Search Strategies and their Success in a Virtual Maze

Simon J. Buechner (buechner@cognition.uni-freiburg.de)
Christoph Hölscher (hoelsch@cognition.uni-freiburg.de)
Jan M. Wiener (mail@jan-wiener.net)

University of Freiburg
Center for Cognitive Science, Friedrichstraße 50
79098 Freiburg, Germany

Abstract

The study investigates strategies for search tasks in a virtual environment. Participants had to find and retrieve an object in an unfamiliar environment. Trajectories as well as search performance were analyzed with respect to three applicable strategies when the target was hidden at two different locations. Prior knowledge about the environment was varied within participants and between environments. The results show a strong preference for the perimeter strategy. Prior knowledge improved search performance only to some extent.

Keywords: spatial cognition, navigation, search strategy, virtual environments

Introduction

Imagine you enter an environment and you are searching for an object. What would be a good strategy to quickly find it? How successful would that strategy be depending on the exact location of the object? The selection of the strategy as well as your search performance will certainly depend on whether you have been to the environment before or not. Although the search for objects is a task we frequently encounter, surprisingly little is known about the strategies people apply when performing tasks like this.

Strategies

A variety of studies investigated the strategies humans apply when navigating between locations (e.g. Golledge, 1995; Hochmair & Frank, 2002; Conroy Dalton, 2003; Wiener, Schnee & Mallot, 2004; Büchner, Hölscher & Strube, 2007). In all these studies participants had some information about the location of the target place. Either they have been there before, they could permanently see it, or they could infer the location from other information such as background knowledge about the functional organization of buildings.

The studies listed above have been conducted in the context of *wayfinding* (in its literal meaning) - participants had to find a way from one location to another having some knowledge about the target location. In a *search* task as we define it here, in contrast, the agent has no information about the location of the target and thus, strategies may play an even more important role for the success and efficiency of the search.

Different scales of space may require different search strategies. When searching for an object in a room sized environment in which all potential target locations are visible, people may apply different strategies than in a building in which walls block visibility. Montello (1993) conceptualized space of different scales. He distinguished between vista spaces, environmental spaces, and geographical spaces. A *Vista* space is the space that is visible from a single observation point (e.g. a room). An *Environmental* space is a space in which an agent has to move in order to comprehend its entire structure (e.g., a maze or a building). *Geographical* spaces are even larger and include, for example, landscapes and other spaces that can only be experienced with some transportation device.

Search strategies have been subject to empirical studies in different scales of space. Ruddle, Payne and Jones (1999) asked participants to search for nine objects in a virtual seascape. Participants were 'flying' over the seascape experiencing the world from a bird's eye view. The authors found that in such a *geographical space* people tend to use an *anchor strategy*. They start searching at a visible object (in this case an island) and then move along the edge of that object? The authors also observed the *lawn mower strategy* (searching the area in regular back-and-forth pattern from one end to the other).

In *vista spaces*, search strategies have been investigated for visually impaired people and blindfolded participants (Tellevik, 1992; Hill & Rieser, 1993). Tellevik (1992) identified three major strategies: *grid-line*, *perimeter* and *reference-point.* The *grid-line* strategy is equivalent to the *lawn-mower* strategy, i.e. people followed a regular pattern from one wall to the other, slowly moving from one end of the room to the other. The *perimeter* strategy involves a path along the outside walls of the room and a subsequent exploration of the center. When following the *reference-point* strategy people select a reference-point (often the starting location), explore one part of the space, return to the reference-point, explore the next part of the space and so on. A similar strategy has also been observed for navigation in abstract spaces like the World Wide Web: Users of websites often navigate between 'hub' page and several subordinate pages (Cockburn & McKenzie, 2001).

Most of the aforementioned strategies, however, are only applicable in spaces in which the observer can oversee the whole space. This is different in *environmental spaces* in which walls and other objects limit visibility and impose constraints on the ways people move through the environment (e.g. the *grid-line strategy* wouldn't work in *environmental space*). One strategy that is, in principle, applicable is the *perimeter* strategy, another one is the *reference-point* strategy, that has indeed been observed in *environmental space* by Hölscher, Meilinger, Vrachliotis, Brösamle, and Knauff (2006). Participants' task was to find several locations in a complex building. A frequently chosen strategy was to search and return to the initial starting point.

In addition to these empirical studies Løvås (1998) presented a computational model for an evacuation situation in which an *always-turn-left* strategy (it works equally for *turn-right*) was a successful option to get out of a building. This strategy is functionally equivalent to the *perimeter* strategy identified by Tellevik (1992) as long as the initially selected wall is not an 'island' (a convex walling that is not connected to the perimeter). It is subject to empirical investigation whether people actually choose such a perimeter strategy when navigating through a maze (*environmental space*). Moreover, evacuation situations are a special case of search tasks in which people are stressed and might behave differently than when searching for an object without time pressure. It therefore remains an open question whether they choose the perimeter strategy for *search* tasks aside from evacuation situations.

Prior Knowledge

Prior knowledge of an environment is important for strategy selection. Hölscher, Meilinger, Vrachliotis, Brösamle, & Knauff (2006) found that people who were unfamiliar with the environment preferred the central point strategy (which is similar to the reference-point strategy described above) over hierarchical strategies that were preferred by people who were familiar with the environment (first going to the target region and then looking for the goal). Thus, it is important to distinguish between *uninformed* and *informed search* (cf. Ruddle, Payne & Jones, 1999). In *uninformed* search participants are naïve with respect to the search space (i.e. they have no knowledge about the environment). In contrast, in *informed* search they have at least some knowledge about the search space. It is not surprising that prior knowledge also increases search efficiency (e.g. Gärling, Lindberg & Mäntylä, 1983; Hölscher, Büchner, Brösamle & Strube, 2007). For an extended taxonomy of wayfinding that distinguishes tasks by different levels of knowledge that is available to the navigator, see Wiener, Büchner & Hölscher (2009).

Research Questions & Hypotheses

In the current study we investigated the effect of prior knowledge on search performance as well as the effect on the dynamics of the learning process. In general, the availability of prior knowledge should accelerate the learning process if asked to repeatedly search for the same target location as participants are then able to directly access knowledge about the environment (e.g. which path choice options are dead-ends). In addition a lack of prior knowledge may require to allocate more cognitive resources to the process of spatial updating (keeping track of the relative location of surrounding objects) in order to remain oriented (cf. Yardley & Higgins, 1998). This could lead to a decreased search performance.

Which strategy do people choose when searching for an object in a maze? We identified three strategies that are applicable in environmental spaces:

1. <u>Directed Random Search</u>. People just start walking. At choice points they randomly decide between path alternatives with a bias to avoid going back. While this is certainly not the most efficient strategy to find a specific location in the environment, it constitutes a conceivable option. The trajectory should show no regular pattern.
2. <u>Reference-point.</u> People try to remain oriented in order to systematically explore the environment. Thus, they select a reference point from which they start short explorations in different directions and to which they return regularly. This should result in a radiating trajectory around one or more reference points.
3. <u>Perimeter strategy.</u> People follow the perimeter of the maze. One advantage of this strategy is that it guarantees that one does not explore one branch of the maze twice. In addition, one easily finds back to the starting point simply by turning around and following the perimeter in the other direction.

We cannot claim that this is a complete list of applicable strategies. The reference-point strategy and the perimeter strategy have been observed before and are described in the literature, but in different contexts. It is reasonable to assume that these strategies are most likely selected in an uninformed search task (i.e. a search in an unfamiliar environment).

How successful are the strategies for different target locations? When applying *directed random* search, target locations close to the starting point should, on average, be found faster and with less detours than locations further away from the starting point. Predictions for search performance for the *reference point* strategy are difficult, because the success of the strategy depends on two factors: (1) the selection of the reference point, (2) the order in which multiple options are explored. However, if participants choose the starting point as the reference they should also find a location close to the starting point faster than further away. The *perimeter* strategy should lead to good search performance for targets located along the perimeter of the environment and to worse performance for targets located in the center of the environment. Although directed random search and reference point should be distinguishable on the level of the trajectory, they are not necessarily distinguishable with respect to search performance/efficiency. The perimeter strategy should lead to different trajectories and different search performance.

How do the learning curves of a search-and-recover task (i.e., find a target return to the start location and find it

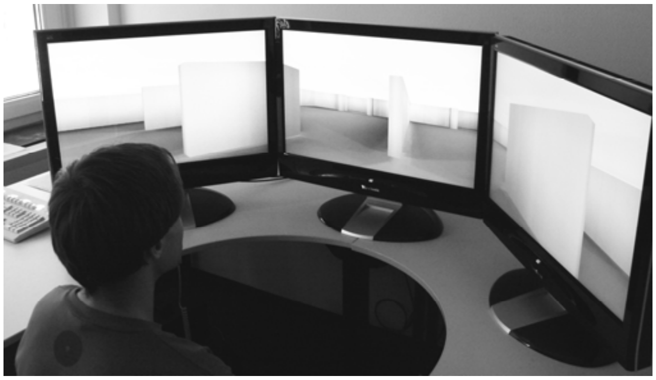

Figure 1: Setup of the immersive virtual environment with three screens. The horizontal field of view of about 170° allows peripheral vision.

again) differ when the target location is compatible or incompatible with the applied strategy? Compatible target locations should be found quicker with fewer detours which should also affect the learning curves. People should learn faster if they can simply apply the initially chosen strategy again.

When the target location is incompatible with the strategy, participants have to deviate from the initially chosen strategy and select a different strategy. The additional processing load should lead to shallower learning curves.

Method

Participants

Thirty-eight participants (16 of them female) finished the experiment without suffering from motion sickness. Participants were between the ages of 19 and 30 years (*M*= 22.8, *SD*= 2.7). They were recruited through postings on campus and e-mailing lists. Most of them were students from a variety of subjects. Participants were paid (7.50 Euro/h) or received course credit for participation.

Setting

Participants had to perform a search-and-recover task in each of two mazes (details see below). The order of presentation of the mazes was balanced between participants. The tasks were presented in an immersive virtual environment on three 28'' screens at 60Hz and a total resolution of 768×4080 px. The screens were set up in a semi-circle, participants were seated about 50cm in front of the center screen, providing a horizontal field of view of about 170° and a vertical field of view of about 60° (Figure 1). The Vizard 3.0 Virtual Environment software was used for the presentation of the virtual mazes. Participants used a Logitech Rumble Pad controller with two joysticks in order to steer around the virtual environment.

Two different virtual mazes were constructed in order to ensure that the results are not tied to the spatial structure of one particular maze. Both mazes were constructed in a way that people encounter many two-option situations when they move through the maze. The mazes were deliberately left

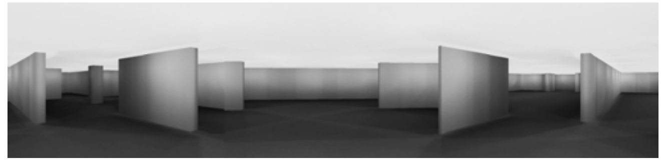

Figure 2: Screenshot of one of the virtual mazes. The mazes were presented on a setup with three screens. The distortions in this picture were not present in the experiment.

empty with no objects in order to avoid that people use objects as landmarks or orientation aids. In order to remain oriented, participants therefore had to refer to geometrical features of the environment (Fig. 2 shows a screenshot of the environment, Fig. 3 the layouts and target locations).

Experimental Design and Procedure

Before performing the experimental tasks participants were instructed how to use the controller and had to perform a practice run. They had to follow a red line in a virtual training maze without running into a wall. All participants managed to complete the task easily.

Participants then were instructed in written form about the experimental tasks. They had to perform a search-and-recover task, i.e. they had to find a red ball hidden in the environment, without being given any information about its location. Once they had found it, they had to estimate the direction to the starting point by turning in the assumed direction and pressing a button. These pointing data exceed the scope of this paper and will be presented elsewhere. Participants were then put back to the starting point and the ball was placed at the same location as before (participants were informed about this). Participants were then again asked to search for the ball. This procedure was repeated until the participant found the ball three times without error.

There were two mazes and two different ball locations (center, edge) resulting in four possible combinations of maze and location. Each participant had to perform the task in two of these combinations. The combinations of maze and location as well as the order of the tasks were balanced across all participants.

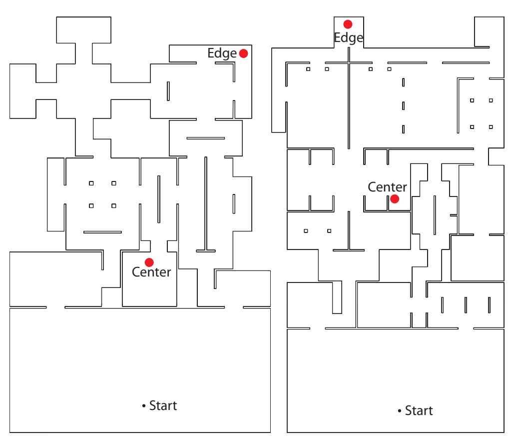

Figure 3: Layouts of the two virtual mazes with the locations of the target objects.

Each participant searched both mazes, maze 1 and maze 2. Prior knowledge was varied between mazes: For one maze participants explored the environment for eight minutes before they had to perform the search task, for the other maze they started with the search task right away. Half of the participants started with maze 1, the other half with maze 2. Within each of these two groups, 50% first performed a task version with prior exploration, and did their second maze without exploration. This was reversed for the other 50% in each group.

The whole experiment lasted no longer than one hour.

Dependent Variables

Since the target's distance from the starting point varied between conditions, the time participants required to find the target was not a feasible measure to judge search performance. Instead we calculated a standardized distance measure, percentage above optimal (PAO). It is calculated by comparing the walked distance with the shortest possible distance as follows:

$$PAO = ((d_{walked} - d_{shortest}) / d_{shortest}) \times 100$$

It allows comparing distances between conditions and mazes as it standardizes the walked distance by the shortest possible distance. The measure provides the relative amount of distance (compared to the shortest possible distance) that people walk in addition to the shortest distance. PAO has also been used by other authors (e.g. Ruddle, Payne & Jones, 1997; Wiener, Schnee & Mallot, 2004) in order to compare performance for search tasks with different lengths. The number and length of stops were recorded as well.

Results

Trajectories

Visual inspection of the trajectories and informal comments by the participants suggest that participants did not perform a random search. We also found little evidence that they applied a reference point strategy. The majority of the trajectories were shaped as those depicted in Figure 4. Participants tended to follow the perimeter of the maze. When they did this they did not necessarily *walk* exactly along the wall, instead they entered rooms shortly and when they saw that there was no path choice option they moved on. In general, participants mostly walked along the perimeter and only if they didn't find the target by the time they returned to the starting point, they changed the strategy and also started searching for the target in the center.

Performance

For the analysis of search performance we considered the first five trials of each participant as the majority of the participants had found the optimal path at this time.

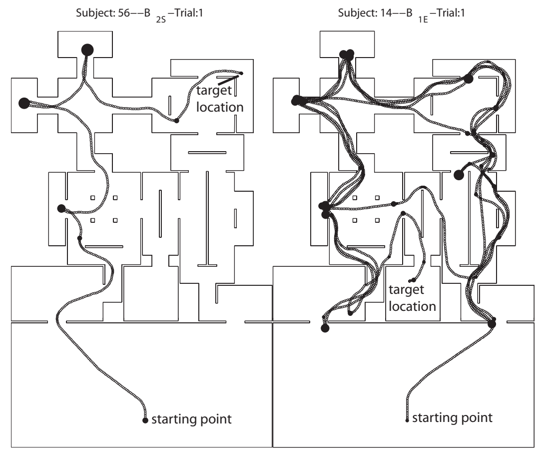

Figure 4: *Left:* Sample trajectory of one participant who almost pristinely applied the *perimeter* strategy and finds the target at the edge location quickly. *Right:* A different participant in the same maze also follows the perimeter strategy. He even searches at locations where he had already been before, before he finds the target at the center location. Circles display stops; the size of the circles is proportional to the stop time.

We didn't expect any performance differences between the two mazes. An ANOVA with factor maze yielded no significant difference in PAO between the two mazes [$F(1,388)= .007$, $p= .931$; $M_{maze1}= 127$, $SD_{maze1}= 275$, $M_{maze2}= 125$, $SD_{maze2}= 276$]. We can conclude that any observed effects are not tied to the particular spatial structure of one maze. Thus, the following analyses are based on the combined data from both mazes.

If participants choose the perimeter strategy they should perform better for the edge location, if they choose one of the other two strategies they should perform better for the center location. An ANOVA for PAO with factors target location (center, edge) and prior knowledge (prior exploration, no exploration) yielded a main effect of target location [$F(1,386)= 46.633$, $p < .001$] and an interaction of both factors [$F(1,386)= 4.020$, $p= .046$]. There was no main effect of prior knowledge [$F(1,386)= 2.272$, $p= .133$]. Participants deviated less from the optimal path when the target object was located at the edge of the maze ($M= 40$, $SD= 78$) than when the object was located in the center ($M= 212$, $SD= 362$). For separate means of each level of the factor refer to Figure 5. The results provide additional evidence for the perimeter strategy, but not for any of the other strategies. In addition they show that prior knowledge has a different effect on search performance depending on the exact location of the target.

Based on the observed main effect for target location the following analyses were conducted separately for each target location. In order to investigate the observed interaction of target location and prior knowledge in detail an ANOVA for PAO with factors prior knowledge (with exploration, without) and trial (one to five) was conducted.

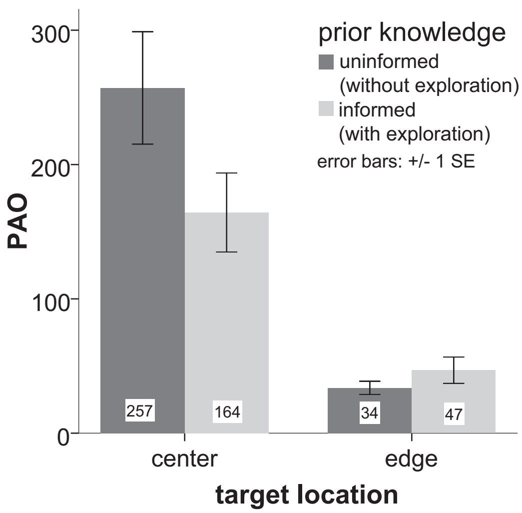

Figure 5: Main effect for target location as well as an interaction of target location and prior knowledge for PAO. Numbers in the white boxes on the bars denote means.

For the *center location* there were main effects for trial [$F(4,185) = 19.401$, $p < .001$] and prior knowledge [$F(1,185) = 4.447$, $p = .036$], though no interaction. For the *edge location* there was only a main effect for trial [$F(4,185) = 5.270$, $p < .001$], but no main effect for prior knowledge and no interaction (cf. Figure 6).

Participants' search performance increased from trial to trial and they approached almost optimal search performance in the fifth trial. They became more and more familiar with the environment and made fewer errors when walking the path from the starting point to the target location. Note, that across all five trials prior exploration only had an effect on search performance if the target was located on the edge.

If only trial one is considered (in which participants have no knowledge about the target's exact location), there is no effect of prior knowledge for both, the center location [$F(1,185) = .120$, $p = .729$] and the edge location [$F(1,185) = 1.206$, $p = .274$]. Figure 6 shows that in trial two, however, the learning curves split for the center location but not for the edge location.

A direct comparison of the two target locations for trial one shows that the PAO value for the edge location is less than 1/6 ($M = 84$, $SD = 56$) of the value for the center location ($M = 576$, $SD = 357$). This difference is highly significant [$t(39.896) = 8.508$, $p < .001$]. The large advantage of the edge location over the center location can only be explained by the application of the perimeter strategy.

Discussion

The main goal of this study was to investigate which of three strategies people choose when they search for a target in an unfamiliar environment. People clearly did not perform a random search. If this was the case, they would have found the target in the center location with fewer detours than when it was located at the edge location.

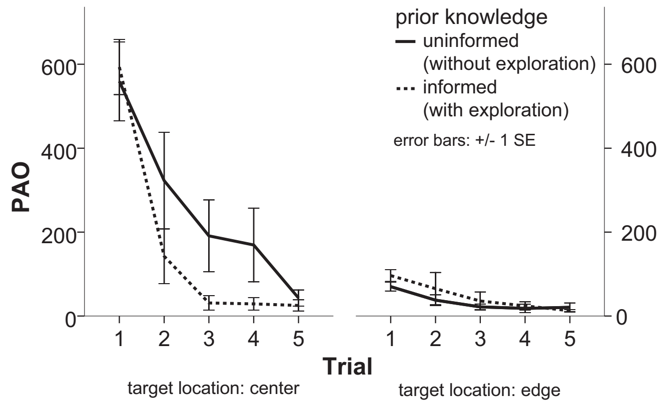

Figure 6: PAO decreases from trial to trial. Missing prior knowledge of the setting only affects search performance for the center location, not for the edge location.

Furthermore, no evidence in favor of a reference-point strategy was observed. The selection of such a strategy would have involved a radiant pattern of trajectories around one or more reference points. Trajectories, however, were mostly tied to the perimeter of the maze. The search performance provides additional evidence for the perimeter strategy. Participants walked fewer detours if the target object was located on the edge of the maze rather than in the center (location effect) although the center location was closer to the starting point than the edge location.

There was no effect of prior knowledge in the first search trial. Participants who had explored the maze before the search-and-recover task did not benefit from that exposure to the environment immediately. Together with the qualitative trajectory analysis, this result suggests that the choice of the ad-hoc strategy is independent of prior knowledge. However, they do benefit in trials two to five, at least when the target is located in the center. Their learning curve is steeper than in the no-prior knowledge condition.

Why does the effect of prior knowledge only show up for the center location? In general, PAO values are much lower for the edge location. This is, presumably, due to the selection of the perimeter strategy which supports finding the target if it is located at the edge. Prior knowledge cannot substantially improve search performance as it is already close to optimal, even in the first trial. If the target was located in the center, participants still selected the perimeter strategy. In the first trial, prior knowledge did not affect search performance, as participants had no information about the target's location. In the subsequent trials, however, prior knowledge facilitated learning and performance approached optimal search performance earlier than for the condition in which people had no prior knowledge. We take this for evidence that prior knowledge about an environment allows for easier disengagement from the initially selected perimeter strategy. They adapt quicker to the fact that this strategy is not successful for the center location. This is compatible with observations of adaptive use of path-choice strategies in real-world wayfinding (Hölscher et al., in press).

The lack of prior knowledge does not necessarily lead to worse performance. If the target is located on the edge, the selection of the strategy compensates for the lack of knowledge. When the target is compatible with strategy selection no additional knowledge is needed to quickly find the target.

Conclusion

We have found evidence that people do not choose their paths randomly when they search for an object in an unfamiliar environment. There is strong evidence from both, trajectories and search performance measures that people tend to use a perimeter strategy. We did not find evidence for a reference point strategy although the selection of this strategy has been observed in environments of the same scale but of different kind (Hölscher et al., 2006). We conclude that the selection of a strategy does not only depend on the knowledge people have about the environment, but also on the environment itself. It is subject to future research to identify features of the environment that determine the selection of a particular strategy.

The results also suggest that prior knowledge about the search space facilitates the disengagement from the initially selected strategy and that the lack of prior knowledge can be compensated by an adequate search strategy.

Future Work

We are currently working on a computational model that implements the perimeter strategy and compares it to a random walker. Preliminary results show that the random walker finds the center location quicker (and with less detours) than the edge location. By contrast, a simulation model that weighs path choice options along the perimeter finds the edge location faster than the center location. Future work will fine-tune the model and integrate the interaction of prior knowledge and the use of the perimeter strategy.

Acknowledgments

The research has been supported by the German Research Foundation (DFG) in the Transregional Collaborative Research Center “Spatial Cognition” (SFB/TR8). We thank Razvan Orendovici for programming work as well as Susanne Fischer for help with the data collection.

References

Büchner, S. J., Hölscher, C. & Strube, G. (2007). Path choice heuristics for navigation related to mental representations of a building. In: *Proceedings of the Euro CogSci 2007*, 504-509, Delphi, Greece.

Cockburn, A. &McKenzie, B. (2001). What Do Web Users Do? An Empirical Analysis of Web Use. *International Journal of Human-Computer Studies*, *54*(6):903--922.

Conroy Dalton, R. (2003). The secret is to follow your nose: Route path selection and angularity. Environment and Behavior, *35*(1):107–131.

Gärling, T., Lindberg, E., & Mäntylä, T. (1983). Orientation in buildings: Effects of familiarity, visual access, and orientation aids. *Journal of Applied Psychology*, *68*(1), 177-186.

Golledge, R. (1995). Path selection and route preference in human navigation: A progress report. In Frank, A. and Kuhn, W., (Eds.), Conference on Spatial Information Theory - LNCS, Vol. 988, Semmering, Austria.

Hill, E. W., & Rieser, J. J. (1993). How persons with visual impairments explore novel spaces: Strategies of good and poor performers. *Journal of Visual Impairment & Blindness*, *87*(8), 295-301.

Hölscher, C., Büchner, S. J., Meilinger, T., Strube, G. (in press). Adaptivity of Wayfinding Strategies in a Multi-Building Ensemble: the effects of spatial structure, task requirements and metric information. *Journal of Environmental Psychology*.

Hölscher, C., Büchner, S., Brösamle, M., Meilinger, T., & Strube, G. (2007). Signs and Maps – Cognitive Economy in the Use of External Aids for Indoor Navigation. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of CogSci 2007*, (377-382). Austin, TX: Cognitive Science Society.

Hölscher, C., Meilinger, T., Vrachliotis, G., Brösamle, M., & Knauff, M. (2006). Up the Down Staircase: Wayfinding Strategies and Multi-Level Buildings. *Journal of Environmental Psychology 26*(4), 284-299.

Løvås, G. G. (1998). Models of Wayfinding in Emergency Evacuations. *European Journal of Operational Research*, *105*(3), 371-389.

Montello, D. R. (1993). Scale and multiple psychologies of space. In A. U. Frank & I. Campari (Eds.), Spatial information theory: A theoretical basis for GIS (pp. 312-321). Proceedings of COSIT '93, Springer, LNCS 716.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1997). Navigating buildings in "desk-top" virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied*, *3*(2), 143-159.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied*, *5*(1), 54-75.

Tellevik, J. (1992). Influence of spatial exploration patterns of cognitive mapping by blindfolded sighted persons. *Journal of Visual Impairment & Blindness*, *92*, 221-224.

Wiener, J., Büchner, S., and Hölscher, C. (2009). Taxonomy of Human Wayfinding Tasks: A Knowledge-Based Approach. *Spatial Cognition & Computation*, *9*(2), 1–14.

Wiener, J. M., Schnee, A., and Mallot, H. A. (2004). Use and interaction of navigation strategies in regionalized environments. *Journal of Environmental Psychology*, 24(4):475–493.

Yardley, L., & Higgins, M. (1998). Spatial updating during rotation: The role of vestibular information and mental activity. *Journal of Vestibular Research: Equilibrium & Orientation*, *8*(6), 435-442.