bias is relatively unimportant.

Bias in the bearing estimates may be less easy to eliminate, because it appears to be less homogeneous and more dependent on the subtask. In order to correct bias, it must be possible for the subject to develop a relatively simple mental model of how the bias depends on the situation. Whereas in the range case, it appears that much of the bias could be removed if the subject merely learned to increase the range estimate by a fixed percentage, no such simple correction rule has yet been demonstrated for the bearing case. Whether or not it is possible to find a rule that is sufficiently simple to be of use in training subjects to eliminate the bearing bias in a wide range of spatial situations is a topic for future study.

Assuming that a subject's response bias in range and bearing can be eliminated or reduced by appropriate training, one is still faced with problems related to a subject's intrinsic response variability in estimating range and bearing, and to intersubject differences of various types. It may be that training can reduce intrasubject and/or intersubject response variability. However, even if it can't, it may be possible, through the use of appropriate spatial-abilities tests, to predict some aspects of a subject's performance in the experimental spatial-behavior tasks. To the extent that this is possible, it too, like the elimination of bias, would enable one to sharpen one's ability to distinguish between different training methods and select the one most likely to be appropriate to the requirements of the specific training situation.

In conclusion, two points need to be stressed. First, to the extent that the issues related to response errors that have been discussed are independent of the choice among RW, VE, NVE, and Mod as training conditions, it doesn't matter which of these conditions are used in future work addressed to these issues. Second, our preliminary experiment was very limited in that only one, relatively simple, space was used; no intrasubject comparisons of training conditions were made; the training experience of the subjects was not varied systematically and in a controlled manner (other than the assignment to one of the conditions RW, VE, NVE, and Mod); and the range of tests used was highly constrained. With respect to this last point, it should be noted not only that attention was confined to configurational knowledge, but also that the tests of configurational knowledge were very limited. No map-drawing tests were included, and there

were no attempts to explore a subject's ability to make spatial estimates that were not anchored at the location at which the subject was standing (e.g., estimating the bearing and range to a target from a different location, or estimating whether or not a clear line of sight exists between two locations, neither of which coincided with the position of the subject). As recently stressed by Jim Templeman, Jack Loomis, and Rudy Darken (personal conversations), the inclusion of such responses would be of interest with respect to both basic research issues and VE training applications.

In future work concerned with training spatial behavior in specific spaces, we intend not only to study more complex spaces, but also to develop methods for decreasing response bias and scatter in the experimental data, to include a wider variety of spatial behavior measures, to explore the possibility of creating a general VE spatial abilities test to help characterize the individual subjects used in the experimental studies, and to create and evaluate a VE training system that combines an immersive walk-through VE (with an improved motion interface) and a miniature-model VE in such a manner that the training provided by this system clearly exceeds that provided by real-world training.

ACKNOWLEDGMENTS

This work was supported by the Office of Naval Research, Grant # N00014-96-1-0379. We are indebted to Terry Allard, Rudy Darken, Jack Loomis, and Jim Templeman for many useful discussions concerning the use of VE for training spatial behavior. We also wish to acknowledge many useful comments by the anonymous reviewers.

REFERENCES

- Aginsky, V., Harris, C., Rensink, R., & Beusmans, J. (1996). Two Strategies for Learning a Route in a Driving Simulator. Cambridge Basic Research Center, Nissan Research and Development, TR 96-6.
- Bliss, J.P., Tidwell, P.D., and Guest, M.A. (1997). The Effectiveness of Virtual Reality for Administering Spatial Navigation Training to Firefighters. *Presence* 6(1), 73-86.
- Brooks, F. P., Jr. (1992). Six Generations of Building Walkthrough: Final Technical Report to the National Science Foundation, TR92-026.
- Chance, S.S., Gaunet, F., Beall, A.C., & Loomis, J.M. (1998). Locomotion Mode Affects the Updating of Objects Encountered During Travel: The Contribution on Vestibular and Proprioceptive Inputs to Path Integration. *Presence* 7(2), 168-178.
- Chase, W.G. (1983). Spatial Representation in Taxi Drivers. In D. R. Rogers and J. A. Sloboda (Eds.), *Acquisition of Symbolic Skills*. New York: Plenum.
- Colle, H.A., & Reid, G.B. (1998). The Room Effect: Metric Spatial Knowledge of Local and Separated Regions." *Presence* 7(2), 116-128.
- Darken, R.P. & Banker, W.P. (1998). Navigating in Natural Environments: A Virtual Environment Training Transfer Study. *Proceedings of VRAIS '98*, 12-19.
- Darken, R.P. & Sibert, J.L. (1993). A Toolset for Navigation in Virtual Environments. *Proceedings of ACM User Interface Software & Technology*, 157-165.
- Darken, R.P. & Sibert, J.L. (1996a). Navigating Large Virtual Spaces. *International Journal of Human-Computer Interaction*, 49-71.
- Darken, R.P. & Sibert, J.L. (1996b). Wayfinding Strategies and Behaviors in Large Virtual Worlds, *Proceedings of Computer Human Interfaces Conference, 1996 (CHI'96)*, 142-149.
- Eckstrom R.B, French, Harmen, & Dermen (1990). Manual for the Kit of Factor-Referenced Cognitive Tests, Technical Report, Office of Naval Research Contract N00014-71-C-0117. Princeton, N.J.: Educational Testing Service.
- Goerger, S.R. (1998). Spatial Knowledge Acquisition and Transfer from Virtual to Natural

- Environments for Dismounted Land Navigation. MS Thesis, Department of Computer Science, Naval Postgraduate School.
- Goerger, S.R., Darken, R.P., Boyd, M.A., Gagnon, T.A., Liles, S.W., Sullivan, J.A., and Lawson, J.P. (1998). Spatial Knowledge Acquisition from Maps and Virtual Environments in Complex Architectural Spaces. Proceedings of the 16th Applied Behavioral Sciences Symposium, 22-23 April, U.S. Air Force Academy, Colorado Springs, CO., 6-10.
- Guilford, J.P., & Zimmerman, W.S. (1947). The Guilford-Zimmerman Aptitude Survey. Palo Alto: Consulting Psychologists Press.
- Henry, D.P. (1992). Spatial Perception in Virtual Environments: Evaluating an Architectural Application. MSEE thesis, University of Washington.
- Loomis, J.M., Golledge R.G., & Klatzky, R.L. (1998). Navigation System for the Blind: Auditory Display Modes and Guidance." Presence 7(2), 193-203.
- May, M., Peruch, P., & Savoyant, A. (1995). "Navigating in a Virtual Environment with Map-Acquired knowledge: Encoding and Alignment Effects. Ecological Psychology 7(1), 21-36.
- Moeser, S. D. (1988). Cognition Mapping in a Complex Building. Environment and Behavior 20, 21-49.
- Pausch, R., Proffitt, D., & Williams, G. (1997). Quantifying Immersion in Virtual Reality. Computer Graphics Proceedings, Annual Conference Series.
- Peruch, P., Vercher, J.-L., & Gauthier, G.M. (1995). Acquisition of Spatial Knowledge through Visual Exploration of Simulated Environments. Ecological Psychology 7(1), 1-20.
- Presson, C.C., & Montello, D.R. (1994). Updating after Rotational and Translational Body Movements: Coordinate Structure of Perspective Space. Perception 23, 1447-1455.
- Rieser, J. J. (1989). Access to Knowledge of Spatial Structure at Novel Points of Observation. Journal of Experimental Psychology: Learning, Memory, and Cognition 15, 1157-1165.
- Ruddle, R.A., Payne, S.J., and Jones, D.M. (1997). Navigating Buildings in 'Desk-top' Virtual Environments: Experimental Investigations Using Extended Navigational Experience. Journal of Experimental Psychology: Applied, 3(2), 143-159.

- Ruddle, R.A., Payne, S.J. & Jones, D.M. (1998). Navigating Large-Scale 'Desk-Top' Virtual Buildings: Effects of Orientation Aids and Familiarity. *Presence* 7(2), 179-192.
- Ruddle, R.A., Payne, S.J., and Jones, D.M. (1999). Navigating Large-Scale Virtual Environments: What Differences Occur between Helmet Mounted and Desk-Top Displays. *Presence* 8(2), 157-168.
- Ruddle, R.A., Randall, S.J., Payne, S.J., and Jones, D.M. (1996). Navigation and Spatial Knowledge Acquisition in Large-Scale Virtual Buildings: An Experimental Comparison of Immersive and Desk-Top Displays. *Proceedings of the Second International FIVE Conference*, 125-136.
- Satalich, G.A. (1995). Navigation and Wayfinding in Virtual Reality: Finding the Proper Tools and Cues to Enhance Navigational Awareness. MSEE thesis, University of Washington.
- Stoakley, R., Conway, M.J., & Pausch, R. (1995). Virtual Reality on a WIM: Interactive Worlds in Miniature. *CHI'95 Mosaic of Creativity*.
- Tate, D.L., Sibert, L., & King, T. (1997). Virtual Environments for Shipboard Firefighting Training.
- Thorndyke, P. W., and Hayes-Roth, B. (1982). Differences in Spatial Knowledge Acquired from Maps and Navigation. *Cognitive Psychology* 14, 560-589.
- Tlauka, M. & Wilson, P.N. (1996). Orientation-Free Representations from Navigation through a computer-simulated Environment. *Environment and Behavior*, 28(5), 647-664.
- Waller, D., Hunt, E., & Knapp, D. (1998). The Transfer of Spatial Knowledge in Virtual Environment Training. *Presence* 7(2), 129-143.
- Witmer, B.G., Bailey, J.H., & Knerr, B.W. (1995). Training Dismounted Soldiers in Virtual Environments: Route Learning and Transfer. *USARI Technical Report 1022*.
- Witmer, B.G., Bailey, J.H., Knerr, B.W., & Parsons, K.C. (1996). Virtual spaces and real world places: transfer of route knowledge. *International Journal of Human-Computer Studies* 45, 413-428.
- Witmer, B.G. & Kline, P.B. (1998). Judging Perceived and Traversed Distance in Virtual Environments. *Presence* 7(2), 144-167.

Appendix A: Analysis of Variance of Navigation Errors

ANOVA analysis was performed in order to test the statistical significance of the trends that were observed in the data. Bearing errors, absolute bearing errors, log range errors, and absolute log range errors were each analyzed in an ANOVA analysis in which training and subtask were the main factors and the subject factor was nested within training. In the bearing and absolute bearing data, 432 data points were included in the analysis. In the range data analysis, 442 data points were included.

Bearing error depended significantly on subtask [$F(12, 385) = 45.93, p < 0.001$] and subject [$F(31, 385) = 3.1426, p < 0.001$], but not on training method [$F(3, 385) = 1.10, p > 0.001$]. This analysis confirms the observation that there was a significant dependence of the bearing error on the subtask and on individual subject, but not on the method of training. Post-hoc tests confirmed that for some subtasks, bearing errors were more negative (to the left) than in other subtasks. The expected bearing errors ranged from $\bar{m}17.4$ deg (in subtask 3) to 14.6 deg (in subtask 5). Scheffe post-hoc tests confirmed that many pairs of the subtasks yielded significantly different mean bearing errors. In general, most of the significant differences arose because responses on subtasks 2, 3, 6, and 7 tended to be to the left of the target compared to responses on subtasks 5, 8, and 12, which tended to be to the right of the target.

Results of analysis of absolute bearing error were similar, however subject differences were not as large. Absolute bearing error depended significantly on subtask [$F(12, 385) = 6.29, p < 0.001$], but not on training method [$F(3, 385) = 2.1247, p > 0.001$] or on subject [$F(31, 385) = 1.38, p > 0.001$]. Scheffe post-hoc analysis indicated that the expected absolute bearing errors were marginally significantly larger in subtask 3 (expected absolute error of 17.5 deg) than in subtask 10 (expected absolute bearing error of 7.1 deg).

The subtasks for which expected bearing error had the largest magnitude were the same subtasks for which the expected absolute errors were largest. Specifically, subtasks 2, 3, 5, 6, 7, 8, and 12 had larger expected absolute errors than any of the remaining 6 subtasks. In other

words, on some tasks large absolute bearing errors occur because of a systematic response bias (either to the left or the right), while on other tasks, bearing errors tend to be less systematic and smaller in magnitude.

Log range errors depended significantly on training method [$F(3, 395) = 9.77, p < 0.001$], subtask [$F(12, 395) = 8.65, p < 0.001$], and subject [$F(31, 395) = 24.68, p < 0.001$]. Scheffe post hoc analysis indicated that the expected log range error in the VE condition (value of -0.022) was significantly larger than the expected log range errors in the Mod and RW conditions (values of -0.089 and -0.083, respectively). In the NVE task, the expected log range error was -0.033, a value that was not significantly different from any of the values in the other three conditions. In other words, across all tasks, subjects tended to underestimate the distance to the target. This tendency was less pronounced in the VE condition compared to the Mod or RW conditions. While this trend was significant, the magnitude of the difference was relatively small. ANOVA analysis also showed that subtask was a significant factor; however, once again, the magnitudes of the differences across subtasks are quite small. The expected log range values varied from -0.176 on subtask 11 to 0.0157 on subtask 12. For most tasks, the expected log range errors were negative (subjects tended to underestimate the range). Only subtasks 5, 8, and 12 had positive expected values. Scheffe post-hoc analysis showed that the only significant differences among all pairwise comparisons arose when comparing subtask 11 with these subtasks.

An ANOVA on the absolute log range errors showed that training condition [$F(3, 395) = 18.69, p < 0.001$] and subject [$F(31, 395) = 6.067, p < 0.001$] were significant factors. However, subtask [$F(12, 395) = 1.75, p > 0.001$] was not a significant factor. Scheffe post hoc analysis indicated that the absolute log range error in the RW condition (expected absolute log range error of 0.224) was significantly larger than the absolute log range error in each of the other three tasks (expected log range errors of 0.157, 0.140, and 0.145 for VE, NVE, and Mod conditions, respectively). Once again, although this effect is significant, the magnitude of the difference in absolute range errors between real world training and the other methods is not large.

Taken together, these results indicate that the differences in range errors across subtask

are relatively unimportant. However, errors in the range judgments do depend on training method. Specifically, there is a general tendency to underestimate target range, but this tendency is least pronounced in the NVE condition. The magnitude of the log range errors is largest in the RW condition (compared to the other 3 conditions). Although robust and statistically significant, neither of these effects is large in an absolute sense.

FIGURE CAPTIONS

- Fig. 1. Overhead view of space used in the experiment with indications of both target and reference locations.
- Fig. 2. Illustrative view of the space in Mod condition: far away.
- Fig. 3. Illustrative view of the space in Mod condition: zoomed in.
- Fig. 4. Portion of space by elevators as seen in VE or NVE conditions (also pictured in lower left corner from a different viewpoint in Fig. 2).
- Fig. 5. Plots in the X,Y plane of the estimated and actual locations of the targets relative to the reference locations for each of the subtasks 1-13 specified in Fig. 1 and Table 1. The different training groups are indicated by different symbols. The distance between adjacent tick marks on both the X and Y axes is 50 ft. The origin (0,0) is located at the intersection of the lines with tick marks. In order to maintain the same coordinate scale for all plots, a few points (less than 2%) exceeded the coordinate ranges used in these plots and are not shown [this occurred only in the plots for subtasks 1, 9, and 12].
- Fig. 6. Histograms of bearing errors. Top row shows signed errors, bottom row absolute errors. Different columns correspond to different training conditions.
- Fig. 7. Same as Fig. 6, except for log range errors.
- Fig. 8. Contour plot of signed bearing errors. Meaning of gray scale in terms of signed bearing errors is specified by column of numbers on right (in degrees). Contours are estimated by using 2-dimensional linear interpolations between measured data points for grid of subject/condition (abscissa) and subtask (ordinate). Outliers and missing data points are plotted using white asterisks. Dominant effect of subtask is evident from primarily horizontal orientation of constant-error contours.
- Fig. 9. Same as Fig. 8 except for log range errors. Dominant effect of subject is evident from primarily vertical orientation of constant-error contours.

LEGEND FOR TABLE 2

The first page of this table gives the results for bearing errors, the second for log range errors (log to the base ten of the ratio of estimated range to actual range). The four blocks of rows on each page give the results for the four training conditions RW, VE, NVE, and Mod. The two blocks of columns on each page give the results for the signed errors and for the absolute errors.

In each of the 16 subtables on a page, the two numbers in square brackets give the grand mean and the grand standard deviation corresponding to that subtable; i.e., the statistics that are obtained when the data are pooled over all subjects and subtasks for the given training condition and the given type of error (signed or absolute) associated with that subtable. (In all cases, the term "standard deviation" refers to the standard deviation about the mean.)

Within each of the 16 subtables on each page, column (a) serves to identify both the subtask (see Table 1) and the subject; columns (b) and (c) give the means and standard deviations, respectively, for each subtask (where the variable averaged across is the subject); and columns (d) and (e) give the means and standard deviations for each subject (where the variable averaged across is subtask). The row labeled "mean" in each subtable gives the mean values for columns (b), (c), (d), and (e).

TABLE 1: Subtasks

	Water Fountain	Microwave Oven	Conference Table	End of Hallway	Elevator Door
Reference Location					
A		1	2	3	
B		4		5	6
C	7	8		9	10
D			11	12	13

TABLE 2: Means and Standard Deviations

BEARING ERROR (DEGREES)

Signed Errors						Absolute Errors					
	Subt/ Subj (a)	Across Subjects Mean (b)	StDev (c)	Across Subtasks Mean (d)	StDev (e)		Subt/ Subj (a)	Across Subjects Mean (b)	StDev (c)	Across Subtasks Mean (d)	StDev (e)
RW	1	-10.20	10.08	-4.92	15.72		1	12.60	6.38	13.54	8.66
[-2.89	2	-11.33	9.01	-4.50	16.80	[12.76	2	12.22	7.60	14.83	8.02
14.98]	3	-16.05	8.47	-7.31	11.94	8.28]	3	16.25	8.04	11.62	7.37
	4	3.07	12.18	-9.58	8.20		4	9.36	7.56	10.58	6.72
	5	14.00	14.40	-7.83	10.35		5	16.60	10.91	10.50	7.33
	6	-17.40	11.45	-6.62	14.07		6	18.20	9.98	13.23	7.47
	7	-13.90	5.49	-5.80	16.63		7	13.90	5.49	15.20	7.57
	8	11.17	12.14	7.15	18.70		8	12.72	10.28	16.77	10.02
	9	6.70	7.54	4.88	13.61		9	7.90	6.12	10.73	9.29
	10	-6.70	6.27	3.92	14.36		10	7.90	4.46	11.00	9.58
	11	-2.78	9.48				11	7.67	5.70		
	12	16.30	9.26				12	16.30	9.26		
	13	-8.10	12.97				13	12.70	7.86		
mean		-2.71	9.90	-3.06	14.04			12.64	7.66	12.80	8.20
VE	1	-6.56	4.64	0.00	20.41		1	7.22	3.35	14.92	13.24
[-1.25	2	-12.50	7.37	0.15	16.31	[11.10	2	13.30	5.60	14.46	6.29
13.74]	3	-21.11	10.65	1.23	9.98	8.14]	3	21.11	10.65	8.00	5.64
	4	8.89	16.37	-4.62	13.65		4	12.44	13.52	11.85	7.58
	5	16.75	7.17	-0.54	11.10		5	16.75	7.17	7.77	7.63
	6	-10.30	6.58	-4.45	12.26		6	11.10	4.93	10.64	6.90
	7	-8.40	11.85	-0.92	17.23		7	11.40	8.63	13.08	10.60
	8	7.22	10.35	0.70	12.45		8	9.44	8.09	10.10	6.51
	9	7.20	4.42	-0.60	11.41		9	7.60	3.60	8.80	6.68
	10	-4.30	6.77	-3.50	11.41		10	6.90	3.67	10.50	4.81
	11	-2.33	7.75				11	6.33	4.58		
	12	16.56	5.50				12	16.56	5.50		
	13	-2.89	8.01				13	5.33	6.44		
mean		-0.91	8.26	-1.25	13.62			11.19	6.59	11.01	7.59
NVE	1	-8.71	5.56	-9.92	11.40		1	8.71	5.56	11.77	9.30
[-2.20	2	-6.43	5.88	-7.33	13.12	[10.43	2	7.57	3.99	12.67	7.23
12.61]	3	-19.80	6.42	-0.82	4.45	7.33]	3	19.80	6.42	3.91	1.92
	4	7.64	7.64	-3.42	14.94		4	8.21	6.92	13.42	6.29
	5	18.60	10.06	-0.85	12.84		5	18.60	10.06	10.08	7.45
	6	-13.00	6.58	-0.12	13.89		6	13.00	6.58	10.81	8.14
	7	-8.64	5.51	5.62	10.76		7	8.64	5.51	10.31	5.89
	8	5.58	7.05				8	6.92	5.44		
	9	4.92	6.04				9	6.58	3.67		
	10	-3.43	6.58				10	6.29	3.30		
	11	-5.00	11.82				11	11.00	5.26		
	12	7.42	21.54				12	18.42	11.09		
	13	-3.29	7.91				13	7.00	4.24		
mean		-1.86	8.35	-2.41	11.63			10.83	6.00	10.42	6.60
Mod	1	-2.44	7.33	-2.31	12.03		1	5.31	5.31	9.38	7.42
[-1.42	2	-12.06	6.95	-7.85	12.23	[10.6	2	12.06	6.95	11.38	8.71
13.06]	3	-12.69	5.93	-3.50	7.37	7.59]	3	12.69	5.93	6.50	4.64
	4	4.93	7.72	-0.38	17.02		4	7.50	4.72	13.00	10.34
	5	7.50	12.99	-1.62	15.33		5	11.75	8.68	12.54	8.21
	6	-10.75	8.07	-3.08	13.59		6	10.75	8.07	10.92	8.11
	7	-9.38	9.24	5.88	10.20		7	11.38	6.16	9.04	7.28
	8	14.94	12.03	1.35	13.12		8	14.94	12.03	11.73	4.99
	9	9.06	6.73				9	9.06	6.73		
	10	-6.75	7.19				10	7.50	6.28		
	11	-1.00	11.34				11	8.50	6.87		
	12	10.00	11.25				12	10.75	10.43		
	13	-9.00	14.95				13	15.25	7.09		
mean		-1.36	9.36	-1.44	12.61			10.57	7.33	10.56	7.46

LOG RANGE ERROR

Signed Errors

Absolute Errors

	Subt/ Subj (a)	Across Subjects		Across Subtasks			Subt/ Subj (a)	Across Subjects		Across Subtasks	
		Mean (b)	StDev (c)	Mean (d)	StDev (e)			Mean (b)	StDev (c)	Mean (d)	StDev (e)
RW											
[-0.09 0.25]	1	-0.13	0.26	-0.34	0.13	[0.22 0.14]	1	0.25	0.13	0.34	0.13
	2	-0.08	0.26	0.06	0.12		2	0.23	0.11	0.11	0.08
	3	-0.06	0.31	0.17	0.16		3	0.27	0.13	0.19	0.12
	4	-0.15	0.27	0.23	0.12		4	0.26	0.15	0.23	0.12
	5	0.06	0.21	-0.39	0.11		5	0.18	0.09	0.39	0.11
	6	-0.13	0.29	-0.01	0.16		6	0.26	0.18	0.12	0.10
	7	-0.09	0.22	0.11	0.10		7	0.18	0.14	0.14	0.07
	8	-0.06	0.23	-0.24	0.19		8	0.18	0.15	0.28	0.12
	9	-0.05	0.30	-0.23	0.11		9	0.25	0.15	0.23	0.11
	10	-0.19	0.17	-0.19	0.14		10	0.19	0.16	0.21	0.10
	11	-0.19	0.24				11	0.26	0.16		
	12	0.02	0.27				12	0.24	0.10		
	13	-0.09	0.19				13	0.17	0.11		
mean		-0.09	0.25	-0.08	0.13			0.22	0.14	0.22	0.11
VE											
[-0.02 0.19]	1	-0.02	0.17	-0.19	0.12	[0.16 0.11]	1	0.13	0.10	0.20	0.12
	2	0.01	0.18	-0.11	0.08		2	0.14	0.11	0.11	0.08
	3	-0.04	0.17	0.00	0.10		3	0.15	0.08	0.08	0.07
	4	0.04	0.23	0.19	0.18		4	0.17	0.15	0.23	0.11
	5	-0.02	0.19	-0.02	0.11		5	0.15	0.11	0.08	0.06
	6	0.00	0.16	-0.17	0.12		6	0.14	0.07	0.17	0.12
	7	-0.08	0.21	-0.17	0.11		7	0.18	0.13	0.18	0.10
	8	0.07	0.29	-0.14	0.07		8	0.24	0.15	0.14	0.07
	9	0.03	0.19	0.15	0.12		9	0.15	0.10	0.15	0.11
	10	-0.07	0.14	0.24	0.13		10	0.13	0.08	0.24	0.13
	11	-0.12	0.20				11	0.20	0.11		
	12	0.01	0.21				12	0.16	0.13		
	13	-0.03	0.14				13	0.12	0.08		
mean		-0.02	0.19	-0.02	0.11			0.16	0.11	0.16	0.10
NVE											
[-0.03 0.17]	1	-0.01	0.16	-0.05	0.11	[0.14 0.10]	1	0.13	0.08	0.09	0.08
	2	0.08	0.17	-0.26	0.09		2	0.14	0.11	0.26	0.09
	3	0.05	0.11	-0.05	0.11		3	0.08	0.09	0.09	0.07
	4	-0.13	0.16	0.10	0.13		4	0.18	0.09	0.13	0.10
	5	-0.03	0.12	0.08	0.15		5	0.09	0.06	0.14	0.09
	6	-0.02	0.16	-0.11	0.13		6	0.12	0.09	0.14	0.09
	7	0.03	0.18	0.03	0.17		7	0.15	0.07	0.14	0.09
	8	0.00	0.17				8	0.12	0.10		
	9	-0.01	0.09				9	0.06	0.07		
	10	-0.07	0.16				10	0.14	0.09		
	11	-0.26	0.09				11	0.26	0.09		
	12	0.08	0.12				12	0.12	0.07		
	13	-0.04	0.19				13	0.14	0.12		
mean		-0.02	0.14	-0.04	0.13			0.13	0.09	0.14	0.09
Möb											
[-0.09 0.15]	1	-0.08	0.13	-0.02	0.16	[0.15 0.10]	1	0.12	0.08	0.14	0.08
	2	-0.07	0.14	-0.25	0.08		2	0.11	0.10	0.25	0.08
	3	-0.10	0.14	-0.17	0.11		3	0.15	0.09	0.18	0.10
	4	-0.15	0.12	-0.07	0.16		4	0.18	0.08	0.13	0.11
	5	0.01	0.13	-0.15	0.07		5	0.10	0.07	0.15	0.07
	6	-0.08	0.14	0.01	0.18		6	0.14	0.08	0.13	0.12
	7	-0.16	0.15	-0.07	0.08		7	0.18	0.12	0.09	0.05
	8	0.07	0.22	0.01	0.11		8	0.18	0.13	0.08	0.07
	9	-0.06	0.08				9	0.07	0.07		
	10	-0.23	0.13				10	0.23	0.13		
	11	-0.17	0.11				11	0.17	0.11		
	12	-0.01	0.13				12	0.11	0.07		
	13	-0.13	0.14				13	0.16	0.09		
mean		-0.09	0.14	-0.09	0.12			0.15	0.09	0.15	0.09

TABLE 3: Grand Means and Standard Deviations

	Bearing		Log Range	
	Signed	Absolute	Signed	Absolute
RW	-2.9, 15.0	12.8, 8.3	-0.09, 0.25	0.22, 0.14
VE	-1.3, 13.7	11.1, 8.1	-0.02, 0.19	0.16, 0.11
NVE	-2.2, 12.6	10.4, 7.3	-0.03, 0.17	0.14, 0.10
Mod	-1.4, 13.1	10.6, 7.6	-0.09, 0.15	0.15, 0.10
Mean	-2.0, 13.6	11.2, 7.8	-0.06, 0.19	0.17, 0.11

TABLE 4: Correlation ρ between pairs of Training Groups for the Dependence of Mean Errors (across subjects) on subtasks

	Bearing		Log Range	
	Signed	Absolute	Signed	Absolute
ρ (RW, VE)	0.95	0.65	0.39	-0.09
ρ (RW, NVE)	0.91	0.70	0.61	0.10
ρ (RW, Mod)	0.94	0.45	0.76	-0.30
ρ (VE, NVE)	0.96	0.82	0.42	0.32
ρ (VE, Mod)	0.89	0.25	0.69	0.26
ρ (NVE, Mod)	0.82	0.16	0.47	0.51
Mean	0.91	0.50	0.56	0.13

TABLE 5: Comparisons of Grand Standard Deviations to Standard Deviations Across Subjects and Across Subtasks

	Bearing		Log Range	
	Signed	Absolute	Signed	Absolute
Grand σ				
RW	15.0	8.3	0.25	0.14
VE	13.7	8.1	0.19	0.11
NVE	12.6	7.3	0.17	0.10
Mod	13.1	7.6	0.15	0.10
Mean	13.6	7.8	0.19	0.11
X subj σ				
RW	9.9	7.6	0.25	0.14
VE	8.3	6.6	0.19	0.11
NVE	8.4	6.0	0.14	0.09
Mod	9.4	7.3	0.14	0.09
Mean	9.0	6.9	0.18	0.11
X subt σ				
RW	14.0	8.2	0.13	0.12
VE	13.6	7.6	0.11	0.10
NVE	11.6	6.6	0.13	0.09
Mod	12.6	7.5	0.12	0.09
Mean	12.9	7.5	0.12	0.10

TABLE 6: Values of $\sigma(m) / m(\sigma)$

Subtasks

	Bearing		Log Range	
	Signed	Absolute	Signed	Absolute
RW	1.18	0.47	0.30	0.28
VE	1.39	0.72	0.28	0.30
NVE	1.23	0.83	0.63	0.55
Mod	1.03	0.40	0.61	0.46
Mean	1.21	0.61	0.30	0.40

Subjects

	Bearing		Log Range	
	Signed	Absolute	Signed	Absolute
RW	0.43	0.27	1.66	0.86
VE	0.16	0.33	1.42	0.58
NVE	0.44	0.47	0.96	0.65
Mod	0.32	0.29	0.79	0.62
Mean	0.34	0.34	1.21	0.68


TABLE 7: Correlations with Subtask 0 (Target Visible)


	Bearing		Log Range	
	Signed	Absolute	Signed	Absolute
RW	0.27	0.17	0.75	0.76
VE	0.40	0.16	0.19	0.35
NVE	0.49	0.01	0.28	0.46
Mod	0.68	-0.09	0.66	0.52
Mean	0.46	0.06	0.47	0.52
Pooled	0.42	-0.09	0.58	0.62


TABLE 8: Results of Normalizing Data According to Estimates Made in Subtask 0


	Bearing				Log Range			
	Signed		Absolute		Signed		Absolute	
	m	σ	m	σ	m	σ	m	σ
RW	-2.9	15.0	12.8	8.3	-0.09	0.25	0.22	0.14
RW*	-2.5	14.9	12.5	8.5	0.11	0.20	0.18	0.13
VE	-1.2	13.7	11.1	8.2	-0.02	0.19	0.16	0.11
VE*	-1.7	13.9	11.1	8.5	0.10	0.21	0.16	0.16
NVE	-2.2	12.7	10.4	7.3	-0.03	0.17	0.14	0.10
NVE*	-2.5	12.4	10.2	7.5	0.05	0.18	0.15	0.11
Mod	-1.4	13.1	10.6	7.7	-0.09	0.15	0.15	0.10
Mod*	-1.1	13.0	10.1	8.2	0.13	0.14	0.15	0.11
Mean	-1.9	13.6	11.2	7.9	-0.06	0.19	0.17	0.11
Mean*	-1.9	13.5	11.0	8.2	0.05	0.18	0.16	0.13


* denotes results for normalized data





















BEARING ERROR (deg)



LOG RANGE ERROR



ATTACHMENT NUMBER 1

REPORTS AND REPORT DISTRIBUTION

REPORT TYPES

- (a) Performance (Technical) Report(s) (Include letter report(s)) Frequency: Annual
- (b) Final Technical Report, issued at completion of Grant.
NOTE: Technical Reports must have a SF-298 accompanying them.
- (c) Final Financial Status Report (SF 269)
- (d) Final Patent Report (DD 882)

REPORTS DISTRIBUTION		
ADDRESSEES	REPORT TYPES	NUMBER OF COPIES
Office of Naval Research Program Officer Terry Allard ONR 342CS Ballston Centre Tower One 800 North Quincy Street Arlington, VA 22217-5660	(a) & (b) w/(SF-298's)	3
Administrative Grants Officer OFFICE OF NAVAL RESEARCH REGIONAL OFFICE BOSTON 495 SUMMER STREET ROOM 103 BOSTON, MA 02210-2109	(c), (d) & SF-298's only for (a) & (b)	1
Director, Naval Research Laboratory Attn: Code 5227 4555 Overlook Drive Washington, DC 20375-5326	(a) & (b) w/(SF-298's)	1
Defense Technical Information Center 8725 John J. Kingman Road STE 0944 Ft. Belvoir, VA 22060-6218	(a) & (b) w/(SF-298's)	2
Office of Naval Research Attn: ONR 00CC1 Ballston Centre Tower One 800 North Quincy Street Arlington, VA 22217-5660	(d)	1

If the Program Officer directs, the Grantee shall make additional distribution of technical reports in accordance with a supplemental distribution list provided by the Program Officer. The supplemental distribution list shall not exceed 250 addressees.

REPORT DOCUMENTATION PAGE

Form Approved

OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 2/25/2000	3. REPORT TYPE AND DATES COVERED 3rd Annual Technical 2/1/98 - 1/31/99	
4. TITLE AND SUBTITLE Training Spatial Knowledge Acquisition			5. FUNDING NUMBERS N00014-96-1-0379	
6. AUTHOR(S) Nathaniel I. Durlach				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Research Laboratory of Electronics Massachusetts Institute of Technology 77 Massachusetts Avenue Cambridge, MA 02139			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Ballston Center Tower One 800 North Quincy Street Arlington, VA 22217-5660			10. SPONSORING/MONITORING AGENCY REPORT NUMBER 96PRO3143-00	
11. SUPPLEMENTARY NOTES The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report summarizes the work done by MIT in Year 3 (1 February 1998 through 31 January 1999) of the ONR Grant N00014-96-1-0937 entitled "Training Spatial Knowledge Acquisition Using Virtual Environments." It has been prepared by Dr. Thomas E. v. Wiegand, Nathaniel Durlach (PI), Sam Madden, and Rebecca Lee Garnett. The main effort during this period has been directed towards the development of a hardware/software system (the "VE Construction System") to facilitate the creation of complex virtual building environments. Virtual environment modeling of complex spaces is presently a manually intensive task in which required labor increases markedly with the level of three-dimensional detail that must be represented. As most real-world spaces worth modeling are of high complexity, the cost of this process can have a significant negative impact on the scope and timeliness of such projects. The tools we have developed during this year of our program greatly reduce the amount of time required to build the complex photorealistic VE models necessary for performing our envisioned VSPAN experiments. Also, of course, these tools should prove useful for other programs conducted at other laboratories.				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	