Title
Feature-Based Navigation Strategy Selection in Complex Maze Environments

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors
xu, dongming
Hanna, Sean

Publication Date
2026

Feature-Based Navigation Strategy Selection in Complex Maze Environments

Dongming Xu and Sean Hanna
University College London
dongming.xu.23@alumni.ucl.ac.uk, s.hanna@ucl.ac.uk

Abstract

Human navigation in maze-like environments can be approximated by distinct heuristic strategies, yet which strategy best reflects human behavior appears to depend on environmental structure rather than global optimality alone. We compare three families of navigation models—exit-oriented (EXO), target-angle (Angular), and target-distance (Distance)—across 25 grid mazes and multiple start locations. Model behavior is evaluated using success rate, path efficiency, and human-likeness defined as the correlation between model and human state-visit profiles. Across mazes, Angular heuristics are more human-like than EXO in the majority of environments, while hybrid combinations offer limited additional benefit. As an exploratory diagnostic test of whether maze structure contains information about relative strategy alignment, we label each maze by the EXO–Angular human-likeness gap (excluding ambiguous ties) and evaluate a simple leave-one-maze-out classifier using maze-level features. Results show weak but non-trivial signal under balanced accuracy, highlighting both the possible relevance of maze features and the current limits imposed by the small, strongly imbalanced dataset. These findings support a feature-based view of navigation strategy selection, in which environmental structure shapes the apparent human-likeness of different heuristics. Code for all simulations and analyses is available at [https://github.com/9nx46r8rzg-web/grid-maze-heuristic-study.git].

Introduction

Understanding how humans navigate complex environments is a central problem in cognitive science, rooted in the foundational theory of cognitive maps (Maguire et al., 1998; O'Keefe & Nadel, 1978). In structured environments such as mazes, humans often rely on 'fast and frugal' heuristics (Gigerenzer & Gaissmaier, 2011) that exploit environmental regularities to bypass the prohibitive costs of global path-planning (Tolman, 1948). Such heuristics capture key regularities of human behavior while remaining computationally lightweight and cognitively plausible.

A persistent challenge, however, is that no single heuristic consistently matches human behavior across all environments. A strategy that appears highly human-like in one maze may perform poorly in another. This observation suggests that human navigation may not be best characterized by a single dominant strategy, but instead by context-sensitive strategy selection shaped by environmental structure.

Rather than asking which navigation strategy is globally optimal, the present work adopts a feature-based perspective on navigation modeling. From this view, the apparent human-likeness of a navigation heuristic is not an intrinsic property of the algorithm itself, but a consequence of how its structural sensitivities align with the environment in which it is deployed. A strategy may appear highly human-like in one class of environments while diverging substantially from human behavior in another, not because human cognition changes, but because different environmental cues are selectively amplified.

This perspective contrasts with approaches that seek to identify a single dominant navigation strategy or to optimize performance metrics such as path length or success rate. Instead, we focus on comparative alignment: given a particular maze structure, which heuristic best captures the spatial patterns observed in human navigation, and why? By framing the problem in this way, we treat navigation models not as competing accounts of human optimality, but as descriptive probes that reveal how environmental structure constrains human-like behavior.

From a resource-rational perspective, the frequent alignment of human behavior with Angular consistency over global goal-seeking (EXO) can be understood as an adaptive trade-off. Rather than continuously recomputing globally optimal paths, human navigation may prioritize maintaining a stable internal heading that reduces cognitive overhead during movement through complex environments. Importantly, this framing does not imply that humans explicitly select between discrete navigation strategies.

Instead, we propose that human navigation behavior reflects the interaction of stable cognitive biases—such as directional persistence and goal salience—with the structural affordances of the environment. Under this view, what appears as "strategy selection" is better characterized as the context-dependent expression of these biases, whose relative dominance is shaped by environmental structure rather than by deliberate algorithmic choice.

Accordingly, our goal is not to introduce a new navigation algorithm, but to characterize when existing simple heuristics approximate human navigation behavior and whether these conditions can be partially inferred from maze-level structural features. By adopting this structure-sensitive, bias-centered perspective, we align with broader accounts of human cognition as adaptive, context-dependent, and constrained by representational affordances of the environment.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

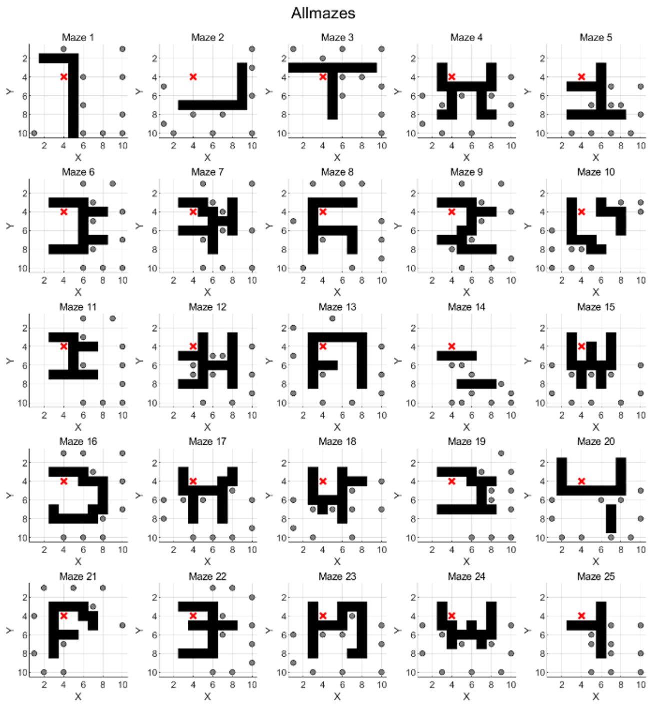

Figure 1: The layout of all maze obstacles, starting points, and endpoints.

Methods

Environment and Human Data

Human navigation behavior is analyzed within a set of 25 grid-based maze environments. Each maze consists of a bounded grid with obstacles forming corridors and junctions, and a single designated exit. The human trajectory data utilized in this work were obtained from the comparative benchmark study by De Cothi et al. (2021), which evaluated human navigation alongside that of rats and reinforcement learning (RL) agents. In these experiments, participants were tasked with navigating from various starting locations to the fixed exit.

For each maze, trajectories from all participants and starting positions were aggregated to form a human state-visit profile, representing the normalized frequency with which each grid cell was visited. This representation abstracts away from individual trajectories while preserving global spatial patterns of exploration and path usage.

Navigation Models

The analysis considers three families of heuristic navigation models.

Exit-Oriented (EXO). The EXO model selects actions that maximize progress toward the exit, using local directional information to favor movements that reduce distance or increase alignment with the exit location. This heuristic is inspired by agent-based visibility analysis (Turner & Penn, 2002), favoring movements that align with the perceptual flow toward the goal.

Angular (Angular). Angular models prioritize minimizing angular deviation between the agent's movement direction and the vector pointing toward the exit. This approach leverages the established human sensitivity to directional cues in spatial orientation (Hegarty et al., 2006; Montello, 1993). Variants differ in discretization and weighting, but all share the core principle of directional alignment.

Distance-Based (Distance). Distance-based models select actions that minimize Euclidean distance to the exit. While computationally simple, this metric is often criticized for ignoring environmental structure (Golledge & Spector, 1978), as it lacks explicit consideration of angular alignment or corridor constraints.

Several hybrid combinations of these heuristics were also evaluated in the original implementation. As these variants did not substantially alter the qualitative pattern of results, detailed comparisons are provided in the accompanying code repository.

Interpretation of Heuristic Biases

Although the navigation models considered here are implemented algorithmically, our analysis treats them primarily as descriptive heuristics with distinct structural sensitivities. Importantly, we do not claim that humans explicitly compute the quantities optimized by any given model. Rather, each heuristic can be interpreted as inducing a characteristic bias over spatial choices, which may align more or less closely with human behavior depending on environmental structure.

The EXO strategy emphasizes global goal salience by prioritizing progress toward the exit, effectively treating the goal location as a dominant attractor throughout navigation. Such a bias may be advantageous in environments where the exit is perceptually or structurally salient and where deviations from a direct route are limited. In contrast, Angular strategies prioritize local directional consistency, favoring movements that maintain alignment with the target direction even when global progress is temporarily impeded. This bias can be interpreted as a tendency to preserve heading commitment, potentially reducing the cognitive cost associated with frequent reorientation in complex environments.

Distance-based strategies, while intuitively appealing, neglect aspects of environmental structure such as corridor constraints and turning affordances. As a result, they often fail to capture the spatial patterns observed in human navigation, despite their simplicity.

From this perspective, differences in human-likeness across models are not interpreted as evidence that humans adopt one strategy over another in a literal sense. Instead, they reflect how different heuristic biases resonate with the spatial cues made available by a given maze. This interpretation motivates our focus on comparative evaluation across environments and underlies the feature-based analysis introduced below.

Evaluation Metrics

We distinguish between three levels of evaluation. First, performance-based evaluation asks whether a model reaches the exit effectively, using success rate and path efficiency. Second, human-based evaluation asks whether a model reproduces the spatial patterns observed in human behavior, using the human-likeness correlation. Third, structure-sensitive evaluation asks whether the relationship between model behavior and human behavior varies systematically across maze environments. These levels are complementary but not equivalent: a model may be successful and efficient while still failing to reproduce human-like spatial visitation patterns. Model behavior was evaluated using three complementary measures. **Success rate** quantifies the proportion of trials in which the agent successfully reaches the exit. **Efficiency** is measured as the number of steps taken relative to the shortest path length. **Human-likeness** is defined as the Pearson correlation between the model's state-visit profile and the corresponding human profile for each maze. This measure captures similarity in spatial visitation patterns rather than exact trajectory matching, providing a robust aggregate comparison across participants and starting locations.

Human-Likeness as a Representational Measure

A central component of our evaluation is the use of human-likeness as a measure of model alignment with human navigation behavior. Rather than comparing individual trajectories or step-by-step action choices, we characterize behavior at the level of spatial visitation patterns. For each maze, human and model trajectories are aggregated into state-visit profiles, which represent the normalized frequency with which each grid cell is visited across trials.

This representation emphasizes where agents tend to spend time and which regions of the environment are preferentially explored, while abstracting away from fine-grained temporal ordering. Such an approach is well suited to environments in which multiple routes can lead to the same outcome and where local detours may not reflect substantive differences in navigational strategy. By focusing on spatial distributions, the human-likeness measure captures global structural similarities between human and model behavior that may be obscured by trajectory-level noise.

This measure aligns with accounts that treat navigation as the accumulation of spatial experience over time, rather than as the execution of an explicitly optimized sequence of actions, emphasizing how space is experienced and utilized rather than how actions are selected step by step.

We quantify human-likeness using the Pearson correlation between human and model state-visit profiles within each maze. High correlation indicates that the model reproduces the spatial patterns observed in human navigation, including preferences for particular corridors, junctions, or regions of the maze. Importantly, this measure does not reward models for strictly following shortest paths or minimizing steps. Instead, it reflects similarity in how space is utilized during navigation.

This distinction is critical for the present study. Models that achieve high success rates or efficient paths may nonetheless diverge substantially from human behavior if they exploit structural shortcuts or adopt overly deterministic policies. Conversely, a model may appear inefficient by traditional metrics while still capturing human-like exploration patterns. By prioritizing representational similarity over performance optimization, the human-likeness measure aligns directly with our goal of comparing heuristic biases as descriptive accounts of human navigation.

Finally, by computing human-likeness separately for each maze, this measure enables direct comparison of model alignment across environments. Variation in human-likeness across mazes can therefore be interpreted as evidence that environmental structure modulates the degree to which a given heuristic resonates with human spatial behavior.

Exploratory Feature-Based Analysis (E1)

To explore whether maze structure contains information about relative human-likeness, the analysis focuses on the comparison between EXO and Angular strategies. For each maze i, we compute

$$\Delta_i = H_{\text{Angular}}(i) - H_{\text{EXO}}(i),$$

where H denotes human-likeness. Mazes with small absolute differences are treated as ambiguous and excluded using a tie-removal threshold ε, set to the 10th percentile of $|\Delta|$ across mazes. Remaining mazes are labeled as *Angular-more-human-like* or *EXO-more-human-like*.

Maze-level structural features (e.g., connectivity, branching, and corridor statistics), along with path-based features derived from shortest paths, are extracted. These features are used in a simple logistic regression classifier as an exploratory diagnostic probe, rather than as a robust predictive model, to assess whether coarse maze-level structure carries information about relative strategy alignment. The classifier is evaluated using leave-one-maze-out cross-validation. Given the strong class imbalance, balanced accuracy is reported alongside standard accuracy, and results are compared to a majority-class baseline. Because the number of mazes is small and the resulting labels are strongly imbalanced, this analysis is intended to test for preliminary structural signal rather than to establish a generalizable classifier.

Results

Overall Model Performance Across Mazes

Row	EXO	target_angle
Maze_1	0.30342	0.733706727
Maze_2	0.39131	0.680831356
Maze_3	0.44029	0.358051374
Maze_4	0.51078	0.644787234
Maze_5	0.43354	0.753965829
Maze_6	0.45513	0.653059798
Maze_7	0.42675	0.744549953
Maze_8	0.61622	0.466278661
Maze_9	0.45532	0.649647999
Maze_10	0.49125	0.802429708
Maze_11	0.36958	0.65452647
Maze_12	0.46205	0.73309652
Maze_13	0.70964	0.404268695
Maze_14	0.52391	0.737035879
Maze_15	0.53667	0.742024728
Maze_16	0.55158	0.801347146
Maze_17	0.45555	0.514255765
Maze_18	0.37108	0.610341644
Maze_19	0.48018	0.768292343
Maze_20	0.00481	0.548937231
Maze_21	0.75707	0.612995163
Maze_22	0.36151	0.701281116
Maze_23	0.32349	0.344971988
Maze_24	0.3867	0.650253404
Maze_25	0.41028	0.74948401

Figure 2: Human-likeness correlation scores for EXO and Angular across all mazes.

Navigation models are first compared using success rate and efficiency across all maze environments. All three strategy families—EXO, Angular, and Distance—are capable of reaching the exit in the majority of mazes, though they differ in path length and robustness across starting locations. Distance-based strategies generally achieve high success rates but often produce trajectories that deviate from human spatial visitation patterns. Angular strategies tend to balance success and efficiency, while EXO strategies exhibit greater variability across environments.

Across the full set of mazes, hybrid combinations of heuristics yield only marginal improvements in success rate or efficiency relative to their base components. In particular, combining angular and distance cues does not consistently outperform the pure Angular strategy. As these effects are small and do not alter the qualitative pattern of results, we focus subsequent analyses on the three primary strategy families. Detailed performance comparisons for hybrid variants are provided in the accompanying code repository.

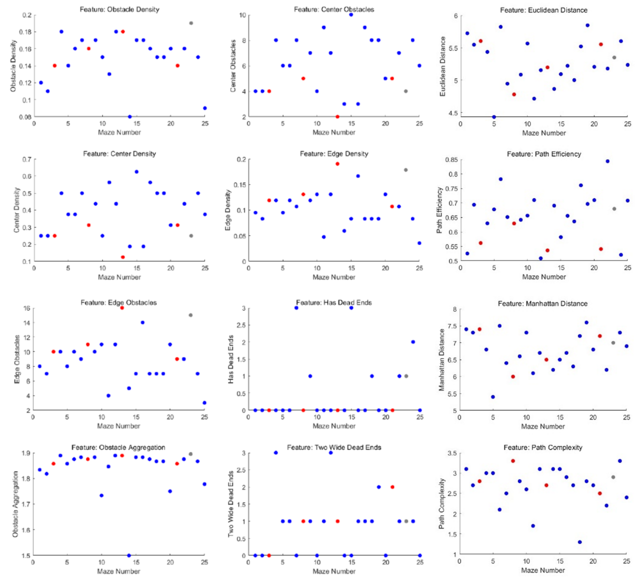

Figure 3: Maze-level structural and path features.

Human-Likeness Varies Systematically Across Environments

Human-likeness, quantified as the correlation between model and human state-visit profiles, reveals a markedly different pattern from success rate alone. While Angular strategies achieve higher human-likeness scores than EXO in the majority of mazes, this advantage is not uniform. In a subset of environments, EXO produces visitation patterns more closely aligned with human behavior.

Importantly, these differences are environment-specific rather than global. The relative ranking of strategies varies across mazes, indicating that no single heuristic dominates in all settings. This result motivates an environment-conditioned analysis of strategy alignment rather than a single aggregate ranking.

Exploratory Feature-Based Analysis of Relative Human-Likeness (E1)

To examine whether maze structure contains preliminary information about relative strategy alignment, EXO and Angular strategies are compared using the exploratory feature-based analysis described above. For each maze, the difference in human-likeness between Angular and EXO strategies is computed, and ambiguous cases are excluded using a conservative tie-removal threshold. This procedure yields 22 labeled mazes, with Angular favored in 19 cases and EXO favored in 3.

Using maze-level structural and path-based features, a simple logistic regression classifier evaluated under leave-one-maze-out cross-validation yields an accuracy of 0.77. Given the strong class imbalance, balanced accuracy (0.59) and area under the ROC curve ($AUC = 0.61$) are additionally reported. Rather than demonstrating practical predictive performance,

	EXO	target_angle
Obstacle_Density	0.224000405	-0.272139141
Center_Obstacles	-0.155469112	0.066851718
Edge_Obstacles	0.235424601	-0.31208352
Has_DeadEnds	-0.151306883	0.196698948
Center_Density	-0.155469112	0.066851718
Edge_Density	0.235424601	-0.31208352
Obstacle_Aggregation	0.26316294	-0.279804126
Two_Wide_DeadEnds	0.302433746	-0.198131803
Euclidean Distance	-0.314615385	0.423846154
Manhattan Distance	-0.360694311	0.336031452
Path Efficiency	0.096153846	0.293076923
Path Complexity	-0.052204941	-0.30510888

Figure 4: Human-likeness correlation scores for EXO and Angular models in each maze.

these results suggest that coarse geometric summaries of maze structure may carry limited but non-trivial information related to relative strategy alignment.

Classification performance is asymmetric across labels. The analysis reliably identifies mazes in which Angular strategies are more human-like, but shows limited sensitivity for EXO-favoring environments. This asymmetry reflects both the scarcity of EXO-labeled mazes and the conservative nature of the tie-removal procedure. Together, these results suggest that maze features contain informative but constrained signals regarding relative human-likeness.

Discussion

The central claim of this work is that the apparent human-likeness of simple navigation heuristics is systematically shaped by environmental structure rather than by global optimality or algorithmic sophistication. Across a diverse set of maze environments, we find that no single heuristic uniformly captures human navigation behavior. Instead, different strategies align more closely with human spatial patterns under different structural conditions (Werner et al., 2000). This finding supports a view of human navigation as context-sensitive and adaptive, in which environmental affordances selectively amplify particular heuristic biases.

From a *resource-rational* perspective, the predominance of Angular strategies in many mazes suggests that maintaining local directional consistency is a key factor in human navigation through structurally complex environments. We interpret this angular consistency as a mechanism to reduce the cognitive demands associated with frequent reorientation. In mazes with long corridors and multiple junctions, preserving a stable heading may provide a robust local cue (Hochmair & Frank, 2002) that maintains progress even when a direct path to the exit is temporarily obstructed. Notably, the robust alignment of the Angular heuristic with human paths bears a functional resemblance to the predictive properties of the Successor Representation (SR) framework (Stachenfeld et al., 2017). This correspondence suggests that simple directional heuristics may act as computationally efficient approximations to richer predictive maps, without requiring explicit state–state transition learning.

Conversely, the relative success of exit-oriented (EXO) strategies in a subset of mazes indicates that global goal salience can dominate when environmental structure affords direct and stable progress toward the exit. Grounding the EXO model in Gibsonian affordances (Gibson, 1979), our results suggest that visual salience can override directional commitment, consistent with theories of visually controlled locomotion (Warren, 1998). Consequently, in environments where the exit remains perceptually salient and deviations are minimal, the EXO bias effectively aligns with human intent.

Importantly, these differences do not imply that humans switch between explicit, discrete algorithms. Rather, they reflect how stable aspects of human spatial behavior resonate differently with the representational demands imposed by different structural contexts.

In additional analyses, we examined whether the relative human-likeness ranking between EXO and Angular strategies within each maze was stable across different starting locations and stochastic rollouts. We found that in mazes favoring EXO strategies, this advantage was consistently observed across starting positions, rather than being driven by a small subset of favorable initial conditions. This stability suggests that the observed differences reflect environment-conditioned bias dominance, rather than noise-driven effects or implicit strategy switching at the level of individual trajectories.

The exploratory feature-based analysis provides preliminary support for this interpretation while also revealing the nuances of strategy alignment. While the logistic regression classifier achieved a modest balanced accuracy, its deviation from the chance baseline suggests a systematic relationship between maze structure and strategy alignment. Rather than viewing this as a failure, we interpret it as evidence that coarse, geometric structural summaries capture only part of the information relevant to human navigation. Humans likely rely on finer-grained spatial cues, such as perceived corridor continuity or local visibility salience, which are not fully reflected in aggregated maze features. In this sense, the present features primarily capture global structural effects, such as overall connectivity, branching, and shortest-path properties. By contrast, local structural effects refer to the immediate decision context experienced by the navigator, including corridor continuity, junction geometry, and moment-to-moment visibility from the current position.

Several limitations of the present analysis offer avenues for future research. First, the maze-level features used here necessarily compress rich spatial structure into a small set of summary statistics. These geometric features may fail to represent the *local affordances* that strongly influence human decision-making.

Second, the current analysis treats each maze as a static environment and evaluates human-likeness in aggregate. This approach does not capture potential learning effects or dynamic strategy adaptation over time. Humans may initially rely on exploratory cues resembling EXO-like behavior before transitioning to more efficient directional heuristics as they construct a mental map.

Finally, while the heuristics examined here are intentionally simple, real human navigation is likely supported by a mixture of biases rather than a single dominant strategy. From this perspective, the observed alignment between humans and particular heuristics should be interpreted as partial overlap in representational sensitivity (Gallistel, 1990), not as evidence that humans implement any one model in isolation. Future work should transition from static geometric summaries to dynamic, agent-centric representations—such as isovist-based visibility graphs (Turner, 2001)—to better mirror the navigator's first-person perspective. Operationally, this could involve computing visibility fields at each visited state, extracting features such as visible corridor length, junction openness, goal visibility, and changes in visual access over time, and then relating these local perceptual features to moment-by-moment model–human alignment.

In conclusion, this work reframes heuristic comparison as a problem of structure-dependent alignment rather than strategy superiority. By evaluating simple heuristics across multiple environments and prioritizing representational similarity over performance metrics, we show how environmental structure constrains which aspects of human navigation behavior are most salient. This approach positions navigation models as descriptive probes of environmental sensitivity, offering a principled foundation for future work on adaptive strategy expression and structure-aware modeling in human navigation.

References

De Cothi, W., Nyberg, N., Griesbauer, E. M., Duvelle, É., Barry, C., & Spiers, H. J. (2021). Predictive maps in rats and humans for spatial navigation. *Current Biology*, *31*(10), 1221–1233.

Gallistel, C. R. (1990). *The organization of learning*. MIT Press.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–482.

Golledge, R. G., & Spector, A. N. (1978). Comprehending the urban environment: Theory and practice. *Geographical Analysis*, *10*(4), 403–426.

Hegarty, M., Montello, D. R., Richardson, A. E., Ishikawa, T., & Lovelace, K. (2006). Spatial abilities at different scales: Individual differences in aptitude-test performance and spatial-layout learning. *Intelligence*, *34*(2), 151–176.

Hochmair, H. H., & Frank, A. U. (2002). Influence of estimation errors on wayfinding-decisions in unknown street networks—analyzing the least-angle strategy. *Spatial Cognition and Computation*, *2*(4), 283–313.

Maguire, E. A., Frith, C. D., Burgess, N., Donnett, J. G., & O'Keefe, J. (1998). Knowing where things are: Parahippocampal involvement in encoding object locations in virtual large-scale space. *Journal of Cognitive Neuroscience*, *10*(1), 61–76.

Montello, D. R. (1993). Scale and multiple psychologies of space. In *Spatial information theory: A theoretical basis for gis* (pp. 312–321). Springer.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, *20*(11), 1643–1653.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208.

Turner, A., & Penn, A. (2002). Encoding natural movement as an agent-based system: An investigation into human pedestrian behaviour in the built environment. *Environment and Planning B: Planning and Design*, *29*(4), 473–490.

Warren, W. H. (1998). Visually controlled locomotion: 40 years later. *Ecological Psychology*, *10*(3-4), 177–219.

Werner, S., Krieg-Brückner, B., Mallot, H. A., Schweizer, K., & Freksa, C. (2000). Spatial cognition: The role of landmark, route, and survey knowledge in human and robot navigation. *Informatik-Forschung und Entwicklung*, *15*(4), 230–240.