

Spatial Navigation in Virtual Reality Environments: An EEG Analysis

WALTER F. BISCHOF, Ph.D., and PIERRE BOULANGER, Ph.D.

ABSTRACT

Past research has linked theta oscillations (electroencephalographic activity in the 4–8-Hz range) to spatial navigation in rodents and humans, and to the encoding and retrieval of spatial information in rodents. In the present study, electroencephalographic activity was measured while humans navigated through virtual mazes. Results confirmed previous findings that the frequency of theta episodes is directly related to the difficulty of maze navigation. We were also able to show that theta episodes occur most likely at points in a maze where new hallways come into view, or after navigational mistakes have been realized and are being corrected. This indicates that, just as in rodents, theta episodes in humans are related to the encoding and retrieval of spatial information.

INTRODUCTION

OSCILLATORY PATTERNS OF ACTIVITY in the electroencephalogram (EEG) of animals and humans have long been linked to a number of cognitive functions. Theta oscillations (4–12 Hz in rodents) have been implicated in spatial navigation in animals,^{1–4} and in particular have been linked in to the encoding and retrieval of spatial information in the hippocampus of rodents.² More recently, task-related characteristics of theta oscillations (4–8 Hz in humans) have also been studied in humans. Kahana et al.⁵ recorded subdurally from epileptic patients while they were navigating virtual mazes presented on a computer screen, and found multiple episodes of high-amplitude theta activity in a number of cortical areas including the frontal and temporal cortex. They further found that the frequency of theta-wave episodes occurred more frequently in complex mazes and that they were more frequent in recall trials than in learning trials. These results were confirmed in a further study.⁶

Using scalp EEGs, Nishiyama and Yamaguchi⁷ identified two regions with bursts of high theta activity during virtual maze navigation, one in the

frontal region, possibly associated with spatial working memory, and one in the parietal-temporal region. The sequential theta activity of these areas suggests that information flow in the theta range reflects connections between hippocampus, prefrontal and parietal cortex, connections that are establishing the cortical network involved in human spatial cognition. Using magnetoencephalography (MEG), Araujo et al.⁸ found a strong increase of theta activity in temporal regions for virtual maze navigation, but not for either mental calculations or passive viewing of maze navigation, supporting the conclusion that theta activity is closely tied to spatial navigation.

Further evidence on the neural basis of navigation in humans comes from a functional brain imaging (PET) study during navigation of virtual environments.⁹ Activation of the right hippocampus was strongly associated with knowing accurately where places were located and navigating accurately between them. Getting to those places quickly was strongly associated with activation of the right caudate nucleus, while activity in other regions was probably associated with non-spatial aspects of navigation.

In summary, these studies provide converging evidence on the structure and function of the human cortical network involved in spatial navigation, and on the role of theta oscillations in linking the activity of the cortical regions involved in this network. Further, they indicate that the occurrence of theta episodes can be linked to characteristics of the navigational tasks, more specifically, that the frequency of theta episodes as well as amplitude of theta activity may be taken as an indicator of difficulty of maze navigation. In these studies, virtual reality (VR) environments were created and navigated in order to study the structure and function of cortical navigational mechanisms. The correlation between theta activity and difficulty of maze navigation also permits inferences in the other direction, namely to use theta activity to assess ease of navigation through virtual reality environments. Hence, theta activity can be used as another measure for evaluating VR environments, in addition to standard behavioral measures such as number of errors, learning time, etc.

In navigating through the environment, we rely on a multitude of cues for planning large-scale movements (such as walking through the environment), for combining trajectories of previously traveled paths into a mental representation of the environment, and for determining heading. These cues include both static cues (e.g., landmarks, displacement information, depth information) as well as dynamic cues (e.g., optic flow patterns). Recent behavioral studies indicate that the difficulty of navigating through virtual environments is closely related to the number and types of visual cues that are available for guiding navigation. For example, Kirschen et al.¹⁰ found that landmark information and optic flow information significantly reduced the time participants took to navigate a virtual maze. These studies suggest that rendering of perceptually rich VR environments can make navigation much easier than perceptually impoverished renderings, and thus can help to reduce the cognitive load in navigation tasks. However, perceptual enrichment of VR environments may not necessarily lead to easier navigation. In a perceptually impoverished VR environment, we may rely much more on internally represented, cognitive maps of the environment, whereas in perceptually rich environments one can rely more on landmark-based guidance. This, in turn, may prevent the building of an efficient and easily accessible representation of the environment.

Many of the past EEG studies on navigation in VR environments have imposed fairly severe con-

straints on the movements that could be made. In the studies of Kahana et al.⁵ and Nishiyama and Yamaguchi,⁷ for example, participants were transported at fixed velocity through maze hallways and viewpoint changes were instantaneous and limited to 90 degrees. This permitted equating movement patterns across different participants, but it also led to highly artificial movement patterns that were not under control of the observers. The present study was aimed at overcoming this limitation by studying navigation with more natural movement patterns.

MATERIALS AND METHODS

The mazes used in the present experiment were created using the Mandala system,¹¹ a reconfigurable VR tool that allows the generation of complex VR environments using a simple markup language and their display on range of VR hardware, including simple screens (as used in the present experiment), stereoscopic display systems, and VR caves. In the Mandala system, a VR environment is made up of a large number of cubes ("cells"), and each cell can have a floor, a ceiling, visible and invisible walls, and can contain arbitrary objects. Further, one can define a number of events for each cell (e.g., the avatar entering or leaving the cell, the avatar colliding with one of the walls), and each of these events can trigger some action (e.g., writing a message into a log file, sending a message to a sound server, or sending a signal to the EEG computer). Many of the features of the Mandala system were developed for the present study, and precautions were taken to ensure that display updates and refreshes were as efficient as possible, so that temporal aspects of navigation in complex mazes could be studied.

Participants

Fourteen graduate students (11 male, three female) from the Department of Computing Science participated in this study. All had normal or corrected-to-normal vision.

Maze stimuli

All mazes consisted of a set of T-junctions arranged in such a way that the goal position could be reached from the starting position in a sequence of 10 left or right turns. The layout of one maze is shown in Figure 1, with cell 0 (on the right




FIG. 1. Layout of a maze used in the experiments. The maze consists of a number of T-junctions, arranged such that the goal position (cell 94g) can be reached from the starting position (cell 0) in a sequence of 10 left/right turns LLRRLRLRRL. Each T-junction consists of six or seven cells: one or two cells before the junction (e.g., cells 1 and 2), the junction cell (e.g., cell 3), a correct arm consisting of two cells (e.g., cells 5 and 5), and an incorrect arm consisting of two cells (e.g., cells 127 and 128f). The path beyond an incorrect arm is blocked by a transparent screen (e.g., between cells 128f and 129). The paths between other cells (e.g., cells 8 and 114) are blocked by visible walls.

hand side) indicating the starting position and cell 94g (at the top) indicating the goal position. In this case, a sequence of left/right turns LLRRLRLRRL leads from the starting position to the goal position. A set of 10 different mazes was created, and random subsets of these mazes were presented in the experiments.

Each T-junction consisted of either six or seven cells, a junction cell (e.g., cell 3), two cells each for the left and right arm (e.g., cells 4/5 and 127/128f), and either one or two cells before the intersection (e.g., cells 1/2). From every junction, another junction was visible along both, the left arm and the right arm. An invisible barrier prevented navigation in the incorrect direction. For example, in the

first T-junction of Figure 1, a barrier was placed between cells 128f and 129, so that the T-junction consisting of cells 129–135 could not be reached.

Several maze views are shown in Figure 2. Depending on the experimental condition, the maze walls could be either plain (Fig. 2a) or colored (Fig. 2b), and could either have an arrow indicating the correct direction (Fig. 2c) or not (Fig. 2b). Finally, a message on the wall at the end of a hallway indicated the goal position (Fig. 2d). Upon entering a new maze, a participant started in the starting cell (cell 0 in Fig. 1) facing the first intersection (cell 3). As soon as the goal cell (cell 94G in Fig. 1) was entered, the maze was exited and a new maze was entered.




FIG. 2. Four different maze views: (a) View of a hallway in a “plain” maze. (b) View of a hallway in a “color” maze. (c) View of a hallway in a “plain” training maze, with arrows indicating the correct direction of navigation. (d) View of the goal position in a “color” maze.

Maze navigation

To travel through a maze, participants used the arrow keys of the keyboard, namely forward (\uparrow), backward (\downarrow), left (\leftarrow), and right (\rightarrow). A single forward- or backward-key press changed the position by about 0.05 cell width, and a single left- or right-key press changed the orientation by about 3 degrees. By pressing the forward-key continuously, one could move forward at a velocity of about 1.2 cells/sec, and by pressing the left- or right-key continuously, one could rotate at a velocity of about 60 deg/sec. Under optimal conditions, a maze could thus be traversed in about 40 sec.

Response recording

During maze navigation, the following events were recorded: entering a cell, leaving a cell, colliding with a visible obstacle (i.e., walking into a wall), and colliding with an invisible obstacle (i.e., walking

into an invisible barrier). All events were recorded using a time-stamp. It was thus possible to determine in which cell the participant was located at any moment in time, as well as how many errors were made (i.e., how often collisions with an invisible barrier occurred), how often a path was backtracked, and how long it took to traverse a whole maze.

Maze variations

Two different types of mazes were presented, plain mazes and colored mazes. In plain mazes, all walls had the same color (Fig. 2a), whereas in the color mazes, the color of each T-junction was chosen randomly from the set {blue, cyan, green, red, magenta, yellow}. All walls of a single T-maze had the same color (Fig. 2b). The texture of all walls, as well as ceilings and floors did not change.

All mazes were generated in two versions, as training mazes and as test mazes. In training mazes,

a large yellow arrow indicated the correct direction at each intersection (Fig. 2c), whereas the arrows were absent in the test mazes.

Apparatus

All stimuli were presented on a NEC MultiSync LCD 1850E Screen with a refresh rate of 60 Hz and a display area of approximately 36 cm by 28.5 cm, or 35 arc deg \times 28 arc deg at the viewing distance of approximately 57 cm. All manual responses were entered on the keyboard.

EEG methods

EEG signals were collected from 38 gold electrodes embedded in an electrode cap and amplified using a Neuroscan NuAmp amplifier. Recording locations were based on the electrode placement system of the American Electroencephalographic Society¹² and included the following electrodes: FP1, FP2, F7, F3, FZ, F4, F8, FT9, FT7, FC3, FCZ, FC4, FT8, FT10, T7, C3, CZ, C4, T8, TP7, CP3, CPZ, CP4, TP8, P7, P3, PZ, P4, P8, PO1, PO2, O1, OZ, and O2. All channels were referenced to linked mastoids with a forehead ground electrode (AFZ). Linked mastoids served as a reference. Horizontal eye movements were monitored with bipolar electrodes on the outer canthus of each eye, and vertical movements were monitored from electrodes placed above and below the left eye. EEG and EOG were recorded at a sampling rate of 500 Hz for the duration of the whole experiment.

Procedure

In each experimental condition, a participant learned to walk through five randomly selected T-mazes. Each maze was presented once as a training maze, followed by a number of presentations as test mazes without arrows. A test maze was presented repeatedly until the participant traversed it without error (i.e., without walking into an invisible barrier) or until the test maze had been presented eight times, whichever came first.

Upon entering a maze, participants were allowed to rest, but were instructed to traverse that maze without pause once they had left the starting cell. As soon as the goal cell was entered, the maze was exited and the next maze was entered. No feedback about performance was given, but participants could easily determine whether they had successfully traversed a maze without error, since only after an error-free traversal was the next training

maze (with arrows) presented. Participants were instructed to remain still and to not talk during the experiment.

There were two experimental conditions, a "plain" condition and a "color" condition. In the plain condition, all maze walls were plain (Fig. 2a), whereas in the "color" condition, all walls were colored (Fig. 2b). Each participant traversed five different mazes, once with arrows in a training trial, followed by (at most eight) test trials, so each participant traversed at least 10 and most 45 mazes. These conditions were presented in a between-subject design; that is, each participant traversed either plain or colored mazes.

RESULTS

Behavioral measures

The two experimental groups, color and plain, were compared with respect to a number of behavioral measures. For each of the five presented mazes, participants in the color group required on average 3.18 test maze traversals, which was slightly higher than the number of test maze traversals of 2.53 for the plain group ($t(13) = 0.727, p > 0.1$). The average number of errors (i.e., collisions with an invisible barrier after taking a wrong turn) was 4.59 for the color group, which is slightly higher than the average number of errors of 2.66 for the plain group ($t(13) = 1.00, p > 0.1$). Finally, the average duration spent on traversing a single test maze was 70.5 sec for the color group, whereas it was 53.1 sec for the plain group ($t(13) = 2.81, p < 0.05$).

In summary, the participants who traversed the colored mazes, spent longer time on traversing the mazes, and, at least in tendency, made more navigation errors and had to traverse mazes more often before making an error-free traversal than did the participants in the plain condition.

Electroencephalographic measures

An EEG was recorded continuously for a complete experimental session. These data were segmented into episodes corresponding to the traversal of a single maze, beginning at the time the starting position (e.g., cell 0 in Fig. 1) was left, and ending at the time the goal position (e.g., cell 94g in Fig. 1) was entered. Each of these maze episodes was analyzed independently.

For a given EEG channel, short-term spectrograms were computed for each maze episode. To this end, discrete Fourier transforms were computed

for Hanning windows of 1000 msec in duration, with successive windows overlapping by 900 msec. The spectrograms for all 17 test maze traversals of one participant and one electrode (Fz) are shown in Figure 3. For plotting purposes, the spectrograms were constrained to a frequency range of 4–12 Hz. Brightness in each spectrogram is inversely related to the log power of the short-term Fourier spectrum, normalized over a whole episode; that is, the darker a spectrogram at a given time point and frequency, the higher is the power of the corresponding Fourier component. Theta episodes should thus show as dark bands in the 4–8-Hz range.

The analysis of the EEG spectra related theta-episodes during a maze traversal to locations in a maze. An example of this analysis is shown in Figures 4 and 5. The spectrogram in Figure 4a reproduces the last but one spectrogram of Figure 3. In this spectrogram, 10 theta episodes are identifiable if one ignores the very last episode that may be related to exiting the maze. Figure 4b shows the corresponding average power of the theta-band (4–8 Hz) over time. The time points of the theta episodes allow determining where in the maze the participant was at the time. This is shown in Figure 5. As in Figure 1, cell 0 is the starting position, and cell 94g is the goal position, hence the goal position

could be reached from the starting position in a sequence of turns LLRRLRLRRL. The arrows indicate the positions in the maze where a theta-episode occurred. The first five episodes are indicated by single arrows, and occur in cells 4, 22, 36, 42, and 54, immediately after the participant had turned into one arm of a T-junction and saw a new hallway segment. The last theta-episode occurred in cell 89, shortly before the last T-junction of the maze. The other four theta-episodes occurred in cells 65f/64 and 104f/103, respectively. In both these cases, the participant had made an incorrect navigational decision and walked into an invisible barrier. In each pair, the first theta episode occurred at the barrier (cells 65f and 104f, respectively), and the second episode occurred in the neighboring cell (cells 64 and 103, respectively), on the way out from the incorrect T-junction arm, and with the correct T-junction arm in view.

The example in Figures 4 and 5 indicates that theta episodes are more likely to occur when a new hallway gets into view or, after a participant has realized that a navigational error has been made (after walking into a barrier). If this is so, then the average power of theta episodes should be higher immediately after a turn in a T-junction than immediately before the junction. In Figures 4c, a T-junction is shown together with a definition of cell types (cells 1/2 immediately before a junction, cell 3 the junction, cells 4/5 immediately after a junction in the correct direction, and cells 7/8 immediately after a junction in the incorrect direction). Figure 4d shows the average power of the theta-band for each of these cell types, for the spectrogram in Figure 4a. It can be seen that the power in the theta-band is higher in cells 4–6, immediately after a turn in a junction.

A statistical analysis of *all* maze traversals confirmed the above analysis. The average power in the theta band was analyzed in a two-factor mixed factorial analysis of variance, with experimental group (color and plain) as a between-subjects factor, and cell type (cells 1–7) as a within-subjects factor. The results show a non-significant effect of experimental group ($F(1,124) = 0.239, p > 0.1$), a significant effect of cell type ($F(6,124) = 6.836, p < 0.01$), and a non-significant interaction ($F(6,124) = 1.1, p > 0.1$). The means are shown in Figure 6, indicating that the power in the theta band is higher immediately after a junction, when a new hallway is visible, than immediately before a junction. The results also show that the experimental condition (color and plain) had a consistent, although not significant effect; that is, the power of theta-episodes tended to be higher in the color condition.




FIG. 3. EEG spectrograms of 17 maze traversals of one participant. Each spectrogram shows the power of the frequency components in the range of 4–12 Hz at different time points in the maze traversal. Brightness of the spectrograms is inversely proportional to the log power of the short-term Fourier spectrum, normalized over the whole episode. Theta episodes are identified by dark bands in the lower half of the spectrograms, in the 4–8 Hz frequency range.

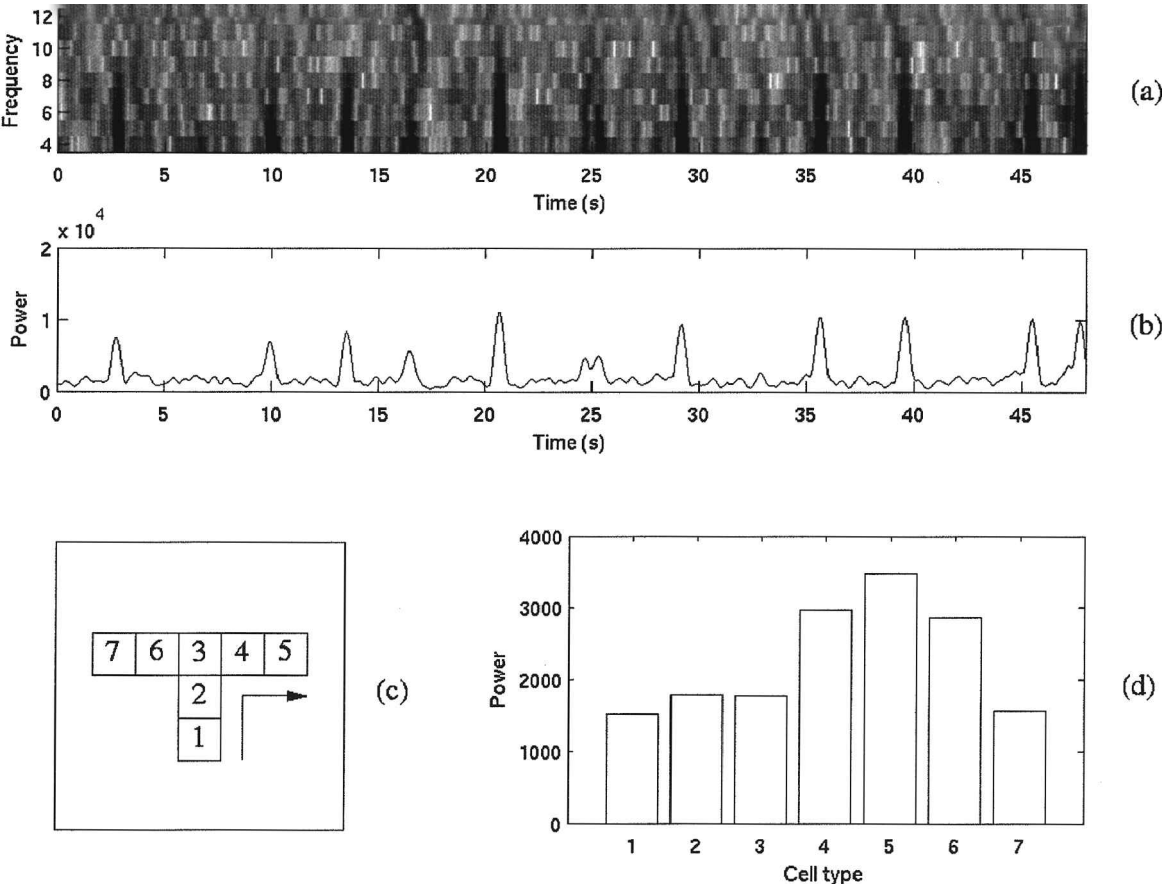


FIG. 4. The top panel shows a magnified view of the spectrogram of a single maze traversal, namely the last but one in Figure 3. The middle panel shows the power of the theta component (integrated over the range of 4–8 Hz) of the spectrogram in the top panel. Ten theta episodes can be identified, if one ignored the very last one. The bottom panel shows a T-junction on the left, with the direction of correct navigation indicated by the arrow. Within the T-junction, cells types are numbered in the range of 1–7, with cell types 1–2 before the junction, cell type 3 corresponding to the junction, cell type 4–5 after the junction in the correct navigation direction, and cell types 6–7 after the junction in the incorrect navigation direction. The histogram on the right shows the average power of the theta components for each of the cell types. This histogram was obtained by determining every 50 msec, in what cell the participant was located, and adding up the power of the theta band correspondingly.

DISCUSSION

The behavioral results show that participants who traversed the colored mazes, spent more time on traversing the mazes, and, at least in tendency, made more navigation errors and required more attempts before managing to traverse a maze without error, than did the participants who traversed the plain mazes. This result stands in contrast to the fact that the participants in the color condition had actually more information available for guiding their navigation than did the participants in the plain condition. This result is consistent with navigation strategies reported by the participants. In the plain

condition, all hallways looked the same, and the only way to remember a path was to rely on learning a sequence of left-right turns. Once participants realized and used this strategy, the number of navigation errors dropped sharply. In contrast, only few of the participants in the color condition reported adopting this strategy. Instead, they relied on navigation rules that took color cues into account and appeared to encode a number of local navigational decision rules (e.g., to turn left at the second magenta intersection, or to turn right after to green intersections in a row). As the behavioral results show, this strategy turned out to take more time and to be more prone to errors. The EEG results were




FIG. 5. Layout of the maze that was traversed in the maze traversal of Figure 4. The arrows indicate where in the maze each of the 10 theta episodes in Figure 4 occurred. The first five occur immediately after a T-junction, when a new hallway segment is in view, and the last episode occurs a little bit further down the hallway. The other four episodes occur in pairs, in cells 103/104f and 64/65f, respectively. In both cases, the participant had taken a wrong turn, walked into a transparent barrier (between cells 65f/66 and 104f/105, respectively) and had to turn around to walk into the other arm of the T-junction.

consistent with the behavioral results: The power of theta episodes tended to be higher in the in the color condition than in the plain condition.

The most interesting results were obtained in the detailed analysis of the maze points where theta episodes occurred. It was found that these episodes did not occur either uniformly or at random points in the mazes. Rather, theta episodes occurred more likely and the power of these episodes was higher immediately after participants had made a turn in a junction and a new hallway came into view. Second, they were also more likely and more powerful after participants had walked into an invisible barrier and had to revise their navigational rules. In contrast, theta episodes were less likely when the participants were simply moving down a hallway. In

the first case, participants were probably retrieving a stored view that could guide their further navigation. In the second case, participants were most likely storing the view with a revised navigation rule. These results provide further evidence on the relation between theta waves and spatial navigation in humans: Theta waves may be directly related to the storage and retrieval of spatial information for navigation, not just in rodents,² but also in humans.

ACKNOWLEDGMENTS

This project was supported by grants from the National Engineering and Research Council of Canada. We would like to thank Reiko Graham,




FIG. 6. Average power in the theta band, separately for each cell type, averaged over all participants and all test maze traversals, graphed separately for the plain and the color condition. The results indicate that the power of the theta band is higher for cells immediately after a junction than for cells before the junction, suggesting that theta episodes are more likely and stronger when a new hallway segment becomes visible, or when an incorrect decision has been made, than when simply moving down a hallway.

Sherry Zsoldos, and Rachel Thomas for their help in running the experiments.

REFERENCES

1. Winson, J. (1978). Loss of hippocampal theta rhythms results in spatial memory deficits in the rat. *Science* 201:160–163.
2. O'Keefe, J., & Recce, M. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.
3. Skaggs, W.E., McNaughton, B.L., Wilson, M.A., et al. (1996). Theta phase precision in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6:149–172.
4. Gavrilov, V.V., Wiener, S.I., & Berthoz, A. (1995). Enhanced hippocampal theta EEG during whole body rotation in awake restrained rats. *Neuroscience Letters* 197:239–241.
5. Kahana, M.J., Sekuler, R., Caplan, J.B., et al. (1999). Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature* 399:781–784.
6. Caplan, J.B., Kahana, M.J., Sekuler, R., et al. (2000). Task dependence of human theta: the case for multiple cognitive functions. *Neurocomputing* 32–33:659–665.
7. Nishiyama, N., & Yamaguchi, Y. (2001). Human EEG theta in the spatial recognition task. Presented at the 5th World Multi-Conference on Systemics, Cybernetics and Informatics, Orlando, Florida.
8. De Araujo, D.B., Baffa, O., & Wakai, R.T. (2001). Theta and alpha oscillations: Dependency on navigation tasks. In: Nenonen, J., Ilmoniemi, R.J., Katila, T., (eds.), *Biomag2000. Proceedings of the 12th International Conference on Biomagnetism, Helsinki University of Technology, Espoo, Finland*. pp. 343–346. biomag2000.hut.fi/proceedings.html
9. McGuire, E.A., Burgess, N., Donnett, J.G., et al. (1998). Knowing where and getting there: a human navigation network. *Science* 280:921–924.
10. Kirschen, M.P., Kahana, M.J., Sekuler, R., et al. (2000). Optic flow helps human learn to navigate through synthetic environments. *Perception* 29:801–818.
11. Torres, D. (2003). A reconfigurable VR tool for spatial navigation (submitted).
12. American Electroencephalographic Society. (1991). Guidelines for standard electrode position nomenclature. *Journal of Clinical Neurophysiology* 8:200–202.

Address reprint requests to:

Walter F. Bischof, xx,
Department of Computing Science
University of Alberta
Edmonton, Alberta, T6G 2E9, Canada

E-mail: wfb@cs.ualberta.ca

