

Dynamic Cognitive Maps for Agent Landmark Navigation in Unseen Environments

Samuel S. Sohn*
Rutgers University
samsksohn@gmail.com

Serena DeStefani
Rutgers University
serena.destefani@rutgers.edu

Mubbasir Kapadia
Rutgers University
mk1353@cs.rutgers.edu

ABSTRACT

The development of autonomous agents for wayfinding tasks has long maintained the usage of naive, omniscient models for navigation. The simplicity of these models improves the scalability of crowd simulations, but limits the utility of such simulations to the visualization of general behaviors. This restricted scope does not allow for the observation of more nuanced, individualized behaviors. In this paper, we demonstrate a novel framework for agent simulations that does not rely on omniscience. Instead, each agent is equipped with a memory architecture that enables wayfinding by maintaining a cognitive map of the space explored by the agent. Based on findings from simulation studies, cognitive science, and psychology, we describe a wayfinding procedure that simulates human behavior and human cognitive processes, incorporating landmark navigation, path integration, and memory. This cognitive approach makes observations of agent behavior more comparable to those of human behavior.

CCS CONCEPTS

• **Computing methodologies** → **Real-time simulation**;

KEYWORDS

Cognitive Model, Wayfinding, Memory Decay, Cognitive Map, Percolation Theory

ACM Reference Format:

Samuel S. Sohn, Serena DeStefani, and Mubbasir Kapadia. 2018. Dynamic Cognitive Maps for Agent Landmark Navigation in Unseen Environments. In *MIG '18: Motion, Interaction and Games (MIG '18)*, November 8–10, 2018, Limassol, Cyprus. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3274247.3274518>

1 INTRODUCTION

Effective navigation is a fundamental skill of primary importance in biological organisms' survival. Our aim is to simulate agent navigation in an unknown or partially known environment. We present a path-finding algorithm that is modeled after human cognitive processes, incorporating landmark navigation, path integration,

and memory. We want our agents to be as similar to humans as possible, so we make their memory imperfect (i.e., subject to decay).

Our approach stands in contrast to models of agent simulation that have been built upon the assumption of omniscient knowledge [Genesereth and Nilsson 1987] of the surrounding environment. Although this approach simplifies the problem of making the models scalable, it limits their application to real-world contexts. In addition, such an approach ignores the nuanced behavior of individuals. In this paper, we propose a agent simulation model not based on omniscience. Each agent relies on a complex cognitive architecture comprised of distinct memory layers [Wang et al. 2016] that sustain a personal cognitive map of the portion of the environment that has already been explored.

Landmarks are essential to navigation. We consider a landmark anything that captures our visual attention [Hunt and Waller 1999], including naturally salient reference points and any location in the architectural layout that can serve as the object of a decision-making process (e.g., the entrance of a room or a fork in a hallway). The way that memories of close landmarks interact with and modify each other was modeled by a spring-mass system. This is an intuitive model that accurately parallels recent accounts of cognitive maps [Couclelis et al. 1987] [Schmidt et al. 2003]. In addition, using this model facilitates the modeling of memory decay. This cognitive approach is adopted in order to base our agent simulation on cognitive psychology, making observations of agent behavior comparable to those of human behavior. The advantages of observing human behavior through agents are (a) that accurate agent behavior can inform internal human mechanisms and (b) that the wayfinding model manifested by agent behavior can be intuitively visualized.

The findings thus have the potential to be applied to different contexts. We can tailor online location services to human navigation by highlighting only those landmarks relevant to a particular user's path-finding needs. A landmark-based location service will help the user to orient himself in the new environment more efficiently and learn its layout more quickly. We can also use a virtual recreation of a public space to study and improve the strategic location of landmarks.

2 RELATED WORK

2.1 Navigation

There are various types of navigation, ranging from the simplest to the most computationally expensive, at the memory and cognitive level. The most basic is locomotor guidance, in which the subject simply travels toward a beacon that is perceptually unobstructed. Next is landmark navigation, which requires to store some landmarks [Hunt and Waller 1999] in memory. Then we have dead reckoning or path integration [Loomis et al. 1999], which represents

*This is the corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MIG '18, November 8–10, 2018, Limassol, Cyprus

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6015-9/18/11...\$15.00

<https://doi.org/10.1145/3274247.3274518>

the ability to update and store the current position with respect to starting position. Route-based navigation refers to the ability to store different positions, whereas map-based navigation relies on survey knowledge that has been acquired while exploring the environment [Foo et al. 2005].

Some researchers believe that animals and humans form a cognitive map which is perfectly Euclidean [Gallistel 1990], while others claim that the evidence in favor of the presence of a cognitive map is not conclusive [Foo et al. 2005].

The latest research shows that humans are not particularly good at path-integration [Foo et al. 2005]. Their preferred navigation method is landmark navigation, even though they are able orient themselves in the absence of landmarks [Riecke et al. 2000]. Overall, research shows that human beings are able to apply flexible strategies depending on the circumstances [Hunt and Waller 1999].

2.2 Memory Decay

The general mechanism of memory decay is not exponential; memory does not fade at a constant rate [Staddon et al. 2002]. If memory decay is not exponential, what is its nature? Empirical data shows that the rate of decay is fast at first but slows after a certain time; a sum of two exponential curves models this data well [Staddon et al. 2002]. This mechanism is necessary to account for Jost's law: given two memories, one recent and one old, the older one will decay more slowly.

What happens then to our memory of spaces as we move away from them and time passes? Distinct models of memory's spatial distortions have been proposed [Schmidt et al. 2003], all considering the presence of a target landmark and of another additional landmark, or of multiple additional landmarks.

According to the Nelson-Chaiklin model, the nearest landmark undergoes distortion, and the distance between the target and the landmark is underestimated according to a power function. If there is more than one landmark, each contributes to the distortion in proportion to its distance from the target.

According to the Huttenlocher model, the fact that the target location belongs to a certain category will have an influence on the recollection of its position. Schmidt and his colleagues [Schmidt et al. 2003] improve on the Huttenlocher's model by proposing a partition model, in which a space is subdivided in zones organized around single landmarks; the midpoint between landmarks may become an additional "attraction point" originating memory distortions.

2.3 Computational Models of Wayfinding

How can we build a computational model of wayfinding as a *human process*? The first comprehensive computational account of wayfinding, the TOUR model, was published by Kuipers in 1978 [Kuipers 1978]. This model's goal is to model a cognitive map of a large-scale space using a hierarchical networks of links between locations, integrated with a collection of available routes. Each route is defined as a sequence of possible actions. An additional level in the hierarchy is due to the presence of overarching regions that segment the space in question.

A hierarchical structure is also used by ELMER [McCalla et al. 1982] to integrate routes elaborated on the base of previous knowledge and navigation plans. Instead *The Traveler* [Leiser and Zilbershatz 1989] is built around a cognitive agent that tries to build a feasible path to a target location by connecting nodes in a network. The model SPAM [McDermott and Davis 1984] uses a fuzzy map that relies on spatial values that are allowed to vary within a pre-defined range. Finally NAVIGATOR [Gopal et al. 1989] relies on route-based knowledge, using a city-block metric. All these models avoid the use of a purely Euclidean map.

3 GEOMETRIC ANALYSIS OF BUILDING ENVIRONMENTS FOR NAVIGATION

The virtual environment is assumed to be a one-story building that is composed of rectangular rooms adjoined to connected rectangular hallways. A geometric analysis of building environments is needed in order to automatically generate a set of standard landmarks for any floor plan that fits certain assumptions. The assumptions are: the single entrance of a room is assumed to be smaller than the length of any of its walls (Assumption *R*), and the entrance of a hallway is assumed to be as wide as the inner width of the hallway (Assumption *H*). The geometric analysis will leverage a binary floor plan image, which can be generated from the navigation mesh [Kallmann and Kapadia 2016] of the virtual environment (Fig. 1.a).

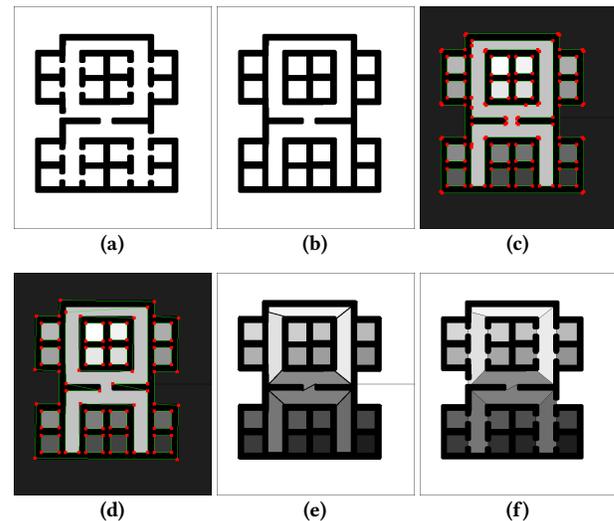


Figure 1: (a) is a binary floor plan image generated from a navigation mesh. (b) shows the result of the first segmentation phase. (c) shows the unsmoothed contours of the connected component of hallways. (d) shows the smoothed contours. (e) shows the result of the second segmentation phase. (f) shows the result of the watershed transform.

3.1 Floor Plan Segmentation

The floor plan image is segmented in two phases. In the first phase, a morphological closing operation [Serra 1983] is applied to the image to separate rooms from hallways (Fig. 1.b). Based on Assumption *R*, the closing operation can be parameterized to close the entrance of a room without filling its entirety, but Assumption *H*

makes this approach inapplicable to hallways. This phase leaves the hallways as one connected component (HCC), which the second phase segments coherently.

During the second phase, the contours of the pixels belonging to the HCC are first smoothed (Figs. 1.c, 1.d). For each point on the smoothed contours, a line is drawn to the closest other point, cutting the HCC (Fig. 1.e). A final labeling of connected components will properly distinguish between coherent hallway segments (henceforth referred to as hallways) given the following constraint. Suppose that for two hallways h_1 and h_2 adjoined into a bend, l_i represents length, w_i represents width, and $m = \sqrt{w_1^2 + w_2^2}$. Assuming that l_1 and l_2 are respectively at least $w_2 + m$ and $w_1 + m$, the pair of points defining the bend will always be connected after the second phase. This constraint also applies to T-junctions and 4-way intersections, which can both be decomposed into multiple overlaid bends.

After the two phases, a watershed transform is applied to the labelled image to determine which segments (rooms and hallways) are connected in the virtual environment, which is used to form a cell-and-portal graph [Haumont et al. 2003] (Fig. 1.f). (Our use of the watershed transform was inspired by Haumont et al. (2003) [Haumont et al. 2003].)

4 MEMORY-ENABLED AUTONOMOUS AGENTS

4.1 Landmarks (Visual Attention)

The visual attention of the agent captures non-occluded landmarks in the agent’s visual field up to configurable capacity [Habekost and Starrfelt 2008]. Landmarks consist of manually curated objects in the virtual environment and procedurally generated “reference points in the environment [...] where navigational decisions have to be made” [Millonig and Schechtner 2007]. Such reference points are located at the portals of the prior cell-and-portal graph. Each landmark’s priority in the agent’s visual attention is determined by its salience, which is composed of allocentric (i.e., inherent to the landmark) and egocentric (i.e., inherent or intentional to the agent) components. These components can be parameterized for interesting crowd scenarios. For example, simulating agents attuned to either object landmarks or reference point landmarks can give more nuanced distinction to the effect of landmark arrangements on navigation. In order to simplify our evaluation (Sec. 5), the visual attention’s capacity was uncapped.

4.2 Route Formation (Short-Term Memory)

The purpose of the short-term memory (STM) is to convert the instantaneous perception information from the visual attention into temporal position information (i.e., routes travelled between landmarks by the agent). As landmarks are captured by the visual attention, the short-term memory stores them with their priorities from the visual attention. While a landmark remains in the visual attention, its STM priority does not decay. Otherwise, its STM priority decays by a set function until it reaches zero, at which point it is removed from the short-term memory. The retention of the short-term memory is limited to a configurable capacity that the sum of the STM priorities cannot exceed. If this capacity does not

exceed that of the visual attention, the agent may have no memory of past landmarks, possibly resulting in a degenerate cognitive map. For the purpose of our evaluation (Sec. 5), the short-term memory capacity was uncapped and the decay function was made linear.

If an agent navigates by steering in the direction of its own perception, it will accrue temporal perceptual information containing surrounding landmarks, which the agent can use to determine its general position via the closest landmark (Fig. 2). When the agent becomes nearer to a different landmark, the route (i.e., the change in landmark) is stored in its long-term memory.

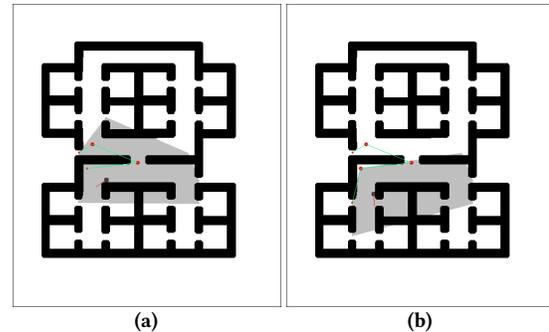


Figure 2: The above images show the extent of the short-term memory in light-gray, the cognitive map’s masses and springs as red circles adjoined by green lines, and the agent’s location and heading as a dark-gray circle with a red radius.

4.3 Cognitive Map (Long-Term Memory)

The cognitive map, which serves as an agent’s long-term memory, is modeled as a spring-mass system in the plane of the virtual environment. In the cognitive map, the masses represent landmarks, the values of the masses represent the landmarks’ memorability, and the springs represent traversed routes (Fig. 2).

When a route is first formed, its landmarks become masses connected by a spring. While a landmark remains in the short-term memory, its position in the cognitive map is the same as in the virtual environment. Otherwise, its position in the cognitive map decays according to the physics of the cognitive map, which is separate from that of the virtual environment, meaning that, for instance, masses will not collide with walls. This decay changes where the agent recalls the landmark to be located.

Initially, a landmark’s mass value in the cognitive map is its priority from the visual attention. While a landmark is either closest to the agent or in the visual attention, its mass in the cognitive map does not decay. Otherwise, the agent is decreasingly able to recall that the landmark exists and the landmark’s position in the cognitive map is subject to accelerating decay. The decrease in mass is limited to m_{min} .

The decay of the cognitive map is based on the Nelson-Chaiklin model of memory distortion, where “distortion is [...] toward the nearest landmark because the distance between landmark and target is underestimated according to a power function” [Schmidt et al. 2003]. Therefore, the equilibrium length of all springs is zero, resulting in a convergent decay. Also, each time a landmark is revisited, its mass value is reset to its original salience and increased by a multiplicative factor. During this process, if a landmark’s mass

value exceeds a set threshold, the cognitive map will set its mass value to m_{max} and stop decaying the mass's value and position. This behavior is consistent with long-term potentiation observed in humans [Scharf et al. 2002].

4.4 Wayfinding

When an agent is tasked with a navigational goal, its behavior reflects its knowledge about the goal. If the goal is in the short-term memory, it is directly navigated to. If the goal is unknown, the agent performs a blind search. Otherwise, if intermediary goals (i.e., subgoals) are unknown, the agent will recall the goal's heading and perform an oriented search [G Golledge 1999] in the virtual environment (implemented as greedy best-first search where the heuristic is the difference between a local heading and the goal heading). If subgoals are known, oriented search is performed for each subgoal in sequence.

To determine potential subgoals, uniform cost search to the goal is performed on the cognitive map by using the masses and springs as vertices and edges in a graph. From a given landmark/vertex, the cost of each neighboring route/edge r is computed by

$$a \times E(r) + (1 - a) \times (1 - M(l)), \quad (1)$$

where $0 \leq a \leq 1$, l is r 's other landmark, $E(r)$ is the normalized Euclidean distance, and $M(l)$ is the normalized mass value of l from the cognitive map.

4.4.1 Edge Cost Example. To illustrate, suppose that a cycle of edges forms a square with side length s . On this square, the agent is located at one vertex and its goal is located at an adjacent vertex. The two paths from the agent to the goal will be referred to as P_1 and P_2 , where P_1 has a length of s and P_2 has a length of $3s$. Both paths have v equidistant subvertices each, meaning that the edge length ratio of P_2 to P_1 is 3:1. Suppose that the masses of P_2 are three times larger than those of P_1 . For $a = \frac{1}{2}$, the sum of the edge costs (according to the above Expression 1) of P_1 and P_2 are both equal to $\frac{1}{2}s$ (Fig. 3.a).

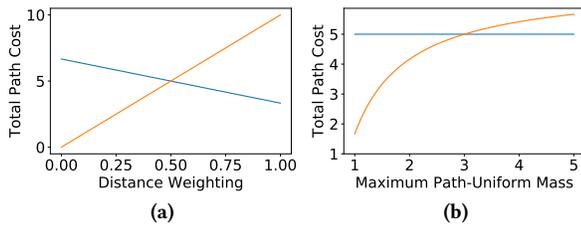


Figure 3: (a) shows the relationship between total path cost and the distance weighting (a) for P_1 (orange) and P_2 (blue). This graph assumes that s is 10 units long. (b) shows (for $a = 0.5$) how strengthening or weakening the memory of the P_2 (blue) (i.e., the maximum path-uniform mass) affects the total path cost of P_1 (orange). When the maximum path-uniform mass is less than its initial mass of 3, the agent favors the orange path for its increased relative salience.

4.4.2 Distance Weighting. The parameter a (i.e., the distance weighting) is responsible for the agent's preference between a shorter, unfamiliar path and a longer, familiar path. When a is increased in the prior example, P_2 is favored for its higher salience

in memory. When a is decreased, P_1 is favored for its recalled shortness (Fig. 3.a). The parameter a is not dynamic, so it should be preset based on experimentation. On the other hand, the masses are dynamic. Given a fixed $a > 0$, if the agent stops sensing, the decay of all masses (less than m_{max}) will converge to m_{min} (Fig. 3.b), and the path-finding will return to uniform cost search using Euclidean distances. Likewise, if the agent continues to explore, all masses will converge to m_{max} , resulting in the same uniform cost search. Both cases will return the same optimal path, because the route costs preserve relative distances when a is greater than zero and masses are equivalent. However, because the recall probabilities of the masses in the two cases are opposite, the former case is expected to return no landmarks and the latter case is expected to return all landmarks.

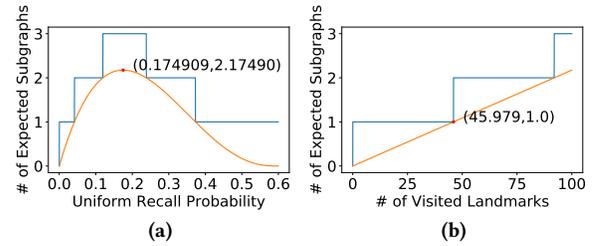


Figure 4: (a) shows the relationship between uniform recall probability and the expected number of subgraphs when 100 landmarks have been visited. (b) shows the relationship between the number of visited landmarks and the expected number of subgraphs when the uniform recall probability is 0.174909.

4.4.3 Path-finding Preceding Recall. To determine whether a subgoal is recalled, its corresponding landmark's mass from $[m_{min}, m_{max}]$ is converted to a probability and sampled. A complete path is calculated in memory before determining the recalled landmarks, because path-finding on a disconnected graph is nonsensical and worsens as either recall probability decreases or the number of landmarks increases. To illustrate, suppose that the graph of the cognitive map is arranged in a square lattice and that the recall probability for each landmark is p . In an instance of computing the path, the landmarks that are not recalled have their neighboring routes removed from the graph, resulting in a set S of disjoint subgraphs. If $p = 0$, S will be empty. If $p = 1$, S will only contain the original parent graph. For $0 < p < 1$, the expectation is not readily clear, so percolation theory can be used to approximate the expected size of S . In general terms, percolation theory facilitates the understanding of connected components that form when edges are randomly chosen from different types of graphs [Stauffer and Aharony 1994]. On an infinite square lattice, there exists a site percolation threshold p_c at which the expected size of a subgraph becomes infinite. According to Newman and Ziff (2000) [Newman and Ziff 2000], this occurs at $p_c \approx 0.59274621$. Therefore, for $p \geq p_c$, the expected size of the set of subgraphs is one. For $p < p_c$, the expression

$$|p - p_c|^{-\frac{43}{18}} \quad (2)$$

is used to approximate the expected size of a subgraph [Smirnov and Werner 2001]. In practice, Expression 2 overestimates the expected size of the subgraph, resulting in an underestimation of the number of subgraphs.

A square lattice was chosen because according to the assumptions about the virtual environment (i.e., its composition of axis-aligned rectangular cells), the maximum degree of a landmark is four. This intuitively represents a building composed of hallway intersections. Compared to a honeycomb lattice (which represents a building composed of hallway T-junctions), the square lattice results in fewer disjoint subgraphs on average, because the honeycomb lattice's site percolation threshold is higher ($p_c \approx 0.6970402$ [Feng et al. 2008]). Therefore, the using a square lattice yields a best-case analysis.

4.4.4 Percolation Example. Suppose that $p = 0.1$ and the number of landmarks in the parent graph is 100. The expected number of recalled landmarks $E(l)$ is exactly 10 and the expected size of a subgraph $E(s)$ is approximately 5.4, so the expected size of S (the set of disjoint subgraphs) is the quotient of $E(l)$ and $E(s)$; i.e., approximately 2 (Fig. 4.a). This indicates that when uniform recall probability is greater than or equal to p_c , the path is expected to be computed on one subgraph. As uniform recall probability decreases below p_c (with memory decay) and the number of visited landmarks increases (with exploration), the path is expected to span across an increasing number of disjoint subgraphs.

Figure 4.b shows that for $p = 0.174909$ (Fig. 4.a) and less than 46 visited landmarks, the path is expected to be computed on one subgraph. This should not be confused with the expected number of recalled landmarks. Only about 8 landmarks are recalled from 46 landmarks. For the given p , 46 recalled landmarks is expected to result in 6 disjoint subgraphs.

5 EVALUATION

For this evaluation, we compare the performance of our agents with that of human subjects from the experiment conducted by Buechner et al. [Buechner et al. 2009]. We use the percentage above optimum (PAO) metric in order to compare our results to those of Buechner et al.

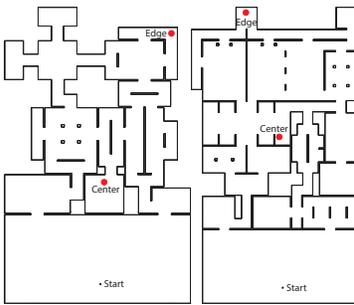


Figure 5: The floor plans above are the maze environments used in Buechner et al. (2009) [Buechner et al. 2009].

5.1 Experiment

We closely replicated the experimental procedure and virtual environments (Fig. 5) from Buechner et al. (2009) [Buechner et al. 2009], in which human subjects were given four types of navigational tasks (i.e., combinations of informed/uninformed tasks and center/edge goals). In each of the two maze environments, a single starting location was paired with a goal either in the center of the

maze or on the edge of the maze. Each task used the same start-goal pair for its trials. After every time the goal was reached (i.e., a trial was completed), the subject would be relocated back to the start. Although Buechner et al. had subjects repeat trials until the goals were reached near-optimally three consecutive times, they found that within the first five trials, the optimal path was found by a majority of the human subjects. Therefore, they have only provided data for the first five trials, which we will be comparing to. Our agents were only required to complete five trials per task and upon completing a trial, they were relocated to the start with their memories preserved. In their experiment, Buechner et al. also had half of the subjects for each type of goal (i.e., center or edge) start their first trial uninformed and had the other half start their first trial after 8 minutes of exploration.

5.2 Blind Search

For the purpose of this experiment, we adapted the agents' blind search to a stochastic perimeter strategy, which resembles observations of human subjects made by Buechner et al. [Buechner et al. 2009]. Initially, the agent begins searching the cells along the perimeter of the maze until it returns to the starting cell. During this initial search, it has a random chance of changing its strategy to a greedy best-first search, where the agent will prioritize cells in ascending degree-order. This will result in the agent visiting all adjacent rooms before advancing. Blind search is used until the goal is able to be recalled, at which point, the different navigational strategies described in Section 4.4 would be used.

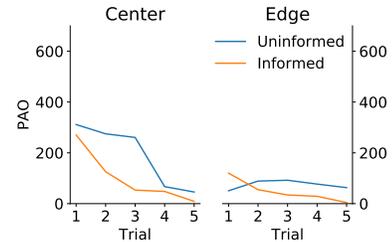


Figure 6: The above graph shows the average PAO of 20 agents for each class of navigational tasks.

5.3 Results

Twenty agents were simulated in real-time for each task while uninformed and again while informed. Figure 6 shows the average percentage-above-optimal (PAO) values resulting from the simulations.

$$PAO = \frac{d_{walked} - d_{shortest}}{d_{shortest}} \times 100 \quad (3)$$

The PAO metric is a measure of inefficiency when navigating to a destination that is calculated by Equation 3, where $d_{shortest}$ is the shortest distance to the goal and d_{walked} is the actual distance that the agent travelled to the goal.

Based on the PAO metric, we compared (Fig. 7) the model's average PAO across all five trials with Buechner's results (ground truth) in all conditions: in the uninformed condition with target at the center ($M= 191.92, SD= 125.32$, Buechner's $M= 257, SD= 362$) and at the edge ($M= 74.30, SD= 17.49$, Buechner's $M= 34, SD= 78$) and in the informed condition with target at the center ($M= 101.26,$

$SD= 103.47$, Buechner's $M= 164$, $SD= 362$) and at the edge ($M= 48.46$, $SD= 43.91$, Buechner's $M= 47$, $SD= 78$). Our model appeared to deviate less from the optimal path when the target was at the center, and slightly more when the target was at the edge in the uninformed condition.

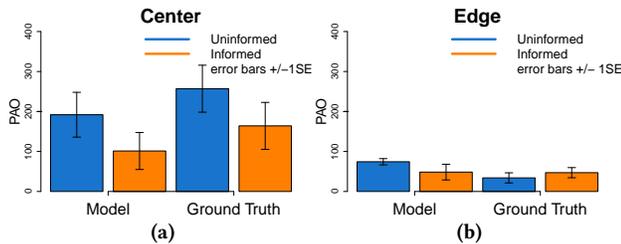


Figure 7: The above barplots show average PAO values across trials in the uninformed and informed conditions, with targets at the center (a) and edge (b), comparing the model with Buechner's results (ground truth).

5.3.1 Key Observations.

- (1) For all informed navigation tasks, a monotonic decrease in inefficiency is observed in the average PAO values.
- (2) For all navigation tasks, there is no apparent advantage to prior exploration in the first trial.
- (3) The mean of the average PAO values of center-goal tasks across all trials is higher than that of edge-goal tasks.

These observations are consistent with those made by Buechner et al. [Buechner et al. 2009]. The one inconsistency in our results is the unimodality of the uninformed edge-goal tasks, which we posit is a result of the stochastic perimeter strategy (Sec. 5.2). Although we hypothesize that a decrease in the likelihood of changing strategy during the initial search of the perimeter strategy would improve our results, we believe that it is first necessary to evaluate the PAO metric.

The PAO metric lends itself to a straightforward interpretation which cannot be discounted, but it does not provide much actionable insight on the quality of a cognitive model for navigation. We postulate that a naive model relying on omniscience could successfully replicate the average PAO fingerprint of humans without capturing the nuances of human decision making.

6 CONCLUSION AND FUTURE WORK

This novel framework for landmark navigation in crowd simulation has the potential to give new nuances to the navigational behaviors of agents. and its intuitiveness makes it accessible for exploring new crowd scenarios with no sacrifice in the findings of simulation studies, cognitive science, and psychology. The comparatively positive outcome of the evaluation evidences this claim.

We are currently integrating communication into this framework, and we hope to simulate rich, multilingual crowd scenarios to further demonstrate the flexibility of our framework. Because agents already use landmark navigation for wayfinding, the communication of routes follows intuitively.

ACKNOWLEDGMENTS

This work has been funded in part by NSF IIS-1703883, NSF SAS-1723869, and DARPA SocialSim-W911NF-17-C-0098.

REFERENCES

- Simon J. Buechner, Christoph Hölscher, and Jan Wiener. 2009. Search Strategies and Their Success in a Virtual Maze. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 31.
- H. Couclelis, R.G. Golledge, N. Gale, and W. Tobler. 1987. Exploring the Anchor-Point Hypothesis of Spatial Cognition. *Journal of Environmental Psychology* 7, 2 (June 1987), 99–122. [https://doi.org/10.1016/S0272-4944\(87\)80020-8](https://doi.org/10.1016/S0272-4944(87)80020-8)
- Xiaomei Feng, Youjin Deng, and Henk W. J. Blöte. 2008. Percolation transitions in two dimensions. *Phys. Rev. E* 78 (Sep 2008), 031136. Issue 3. <https://doi.org/10.1103/PhysRevE.78.031136>
- Patrick Foo, William H. Warren, Andrew Duchon, and Michael J. Tarr. 2005. Do Humans Integrate Routes Into a Cognitive Map? Map- Versus Landmark-Based Navigation of Novel Shortcuts. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31, 2 (2005), 195–215. <https://doi.org/10.1037/0278-7393.31.2.195>
- Reginald G Golledge. 1999. *Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes*. Vol. 10.
- Charles R. Gallistel. 1990. *The Organization of Learning*. Learning, development, and conceptual change., Vol. viii. The MIT Press, Cambridge, MA, US.
- Michael R Genesereth and Nils J Nilsson. 1987. *Logical Foundations of Artificial Intelligence*. Morgan Kaufmann 2 (1987).
- Sucharita Gopal, Roberta L. Klatzky, and Terence R. Smith. 1989. Navigator: A Psychologically Based Model of Environmental Learning through Navigation. *Journal of environmental Psychology* 9, 4 (1989), 309–331. 00126.
- Thomas Habekost and Randi Starrfelt. 2008. Visual attention capacity: A review of TVA-based patient studies. 50 (10 2008), 23–32.
- Dominique Haumont, Olivier Debeir, and François X. Sillion. 2003. Volumetric Cell-and-Portal Generation. 22 (09 2003), 303–312.
- Earl Hunt and David Waller. 1999. Orientation and Wayfinding: A Review. (1999).
- Marcelo Kallmann and Mubbasir Kapadia. 2016. Geometric and Discrete Path Planning for Interactive Virtual Worlds. *Synthesis Lectures on Visual Computing* 8, 1 (Jan. 2016), 1–201. <https://doi.org/10.2200/S00687ED1V01Y201512VCP023>
- Benjamin Kuipers. 1978. Modeling Spatial Knowledge. *Cognitive science* 2, 2 (1978), 129–153. 01107.
- David Leiser and Avishai Zilbershatz. 1989. The Traveller: A Computational Model of Spatial Network Learning. *Environment and behavior* 21, 4 (1989), 435–463. 00127.
- Jack M. Loomis, Roberta L. Klatzky, Reginald G. Golledge, and John W. Philbeck. 1999. Human Navigation by Path Integration. *Wayfinding behavior: Cognitive mapping and other spatial processes* (1999), 125–151.
- Gordon I. McCalla, Larry Reid, and Peter F. Schneider. 1982. Plan Creation, Plan Execution and Knowledge Acquisition in a Dynamic Microworld. *International Journal of Man-Machine Studies* 16, 1 (1982), 89–112. 00050.
- Drew McDermott and Ernest Davis. 1984. Planning Routes through Uncertain Territory. *Artificial intelligence* 22, 2 (1984), 107–156. 00248.
- Alexandra Millonig and Katja Schechtner. 2007. Developing Landmark-Based Pedestrian-Navigation Systems. *IEEE Transactions on Intelligent Transportation Systems* 8, 1 (March 2007), 43–49. <https://doi.org/10.1109/TITS.2006.889439>
- M. E. J. Newman and R. M. Ziff. 2000. Efficient Monte Carlo algorithm and high-precision results for percolation. *Physical Review Letters* 85, 19 (Nov 2000), 4104–4107. <https://doi.org/10.1103/PhysRevLett.85.4104> arXiv: cond-mat/0005264.
- Bernhard E. Riecke, HAHC van Veen, and H. H. Bulthoff. 2000. *Visual Homing Is Possible without Landmarks*. Technical Report. Tech. Rep.
- Matthew T. Scharf, Newton H. Woo, K. Matthew Lattal, Jennie Z. Young, Peter V. Nguyen, and Ted Abel. 2002. Protein Synthesis Is Required for the Enhancement of Long-Term Potentiation and Long-Term Memory by Spaced Training. *Journal of Neurophysiology* 87, 6 (2002), 2770–2777. <https://doi.org/10.1152/jn.2002.87.6.2770> arXiv:https://doi.org/10.1152/jn.2002.87.6.2770 PMID: 12037179.
- Thomas Schmidt, Steffen Werner, and Jörn Diedrichsen. 2003. Spatial Distortions Induced by Multiple Visual Landmarks: How Local Distortions Combine to Produce Complex Distortion Patterns. *Attention, Perception, & Psychophysics* 65, 6 (2003), 861–873.
- Jean Serra. 1983. *Image Analysis and Mathematical Morphology*. Academic Press, Inc., Orlando, FL, USA.
- Stanislav Smirnov and Wendelin Werner. 2001. Critical exponents for two-dimensional percolation. *Mathematical Research Letters* 8, 6 (2001), 729–744. <https://doi.org/10.4310/MRL.2001.v8.n6.a4>
- J. E. R. Staddon, I. M. Chelaru, and J. J. Higa. 2002. Habituation, Memory and the Brain: The Dynamics of Interval Timing. *Behavioural Processes* 57, 2 (2002), 71–88.
- D. Stauffer and A. Aharony. 1994. *Introduction To Percolation Theory*. Taylor Francis.
- Di Wang, Ah-Hwee Tan, and Chunyan Miao. 2016. Modeling Autobiographical Memory in Human-like Autonomous Agents. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 845–853.