

# The Role of Geographical Slant in Virtual Environment Navigation

Sibylle D. Steck<sup>1</sup>, Horst F. Mochntzki<sup>2</sup>, and Hanspeter A. Mallot<sup>2</sup>

<sup>1</sup> DaimlerChrysler Research & Technology, Ulm, Germany

<sup>2</sup> Dept of Zoology, University of Tübingen, Germany.

**Abstract.** We investigated the role of geographical slant in simple navigation and spatial memory tasks, using an outdoor virtual environment. The whole environment could be slanted by an angle of  $4^\circ$ . Subjects could interact with the virtual environment by pedaling with force-feedback on a bicycle simulator (translation) or by hitting buttons (discrete rotations in  $60^\circ$  steps). After memory acquisition, spatial knowledge was accessed by three tasks: (i) pointing from various positions to the learned goals; (ii) choosing the more elevated of two landmarks from memory; (iii) drawing a sketch map of the environment. The number of navigation errors (wrong motion decisions with respect to the goal) was significantly reduced in the slanted conditions. Furthermore, we found that subjects were able to point to currently invisible targets in virtual environments. Adding a geographical slant improves this performance. We conclude that geographical slant plays a role either in the construction of a spatial memory, or in its readout, or in both.

## 1 Introduction

When we find our way in a familiar environment, we use various cues or types of information to find out where we are and, more importantly, where we should head from there. Besides egomotion information, which can be used for path integration, objects and landscape configurations are the most important sources of information. Places can be characterized by recognized objects (local landmarks) or by geometrical peculiarities such as the angle under which two streets meet (cf. Gouteux & Spelke 2001). A mixture of place and geocentric direction information is provided by distant or global landmarks (cf. Steck and Mallot 2000 for a discussion of local and global landmarks). Finally, true geocentric direction (or compass) information is conveyed by cues like the azimuth of the sun (in connection with the time of day) or the slant direction of a ramp-like terrain.

So far, the role of geographical slant and elevation in navigation is only poorly understood. Creem and Proffitt (1998) asked subjects to adjust the slant of a board to previously seen slants of terrain and found that slants as low as 4 degrees are well perceived. In an earlier study, Proffitt et al. (1995) showed that in virtual environments, subjects are also able to reproduce accurately geographical slant ( $5^\circ$  to  $60^\circ$  in  $5^\circ$  steps) on a tilt board. Further, the judgments in the virtual

environments and the naturally presented slants do not differ significantly. This result was confirmed in a study by Proffitt et al. (2001). The memory for elevation of places was studied by Gärling et al. (1990) who showed that subjects were able to judge from memory which of two places in a familiar environment was higher elevated. Subjects who had less experience with the environment tended to exaggerate the elevation differences. Evidence for the use of slant, i.e. the elevation gradient, in human spatial cognition comes from linguistic studies in people living in landscapes with conspicuous slants. Brown and Levinson (1993) and Levinson (1996) report that the Tzeltal language spoken in parts of Mexico uses an uphill/downhill reference frame even in contexts where English or other languages employ a left/right scheme.


In rats, a direct demonstration for the use of slant as a cue to navigation has been provided by Moghaddam et al. (1996). When searching a food source on top of an elevated cone, rats were able to navigate a more direct path than on a flat surface.

Theoretically, there are good reasons to expect that geographical slant should be used in navigation. First, some important navigation tasks such as “find water” can be solved by simply walking downhill. Note that no self-localization is required in this case. Second, geographical slant can provide geocentric<sup>1</sup> compass information which is known to be of great importance in path integration (see Maurer and Séguinot 1995, Mallot 2000). While path integration is principally possible by pure vector summation without any compass, error accumulation is greatly reduced if independent compass information is available. Insects which make extensive use of path integration (Müller and Wehner 1988) obtain compass information from the polarization pattern of the sky light (Rossel 1993). Finally, geographical slant might also act as a local cue characterizing a place. Indeed, it seems quite likely that the same landmark appearing on top of a mountain or halfway along the ascent are readily distinguished. Again, it has been shown in insects, that the so-called snapshot, a view of the environment characteristic of the location is viewed from, is registered to a compass direction (Cartwright and Collett 1982).

In this paper, we address the question whether global geographical slant can be used by human navigators to improve their performance. Three versions of a virtual environment differing only in the overall slant of the terrain were generated. After exploring one of these environments, subjects’ performance and spatial representation was assessed by measuring the overall navigation performance, the quality of pointing to remembered targets, the quality of judging which of two remembered places was higher in elevation, and the orientation of sketch map drawings.

---

<sup>1</sup> The term “geocentric” is used to indicate that some information is given in an observer independent coordinate system, fixed to some anchor point in the world. In contrast, the term “geographical” is used only in connection with the word “slant” to indicate that we are talking about the slant of landscapes rather than the slant of object surfaces. Finally, the term “geometrical” refers to depth as local position information, e.g. “a junction where streets meet at an angle of 45 degrees”.



**Fig. 1.** Virtual Environments Lab with  $180^\circ$  projection screen showing the Hexatown simulation. The subject was seated on a virtual reality bicycle in the center of the half cylinder.


## 2 Method

### 2.1 Subjects

A total of 36 subjects (18 male and 18 female, aged 15–31 years) took part in the experiment. Participation in this experiment was voluntarily and a honorarium was paid for participation.

### 2.2 Virtual Environment

**Graphical Apparatus.** The experiment was performed on a high end graphics computer (Silicon Graphics Inc. ONYX2 3-pipe Infinite Reality), running a C-Performer application that we designed and programmed. The simulation was displayed non-stereoscopically, with an update rate of 36 Hz, on a half-cylindrical projection screen (7m diameter and 3.15m height, Fig. 1). The computer rendered three  $1280 \times 1024$  pixel color-images projected side by side with a small overlap. Images were corrected for the curved surface by the projectors to form a  $3500 \times 1000$  pixel display. For an observer seated in the center of the cylinder (eye height 1.25m), this display covered a field of view of  $180^\circ$  horizontally times  $50^\circ$  vertically. The field of view of the observer was identical to the field of view used for the image calculations. A detailed description of the setup can be found in Veen et al. (1998).



**Fig. 2.** Overview of the three conditions. **Left:** map of the environments. Landmarks indicated by numbers have been used as goals in the exploration phase and as targets in the pointing phase. **Right:** subjects perspective. Each row shows the three pictures projected on the  $180^\circ$  screen. The images are projected with a small overlap; therefore the discontinuities visible here are not present in the actual experiment. The picture shows the view from the place with object 5 in the direction of the street towards the only adjacent place. **Top row** shows the *Flat* slant condition. **Middle row** shows the *Northeast* slant condition. **Bottom row** shows *Northwest*.

**Scenery.** In this experiment, we used three similar environments varying only in geographical slant (Fig. 2). In the control condition, the environment was on a flat plane (*Flat*). In the two other conditions, the environment had a global geographical slant with a slant angle of  $4^\circ$ . (The slant angle is the angle between the surface normal and the vertical; a slant angle of  $4^\circ$  is equivalent to an inclination of 7%.) In pilot studies (Mochnatzki 1999), we found that the simulated slant was well above the detection threshold for geographic slant in the same experimental setup. The slanted environments differed in the orientation of the slant with respect to an arbitrarily chosen “North” direction. In one condition, the geographical slant was oriented in the direction of Northeast (*NE*). In a further condition, the slant was to the Northwest (*NW*). The reasons for using two slanted environments are the following: First, the street raster of the virtual town is not completely isotropic, so different slant directions might have different effects. Second, the sequence of learning tasks used in the exploration phase (see below) introduces an additional anisotropy which makes it necessary to use at least two slant directions.

The model of the environment was generated using MultiGen 3-D modeling software. The environment consisted of an octagonal ground plane surrounded by a flat background showing a regular mountain range. The buildings were constructed using Medit 3-D modeling software. Schematic maps of the town appear in the left side column of Fig. 2. Maps or aerial views were not shown to the subjects. The virtual environment (called “Hexatown”, see Gillner and Mallot, 1998, Steck and Mallot 2000, and Mallot and Gillner 2000) consisted of a hexagonal raster of streets with a distance of 100 meters between adjacent junctions. A junction was built of three adjoining streets forming  $120^\circ$  corners. In each corner, an object (building, gas station, etc.) was placed, see Fig. 2. At the periphery of Hexatown, streets ended blindly. These dead-ends were marked by barriers 50 meters from the junction. A circular hedge or row of trees was placed around each junction with an opening for each of the three streets (or dead ends) connected to that junction. This hedge looked the same for all junctions and prevented subjects from seeing the objects at distant junctions.

The usage of geometrical cues, as demonstrated, e.g., by Hermer and Spelke (1994) and Janzen et al. (2000) is not possible in Hexatown. All junctions are identical and symmetrical, so that when approaching a junction, one cannot infer the approach direction nor the approached place from the geometrical layout. As compared to rectangular city rasters, which are also symmetrical, the hexagonal layout has the advantage that there is no straight-on direction that might be preferred over the branching streets.

**Interaction.** Subjects navigated through Hexatown using a virtual reality bicycle (a modified versions of a training bicycle from *CyberGear<sup>TM</sup>*) which can be seen in Fig. 1 (for details see van Veen et al. 1998). The bicycle has force-feedback, i.e. when pedaling uphill, the subjects have to exert more force than when cycling downhill. The setup thus provides both visual and proprioceptive slant information.

At the junctions,  $60^\circ$  turns could be performed by pressing one of two buttons (left or right) fixed to the bicycle. Once the simulated turn movement was initiated, it followed a predefined velocity profile: turns took 4 seconds, with a maximum speed of  $30^\circ$  per second and symmetric acceleration and deceleration (ballistic movement). The smooth profiles for rotation were chosen to minimize simulator sickness. Translations on the street were initiated by pressing an additional button. Translations were not ballistic; translation velocity was controlled by the pedal revolution, using a mechanical motion model that took into account the current geographical slant. Subjects could only influence the speed, but were not able to change the direction, i.e. they were restricted to the streets.

In the sequel, the motion initiated by pressing the buttons will be referred to as motion decision.

### 2.3 Procedure

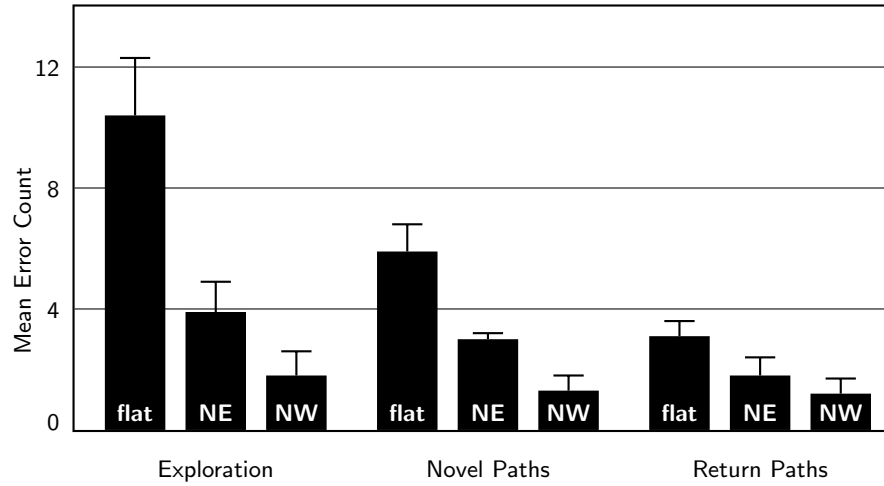
The three experimental conditions were tested in an between-subject design, using 12 subjects per condition. Subjects were run through the experiment individually.

The experiment had four different phases: a navigation phase, pointing judgments, elevation comparison, and map drawing. In the navigation task, the subjects had to find a previously shown goal using the shortest possible path. The navigation phase consisted of 15 search tasks. In the pointing judgment, subjects were asked to carry out directional judgments to previously learned goals. In the elevation judgments, subjects had to choose which learned goal was higher up in the environment. This part was omitted in the *Flat* condition. Finally, subjects had to draw a map from the learned environment. For each part, subjects were instructed separately. Therefore, they were uninformed of all tasks in advance. On average, subjects needed 90 min for all tasks.

**Navigation Phase.** First, the subjects had to solve 15 search tasks in the primarily unknown environment (Fig 2). Before each trial, a full 180° panoramic-view at the goal location was shown. By pressing a button on the handles of the VR-bicycle, the goal presentation was terminated and subjects were positioned at the current starting position. When they had reached their goal, a message was displayed, indicating whether they had used the path with the least number of motion decisions (“fastest” path), or not. The task was repeated until it was first completed without mistakes. During the entire navigation phase, the subjects had the possibility to expose a small picture of the current goal object on a gray background in the bottom left corner of the middle screen by pressing a special button. The starting point of the first five tasks was landmark 15 (*home*). The solutions of the first five search tasks covered the entire maze; we therefore call these phase *exploration*. The next ten routes were either *return paths* from the previously learned goals to the landmark 15, or *novel paths* between the goals, which were learned in the exploration phase. Search tasks involving a return and a novel path were carried out in alternation. The navigation phase ensured that all subjects reached a common fixed performance level for the subsequent pointing judgments.

**Pointing Judgments.** Pointing judgments were made to evaluate the internal representation of the learned environment. The subjects were placed in front of a learned goal, which was randomly chosen. They were asked to orient themselves towards one of four other goals (except *home*) by continuously turning the simulated environment. A fixed pointer (fixed with respect to the screen) was superimposed on the turning image to mark the forward direction to which the goal had to be aligned. Note that this procedure differs considerably from pointing in real environments or in virtual environments presented using a head-mounted display, in that the observer’s arm or body need not move. All that moves during the “pointing” judgment is the image of the simulated environment. For a discussion of pointing procedures, see Montello et al. (1999). Altogether, the subjects had to point to twenty goals. One of these goals was directly visible from one of the reference points. This pointing task was therefore excluded from further analysis.

**Elevation Judgments.** In order to test whether elevation information was also stored, elevation judgments were collected in the *Northeast* and *Northwest* conditions. Pictures of two goals of different elevation were presented in isolation



**Fig. 3.** Mean error count in the navigation phase. Mean number of errors for the three route types (*exploration*, *novel paths*, and *return paths*) and the three slant conditions, *flat*, *slanted NE*, and *slanted NW*. The error bars present one standard error of the mean.

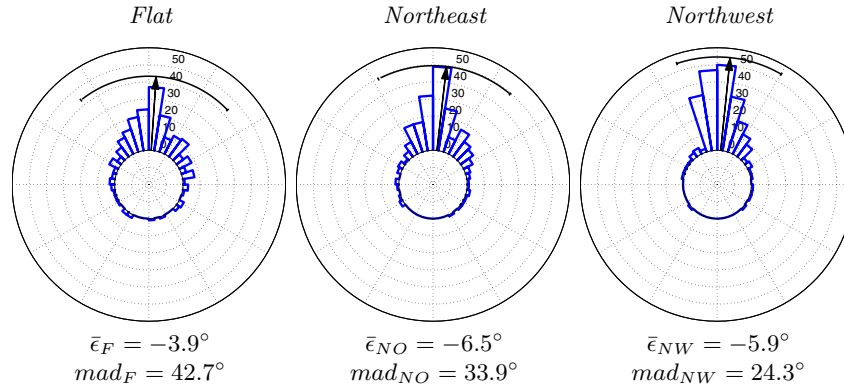
on a gray screen and the subjects had to decide as accurately and as quickly as possible, which goal had appeared at higher elevation in the training environment. For each of the two slant conditions, ten pairs of goals were selected and tested.

**Map Drawing.** In the final phase of the experiment, subjects were asked to draw by hand as detailed a map of the test environment as possible. They were given a pen and a paper. The paper had a printed frame to restrict their drawings. There was no time limit for the subjects.

### 3 Results

#### 3.1 Errors in the Navigation Phase

In the navigation phase, the trajectories of the subjects for every search task were recorded. Every movement decision that did not reduce the distance to the goal was counted as an error. Figure 3 shows the mean number of errors per path type (*exploration*, *return paths*, and *novel paths*) and per slant condition. A three-way ANOVA (3 path types  $\times$  3 slant conditions  $\times$  gender) shows a significant main effect of slant condition ( $F(2, 30) = 5.78$ ,  $p = 0.008^{**}$ ). As figure 3 shows, more errors were made in the *Flat* condition than in the *Northeast* condition. In the *Northwest* slant condition, the least amount of error was made. Further, there was a highly significant main effect of the path type ( $F(2, 60) = 27.69$ ,  $p < 0.001^{***}$ ). In all three slant conditions, the largest number of errors occurred in the *exploration* phase (first five paths, all starting from home). The second



**Fig. 4.** Pointing Error. Circular plots for the slant condition Flat, Northeast and Northwest.  $\bar{\epsilon}$ : circular mean of the error (arrow).  $mad$ : mean angular deviation (segment).

largest number of errors was made for the *novel paths* (connection paths between goals, none of which was the home), while the *return paths* were navigated with the smallest number of errors. Note that the *return paths* alternated with the *novel paths* in the task sequence; therefore the difference in the number of errors of these two path types cannot be explained by differences in the time spent in the environment before each task.

A significant interaction between slant condition and path type was also found ( $F(4, 60) = 4.37, p = 0.004^{**}$ ). It may reflect a floor effect for the *Northwest* slant condition. Since the number of errors was very small in this condition anyway, the effects of condition and path type do not completely superimpose. No difference in the mean number of errors was found between male and female subjects (men:  $11.5 \pm 1.9$ , women:  $10.2 \pm 1.6$ ,  $F(1, 30) = 0.300, p = 0.59$  n.s).

### 3.2 Pointing Judgments

The pointing judgments were stored as angles in degrees with respect to the arbitrarily chosen North direction. Since pointing judgments are periodic data (e.g.,  $181^\circ$  is the same direction as  $-179^\circ$ ), we used circular statistics (see Batschelet 1981) to analyze pointing judgments. The circular means ( $\bar{\alpha}$ ) were calculated by summing the unit vectors in the direction of the pointings. The resultant vector was divided by the number of averaged vectors. The length of the mean vector is a measure for the variability of the data.

To compare the different slant conditions, the deviations from the correct values were averaged over all tasks. Figure 4 shows the deviation from the correct values for all tasks and all subjects. The measured values were distributed in  $9^\circ$  bins. The arrow shows the direction of the circular mean of the errors. The length of the mean vectors is inversely proportional to the “mean angular deviation” shown as circular arc in Fig. 4. The mean vectors are close to zero for all conditions, as is to be expected since we plotted the pointing error.

**Table 1.** Comparison of the variances of pointing in the three slant conditions.

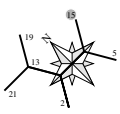
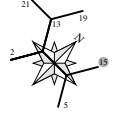
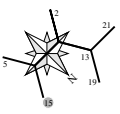
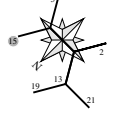
Comparison	$F(198,198) = \frac{mad_1^2}{mad_2^2}$	p
Flat– Northeast	1.72	< 0.001 ***
Flat – Northwest	3.5257	< 0.001 ***
Northeast – Northwest	2.0408	< 0.001 ***

For comparing the variances of the different slant conditions, we compared the arithmetic mean of the squares of the mean angular deviation of each subject using the circular F-test (Batschelet 1981, chap.6.9). There is a highly significant difference between all conditions, see Table 1.

### 3.3 Elevation Judgment

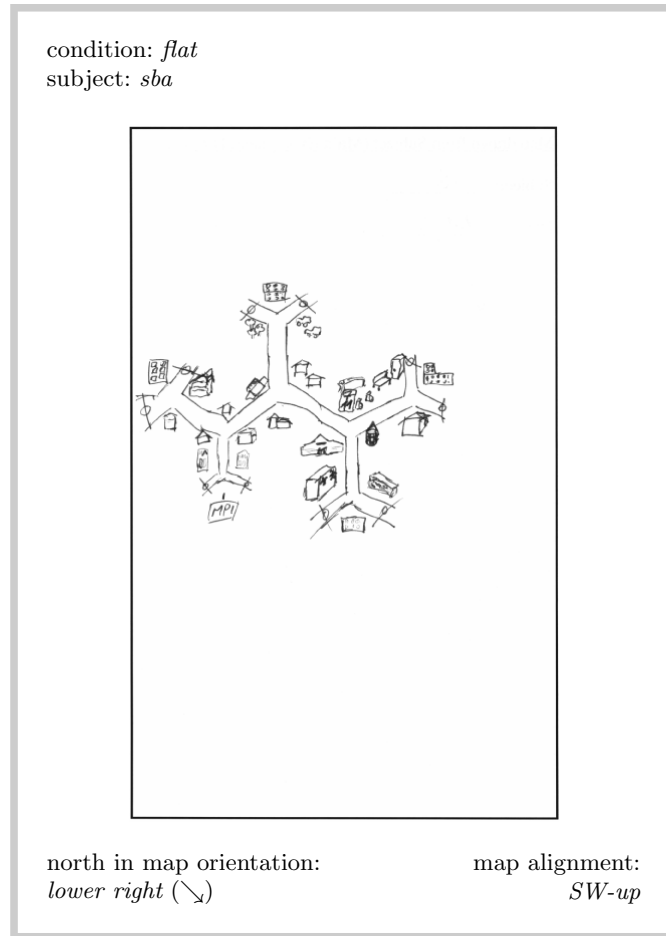
In this part, subjects in the *slanted NE* and *slanted NW* conditions were tested to determine, if they stored the relative elevations of the objects. The subjects in the *Northwest* slant condition gave 109 correct answers out of 120, 90.8%, and the subjects in *Northeast* 94 correct answers out of 120, 78.3%. The answers of the subjects differed significantly from a binomial distribution with  $p = 50\%$  which would imply pure guessing ( $\chi_{NE}^2(10) = 492.0$ ,  $p < 0.001^{***}$ ,  $\chi_{NW}^2(10) = 3838.9$ ,  $p < 0.001^{***}$ ). Therefore, we conclude that the subjects were able to differentiate object elevation. The percentage correct of the *Northwest* condition was significantly higher than the percentage *Northeast* (U-Test after Mann and Whitney  $U(12, 12, p = 0.05) = 37$ ,  $p \leq 0.05^*$ ).

**Table 2.** Alignment of sketch maps in the three slant conditions

	Alignment				ambiguous
	NE–up	NW–up	SW–up	SE–up	
<i>flat</i>					
<i>slanted NE</i>	3	0	6	0	3
<i>slanted NW</i>	6	0	2	0	4
<i>slanted NW</i>	2	5	1	0	4

### 3.4 Map Drawings

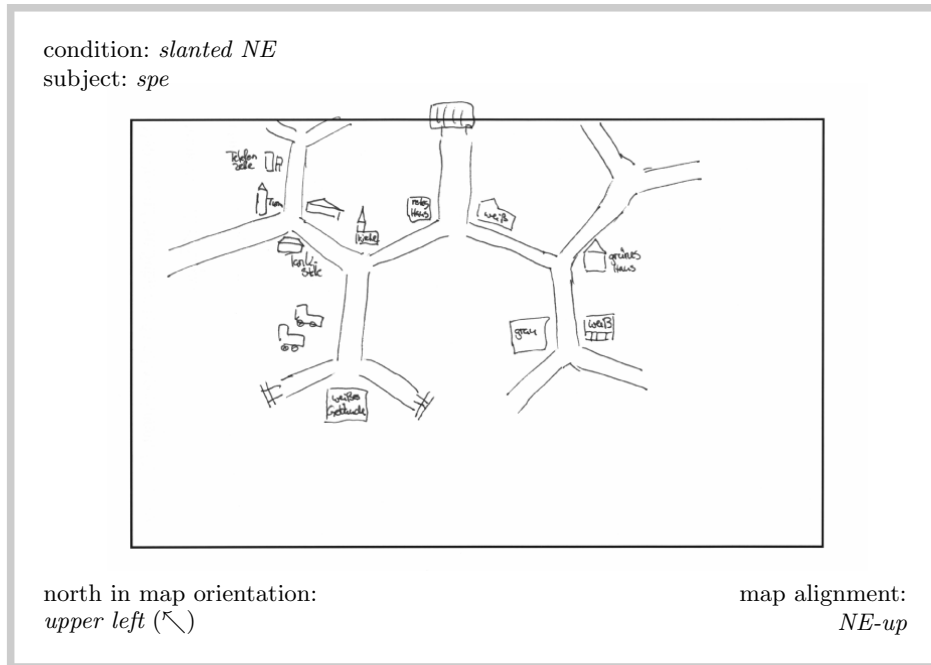
The map drawings were used to study how subjects implemented the geographical slant in their representation. Single maps were mostly quite good, since the



**Fig. 5.** Sketch map drawn by subjects *sba* in condition *flat*. The drawing is aligned in the sense that all buildings are given in perspective with the same vantage point. The top of the page corresponds to Southwest. The bold gray box indicates the size of the sketching paper ( $A4 = 21 \text{ cm} \times 29.7 \text{ cm}$ ). The thin black box is the frame printed on the sketching paper to prevent subjects from starting their drawings too closely to the edge of the paper.

geometry of the junction was often correctly depicted. Only three out of thirty-six subjects drew all junctions as right angle junctions. Four further subjects drew right angles at some junctions. All except one very sparse map, contained object 15, which was the start point of the first five routes.

We were interested in whether the slant conditions influenced the map drawings. Therefore, all maps were examined for alignment. A map was considered “aligned”, if either a uniform orientation of lettering (e.g., Fig. 7) or a perspective of the drawn objects (e.g., Fig. 5) was apparent to the authors. Judgments



**Fig. 6.** Sketch map drawn by subject *spe* in condition *Northeast*. The alignment is apparent from the drawings of the houses. The top of the page corresponds to Northeast, i.e. the more elevated locations are drawn more towards the top of the page. The boxes represent the margin and inner frame of the sketching paper (cf. Fig. 5).

of alignment were carried out independently and maps judged differently are labeled “ambiguous” in Table 2.

The maps were categorized in four groups: NE-up, SE-up, SW-up and NW-up. Table 2 lists the number of drawn maps for all alignment categories for the three different slant conditions (*flat*, *slanted NE*, and *slanted NW*). In the *flat* slant condition, the SW-up alignment was found six times. In this alignment category, object 15 is at the lower edge of the map, and the street, which leads to the next junction, points to the top (cf. Figure 5). Further, the category NE-up (in which the object 15 is at the top edge of the map, and the street, which leads to the next junction, points to the bottom) occurred three times. In the *Northeast* slant condition, the alignment category NE-up occurred six times and SW-up two times. In both cases (NE-up, SW-up), the maps were aligned with the gradient along the geographical slant, with the majority of the maps aligned to the uphill gradient (see Figure 6). In the *Northwest* slant condition, the alignment category NW-up (i.e., uphill along the gradient) occurred five times (cf. Figure 7). There were two maps of the category NE-up and one map of the category SW-up. The distributions of the maps in the alignment categories differ significantly ( $\chi^2(\text{slanted NW}/\text{flat}) = 30.5$ ,



occurred for each route type (exploration, return, novel route) individually, not just for routes leading to a goal uphill in the environment. It appears therefore that slant information is used to improve spatial knowledge in general. In contrast, in the study by Moghaddam et al. (1996), only the navigation to targets on top of a hill was addressed.

A surprising result is the difference between the two slant conditions, which differ only in the direction of the slant relative to the maze layout. We speculate that the difference is related to the fact that in the *slanted NE* condition, the longest route (four segments) is running in zigzag pattern up and down the slope, whereas in the *slanted NW* condition, the longest route is constantly going uphill or downhill. Therefore, the slant information is ambiguous in the *slanted NE* condition.

The results from the navigation part of the experiment are well in line with the pointing judgments. Again, pointing is better for the slanted conditions, and it is also better for the *slanted NW* than for the *slanted NE* condition. We found no difference in judgment accuracy between pointings parallel to the slant and pointings perpendicular to the slant.

Improved pointing in slanted environments is to be expected if slant is used as a compass in a path integration scheme. However, this mechanism does not explain the difference found between the two slant conditions.

## 4.2 Slant Information and Spatial Memory

The results from the elevation judgment part of the experiment show that the subjects remember the relative elevation of the various objects in the maze. This finding is well in line with the results of Gärling et al. (1990). In graph-like, or topological models of spatial memory (Kuipers 2000, Gillner & Mallot 1998), elevation may be attached as a label to each node of the graph. Alternatively, local slant information could be attached to the edges of the graph. In a recent model by Hübner & Mallot (2002), the graph contains local metric information, including distances between adjacent nodes and angles between adjacent edges. A generalized multi-dimensional scaling algorithm is then used to estimate the metric 2D coordinates of each node. This scheme can be generalized to account for slant data. Local slant data could then be used to generate elevation estimates per node, by some sort of 3D path integration. Evidence for 3D path integration in insects has recently been provided by Wohlgemuth et al. (2001).

Overall slant direction, as seems to be present in the map drawings, is not easily represented in pure graph models. Indeed, some metric structure (as opposed to mere neighborhood information) is necessary to represent a global slant direction in a cognitive map. Further experiments with irregular slants will be needed to access the roles of global and local slant directions.

## Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft, Grant Numbers MA 1038/6-1, MA 1038/7-1 and by the Office of Naval Research Grant

N00014-95-1-0573 award to Jack Loomis. We are grateful to Silicon Graphics Inc., Prof. F. Leberl (Univ. Graz), and the Salford University, UK, for providing the VR-models used in this experiments. The authors thank Jan Restat and Pavel Zahorik for comments and suggestions on an earlier draft of this manuscript. We are grateful to Scott Yu for providing the 3D model of our virtual environments lab shown in Fig. 1.

## References

- Batschelet, E. (1981). *Circular Statistics in Biology*. Academic Press, London.
- Brown, P. and Levinson, S. C. (1993). "Uphill" and "Downhill" in Tzeltal. *Journal of Linguistic Anthropology*, 3(1):46 – 74.
- Cartwright, B. A. and Collett, T. S. (1982). How honey bees use landmarks to guide their return to a food source. *Nature*, 295:560 – 564.
- Creem, S. H. and Proffitt, D. R. (1998). Two memories for geographical slant: Separation and interdependence of action and awareness. *Psychonomic Bulletin & Review*, 5:22 – 36.
- Gärbling, T., Böök, A., Lindberg, E., and Arce, C. (1990). Is elevation encoded in cognitive maps. *Journal of Environmental Psychology*, 10:341 – 351.
- Gillner, S. and Mallot, H. A. (1998). Navigation and acquisition of spatial knowledge in a virtual maze. *Journal of Cognitive Neuroscience*, 10:445 – 463.
- Gouteux, S. and Spelke, E. S. (2001). Children's use of geometry and landmarks to reorient in an open space. *Cognition*, 81:119 – 148.
- Hermer, L. and Spelke, E. S. (1994). A geometric process for spatial reorientation in young children. *Nature*, 370:57 – 59.
- Hübner, W. and Mallot, H. A. (2002). Integration of metric place relations in a landmark graph. In Dorronsoro, J. R., editor, *International Conference on Artificial Neural Networks (ICANN 2002)*, Lecture Notes in Computer Science. Springer Verlag.
- Janzen, G., Herrmann, T., Katz, S., and Schweizer, K. (2000). Oblique angled intersections and barriers: Navigating through a virtual maze. *Lecture Notes in Computer Science*, 1849:277–294.
- Kuipers, B. (2000). The spatial semantic hierarchy. *Artificial Intelligence*, 119:191 – 233.
- Levinson, S. C. (1996). Frames of reference and Molyneux's question: Crosslinguistic studies. In Bloom, P., Peterson, M. A., Nadel, L., and Garrett, M. F., editors, *Language and Space*, pages 109 – 169. The MIT Press, Cambridge, MA.
- Mallot, H. (2000). *Computational Vision. Information Processing in Perception and Visual Behavior*, chapter Visual Navigation. The MIT Press, Cambridge, MA.
- Mallot, H. A. and Gillner, S. (2000). Route navigation without place recognition: what is recognized in recognition-triggered responses? *Perception*, 29:43 – 55.
- Maurer, R. and Séguinot, V. (1995). What is modelling for? A critical review of the models of path integration. *Journal of theoretical Biology*, 175:457 – 475.

- Mochnatzki, H. (1999). Die Rolle von Hangneigungen beim Aufbau eines Ortsgedächtnisses: Verhaltensversuche in Virtuellen Umgebungen. Diploma thesis, Fakultät für Biologie, Univ. Tübingen.
- Moghaddam, M., Kaminsky, Y. L., Zahalka, A., and Bures, J. (1996). Vestibular navigation directed by the slope of terrain. *Proceedings of the National Academy of Sciences, USA*, 93:3439–3443.
- Montello, D. R., Richardson, A. E., Hegarty, M., and Provenza, M. (1999). A comparison of methods for estimating directions in egocentric space. *Perception*, 28:981 – 1000.
- Müller, M. and Wehner, R. (1988). Path integration in desert ants, *cataglyphis fortis*. *Proceedings of the National Academy of Sciences, USA*, 85:5287 – 5290.
- Proffitt, D. R., Bhalla, M., Gossweiler, R., and Midgett, J. (1995). Perceiving geographical slant. *Psychonomic Bulletin & Review*, 2:409 – 428.
- Proffitt, D. R., Creem, S. H., and Zosh, W. D. (2001). Seeing mountains in mole hills: geographical-slant perception. *Psychological Science*, 12:418 – 423.
- Rossel, S. (1993). Navigation by bees using polarized skylight. *Comparative Biochemistry & Physiology*, 104A:695 – 708.
- Steck, S. D. and Mallot, H. A. (2000). The role of global and local landmarks in virtual environment navigation. *Presence. Teleoperators and Virtual Environments*, 9:69 – 83.
- Veen, H. A. H. C. v., Distler, H. K., Braun, S. J., and Bühlhoff, H. H. (1998). Navigating through a virtual city: Using virtual reality technology to study human action and perception. *Future Generation Computer Systems*, 14:231 – 242.
- Wohlgenuth, S., Ronacher, R., and Wehner, R. (2001). Ant odometry in the third dimension. *Nature*, 411:795–798.