

Behavioral experiments in spatial cognition using virtual reality

Hanspeter A. Mallot, Sabine Gillner, Hendrik A.H.C. van Veen,
and Heinrich H. Bülthoff

Max-Planck-Institut für biologische Kybernetik
Spemannstr. 38, D-72076 Tübingen, Germany
<http://www.kyb.tuebingen.mpg.de>

Abstract. Virtual reality is used as a novel tool for behavioral experiments on humans. Two environments, Hexatown and Virtual Tübingen, are presented. Experiments on cognitive maps carried out in the Hexatown environment are reported in this paper. Results indicate that subjects are able to acquire configuration knowledge of the virtual town even in the absence of physical movement. Simpler mechanisms such as associations of views with movements are also present. We discuss the results in relation to a graph-theoretic approach to cognitive maps.

1 General Introduction

1.1 Mechanisms and competences

In animal navigation, at least three basic mechanisms of spatial memory have been identified which we will refer to as path integration, guidance, and direction. Guidance and direction use local position information, i.e., sensory input characteristic of a given place, but use this information in different ways. Path integration uses egomotion data and can function even in the absence of local position information. In more detail, the characteristics of the three mechanisms are as follows:

- *Path integration* or dead reckoning is the continuous update of the egocentric coordinates of the starting position based on instantaneous displacement and rotation data (see Maurer & Séguinot 1995 for review). Odometry data are often taken from optic flow but other modalities such as proprioception (e.g., counting steps) may be involved as well. Since error accumulation is a problem, the use of global orientation information (“compasses”, e.g., distant landmarks or the polarization pattern of the skylight) is advantageous. Path integration involves some kind of working memory in which only the current “home-vector” (coordinates of the starting point) is represented, not the entire path.
- *Guidance* is a mechanism in which the navigator keeps a fixed or otherwise well-defined orientation with respect to a landmark or a group of landmarks (see O’Keefe & Nadel, 1987, p. 82). A well studied guidance mechanism

in insect navigation is snapshot-based homing, i.e. the approach of a place whose local view matches a stored “snapshot” (Cartwright & Collett 1982). This mechanism requires long-term storage of the view or snapshot visible at that point. From a comparison of the stored view with the current view, an approach direction can be derived. Moving in this direction will lead to a closer match between the two views (Cartwright & Collett 1982, Mallot et al. 1997). More generally, “view” and “snapshot” may be replaced by local position information from any sensory modality.

- *Direction* is the association of a recognized view (local position information) to a movement. As for guidance, long-term memory of the local position information (view) is required. In addition to that, a movement direction is stored, i.e., the recognized view acts as a pointer to some other goal. The existence of such associations has been shown in bees (Collett & Baron 1995) and humans (Mallot & Gillner 1997). The direction mechanism can be generalized to associate views with more complex behaviors such as wall following or passing through a door (Kuipers & Byun, 1988).

For a more comprehensive discussion of these and other mechanisms, see O’Keefe & Nadel (1978; page 80ff) and Trullier et al. (1997).

Distinguishing between these mechanisms leads to a related distinction between different types of landmarks. We use the term “global landmarks” for distant markers serving as a compass system in path integration. In contrast, the terms “local landmark” or “view” refer to local position information as used for guidance and direction. Note that the same object can have both roles in a given scene.

Using these basic mechanisms, different levels of complexity of spatial knowledge and behavior can be formulated. Concatenating individual steps of either guidance or direction results in routes. These routes will be stereotyped and could be learnt in a reinforcement scheme. More biologically plausible, however, is instrumental learning, i.e., the learning of associations of actions with their expected results. This can be done step-by-step without pursuing a particular goal (latent learning). Instrumental learning entails an important extension of the two view-based mechanisms in that the respective consequences of each of a number of possible choices (either movements or snapshots to home to) are learnt. This offers the possibility of dealing with bifurcations and choosing among alternative actions. Thus, the routes or chains of steps can be extended to actual graphs which are a more complete representation of space, or cognitive map (Schölkopf & Mallot 1995, Mallot et al. 1997). The overall behavior is no longer stereotyped but can be planned and adapted to different goals.

1.2 The view-graph approach to cognitive maps

We discuss the representation of spatial configurations in graph structures for the direction mechanism. Similar schemes can be developed for the other mechanisms or combinations thereof. As the basic element, let us consider the association of

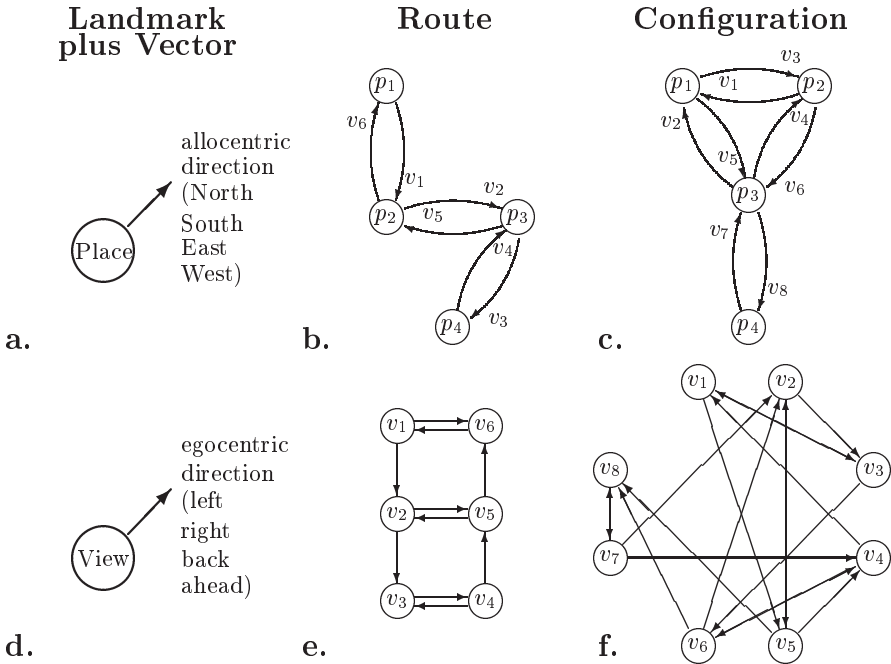


Fig. 1. The graph approach to space representation. Top row (a.–c.): Place-graphs. The nodes are places recognized irrespective of body orientation, the links (arrows) between them carry allocentric direction information. Bottom row (d.–f.): View-graphs. The nodes are recognizable views or other positional information (i.e., depend on the observer's viewing direction) and the arrows carry directional information relative to gaze. Each view v_i in Figs. e. and f. corresponds to a directed connection in Fig. b. and c. From left to right, increasingly complicated spatial layouts are shown.

a recognized view with a movement direction. Additionally, we store the view expected when performing this movement:

$$(\text{current view}, (\text{movement direction}, \text{expected next view})). \quad (1)$$

This association is illustrated in Fig. 1d. When going from one view to the next, navigation can initially follow the movement direction associated with the present view. The additional information on what view to expect next can be used to improve view recognition at the next step. If direction and snapshot-based homing are to be combined, the expected next view is the one to home to.

Chains of such association structures implement a route memory. If different routes are to be learned that share some common section, the decision at the crossroads requires more complicated memory. One way to think of this memory

is to store all possible connections

$$\begin{aligned} &(\text{current view, (movement direction 1, expected next view)}, & (2) \\ &\quad \vdots \\ &(\text{movement direction } n, \text{ expected next view})), \end{aligned}$$

and have a separate planning device select one of the possible movements. The resulting memory structure is a graph of views and movement decisions as shown in Fig. 1e,f. A neural network theory for storing the required information in the form of a labelled graph has been presented by Schölkopf and Mallot (1995). For related approaches including hippocampal modelling, see Touretzky and Redish (1996), Muller, Stead & Pach (1996), and Prescott (1996).

1.3 Behavioral experiments in virtual reality

In the work reported in this paper, we chose interactive computer graphics, or virtual reality (VR), as our experimental method. Previous studies using virtual reality have focussed on the transfer of knowledge between different media used for acquisition and testing. May, Péruch & Savoyant (1995) and Tlauka & Wilson (1996), for example, have tested map-acquired knowledge in a pointing task performed in virtual reality. Tong, Marlin & Frost (1995), using a VR bicycle, showed that active exploration leads to better spatial knowledge than passive stimulus presentation. Sketch maps produced after exploration of various virtual environments have been studied by Billinghamurst & Weghorst (1995). Design principles for constructing easy-to-navigate virtual environments have been studied by Darken & Sibert (1996). Ruddle et al. (1997), using a VR-setup similar to the one used in this study, showed that navigation performance in a simulated indoor environment is essentially as good as found in experiments carried out in real buildings. In the present paper, we use virtual reality to isolate the various cues used for the build-up of spatial knowledge and to study the underlying mechanisms. The advantages of virtual reality for this application are (i) the high controllability of computer graphics stimuli, and (ii) the easy access to behavioral data, such as the subject's movement decisions.

Stimulus control. When investigating the information sources used in navigation, it is advantageous to be aware of the exact movement trajectories of the subjects and the visual information available along these trajectories. This can easily be achieved with interactive computer graphic (see Section "Methods"). The various parameters of the sensory input can be easily separated. For instance, in our experiments, we varied the number of buildings visible simultaneously in one view without changing the illumination etc. In real world experiments, such separate stimulus conditions are much harder to realize. Another interesting experimental paradigm is the modification or exchange of various features of the environment after learning. Aginsky, Harris, Rensink & Beusmans (1996) exchanged landmarks after training in a route-learning task. The effects of landmark exchange on navigation have been addressed by Mallot & Gillner (1997).

In principle, the method also allows complete control over vestibular and proprioceptive feedback. In our experiments, for instance, both were completely absent allowing the effects of visual input to be studied in isolation.

Measuring behavior. Navigation performance can be accessed most directly by the paths or trajectories that the subjects take during the exploration. In virtual reality experiments, egomotion is very simple to record, since it is equivalent to the course of the “simulated observer” used for rendering the computer graphics. In this paper, we present a number of novel techniques for data evaluation that are particularly suited for the virtual reality experiments described.

2 Virtual environments

The experiments which we report in this paper have all been conducted in a specially designed virtual village called “Hexatown”. This section describes the layout and construction of this village and explains how subjects interact with it. We also list and motivate the different stimulus conditions that are used in the experiments. A valid question that one can ask is whether results obtained using this kind of virtual environments can be transferred to navigation in the natural environment. The second part of this section describes another virtual environment that we are developing and that hopefully will help to address this difficult question.

The virtual environments were constructed using Medit 3D-modelling software and animated with a framerate of 36 *Hz* on a SGI Onyx RealityEngine² using IRIX Performer software.

2.1 Hexatown

Geometry. A schematic map of Hexatown appears in Fig. 2a. It is built on a hexagonal raster with a distance between two places of 100 meters. At each junction, one object, normally a building, was located in each of the 120 degree angles between the streets; thus each place consisted of three objects. In the places with less than three incoming streets, dead ends were added instead, ending with a barrier at about 50 meters. The whole town was surrounded by a distant circular mountain ridge showing no salient features. The mountains were constructed from a small model which was repeated periodically every 20 degrees.

An aerial view of Hexatown is shown in Fig. 3¹. It gives an impression of the objects used. The spacing and position of the trees corresponds to viewing condition 1 and 2 (see below).

The hexagonal layout was chosen to make all junctions look alike, such that no information is contained in the geometry of the junctions themselves. In rats

¹ Color versions of Figures 3 and 4 are available from the world wide web, <http://www.kyb.tuebingen.mpg.de/links/hexatown.html>.




Fig. 2. a. Street map of the virtual maze with 7 places numbered 0 – 6 and 21 views numbered 0 – 20. The ring around each place indicates the hedges used in viewing conditions 1 – 3 to occlude distant places. b. Possible movement decisions when facing the view marked *a*. *L*: turn left 60 degrees. *R*: turn right 60 degrees. *G*: go ahead to next place.

navigating simple cross-like mazes with non-orthogonal arms, it has been shown that geometry may be more important than landmarks presented as cues. When a conflict between local landmarks placed within the arms and geometry information was introduced, rats behaved as if they would ignore the landmarks and followed the geometry information (Cheng 1986). Since in this study, we want to limit the available information to local landmarks (“views”), the hexagonal layout seems appropriate. An additional problem arising in Cartesian grids (city-block raster) is the fact that long corridors are visible at all times and the possible decisions at a junction are highly unequal: going straight to a visible target or turning to something not presently visible.

Simulated movements. Subjects could move about the town using a computer mouse. In order to have controlled visual input and not to distract subject’s attention too much, movements were restricted in the following way. Subjects could move along the street on an invisible rail right in the middle of each street. This movement was initiated by hitting the middle mouse button and was then carried out with a predefined velocity profile without further possibilities for the subject to interact. The translation took 8.4 seconds with a fast acceleration to the maximum speed of 17 meters per second and a slow deceleration. The movement ended at the next junction, in front of the object facing the incoming street. 60 degree turns could be performed similarly by pressing the left or right mouse button. Again, the simulated movement was “ballistic”, i.e., following a predefined velocity profile. Turns took 1.7 seconds with a maximum speed of 70 degrees per second and symmetric acceleration and deceleration.



Fig. 3. Aerial view of Hexatown. Orientation as in Fig. 2a. The white rectangle in the left foreground is view 15, used as “home”-position in our experiments. The aerial view was not available to the subjects. Object models are courtesy of Silicon Graphics, Inc., and Prof. F. Leberl, Graz.

Fig. 2 shows the movement decisions that subjects could choose from. Each transition between two views is mediated by two movement decisions. When facing an object (e.g., the one marked “a” in Fig. 2), 60 degrees turns left or right (marked “L”, “R”) can be performed which will lead to a view down a street. If this is not a dead end, three decisions are possible: the middle mouse button triggers a translation down the street (marked “G” for go), while the left and right buttons lead to 60 degrees turns. If the street is a dead end, turns are the only possible decision. In any case, the second movement will end in front of another object.

Viewing conditions. We used four stimulus conditions with varying degrees of visibility of the environment. In condition 1, navigation is view-based in the strictest sense, i.e. subjects have to rely on local views only. The other conditions contain increasingly more non-local information. For example views occurring in the four conditions, see Fig. 4.

In condition 1, 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 more distant junctions. The objects were placed 22 meters away from the center of the junction. The arrangement was such that when entering the circular hedge, the buildings to the left and right were already outside the observer’s field of view (60 degrees). Thus, the three

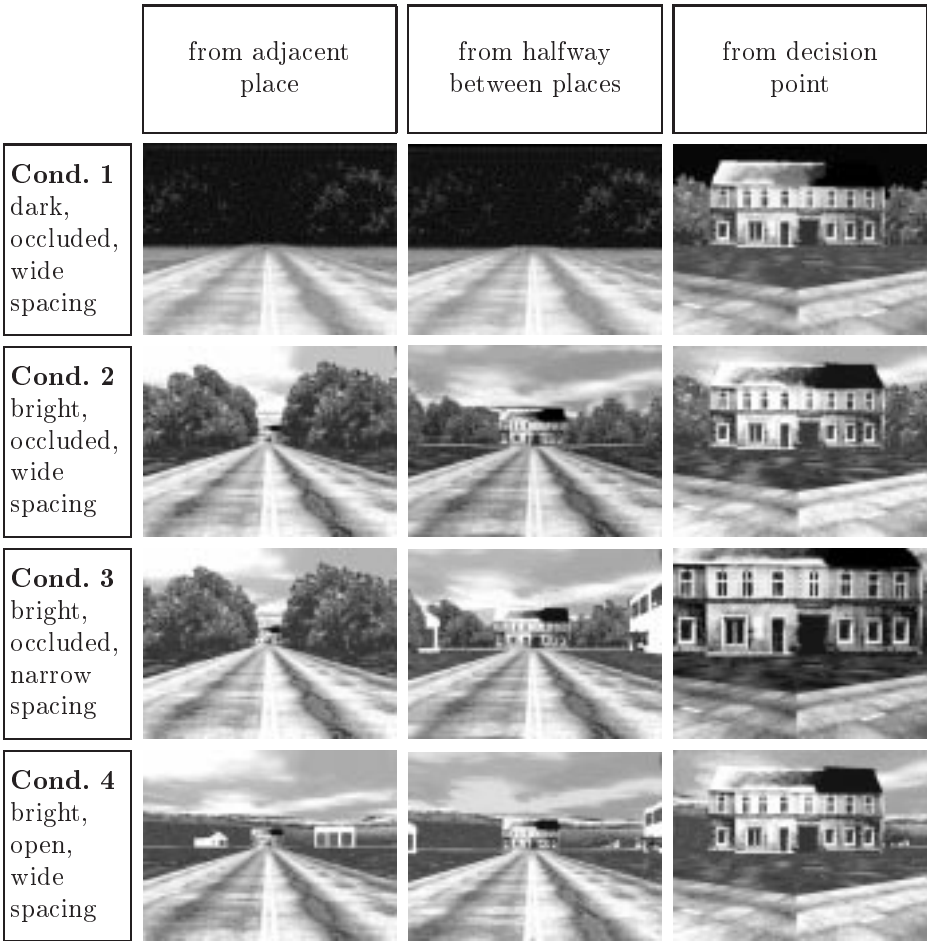


Fig. 4. Viewing conditions used in the experiments. For each condition, three views are shown. The views in the left column occur when looking from a place (No. 5) into a street. The views in the middle column can be seen during the motion along a street, in this example from place 5 to place 3. The views in the right column show an object (view 11) as seen from the corresponding junction (place 3).

buildings at one junction could never be seen together in these conditions. The sky was dark, as if navigating the town at night, and illumination was as with a torch or the headlights of a car reaching about 60 meters. Thus, the building at the far end of a street was not visible in this condition.

Condition 2 was the same as condition 1, except that illumination now came from the bright sky. While this changed the overall impression considerably, the only additional information provided in condition 2 was the view at the far end of the street, which was now visible from the adjacent place.




Fig. 5. View of a preliminary version of the Virtual Tübingen model. The buildings appearing in white are geometry models without textures being added. The other buildings show complete geometry and texture mapping.

In condition 3, the buildings were placed closer to the junction point, at a distance of 15 meters. This had the effect that all three buildings of a place could now be seen at a glance when entering the circular hedge. Condition 3 thus provided place information, rather than mere local views.

Condition 4 was the same as condition 2, but now the hedges were removed. Therefore, subjects could not only see all buildings from the place being currently approached, but also more distant places. In this condition, landmark information is no longer locally restricted.

2.2 Virtual Tübingen

An artificial virtual environment can be designed to fit the requirements of the experiments that will be conducted in it. In the case of Hexatown, for instance, the hexagonal grid of roadways was explicitly constructed to force movement decisions at every junction. If, however, one would like to verify the applicability of experimental results obtained with virtual environments like Hexatown for navigation in the real world, one should conduct a comparative study of navigation behavior in real and virtual environments. This notion has motivated us to create a virtual copy of an existing environment: Virtual Tübingen.

We have started to construct a virtual model of the historic centre of Tübingen, a 600 by 400 m area densely packed with approximately 600 buildings. The richness and complexity of the buildings and street network in this part of the city provide a very interesting environment for navigation studies. The shapes and

façades of the buildings are all different, and streets often tend to be curved and have varying width. They form an irregular network reflecting the topographical situation of Tübingen on a ridge rising between two valleys. Correspondingly, altitudes vary considerably and the resulting changes in slope add to the complexity of the navigation cues. Clearly, the construction of such a model is a long-term project and the current state of the model is far from being complete. Figure 5 shows an aerial view of an early version of Virtual Tübingen, extending from the city hall and the market place (on the left) to the north side of the “Holzmarkt” square. More information about the construction of this model can be found in Van Veen et al. (in press). Once this model has been completed, experiments on spatial behavior will be conducted both in real Tübingen and in its virtual counterpart. If the results obtained in both environments are consistent with each other, further experiments can be performed in the virtual environment taking advantage of the advanced features (such as the ability to manipulate the structure and appearance of the city).

3 Acquisition of spatial knowledge in Hexatown

3.1 Methods

Setup. Experiments were performed using a standard 19-inch SGI monitor. Subjects were seated comfortably in front of the screen and no chin-rest was used. They moved their heads in a range of about 40 to 60 cm in front of the screen which results in a viewing angle of about 35 – 50 degrees.

Procedure. In the experiment, subjects found themselves facing some view v_1 . They were then presented with a target view v_2 printed out on a sheet of paper and asked to find this view in the virtual town (task $v_1 \rightarrow v_2$). When they found the view, feedback was given in the form of a little sign appearing on the screen. If they got lost in the maze, i.e., if they deviated from the shortest possible path by more than one segment, the trial was stopped and another sign announced the failure. The sequence was terminated when the shortest possible way, i.e., the way involving the minimal number of decisions (mouse clicks) was found. The whole exploration phase contained 12 such search tasks, or ways to be found. The first four ways were excursions from view no. 15, which served as a “home” position. View 15 showed a poster wall saying “Max-Planck-Institut für biologische Kybernetik”. The following 8 searches were either returns to home or novel paths not touching on view 15. The return and novel path tasks were presented alternately in two sequence conditions: in the **returns-first condition**, the first task was a return, whereas in the **novel-first condition**, the sequence started with a novel path. In both conditions, the four excursions were performed prior to both returns and novel paths.

After the exploration phase described above, subjects were asked to rate distances in the maze and to produce a sketch map. The results from these parts of the experiments have been published elsewhere (Gillner & Mallot 1998).

Subjects. Eighty paid volunteers participated in the experiment, 40 of which were male and 40 female, aged 15 – 38. Twenty subjects (10 male, 10 female) took part in each of the four viewing conditions (Fig. 4). Within each viewing condition, the group of subjects was split equally (10 to 10) between the two sequence conditions (returns–first and novel–first), as well as the two instructions for the distance estimation (“distance” and “airline–distance”).

3.2 Exploration performance

Figure 6 shows an example for the cumulative trajectories taken by a single subject in the twelve search tasks. Paths 1 – 4 are excursions, 5, 7, 9, and 11 are returns and 6, 8, 10, and 12 are novel routes. Overall, there is a tendency for lower error rates in the search tasks performed later. That is to say, there is a transfer of knowledge obtained in earlier searches to the later searches. The decrease of errors is not monotonic and is missing in some subjects.

For a quantitative analysis, errors were defined locally as movement decisions that do not reduce the distance of the goal. Each movement decision equals clicking the mouse buttons twice (cf. Fig. 2b). Distance to the goal is measured as the minimum number of decisions needed to reach it (“decision–distance²”). Thus, if a subject enters a street leading away from the goal, the return from that street will be counted as a correct decision even though the current position is not part of the shortest path. This is due to the fact that the return from the false street is again an approach to the goal. In cases where the correct decision is a 60 degrees turn left followed by a “go”, the 120 degrees turn left would leave the decision–distance to the goal unchanged. This decision (and the mirror–symmetric case) is also counted as an error, in accordance with the above definition of local error. In cases where the subject has to perform a full 180 degrees turn, the initial 120 degrees turning step may be either left or right. Since both turns lead to a reduction in the number of steps remaining to the goal, either decision would be considered correct.

Average error rates for each path type are shown in Fig. 7. For each viewing condition (1 – 4; see Methods section), the excursions, returns, and novel paths were lumped into groups of four. As mentioned above, the excursions were performed first, while the novel and return paths were performed alternately, starting with a return in one group of subjects and starting with a novel route in a second group. The data show a learning effect in the sense that excursions take more errors than the later paths. They also show a clear effect of condition: higher visibility results in lower error rates. This general relation does not hold

² One could argue that false turns should not be counted as errors at all, since in a more realistic setup, they correspond to simple head turns rather than movements of the body. In fact, when accessing subjects’ sense of distance, we did find in additional measurements that turns are ignored in distance estimation (Gillner 1997, Gillner & Mallot 1998). For the specification of the task, i.e., finding the shortest possible route to a goal, we think, however, that erroneous turns should be considered. Turns do take time, after all.




Fig. 6. Travelling frequencies for each view-transition for the twelve paths (subject GPK, viewing condition: 1; sequence condition: returns-first). Left column: excursions, middle column: returns, right column: novel paths. The overall number of errors decreases at later stages of exploration.

for the comparison of conditions 1 and 2, however, which differ in the visibility of the neighboring places.

A 3-way analysis of variance (ANOVA, 4 conditions \times 3 path types \times 2 genders) of error-rate as the dependent variable reveals significant effects of condition ($F(3, 72) = 17.31, p < 10^{-4}$) and path type ($F(2, 144) = 60.65, p < 10^{-4}$),




Fig. 7. Total number of wrong movements performed in the different path types (excursion, novel, return). Numbers are averaged over 20 subjects, error bars are standard deviations. 1–4: viewing conditions.

but not of gender ($F(1, 72) = 0.22$, n.s.). Additionally we found an interaction of condition and path type ($F(6, 144) = 2.66$, $p = 0.018$). The error rates of novel paths are slightly higher than those of the returns (see Fig. 7). This effect is not significant, however.

3.3 Knowledge Transfer across Routes

In our procedure, learning occurs on two time scales. During each of the twelve tasks, a route is learned. When switching from one route to the next, part of the knowledge acquired in the earlier routes might be transferred to the new ones. To test this, we define a transfer coefficient τ in the following way:

Let R and N denote two routes, for instance the first return and novel path, respectively. Our group of subjects is divided into two subgroups, one of which explores R first and N second, whereas the second group explores N , then R . Four such pairs of routes have been tested. We accumulate the data from these four tested pairs of returns and novel paths:

- $E_{R,1}$ errors in returns in the returns–first condition
- $E_{N,1}$ errors in novel paths in the novel–first condition
- $E_{R,2}$ errors in returns in the novel–first condition
- $E_{N,2}$ errors in novel paths in the returns–first condition

Thus, $E_{R,1}$ and $E_{N,2}$ refer to the first group of subjects (returns first condition) and $E_{N,1}$, $E_{R,2}$ to the second. If transfer occurs, the route explored first should have higher error rates in both cases. We define:

$$\tau = \frac{E_{R,1} - E_{R,2} + E_{N,1} - E_{N,2}}{E_{R,1} + E_{N,1}} \quad (3)$$

If error rates do not depend on position, τ will be zero; if everything is learned already when exploring the first route, $E_{R,2}$ and $E_{N,2}$ will be zero and τ evaluates to 1.




Fig. 8. Average error numbers for novel paths and returns in the novel-first and returns-first sequence conditions. All subjects from viewing conditions 1 and 2 with an overall error number below the median were included in this plot. ● novel routes; ○ returns. ΔE_R : disadvantage of returns-first group in the return routes. ΔE_N : disadvantage of novel-first group in the novel routes. For both returns and novel routes, error rate drops when other routes are explored before. The transfer-coefficient (Equation 3) is $\tau = 0.4$.

For this evaluation, the subjects from viewing conditions 1 and 2 were pooled, because there were no significant differences between the respective error rates (3 way ANOVA: 2 conditions \times 4 routes \times 2 gender, $F(1, 36) = 0.014, p = 0.9075$). If we take the average over all 40 subjects, no significant effect of transfer is found. If, however, only the 20 subjects with the lowest overall error rate are considered, a transfer effect with $\tau = 0.4$ is found (see Fig. 8). In this case, eleven subjects were from the returns-first condition and nine subjects from the novel-first condition. The result indicates that the good navigators show significant transfer of knowledge even from one route to the next. Transfer across more steps of the exploration procedure is not visible in this evaluation, which does not mean that we exclude such a transfer.

3.4 Persistence

An inspection of Fig. 6 shows that in almost all cases where the subject started from view 15, the first movement decision was LL even though RR would have

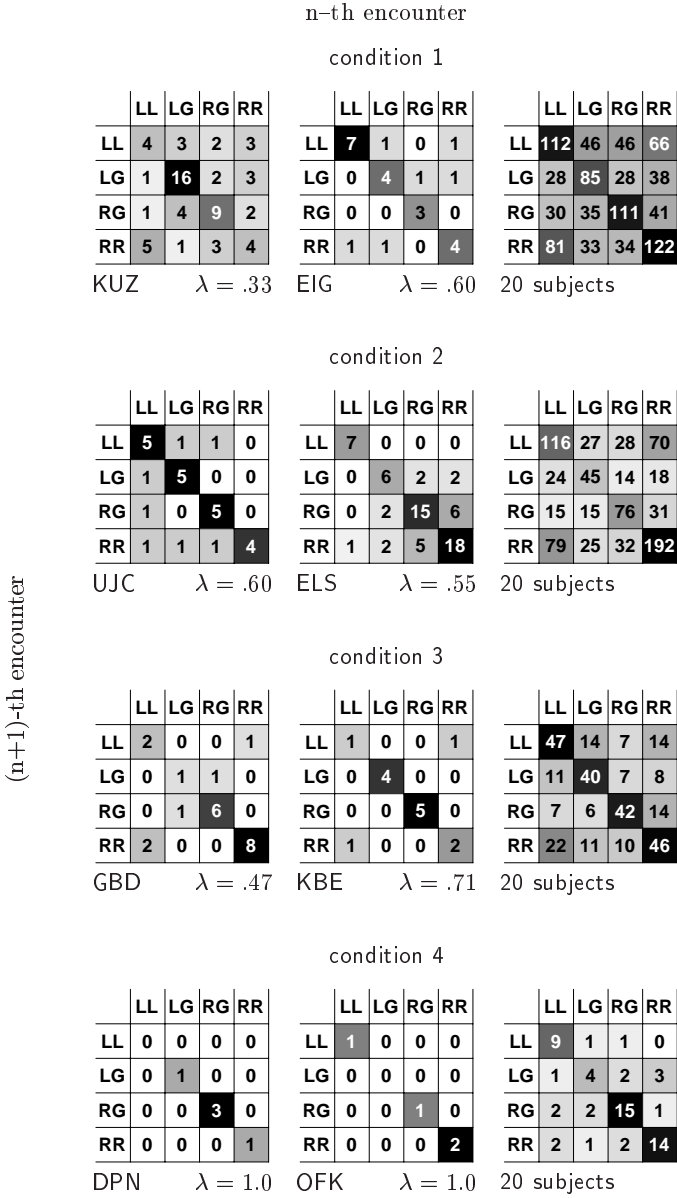


Fig. 9. Examples of return statistics for selected subjects for the four viewing conditions. In the subjects shown for the first three conditions, statistical independence of the decisions taken at the n -th and the $n - 1$ -th encounter could be rejected in all cases, i.e., persistence rate λ was significantly different from zero. For condition 4, where errors were generally quite rare, statistical independence could not be rejected for any subject; still, the matrices show high diagonal entries.

been just as good. Together with similar behaviors observed from other subjects, this lead to the conjecture that at least some movement decisions reflect simple, fixed associations between the current view and some motion that is performed whenever the view occurs. In order to test this in more detail, we analysed the return statistics of the decision sequences.

Let $m_{n,v} \in \{LL, LG, LR, RL, RG, RR\}$ denote the movement decision taken at the n -th encounter of view v (see Fig. 2b for possible movement decisions). We are interested in cases where the movement chosen at the n -th encounter of view v is the same as that taken at the $n - 1$ -th encounter of the same view.. More generally, we count the cases where movement j is taken at encounter $n - 1$ and movement i at encounter n ($i, j \in \{LL, LG, LR, RL, RG, RR\}$):

$$F_{i,j} := \#\{(n, v) | m_{n,v} = i, m_{n-1,v} = j\}. \quad (4)$$

It is important to note two points: First, the two encounters n and $n - 1$ do not occur in subsequent time steps (unless $m_{n-1,v} \in \{RL, LR\}$). Rather, long sequences of other views may occur in between. Second, the frequency $F_{i,j}$ is accumulated over all views. Thus we are looking for an average persistence rate rather than for a view-specific one.

In the experiments, each search task is repeated until the subject finds the shortest possible path. This procedure can in itself produce repetition rates above chance if parts of the path are repeated correctly several times. To exclude this type of error, we restrict our analysis to repetitions where both decisions were false in the sense that they did not lead to an approach to the goal (local definition of errors). Finally, we dropped the cases involving the decisions LR and RL, since these are quite rare.

Example data from individual subjects are shown in Fig. 9. The numbers on the diagonal correspond to cases where the same decision was chosen in two subsequent encounters even though the decision was false in both cases. From these matrices, we can estimate average movement transition probabilities $p_{ij} := P(m_{n,\cdot} = i | m_{n-1,\cdot} = j)$; averaging is performed with respect to the different views involved. A simple statistical model for these transition probabilities is:

$$p_{ij} = \begin{cases} \lambda + p_i & \text{if } i = j \\ p_i & \text{if } i \neq j \end{cases} \quad (5)$$

where $\lambda, 0 \leq \lambda \leq 1$ and $\lambda + \sum_{i=1}^4 p_i = 1$. It states that there is a bias λ for the repetition of the movement chosen at the previous encounter. Other than that, the decisions at subsequent encounters are independent. If $\lambda = 0$, true independence is obtained.

The analysis could be applied to data from 67 out of 80 subjects. For the remaining 13 subjects, the number of total errors was too low to fit the model. Ten of these had been tested in viewing condition 4, where the error rates were lowest. Goodness-of-fit was tested with the χ^2 -Test; choosing a significance level of 5 %, the best fitting model could not be rejected in any of the 67 subjects.

In order to get an impression of the confidence intervals for λ , we repeated the analysis with fixed $\lambda = 0$ in Eq. 5. By testing goodness-of-fit for this model

with the χ^2 -Test, 18 cases could be rejected on the 10 %-level, 9 of which could be rejected on the 1 %-level as well.

Average persistence rate over all subjects was 0.33, indicating that about one third of the erroneous movement decisions were based on persistence. A regression analysis of persistence rate λ with the overall number of errors for each subject did not reveal a significant correlation.

4 Discussion

4.1 Navigation in virtual environments

The results show that spatial relations can be learned from exploration in a virtual environment even under rather restricted viewing conditions. Here, we briefly summarize the most important findings:

Effect of viewing condition. The four viewing conditions differ in the amount of information available to the subjects. Not surprisingly, the number of errors during the search phase decreases as more information is provided. This is in spite of the fact that the field of view was the same in all four conditions. Bringing the objects closer to the places in condition 3 and removing the occluding hedges in condition 4 is reminiscent of zooming out the whole scene with a wide-field lens. Perúch et al. (1997) showed that this zooming does not improve path-integration performance in a triangle-completion task. This discrepancy may characterize a difference between path integration and landmark navigation. Alternatively, his negative result may be due to the marked errors in perspective associated with zooming.

The comparison between conditions 1 and 2 (night and day) does not show an improvement in error rates. This is surprising since more information is available in condition 2 (objects at the far end of the streets become visible). This finding may be related to the fact that the local structure of the maze becomes more complicated in condition 2, where six objects are visible from each place.

Transfer and latent learning. The overall number of errors was smaller for the later search tasks. For the 50 % best subjects, this effect was already clearly visible for the comparison of one search task with the next (Fig. 8). If subjects simply learned a set of independent routes, e.g., by reinforcement learning, each search would be a new task and no such transfer would be expected. The knowledge being transferred from one route to the next is not just a route-memory but involves the recombination of route segments; this is to say, it is of the configuration type. Its acquisition is akin to latent learning, since knowledge obtained during one search can be employed later in other, unrelated search tasks.

As can be seen from Fig. 8, transfer was strong from the novel to the return paths, but not the other way around. One possible explanation of this finding is that the novel paths are more difficult than the returns. When considering the shortest possible paths, the novel paths involve 14 different views, 8 of which also occur in the returns. The returns involve only 9 different views, i.e., almost

all of their views are already known from the novel paths. The only view not occurring in the novel paths is the final goal of the returns, view 15. The transfer asymmetry may thus be due to the fact that the novel routes contained more information about the returns than vice versa.

The occurrence of transfer from one route to another is also evidence for the presence of goal-independent memory of space, i.e., a cognitive map.

Persistence. Along with these arguments for configuration knowledge, evidence for simpler types of spatial learning was also found. The persistence rates presented in Section 3.4 indicate that at least some of the subjects based a considerable part of their movement decisions on simple associations of views with movements. This strategy is efficient for learning non-intersecting routes but will lead to errors for views at a crossroads where the correct motion decision depends on the current goal. We speculate that the persistence rate will decrease if longer training sequences are used.

Subject differences. Subjects differed strongly in terms of the number of errors made when searching a goal as well as in the amount of transfer learning. However, no clustering in different groups can be obtained from our data. In particular, no significant gender differences were found.

4.2 View-based navigation

In viewing conditions 1 and 2, subjects had to rely on local views as their only position information. Their performance and the transfer learning is therefore view-based in an obvious sense. However, this result does not exclude the possibility that some more complicated representation of space is constructed from the local view information. Here we summarize the evidence against such a representation, i.e., evidence for a view-based mechanism of navigation.

Returns aren't easy. After having learned the four excursions, the returns to the starting point along the very same paths are almost as difficult as novel paths (Fig. 7). The advantage on the order of just one error per search task, is not significant ($F(1, 76) = 2.860, p = 0.095$). If the subjects acquired a place-based representation of space, it would be the same for excursions and returns, since the according place-graphs are symmetric (see Fig. 1). In this case, we would therefore expect that returns should be much easier and more reliable than novel paths. The weak difference between the number of errors occurring in returns and novel paths seems to indicate that this is not the case. It is rather more in line with a view-based mechanism, since the views occurring along the return path are as different from the original views as any other views in the maze.

Recognition and action. The average persistence rate of 33% indicates that direct associations of views to movement decisions can be learned. As was pointed out in the introduction, the association pair of view and motion decision is the basic element of a view-based memory of space.

4.3 Local information combined to a graph?

If the representation is in fact view-based, a graph structure is the only representation we can think of that would account for the transfer and planning behavior observed. Independent evidence for a graph-like representation comes mainly from the sketch maps recorded by Gillner & Mallot (1998). Maps are often locally correct but globally inconsistent. Also, places with correct local connectivity have been translocated to erroneous positions. Connectivity can be correct even though metric properties of the sketch maps, such as angles and lengths are grossly mistaken.

The distance estimates do not reflect the decision-distance, which is the graph-distance of the view-graph, but correlate better with walking distance, i.e. the graph distance of the place graph (Gillner & Mallot, 1998). It therefore appears that we cannot decide between the view-graph and the place-graph representations at this point. In ongoing work (Mallot & Gillner, 1997), this question is addressed with additional experiments.

5 Conclusion

In our view, the most important result of this study is the fact that configuration knowledge can be acquired in virtual environments. This is in spite of the fact that the subjects did not actually move, but were interacting with a computer graphics simulation. With respect to the high controllability of visual input, this result may well make virtual reality a valuable addition to more realistic field studies, where stimulus control is often a problem. More realistic simulations such as Virtual Tübingen and large displays are currently being developed and tested in our laboratory.

With respect to our starting point, i.e. view-based navigation, we think that three conclusions can be drawn:

1. *Views suffice.* Map learning is possible if only local, i.e., view-information is provided. In this sense, navigation can be view-based.
2. *Graph vs. view from above.* The representation contains local elements, i.e., a place or view with one or several movement decisions and the respective outcome associated with it. These local elements need not be globally consistent, they need not combine into a metric survey map. Rather, a graph-like representation is sufficient to account for our results.
3. *Places vs. views.* It is not clear from our data, whether the nodes of this graph are places or views. We have not found evidence that the local views are combined into a representation of space independent of the orientation of the viewer. Thus, a view-based representation seems more likely at this point.

References

1. V. Aginsky, C. Harris, R. Rensink, and J. Beusmans. Two strategies for learning a route in a driving simulator. Technical Report CBR TR 96-6, Cambridge Basic Research, 4 Cambridge Center, Cambridge, Massachusetts 02142 U.S.A., 1996.
2. M. Billinghurst and S. J. Weghorst. The use of sketch maps to measure cognitive maps of virtual environments. In *Proc. IEEE 1995 Virtual Reality Annual International Symposium*, pages 40 – 47, Piscataway, NJ, 1995. IEEE Press.
3. B. A. Cartwright and T. S. Collett. How honey bees use landmarks to guide their return to a food source. *Nature*, 295:560 – 564, 1982.
4. K. Cheng. A purely geometric module in the rat's spatial representation. *Cognition*, 23:149 – 178, 1986.
5. T. S. Collett and J. Baron. Learnt sensori-motor mappings in honeybees: interpolation and its possible relevance to navigation. *Journal of Comparative Physiology A*, 177:287 – 298, 1995.
6. R. P. Darken and J. L. Sibert. Navigating large virtual spaces. *International Journal of Human-Computer Interaction*, 8(1):49-71, 1996.
7. M. O. Franz, B. Schölkopf, H. A. Mallot, and H. H. Bülthoff. Learning view graphs for robot navigation. *Autonomous Robots*, in press.
8. S. Gillner. *Untersuchungen zur bildbasierten Navigationsleistung in virtuellen Welten*. Knirsch-Verlag, Kirchentellinsfurt, 1997.
9. S. Gillner and H. A. Mallot. Navigation and acquisition of spatial knowledge in a virtual maze. *Journal of Cognitive Neuroscience*, in press.
10. B. J. Kuipers and Y.-T. Byun. A robust, qualitative approach to a spatial learning mobile robot. In *SPIE Vol. 1003 Sensor Fusion: Spatial Reasoning and Scene Interpretation*. International Society for Optical Engineering (SPIE), 1988.
11. H. A. Mallot, M. Franz, B. Schölkopf, and H. H. Bülthoff. The view-graph approach to visual navigation and spatial memory. In *Artificial Neural Networks — ICANN 97*, 1997.
12. H. A. Mallot and S. Gillner. Psychophysical support for a view-based strategy in navigation. *Investigative Ophthalmology and Visual Science*, 38(Suppl.):4683, 1997. ARVO-Abstract No. 4683.
13. R. Maurer and V. Séguinot. What is modelling for? A critical review of the models of path integration. *Journal of theoretical Biology*, 175:457 – 475, 1995.
14. M. May, P. Péruch, and A. Savoyant. Navigating in a virtual environment with map-acquired knowledge: Encoding and alignment effects. *Ecological Psychology*, 7(1):21-36, 1995.
15. R. U. Muller, M. Stead, and J. Pach. The hippocampus as a cognitive graph. *Journal of General Physiology*, 107:663 – 694, 1996.
16. J. O'Keefe and L. Nadel. *The hippocampus as a cognitive map*. Clarendon, Oxford, England, 1978.
17. P. Péruch, M. May, and F. Wartenberg. Homing in virtual environments: Effects of field of view and path layout. *Perception*, 26:301 – 311, 1997.
18. T. Prescott. Spatial representation for navigation in animals. *Adaptive Behavior*, 4:85 – 123, 1996.
19. R. A. Ruddle, S. J. Payne, and D. M. Jones. Navigating buildings in “desk-top” virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied*, 3:143 – 159, 1997.
20. B. Schölkopf and H. A. Mallot. View-based cognitive mapping and path planning. *Adaptive Behavior*, 3:311 – 348, 1995.

21. M. Tlauka and P. N. Wilson. Orientation-free representations from navigation through a computer-simulated environment. *Environment and Behavior*, 28:647 – 664, 1996.
22. F. H. Tong, S. G. Marlin, and B. J. Frost. Visual-motor integration and spatial representation in a visual virtual environment. *Investigative Ophthalmology and Visual Science*, 36:1679, 1995. ARVO abstract.
23. D. S. Touretzky and A. D. Redish. Theory of rodent navigation based on interacting representations of space. *Hippocampus*, 6:247 – 270, 1996.
24. O. Trullier, S. I. Wiener, A. Berthoz, and J.-A. Meyer. Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51:483 – 544, 1997.
25. H. A. H. C. van Veen, H. K. Distler, S. J. Braun, and H. H. Bühlhoff. Using virtual reality technology to study human action and perception. *Future Generation Computer Systems*, in press. See also: <http://www.kyb.tuebingen.mpg.de/> (Technical Report 57).