This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

Cognitive Path Planning with Spatial Memory Distortion

Rohit K. Dubey, Samuel S. Sohn, Tyler Thrash, Christoph Hölscher, André Borrmann, Mubbasir Kapadia

Abstract—Human path-planning operates differently from deterministic AI-based path-planning algorithms due to the decay and distortion in a human's spatial memory and the lack of complete scene knowledge. Here, we present a cognitive model of path-planning that simulates human-like learning of unfamiliar environments, supports systematic degradation in spatial memory, and distorts spatial recall during path-planning. We propose a Dynamic Hierarchical Cognitive Graph (DHCG) representation to encode the environment structure by incorporating two critical spatial memory biases during exploration: *categorical adjustment* and *sequence order effect*. We then extend the "Fine-To-Coarse" (FTC), the most prevalent path-planning heuristic, to incorporate spatial uncertainty during recall through the DHCG. We conducted a lab-based Virtual Reality (VR) experiment to validate the proposed cognitive path-planning model and made three observations: (1) a statistically significant impact of sequence order effect on participants' route-choices, (2) approximately three hierarchical levels in the DHCG according to participants' recall data, and (3) similar trajectories and significantly similar wayfinding performances between participants and simulated cognitive agents on identical path-planning tasks. Furthermore, we performed two detailed simulation experiments with different FTC variants on a Manhattan-style grid. Experimental results demonstrate that the proposed cognitive path-planning model successfully produces human-like paths and can capture human wayfinding's complex and dynamic nature, which traditional AI-based path-planning algorithms cannot capture.

Index Terms—Cognitive Path-Planning, Human Wayfinding, Fine-To-Course, Spatial Memory, Agglomerative Hierarchical Clustering

1 Introduction

Remarkable strides have been made in computer graphics, immersive game technology, and pedestrian simulation. With the advancement of virtual and augmented reality technology in terms of computation power and reduced cost, computational heavy and realistic games (e.g., Call of Duty, Black Ops Cold War) are getting more realistic with outstanding graphics, surround sound, and in-game physics. Even though such games now deliver a sense of realism that mimics the real world in great detail, the immersive element is still lacking when it comes to the behavior of NPCs. According to [1], the believability of NPC behavior is crucial for immersion in games. Although, it is possible to add randomness to NPC's navigation behavior to simulate variation in decision-making. But, such variation may not look realistic, and it is highly likely, that a human player might understand the underlying random phenomena and get off-put by it.

One of the critical challenges to the modeling and simulation of realistic pedestrian evacuation behavior is developing a cognitive path-planning model that mimics the decision-making process based on incomplete and distorted spatial memory. In the last decade, several works on agent navigation have focused on producing realistic local behaviors [2], [3], [4] by formulating tasks in a biologically plausible way [5]. However, there is still a disconnect between AI-based techniques for path planning [6], [7], which focuses on efficiency and optimality considerations, and cognitive studies [8], [9], [10] which seek to understand how humans navigate in complex spaces fundamentally. This work aims to bridge the gap between AI-based path-planning techniques and cognitive wayfinding techniques to aid in realistic pedestrian simulation and enhance computer games' realistic experiences by leveraging Virtual Reality (VR)-based behavioral study. The three main objectives for conducting the VR experiment are (1) Parameterization of the proposed hierarchical data structure to model a realistic path planning model. (2) To validate, formalize and model the existing theories of spatial distortion during path planning. (3) To validate human and agent wayfinding behaviors for evaluation.

A variety of research in cognitive science, artificial intelligence, and virtual humans has been conducted to investigate the fundamental principles by which humans navigate. The importance of cognitive maps [11], and cognitive graphs [12] for navigation has gained widespread acceptance in the research community. In 1948, Tolman [11] suggested that rodents develop an internal model of their space called a cognitive map. The cognitive map's important role in flexible navigation was also evident under various adaptive situations (e.g., goal, transition, and reward revaluation). Some research also suggests that spatial representations should be characterized as cognitive graphs consisting of locations connected by paths [12]. Cognitive maps and cognitive graphs can coexist without competing against each other. These two forms of spatial representations can operate concurrently or independently and may incorporate both spatial and nonspatial knowledge [13]. Originally, the cognitive map was characterized as geometrical and non-hierarchical, but more recent research suggests that spatial representations are hierarchical and complex, encoding other types of information in addition to simple geometry [8], [10], [14]. Evidence from [15], [16] argues in favor of the importance of nonspatial features (i.e., high-frequency names, number of turns, frequency of visits) for the grouping of landmarks while mentally representing an environment. Investigations using distance judgments [17], direction judgments [18], and the spatial recall of objects' locations [19] suggest that spatial representations are prone to distortions that are systematic rather than random. Such distortions in memory are primarily caused by the perceptual organization that people impose on the environment. Multiple studies

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

have confirmed that classifying stimuli into categories [20] and the sequence in which they are observed [21] can affect the way the stimuli are perceived, thus resulting in biases. Two prominent causes of spatial memory distortion are *Categorical Adjustment* (CA) [22] and *Sequence order Effect* (SE) [21]. The CA model formalizes the tendency to parse spatial layouts into categories and estimate objects' locations within those layouts by combining information from different hierarchical spatial categories. In the SE model, the location of a category's center is influenced by the order with which landmarks were experienced during navigation. According to the SE model, this center location shifts toward landmarks observed earlier in navigation.

To model the fact that mental spatial representation is hierarchical, we propose a Dynamic Hierarchical Cognitive Graph (DHCG) representation based on Euclidean and non-Euclidean landmarks features. We employ Agglomerative Hierarchical Clustering (AHC) and extract hierarchical levels based on empirical evidence (Section 6) to identify distinctive clusters based on the spatiotemporal landmark features observed during the agent's navigation. To model distortions in spatial memory while considering the sequence order of visited landmarks and the systematic bias of recalled landmarks towards the category center, we incorporate *Categorical Adjustment* (CA) [22] and *Sequence order Effect* (SE) [21] models. Finally, we consolidate the above findings into a cognitive path-planning model that is stochastic and subjective by extending the most commonly used Fine-to-Coarse (FTC) strategy (proposed in [23] and supported by [9]). We also performed a lab-based Virtual Reality (VR) experiment to understand the influence of exploration duration and the order of landmark visits on a cognitive graph and to tune the number of hierarchical levels in a cognitive graph.

Findings from the VR experiment revealed a moderate negative correlation between exploration duration and wayfinding performance in terms of distance traveled (i.e., exploration time increased while average path-distance for six wayfinding tasks decreased). We also found a strong positive correlation between the sequence order of landmark visits and route choices at intersections. Similar results for the sequence order effect were also observed with agent simulations that employed the proposed cognitive path-planning model. In these simulations, we observed a reduction of ~21% on Path Edit Distance (PED) for the trajectories generated after 10 and 20 minutes of familiarization. A gradual increase in the agents' (FTC+CA) path distance during path-planning was also observed after 0, 10, and 20 minutes of memory decay, highlighting an increase in distortions caused by CA during spatial recall over an extended time. Finally, multiple simulations on a Manhattan-style grid environment are performed that showcase the difference between the proposed cognitive path-planning algorithm and an AI-based path-planning algorithm.

The primary goal of this paper is to propose a realistic (human-like) path-planning model for a Non-Playable Character (NPC) to exhibit engaging and believable game-play and for advancing pedestrian dynamic simulators and virtual-reality-based games. Adding realistic spatial memory to agents also enhances general circulation and egress simulations, which can assist designers/architects in the early design process of a building. The key contributions of this paper are: (1) The paper proposes a cognitively inspired Dynamic Hierarchical Cognitive Graph (DHCG) to encode an occupant's short-term memory in an unfamiliar environment based on their exploration pattern. (2) The paper improves the most commonly employed human wayfinding heuristic, "Fine-To-Course" by adding distortions to spatial recall during wayfinding. (3) The paper grounds the parameterization of the proposed data structure on behavioral data collected from a VR study. (4) The paper presents simulation results that bear a high resemblance to human participants in terms of path distance, the number of turns in the path, and the landmarks used along the path.

2 Related Work

2.1 Cognitive Models of Navigation

The early traditional view of mental spatial representations is that they take the form of Euclidean non-hierarchical cognitive maps [11], [24]. Here, we highlight some of the non-hierarchical cognitive map-based path-planning models. The Tour model was one of the first computational models to capture the non-hierarchical cognitive map [24]. The locations, order of the places visited, and local intersections' geometry was the Tour model's building blocks. In addition, a real-time, biologically plausible neural network capable of inferring paths to remote places with the cognitive map's assistance was proposed in [25]. Similarly, an abstract representation of allocentric maps of an environment-based computational model for robot exploration was proposed in [26]. These allocentric maps were computed using the shapes of surfaces and boundaries relative to the robot's location. To generate more precise knowledge of the exploration space and improve metrical mapping techniques, a spatial memory model HSSH (Hybrid Spatial Semantic Hierarchy) was proposed in [27], extending the SSH (Spatial Semantic Hierarchy) model proposed in [28]. In contrast, [29] attempts to generate a cognitive-map-like representation using an incomplete plan of the environment. This agent effectively uses distance and orientation information to find the way back home. Recently, [30] proposed a wayfinding model for an autonomous agent based on a dynamic cognitive map using landmarks. This work incorporates memory decay and spatial memory distortion but relies on an outdated distortion model based on a non-hierarchical cognitive map.

Research in spatial cognition and artificial intelligence suggests that cognitive maps can also be hierarchical [8], [31], [32]. Research done in [10], [13], [33] suggests that a human's cognitive map contains three levels of associative memory (i.e., survey, graph, and route) to form a hierarchy in which higher levels encompass the lower levels. Inspired by spatial cognition research, [34] proposes a hierarchical model to solve the traveling salesman problem (TSP). A computational model of spatial navigation based on the hierarchical representation of the space is also proposed in [10]. Their model describes the experimental data collected in [8], which provides strong evidence for the environment's hierarchical organization. Spatial representations such as Geographic Information System tend to fit the hierarchical organization [35] naturally. In a hierarchy, objects are arranged as ordered structures over several distinct levels. Objects at higher levels are more abstract than the detailed lower-level objects. One of the advantages of the hierarchical grouping of objects is that it allows complex tasks to be broken down into multiple independent subtasks, allowing the simulation of parallel distributed architectures. Therefore, we believe that hierarchical representations present some properties that suit human spatial memory modeling and thus cognitive maps. For a complete understanding of the work done in the field of cognitive maps and cognitive graphs, we guide readers towards [12], [36], [37], [38], [39].

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

2.2 AI Techniques for Navigation

Inspired by human cognition, researchers have proposed various hierarchical spatial representations of the environment and hierarchical search algorithms that perform path-planning to overcome the computationally challenging problem of real-time path-planning in a large-scale environment. Multiple hierarchical approaches such as HPA* [6], Hierarchical Path-Finding Theta (HPT*) [7], Coarse-to-Fine (CFA*) [40], and Fine-To-Coarse (FTC-A*) [41] have been proposed to quickly produce abstract paths and to improve the path-planning speed. HPA* partitions the search space into rectangular blocks of large spaces and computes optimal local paths between a limited set of entrances and exits in the blocks. Subsequently, in [7], the authors proposed HPT* by combining HPA* with Theta*. The notable difference between HPA* and HPT* is the replacement of A* with Theta* as a subroutine. These authors found that HPT* (compared to HPA*) resulted in a lower number of visited nodes and lower memory requirements for large open maps with distant nodes having a line of sight. Unlike human navigation behavior, HPA* and CFA* return a complete global path. [23] suggests that human beings plan an incomplete path in the beginning of a route to start moving immediately. In [40], the authors proposed a two-step-based path-planning Coarse-to-Fine A* algorithm. By this method, no new graph is needed to abstract the search space. Instead, a simple grid representation is used to form another coarser-resolution grid map. The authors' goal was to preserve the original grid map's occupancy structure by systematically reducing the grid map resolution. Recently, reinforcement learning (RL) has been employed to train human-like navigational agents [42], [43], [44] by learning a map from egocentric views, employing a Deep-RL framework for robot navigation using depth information from an RGB-D sensor, and learning a partial map using a deep neural network. Notably, this work is able capture several human wayfinding behaviors without using global knowledge. However, these works generalize the agent's ability on pixel-level variations, work only in a small environment, are not biologically plausible, and lack the formation of human-like memory. We argue that a much richer environment involving local landmarks as cues and distortions in spatial memory recall is essential to aid human-like path-planning.

Inspired by the work done in [23] and motivated by the absence of a computational model for FTC wayfinding heuristics, the authors in [41] proposed an FTC-A* robot navigation system that integrates regionalized spatial representations with an efficient path-planner. Fine-grained route planning is performed between neighboring regions, while coarse route planning is performed between distant regions to reduce the search space. Our work extends this work specifically in two ways. First, we produce an experience-driven and dynamic hierarchical spatial representation with time instead of a fixed two-level regionalized spatial representation. Second, we incorporate spatial memory distortions in path-planning that naturally occur due to memory decay. The motivation behind our proposed cognitive path-planning model is to develop a method that can perform human-like path-planning in an unfamiliar environment. Specifically, we want to develop a path-planning model that learns an environment during general exploration and builds a cognitive map without relying on a complete environmental representation as proposed in most prior works.

3 Overview

An overview of the major components of the proposed framework is provided by Figure 1. We propose a dynamic hierarchical cognitive graph representation DHCG of spatial memory to enable human-like path-planning. We first formulate a similarity metric for comparing landmarks, which is supported by studies on human cognition. Next, we apply this metric to a hierarchical clustering algorithm and produce a three-level hierarchy, which has the same number of levels as human spatial memory. Finally, we introduce spatial memory distortions from elapsed time and interfering landmarks, which influence how paths are recalled from the DHCG (Section 4). In Section 5, we adapt FTC, the path-planning heuristic most commonly used by humans, to the DHCG. Computationally, it computes a subgraph from the DHCG, in which landmarks closer to the agent are represented by low-level nodes in the DHCG with fine-grained information and landmarks closer to the goal are represented by high-level nodes with coarse-grained information. This enables the agent to make local decisions using a coarse mental heading toward its goal. Next, we conduct a lab-based Virtual Reality (VR) experiment to investigate and validate our hypotheses on the impact of exploration duration and exploration patterns on the formation of a cognitive map and subsequent wayfinding decisions (Section 6). We validate and test the proposed cognitive path-planning model against participants' data in Section 7. Finally, to showcase the model's generalizability, we simulate cognitive wayfinding behavior on an irregularly shaped real-world building using multiple variations of FTC with and without distortions in Section 8.

4 Dynamic Hierarchical Cognitive Graph (DHCG)

We provide a detailed description of our contributions to cognitive graphs and cognitive path-planning research in the following two sections.

Features	Definition	Weights	Type
Path Distance	Walking distance between two landmarks	0.4	Spatial
Visibility	Visibility between two landmarks (binary)	0.1	Spatial
Recency	Time elapsed from last visit (in seconds)	0.2	Non-spatial
Frequency	Number of times two neighbouring landmarks are visited from each other	0.3	Non-spatial

TABLE 1: Selected features for clustering landmarks using agglomerative hierarchical clustering.

4.1 Clustering Metric for Landmarks

The selection of an appropriate metric is crucial in the development of DHCG. Based on previous research, both spatial and non-spatial factors influence the clustering of regions in the cognitive map. In Table 1, we describe the four features that constitute the comparison metric for Agglomerative Hierarchical Clustering (AHC): path distance, visibility, recency, and frequency. The path distance between two landmarks has been shown to have a stronger correlation to the grouping of landmarks [45] and has thus been

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

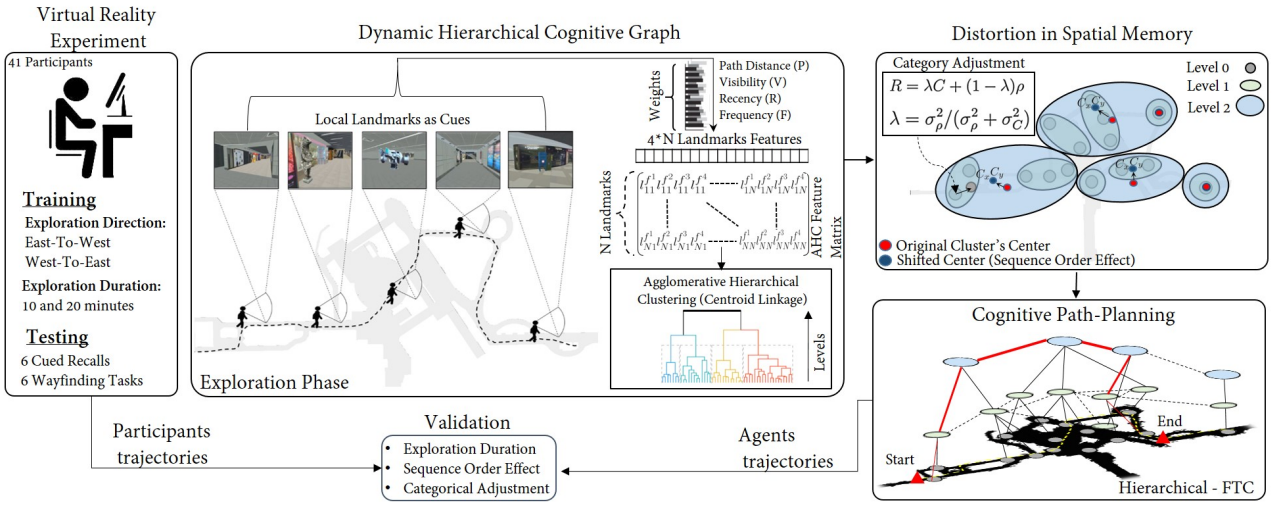

Fig. 1: Framework Overview. We present a cognitive model of agent path-planning that simulates human-like learning of unfamiliar environments: (1) We propose a DHCG representation to encode the environment structure. (2) The influence of the landmarks sequence order and the systematic bias towards the center of the location's cluster is incorporated to model spatial recall errors due to memory decay. (3) The "Fine-To-Coarse" (FTC) path-planning heuristic commonly applied by humans during navigation is extended to incorporate spatial uncertainty during recall. (4) We conduct a VR experiment to validate the proposed path-planning model.

given the highest weight. The visibility between two landmarks enables an association to be formed between them, so landmarks that are visible from each other's locations are grouped in the same cluster [46]. To incorporate the temporal clustering effect, landmarks that are visited immediately after each other tend to be grouped together [47]. Finally, landmarks that are visited frequently from each other (e.g., the entrance of a building and a reception desk) also tend to be grouped [45]. Each of these features is pair-wise (i.e., between two landmarks). For each landmark, these features are used to create a $(4 \cdot n)$-dimensional feature vector that compares the landmark with every other landmark for each of the 4 features, where n is the number of landmarks observed during navigation. Not all features contribute equally to the clustering process, and they vary among individual humans. For this study, we have assigned weights to individual features based on prior studies [45], [46], [47].

4.2 Distortion in Spatial Memory

People tend to cluster their spatial memory according to both spatial and non-spatial properties in a hierarchy [8], [14]. During navigation and cued spatial memory recall, the distance of an object/location which belongs to the same cluster as a cued object is systematically underestimated (i.e., closer than the actual distance between them), and the distance to an object/location from a different cluster from the cued object is overestimated (i.e., further than the actual distance between them) [23], [31], [32].

Category Adjustment. Humans use clusters from their hierarchical cognitive map to estimate a goal location during spatial memory recall, which can become distorted. The Category Adjustment model by [22] offers a computational explanation for reconstructive distortion effects by combining information from two separate spatial memory representations. These two representations are a fine-grained location representing a specific place and a category that stores the center of the cluster to which that place belongs (i.e., the prototypical location of the category). Due to the decay of working memory, the fine-grained location becomes uncertain, forcing the inference of its location from its spatial category. This produces a systematic bias towards the category center, which is modeled by a Bayesian process. According to CA, the estimated location of a place (R) is computed as a weighted combination of its category center (ρ) and the fine-grained location (C). The weight (λ) varies as a function of category dispersion (σ_ρ^2) and the degree of imprecision surrounding the fine-grained location (σ_C^2) as shown in Figure 2.

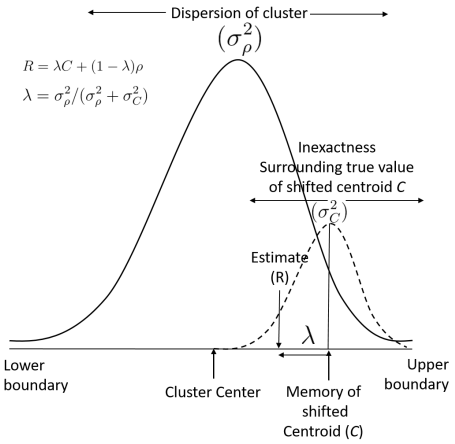

Fig. 2: The category adjustment model used for spatial recall errors during the estimation of goal locations. The Bayesian combination of information about the prior distribution with the present distribution of inexactness surrounding the true value for C results in biased estimates that are more likely to be pulled towards the cluster center.

Sequence Effect on Category Center Location. Findings in [21] provide evidence that the location of a category center is influenced by the order in which locations are visited, and shifts toward the nodes visited earlier. In other words, the temporal order in which the landmarks are visited in a cluster impacts the location

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

of the category center. Thus, we propose to formulate the shifted category center location (C_x, C_y) as shown in Equation 1:

$$C_x = \frac{\sum_{i=1}^{n} L_x^i S^i}{\sum_{i=1}^{n} S^i}; C_y = \frac{\sum_{i=1}^{n} L_y^i S^i}{\sum_{i=1}^{n} S^i}, \quad (1)$$

where the magnitude of sequence order shift for the ith landmark is given by $S^i = 1/i_t \mid 0 \leq t < n$. n represents the total number of landmarks within the cluster. i_t represents the landmark's sequence order during the navigation process in that cluster. L_x^i, L_y^i are the X, Y coordinates of the ith landmark inside the concerned category.

4.3 Memory Decay

Two significant opinions exist in the literature on memory decay. Some researchers provide evidence for memory decay due to elapsed time [48], and others provide evidence for decay due to interference effects [49]. Recently, researchers have converged on the idea that both factors influence memory decay. Based on existing memory theory, the authors in [50] modeled decay due to time and interference, as shown in Equation 2 from their paper. To model memory decay in our framework, we propose replacing the activation of an item with the time of landmark observation during navigation. The probability $p(j)$ of recalling the jth landmark is given by:

$$p(j) = \frac{s \cdot e^{-r_j/m}}{\sum_{i=0}^{N-1} s \cdot e^{-r_i/m}}, \quad (2)$$

where s is the salience of a landmark when it was last observed, and r is the time since the landmark was last observed. m is the memory strength constant that determines the rate of memory decay over time. Based on the trajectory data collected in the VR experiment (Section 6), the observed value of $m = 10^5$. According to the study in [51], the decline in participant's visual working memory between 1 and 4 seconds was small and not statistically significant, but the decline in visual working memory between 4 and 10 seconds was large and significant. Based on the above finding, we introduce a parameter, Memory Decay Rate (β), that is randomly sampled from a range of (4-10) in simulation experiments (Sections 7 and 8).

4.4 Agglomerate Hierarchical Clustering

Given the strong evidence for the hierarchical organization of the human spatial memory, we propose a dynamic hierarchical cognitive graph (DHCG) to encode environmental structure. We analyze navigational landmarks by applying Agglomerative Hierarchical Clustering (AHC), which has been popular for several biological applications. AHC is well-suited for clustering objects into a hierarchy. The resulting dendrogram highlights the progressive grouping of the data that can be leveraged to identify the suitable number of classes into which the data can be grouped. Research into human memory has shown two common types of natural clustering: temporal and semantic. Using AHC wherein no apriori information about the number of clusters is required, we can efficiently parameterize by employing hand-picked features (Table 1). By default, AHC returns a binary tree (known as a dendrogram), resulting in at least $log_2(n)$ hierarchical levels for n-many landmarks. According to our findings from participants' recall data (Section 6) and the three levels of associative memory observed in humans [13], [33], the DHCG is limited to three

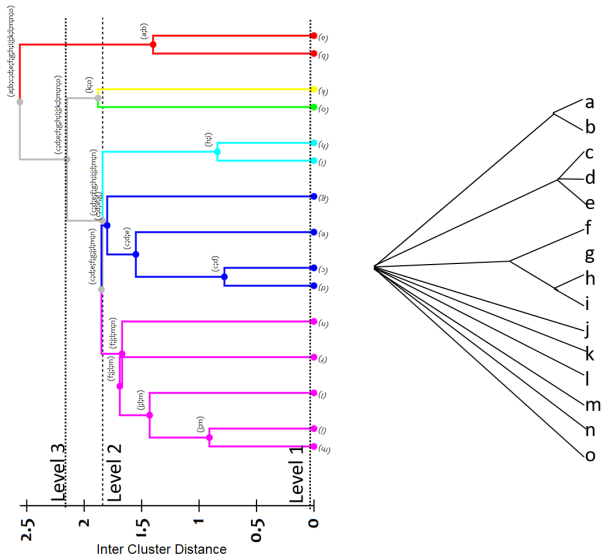

Fig. 3: (left) Dendrogram showcasing Levels 1, 2, and 3 clusters formed after Agglomerative Hierarchical Clustering for an agent after 20 minutes of exploration in East-To-West direction. (right) Ordered tree derived from one participant recall data after 20 minutes of exploration in East-To-West direction.

hierarchical levels by truncating the dendrogram at three levels. The lowest of these levels consists of the individual landmarks, as shown in Figure 3. Each of the other two levels is found by first specifying a threshold on the inter-cluster distance that is stored in each node of the dendrogram. The level at that threshold is the set of all nodes with the highest inter-cluster distance less than the threshold, in which nodes are not ancestors of each other. As a result, each landmark at Level 1 is in its own cluster and has an abstract parent node in Level 2 that clusters the Level 1 node with others. The same principle applies between Level 2 and Level 3 nodes.

The DHCG is formally represented by $G =< L, E >$, where L is an array of 3 sets of nodes $[L_1, L_2, L_3]$, each L_l containing the nodes in level l, and E is the set of edges between any two nodes. Intra-level edges represent routes between landmarks or groups of landmarks, while inter-level edges represent cluster membership.

5 Cognitively-Based Path-Planning

In this section, we utilize the DHCG representation of the environment to extend the well-established Fine-To-Coarse (FTC) wayfinding heuristic [23].

5.1 Abstract Subgraph Construction

FTC uses different hierarchical levels of the DHCG simultaneously during path-planning. This is achieved by computing an abstract subgraph from the DHCG (Algorithm 1) in which locations close to the agent's position are known with fine-grained details and information becomes more abstract with distance. In order to create an abstract subgraph $AG =< N_a, E^{out} >$, we first define the operation $\varphi(n, l)$, which takes a node n and returns the cluster in Level l to which n belongs. Using this formulation, we define several key variables $s_l = \varphi(s_1, l), l \in [1, 3]$, and

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

$d_l = \varphi(d_1, l), l \in [1, 3]$, where s_l and d_l are either the start and destination nodes or their parent nodes in L_2 and L_3 (line 2). Additionally, $\varphi(n, l)$ can be used to determine whether two nodes belong to the same L_2 or L_3 parent node. Within each hierarchical level of the DHCG, nodes should be connected if and only if they are connected by an edge in L_1, which corresponds to a direct physical route (e.g., lines 6 - 10 for Level 1). For Levels 2 and 3, we compute the L_1 edges that are inter-cluster (i.e., belonging to two nodes with different parent nodes in L_2; lines 11 - 19 for Level 2). This implicitly includes edges between different parent nodes in L_3. The edges in the abstract graph are initialized to the L_1 edges that are fully contained within s_2. This represents the fine-grained information close to the start location. The coarse information further from the start is added from the L_1 inter-cluster edges that are either contained in s_3 (between L_2 clusters) or spanning between different L_3 nodes are added to the abstract graph. This ensures that a navigable path exists from the s_1 to the d_3, the L_3 parent node of the destination.

5.2 Abstract Subgraph Pathfinding

With small modifications to their cost values and destinations, existing pathfinding algorithms such as Uniform Cost Search, A*, and D* Lite [52], which typically operate on non-hierarchical graphs, can be used for path-planning on the DHCG using the abstract subgraph, which collapses the three hierarchical levels into one. First, depending on the location of the ultimate destination d_1 with respect to the start s_1, the immediate destination may change. If d_1 and s_1 are in different L_3 clusters, the destination becomes d_3. If they are in different L_2 clusters but are both under s_3, the destination becomes d_2. If they are both in s_2, the destination remains unchanged. Next, the values used to keep track of path costs are redefined to leverage information from the DHCG. Based on the work of [30], the cost function balances between path distance and landmark salience using a parameter $\alpha \in [0, 1]$. For each edge in the abstract subgraph, the distance values are normalized, and for each node, the salience values are normalized. Then, in order to compute the cost $c(u, v)$ from node u to node v, the normalized distance $\delta(u, v)$ and salience $\sigma(v)$ are combined as follows:

$$c(u, v) = (\alpha) \cdot \delta(u, v) + (1 - \alpha) \cdot \big(1 - \sigma(v)\big).$$

When node v represents a cluster of landmarks instead of an individual landmark, its salience value becomes the minimum among its landmarks [53]. To maintain the human-like nature of salience preference, we randomly assign the salience value α from a range of [0,1].

5.3 Introducing Distortions in Fine-To-Coarse Planning

Despite evidence that the original formulation of FTC explains human wayfinding well, FTC was originally ignorant to the distortions that constantly influence human spatial memory. We have identified two notable biases of spatial memory (i.e., category adjustment and sequence order shift) to incorporate into the FTC heuristic. Category Adjustment (CA) biases the recall of a landmark's location towards the centroid of the cluster to which it belongs, and the Sequence order Effect (SE) biases a landmark's recalled position towards landmarks seen earlier. In order to reach its goal, an agent must repeatedly apply the FTC path-planning heuristic. Irrespective of an agent's position, the agent's recalled heading towards higher-level nodes should be distorted by SE.

Algorithm 1 Abstract Subgraph

Input: A navigation graph $G =< N, E >$, where N is a set of visited landmarks and E is a set of corridors/paths connecting those landmarks
Input: $\varphi =< n, L_i >$ A dictionary of landmarks clusters at L_2, and L_3, where n represents a landmark
Input: start nodes s, and destination nodes d
Output: A single-level Abstract Subgraph $AG =< N_a, E^{out} >$, where N_a is the set of abstract nodes

1: $N_a \leftarrow \{\}, E^{inter} \leftarrow \{\}$
2: $L_2 \leftarrow \varphi[s][0], L_3 \leftarrow \varphi[s][1]$
3: **for** $e \in G$ **do**
4: **if** $\varphi[e_s][0] \neq \varphi[e_d][0]$ **then**
5: $E^{inter} \leftarrow e$
  //Add L_1 Edges to Abstract Graph (AG)
6: **for** $e \in G$ **do**
7: **if** $\varphi[e_s][0] = L_2$ & $\varphi[e_d][0] = L_2$ **then**
8: $N_a \leftarrow e_s$
9: $N_a \leftarrow e_d$
10: $E^{out} \leftarrow e$
  //Add L_2 Edges to Abstract Graph (AG)
11: **for** $e \in E^{inter}$ **do**
12: **if** $\varphi[e_s][1] = L_3$ & $\varphi[e_d][1] = L_3$ **then**
13: **if** $\varphi[e_s][0] = L_2$ **then**
14: $e_{new} \leftarrow W_E < e_s, \varphi[e_d][0], 1 >$
15: **else if** $\varphi[e_d][0] = L_2$ **then**
16: $e_{new} \leftarrow W_E < e_s[0], e_d, 1 >$
17: **else**
18: $e_{new} \leftarrow W_E < \varphi[e_s][0], \varphi[e_d][0], 1 >$
19: $E^{out} \leftarrow e_{new}$
  //Add L_3 Edges to Abstract Graph (AG)
20: **for** $e \in E^{inter}$ **do**
21: **if** $\varphi[e_s][1] \neq \varphi[e_d][1]$ **then**
22: **if** $\varphi[e_s][1] = L_3$ **then**
23: **if** $\varphi[e_s][0] = L_2$ **then**
24: $e_{new} \leftarrow W_E < e_s, \varphi[e_d][1], 1 >$
25: **else**
26: $e_{new} \leftarrow W_E < \varphi[e_s][0], \varphi[e_d][1], 1 >$
27: **else if** $\varphi[e_d][1] = L_3$ **then**
28: **if** $\varphi[e_d][0] = L_2$ **then**
29: $e_{new} \leftarrow W_E < \varphi[e_s][1], e_d, 1 >$
30: **else**
31: $e_{new} \leftarrow W_E < \varphi[e_s][1], \varphi[e_d][0], 1 >$
32: **else**
33: $e_{new} \leftarrow W_E < \varphi[e_s][1], \varphi[e_d][1], 1 >$
34: $E^{out} \leftarrow e_{new}$
35: **return** $< N_a, E^{out} >$

This distortion plays a role when the agent performs path-planning on its abstract subgraph and estimates distances. On the other hand, CA plays an increasingly large role as the agent nears its goal. The agent's separation from its goal at L_3 will become a separation at L_2 and then L_1. As the level of the goal's abstraction becomes increasingly fine-grained, the higher-level goal clusters will bias the goal's recalled position.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

5.4 Parameterizing Agglomerative Hierarchical Clustering

Two essential parameters required in the AHC formulation are (1) number of hierarchical levels in a hierarchical cognitive map and (2) the number of nodes at each level. We perform a VR experiment to determine Parameter 1 (Section 6). In this section, we present our method to identify Parameter 2 and provide theoretical proof for its validity. Humans are known to use least-effort heuristics (e.g., FTC) for decision-making [8]. Accordingly, we use this principle to determine our parameters for AHC (i.e., the number of nodes in Levels 2 and 3). We formulate effort/cost as the computational complexity of path-planning on the abstract subgraph, which is a function of the number of nodes n and edges. We compute the expected number of nodes in an abstract subgraph using Eq. 3. We compute the expected number of nodes in an abstract subgraph $y(n, j, k)$ using Eq. 3, where Level 2 has n^j nodes and Level 3 has n^k nodes. We are interested in finding the exponents j and k, which minimize the average size of the abstract graph (i.e., the total number of nodes) and in turn reduces the cost of path-planning. We first solve $dy/dk = 0$ (Eq. 4), which yields $k = j/2$ for $n > 0$, and then $dy/dj = 0$ (Eq. 5), which yields $j = 2/3$ and $k = 1/3$ for $n > 0$. Therefore, we truncate the AHC dendrogram at two inter-cluster distances where the number of clusters becomes $n^{2/3}$ for L_2 and then $n^{1/3}$ for L_3. If the number of levels in the DHCG were to increase, the optimal number of clusters for each additional level would always be the square root of the previous level's number. For instance, the fourth level would have $n^{k/2} = n^{1/6}$ clusters.

$$y(n,j,k) = (|L_3| - 1) + (|L_2|/|L_3| - 1) + n/|L_2|$$

$$y(n,j,k) = n^k + n^{j-k} + n^{1-j} - 2 \quad (3)$$

$$dy/dk = n^{-k}(n^{2k} - n^j)log(n) = 0 \quad (4)$$

$$dy/dj = n^{-j}(n^{1.5j} - n)log(n) = 0 \quad (5)$$

6 Investigating Factors Impacting the Cognitive Map

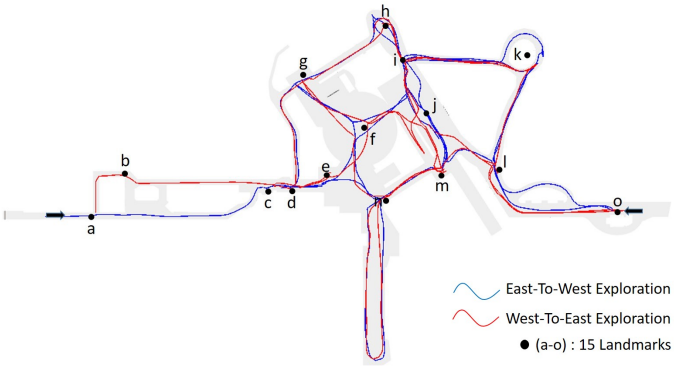

Fig. 4: Top-down view of the floor plan of a real-world shopping mall. Black circles represent the locations of 15 landmarks. Blue and red trajectories represent the exploration directions for East-To-West and West-To-East navigation directions.

6.1 Purpose

We performed a Virtual Reality (VR) experiment to empirically compute the number of hierarchical levels in a cognitive map (Section 6.5.1). During this experiment, participants completed two trials that were each divided into three phases (i.e., exploration, cued-recall of landmark locations, and immersive wayfinding tasks). Participants were randomly assigned to two groups that determined the order with which landmarks were seen during exploration. Participants were exposed to each landmark twice with either exploration pattern (i.e., West-To-East and East-To-West). Moreover, the distribution of participants in two groups was randomized, and the order of wayfinding tasks was also presented in random order. We believe the randomness could have eliminated or minimized the impact of any biases that may arise in the experiment (e.g., laziness and lack of attention) during the experiment. For this human participant data, we investigated two hypotheses. Specifically, we expected exploration duration to be negatively correlated with wayfinding performance (H1) and the order in which landmarks were seen during exploration to bias the order with which landmarks were visited during wayfinding (H2). In Section 7, we also compare the human participants and cognitive agents in terms of path-planning behavior.

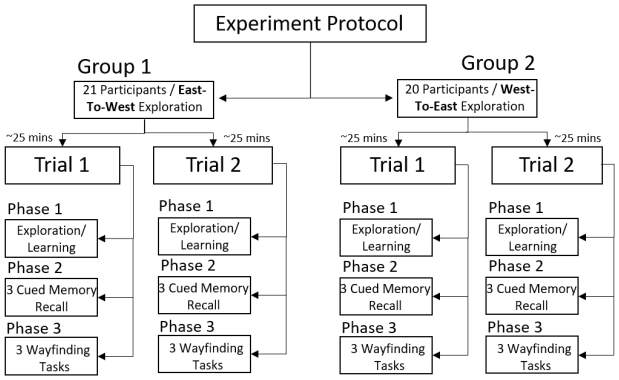

Fig. 5: Behavioral Experiment Protocol

6.2 Participants

Ethics approval was acquired for all experiments from the ethics commission at ETH Zurich (EK 2019-N-79). All participants were university students. Participants were informed that they would receive 5 SGD as a bonus if they performed the experiment thoughtfully and accurately to the best of their capabilities. At the end of the experiment, all participants were paid the additional performance bonus irrespective of their performance. In total, all participants received 25 SGD upon completion of the experiment. No other eligibility criteria based on experience were set. Participant self-reports of gender indicated that 58.53% of them were male and 41.47% of them were female. Participants' ages ranged from 20 to 32 years (Mean = 22.51, Standard Deviation = 1.98). A total of 41 participants' data was collected. A between-subject design was chosen, and participants were divided into two groups (Group 1: 21 participants, Group 2: 20 participants) to learn the environment in two different orders. Group 1 participants were asked to learn the environment in a East-To-West (EW) direction, and Group 2 participants learned the same environment in a West-To-East (WE) direction (see Figure 4 and Table 4).

6.3 Materials

A 3D model of a real-world shopping mall (i.e., ION Orchard in Singapore) served as the virtual environment (Figure 4). Fifteen

Wayfinding Tasks	1a	1b	1c	2a	2b	2c
Origin-Destination	b-h	l-a	o-g	c-k	j-b	i-d

TABLE 2: Six wayfinding tasks used for the VR experiment. Please refer to Figure 4 for the locations of the origin and destination landmarks in the virtual environment.

salient landmarks were placed around intersections that were identified as crucial for wayfinding. The environment did not have any signage or crowd flow during either training or testing phases, so participants had to rely on the shape of the environment and the landmarks themselves to complete the tasks. In the physical environment, participants sat on a chair positioned approximately 1.5 meters from the center of a computer screen (i.e., a Samsung monitor with a 28-inch diagonal and a resolution of 3840 $\times$ 2160 pixels).

6.4 Procedure

The experiment protocol is described in Figure 5. Participants were asked to complete a consent form before starting the experiment. After informed consent, participants were asked to perform two trials (~25 mins. each). Each trial consisted of three phases: (1) exploring the environment passively via video segments, (2) performing three cued recall tasks, and (3) performing three wayfinding tasks.

In **Phase 1**, participants were passively moved along a route by viewing multiple video segments (~1 min. each). The videos were recorded from a first-person perspective (with an eye height of 1.72 m and a walking speed of 1.5 m/s) and explored specific sections of the virtual environment. In each video clip, multiple landmarks were traversed, and each subsequent video continued from the location where the last video ended. All 15 landmarks were seen more than once during this exploration phase. We asked participants to watch videos rather than actively navigate the route to ensure identical stimuli for all participants. The participants were specifically instructed to learn the locations of all landmarks for subsequent cued recall and wayfinding tasks.

In **Phase 2**, participants were asked to perform a cued spatial recall task three times. Three landmark cues were preselected out of 15 landmarks. Landmarks were selected if they were at least two nodes apart, excluding the first and last landmarks shown during phase 1. Participants were instructed to recall the landmarks starting with the cue and including all of the landmarks in the environment. The name of the cued landmark was presented as text on the screen. Participants were then shown a total of 22 landmark names as text in randomized order, including 14 correct landmarks and 8 distractor landmarks. The distractor landmarks were added to deter participants from guessing. At the beginning of the experiment, participants were also informed that they would be incentivized for recalling the landmarks correctly.

Phase 3 was conducted immediately after the cued recall phase. In Phase 3, participants performed three wayfinding tasks (Tasks 1a, 1b, and 1c) as highlighted in Table 2) in the same environment that they learned via video-based exploration in Phase 1. For each wayfinding task, the origin and destination pair was preselected from the set of 15 landmarks. Unique origin-destination pairs were preselected so that multiple alternative paths were possible and included landmarks with high betweenness centrality. In graph theory, $g(v) = \sum_{s \neq v \neq t} \sigma_{st}(v)/\sigma_{st}$, where σ_{st} is the total number of shortest paths from node s to node t and $\sigma_{st}(v)$ is the number of those paths that pass through v) in the paths.

For Trial 2, a new set of videos (~1 min. each and recorded keeping the same navigational direction as Trial 1) were shown to the participants. In Phase 2 of Trial 2, three different landmarks were preselected and used as cues that were not chosen during Trial 1. Finally, in Phase 3, three different sets of origin-destination pairs were employed (i.e., not chosen in Trial 1; Tasks 2a, 2b, and 2c as highlighted in Table 2) using the same preselection criteria as for Trial 1. Overall, each participant performed a total of six cued recall and six wayfinding tasks. The preselected landmark cues and origin-destination pairs were kept the same for the two participant groups.

6.5 Results

6.5.1 Parameter Computation

One essential parameter required in the computational model of the dynamic hierarchical cognitive graph is the **number of hierarchy levels** (L). Below, we present our method to compute the parameter using participants' cued spatial recall data from Trial 2 (i.e., 20 minutes of exploration). We chose not to use the recall data from Trial 1 (i.e., after 10 minutes of environment familiarization) for parameter computation for two reasons. First, not all landmarks were visited twice or more in 10 minutes of exploration, and second, there is a higher likelihood of forming a more substantial mental representation of the environment after 20 minutes of exploration compared to 10 minutes.

As discussed in Section 6.4, participants were asked to perform three cued recall tasks that each began with different landmarks. This method was intended to reveal their individual mental representations of the environment and their basis for clustering the environment into regions at various hierarchical levels. We presume that participants would have imagined moving from the cued recall landmark to another close landmark based on their mental representations recently formulated during the exploration phase. The ordered tree algorithm proposed in [54] and supported in [8] is used to represent clusters in the recall data. The output of the algorithm is an ordered tree (Figure 3 (right)). The tree represents the clusters at multiple levels that the ordered tree algorithm has uncovered from the recall orders. To compute the number of levels, we generate a cumulative distribution of the number of clusters at all levels for four conditions (i.e., two navigation durations and two navigation directions). An average of 95% of the clusters were represented by Levels 1 through 3. To quantitatively evaluate the number of hierarchical levels, we constructed a scree plot with participants' recall data. A scree plot (Figure 6) shows the eigenvalues on the y-axis and the number of levels on the x-axis. At Level 3, we notice that the slope of the curve is clearly leveling off (i.e., forming an elbow), indicating the number of hierarchical levels that includes a representative set of clusters formed in the cognitive map. Thus, we approximate the number of hierarchical level as three in a cognitive map, which is also supported in [10], [13], [33].

6.5.2 Hypothesis 1: Impact of Exploration Duration

To examine the impact of exploration duration on the development of a cognitive map, we examined participants' performance on six wayfinding tasks after exploring the environment for 10 and 20 minutes (i.e., constructing a mental representation of the

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

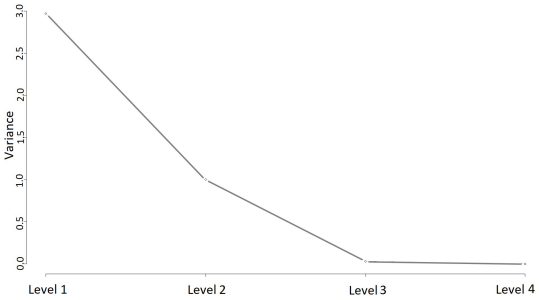

Fig. 6: Scree plot for the cumulative distribution of the number of clusters at each level from participants' recall data.

environment). The average path-distance traveled by participants is presented in Table 3. An independent-samples, two-tailed t-test was conducted to test for a difference between different exploration duration conditions. Participants who learned the environment for 10 minutes with a East-To-West navigation order (M = 248.64, SD = 53.52) compared to the same 15 participants exposed to the environment for an additional 10 minutes (M = 182.19, SD = 46.64) demonstrated a significant reduction in path-distance for the wayfinding tasks, t(28) = 2.29, p = .04. However, we did not notice a similar significant reduction in path-distance for West-To-East navigation order for 10-minute (M = 234.53, SD = 81.88) and 20-minute exploration duration conditions (M = 199.65, SD = 56.48), t(28) = 0.85, p = 0.4. Overall, we notice a moderate negative correlation of exploration duration with wayfinding performance (~21%, p = .04659). Thus, **Hypothesis H1** can be accepted.

	East-To-West (EW)		West-To-East (RW)	
Exploration Duration (mins)	**10**	**20**	**10**	**20**
Task 1a	245.36±42.55	193.93±24.65	205.63±25.87	245.05±118.79
Task 1b	344.15±175.42	221.36±23.01	268.55±68.61	217.8±29.13
Task 1c	190.92±24.16	199.16±37.34	195.17±23.29	172.68±2.91
Task 2a	257.13±173.43	153.98±14.25	271.91±209.61	192.32±70.63
Task 2b	247.86±110.58	222.90±61.27	352.50±303.02	263.65±224.78
Task 2c	206.39±154.11	101.79±16.09	113.42±19.38	106.45±24.32
Average	**248.635±53.52**	**182.186±46.64**	**234.53±81.88**	**199.65±56.48**

TABLE 3: Participants' path-planning data for six wayfinding tasks after familiarizing with the environment for 10 and 20 minutes.

6.5.3 Hypothesis 2: Sequence Order Effect

To investigate the impact of the sequence order in which landmarks were visited during the exploration phase (i.e., independent variable) on the navigation pattern (i.e., dependent variable), we examined participants' performance for six wayfinding tasks under two exploration duration conditions (i.e., 10 and 20 minutes) and two navigational directions. In Group 1 (East-To-West), participants learned the environment starting at the left side of the environment. In Group 2 (West-To-East), participants learned the environment starting at the right side of the environment. In Table 4, we showcase the landmark visit orders for both navigation directions.

Order	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
East-To-West	a	b	c	d	e	g	h	i	f	j	m	n	k	l	o
West-To-East	o	l	m	i	j	k	n	h	f	g	e	d	c	a	b

TABLE 4: Landmark visit order under two navigation conditions

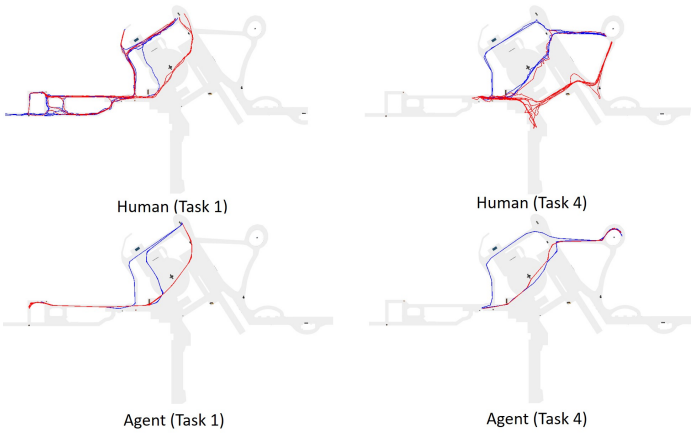

Fig. 7: Visualization of humans' and agent's trajectories to showcase the impact of landmarks sequence order on wayfinding trajectories for two navigation tasks. The trajectories in blue represent wayfinding paths when participants/agent were trained with the East-To-West navigation order. The trajectories in red represent the wayfinding paths when they were trained with the West-To-East navigation order during the exploration phase. We notice similar eastward/rightward shifts in particpiant and agent trajectories when exposed to landmarks in the East-To-West sequence.

For trajectories produced by participants during a wayfinding task, we compute a list of nearest landmark traversed per trajectory (e.g., "abcdgh" for Task 1a trajectory). We then compute the total sum of the landmarks visited using Table 4 for the respective navigation pattern in both groups. For the above example, the distance computed for East-To-West (EW) for the example trajectory "abcdgh" is (a)1+(b)2+(c)3+(d)4+(g)6+(h)7 = 23. We repeat the process for all six wayfinding tasks for all participants in both navigation direction conditions. An independent-samples, two-tailed t-test was performed to test for a difference between navigation directions in terms of navigation pattern. For an exploration duration of 10 minutes, the 15 participants who were training from East-To-West (M = 55.9, SD = 19.80) compared to the 15 participants trained in the West-To-East navigation order (M = 63.92, SD = 21.49) demonstrated significant difference in navigation patterns, t(28) = -1.758, p = .041. Similarly, for the exploration duration of 20 minutes, the 15 participants who were training from East-To-West (M = 50.04, SD = 12.21) compared to the 15 participants trained in the West-To-East navigation order (M = 65.21, SD = 19.76) also demonstrated significant difference in navigation pattern, t(28) = -4.20, $p < .01$. In Figure 7 (top), we notice distinctive eastward (trajectories visualized in blue (Figure 7)) and westward shifts (trajectories visualized in red (Figure 7)) in the trajectories of participants when trained using EW and RW navigation patterns, respectively. Crawford and colleagues [21] provided evidence that the sequence order of presented stimuli has an effect on recalling spatial locations by conducting an experiment using paper packets of 50 mm tall × 280 mm wide

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

(refer to Experiment 1 section in [21] for detailed experiment procedure). Based on the above findings, we believe that a similar sequence order effect can be established in recalling real-world spatial location during human wayfinding. Thus, **Hypothesis H2** can be accepted.

7 Comparison of Human and Cognitive Agent Behavior

To validate and test the proposed cognitive path-planning model, we compare it against participants' data. We begin by replicating the process of cognitive map development in agents similar to the process that human participants presumably followed during the VR experiment. The cognitive agent explored the environment (i.e., the 3D model of ION Orchard in Singapore) using the same trajectories used to train the participants during the exploration phase. A total of four different cognitive map models for the agent were developed (i.e., two exploration directions of East-To-West and West-To-East, and two exploration durations of 10 and 20 minutes). We performed simulations for four different versions of the cognitive path-planning algorithm (i.e., (1) Fine to Coarse (FTC), (2) Fine to Coarse + Categorical Adjustment (FTC+CA), (3) Fine to Coarse + Sequence order Effect (FTC + SE), and (4) Fine to Coarse + Categorical Adjustment + Sequence order Effect (FTC+CA+SE)). Each version of the algorithm employed the specified distortion during the computation of the DHCG and its influence on the cluster's centroid to bias destination location during spatial memory-based recall. Below, we demonstrate two measures of average distance traveled and the average path edit distance for four scenarios on six wayfinding tasks (specified in Table 2 and 3) and compare the results with participants' behavioral data (baseline) by sampling 25 cognitive agents.

7.1 Preliminaries

To realistically simulate the interaction of an agent with the 3D environment, we modeled human-like visual perception. A first-person character-based 3D avatar was designed with an average eye height of 1.72m above the ground. Both horizontal and vertical fields of view (FOV) were modeled (i.e., 120 degrees and 60 degrees, respectively) to account for human neck rotations. The walking speed of 1.5 m/sec was assigned to simulate an average human walking speed. The agents walked the same trajectories as humans to mimic human-like learning of the environment during exploration. Like humans, the agent's visual perception model was able to detect the salient landmarks placed at crucial intersections. After the end of the exploration stage, a navigation graph was formulated with landmarks as nodes and the path/corridors connecting those landmarks as edges. The formulated navigation graph was employed in the construction of a DHCG, as discussed in Section 4. To maintain human-like variability in path-planning, two parameters (i.e., Salience Preference α in the range of 0–1 and Memory Decay Rate β in the range of 4–10) were selected randomly per agent. For each condition and each algorithm version, 25 agents performed the identical six wayfinding tasks that the participants performed.

7.2 Validation of Hierarchical Tree Structures

In Section 4, we proposed a hierarchical representation of spatial memory using AHC, the elected features for clustering, and the identification of the number of levels in a cognitive map. To validate the quality and similarity of the proposed hierarchical structures produced by our algorithm, we compare it with the ordered tree formed from participants' recall data (i.e., baseline). For comparing two hierarchical tree structures, we need a similarity measure. The task is not only to compare the topology but also the content of the clusters formed in the structure. However, to our knowledge, there is very little in the literature about hierarchy comparison [55]. Below, we propose two metrics for comparing trees formed from the proposed algorithm and baseline (i.e., participant recall data).

- **Average Landmarks Depth (ALD)** - We introduce a measure similar to the diversity measure mentioned in [56] that counts the average number of times landmarks appear in the hierarchical tree. If this measure is higher for one hierarchical tree structure than another, then there are more diverse routes in the former tree structure.
- **Average Branching Factor (ABF)** - The branching factor of the hierarchical tree at each level provides a measure of the spatial memory dispersion. If a hierarchical tree has large branching factors on average, then the mental clustering of space in a cognitive map is more dispersed in general.

	Hierarchical Tree Structures				
	Baseline	**Shallow**	**Balanced**	**Deep**	**Proposed Algorithm**
ALD	3.59	1	2	7.5	**3.73**
ABF	3.67	14	3.87	1	**3.5**

TABLE 5: We illustrate the distance between the proposed hierarchical model of the cognitive map with the baseline (i.e., ordered tree generated from participants' recall data), two extreme tree structures (i.e., shallow and deep), and one ideal tree structure in terms of two metrics: Average Landmarks Depth (ALD) and Average Branching Factor (ABF). We report the average value of two exploration directions over 20 minutes of exploration.

In Table 5, we highlight the statistics of the proposed similarity measures for the hierarchical tree structure from the proposed algorithm and ordered trees from participants' recall data. The results highlight the closeness in the hierarchical structure generated from the proposed algorithm with the baseline.

7.3 Sequence Order Effect

To investigate the impact of landmark visit order on the agent's wayfinding performance, we examine the trajectories generated by the FTC model in combination with the distortion caused by the sequence order effect (i.e., FTC+SE). The same data analysis protocol was followed as for the human participant data, as discussed in Section 6.5.3. An independent-samples, two-tailed t-test was conducted to test for differences between 10-minute and 20-minute exploration duration conditions for each navigation direction. Overall, agents who were trained from East-to-West demonstrated a significant difference in the wayfinding trajectories compared to the agents when trained from West-To-East navigation order, t(48) = -2.51, p = .019. In Figure 7 (bottom), we visualize the agent trajectories for two different navigation directions. Similar to humans, we notice a significant rightward and leftward bias in the trajectories generated under East-To-West and West-To-East conditions, respectively. The proposed cognitive

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

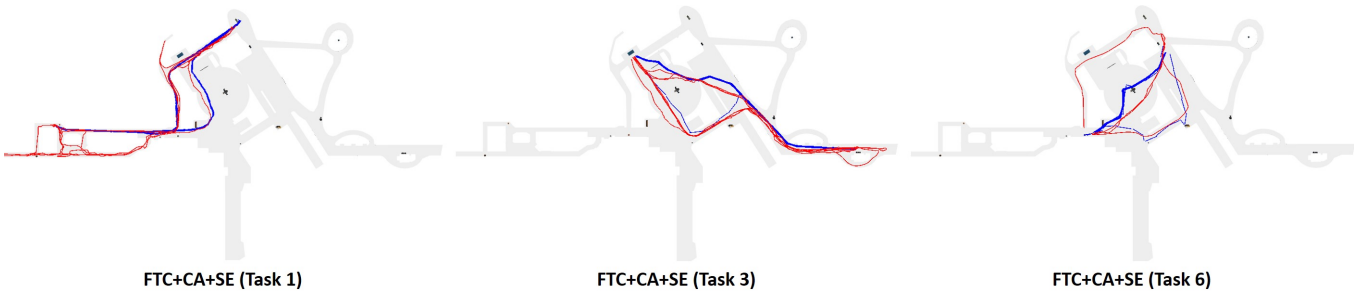

Fig. 8: (a,b,c) Comparison between participants' trajectories (red) and the agent's trajectories (blue) produced using FTC+CA+SE (20 minutes exploration) for three navigation tasks under both (East-To-West and West-To-East exploration direction conditions.

	Exploration Duration	
	10 minutes + 3 Navigation tasks	20 minutes + 3 Navigation tasks
Algorithm	**Avg. Path Edit Distance**	**Avg. Path Edit Distance**
Human Participants	N.A.	N.A.
FTC	2.25	1.80
FTC + CA	2.43	1.97
FTC + SE	**2.23**	**1.76**
FTC + CA + SE	2.29	1.80

TABLE 6: The table showcases results for different combinations of FTC in comparison with participants' data (baseline) collected using the VR data for 10 or 20 minutes exposure to the environment.

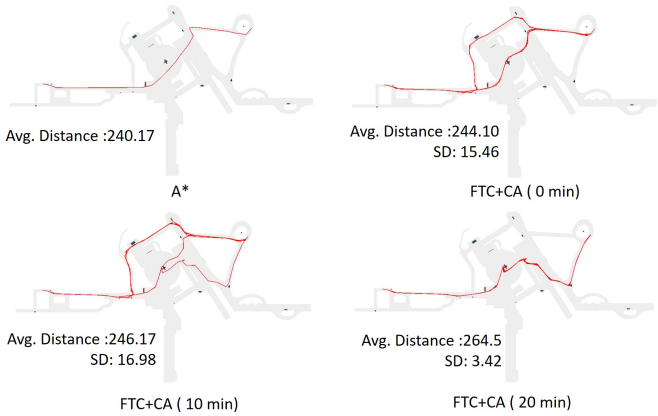

Fig. 9: Visualization of agents' trajectories to showcase the impact of memory decay for three elapsed-time intervals (i.e., 0, 10, and 20 minutes) compared with a standard A* based shortest path.

path-planning model could successfully replicate the sequence order effect as observed in humans and produce a strong positive correlation between the wayfinding trajectories and navigation direction.

7.4 Impact of Exploration Duration

To examine the impact of exploration durations on the development of the DHCG, we examined agents' performance on six wayfinding tasks after 10 and 20 minutes of exploration. The average path edit distance (PED) for different versions of the FTC are presented in Table 6. We notice a significant reduction (˜21%) in the PED between the human trajectories and agents trajectories produced after 10 and 20 minutes of exploration. After longer exploration of twenty minutes, the reduction of PED indicates some convergence between human and agent trajectories. This may be because agents, similar to humans, could manage to find shortcuts between two locations. The closest match to the participants' trajectories was produced using the sequence order effect (FTC+SE).

Trajectories for three randomly selected navigation tasks for both participants and agents, which employed the FTC+CA+SE strategy, after 20 minutes of exploration are visualized in Figure 8 (a-c). The route selections at multiple decision points for all three tasks were similar for both agent and human participants, which is evident by the overlap of blue and red trajectories. We notice that for the same wayfinding task (e.g., Task 1c), two different paths are observed by both human participants and cognitive agents. This is due to the SE distortion caused by the two different navigation directions (EW and WE) during training. For Task 3 in particular (Figure 8 (b)), the average path edit distance between the participant and agent trajectories is 1, which is near perfect. The difference in average traversed distance (13.86 meters) and average travel time (4.92 seconds) over six tasks further evidence the similarities between human and agent trajectories.

7.5 Distortion Effect: Categorical Adjustment

To validate the impact of Categorical Adjustment (CA), we examine the trajectories generated by employing the FTC model in combination with the distortion caused by CA (i.e., FTC+CA) for three elapsed time of 0, 10, and 20 minutes after exploring/learning the environment for 20 mins. At a time interval of 0 mins, there is no memory decay on DHCG, and the landmark salience value is high. The saliency value of landmarks (i.e., nodes in DHCG) reduces with elapsed time as per Section 4.3. We highlight the trajectories generated by employing an FTC+CA in Figure 9 compared to A* trajectories. Trajectories generated by FTC+CA (0 mins) are distinctively closer to the trajectories generated by A* in terms of traversed distance and landmarks traversed at key decision points. Thus, it highlights the fact that in the presence of strong memory (i.e., minimal distortion in landmarks location), the proposed path-planning algorithm manages to find the shortest path between the source and the destination, often observed in human wayfinding under strong spatial memory conditions. Moreover, we notice a gradual increase in traversed distance after 10 and 20 minutes of memory decay. In Figure 9, we distinctly notice variations in trajectories due to the distortion accumulating over time. The CA effect causes these variations in trajectories during the spatial recall. Lower salience values for recalled landmarks are associated with more bias towards the cluster center and thus more

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

changes in the nodes and edges traversed between the source and destination nodes during path-planning on DHCG.

8 Simulation Results on a Manhattan Grid

We elect a uniform grid environment with multiple regular intersections to exemplify the spatial distortion biases in route selection at decision points. A 3D model of a 10×10 Manhattan grid was used to conduct these simulations (Figure 10). One start location and ten destinations were preselected. We compared the path-planning results between A* and FTC and its variants to demonstrate the influence of each type of distortion on agent wayfinding behavior. The parameters $S = 10^5$ (memory strength) and $\alpha = 0.5$ (preference of path distance over memory) were used to generate all simulation results. In Figure 10 and Table 7, we present qualitative and quantitative results respectively.

In the scenario of East-To-West navigation direction by an agent, we observe an increase in average total distance traveled for 10 wayfinding tasks in comparison to the shortest path produced by A* (i.e., 117.56). The different combinations of the proposed cognitive path-planning algorithm, ranked in terms of average total distance and path edit distance, were FTC (avg dist = 131.14, PED = 10.8), FTC+CA (avg dist = 144.23, PED = 10.9), FTC+SE (avg dist = 134.03, PED = 9.3), and FTC+CA+SE (avg dist = 150.38, PED = 10.6). Both CA and SE distortion strategies increased the path distance compared to A* and FTC (i.e., complete knowledge and no distortion). The longest distances traveled were observed in FTC+CA+SE. Moreover, the average total path edit distance of 10.9 is observed the most for the FTC+CA combination. The visualized trajectories in Figure 10 also highlight the variations in route choices qualitatively. The trajectories produced by AI-based path-planning (A* in this case) fail to produce realistic paths and produce fixed and similar-looking paths. Recently, inspired by the natural path-creation process [57] proposed a method to efficiently supervise the path of NPC in an interactive virtual environment such as a game or VR by estimating a weight map and path similarity based on the user's path. The goal was to provide a new route to the NPC by referring to the user's movement trajectory rather than a fixed path. In our model, we achieve similar goals of human-player in VR-like path without relying on the human-player path dataset. In Figure 10 we notice the variation in trajectories with different combinations of FTC and distortions. A game-level designer can employ different distortion types as a plug-and-play feature to FTC to mimic varied human movement behavior. Moreover, more distortions can be investigated and formulated to extend our proposed path-planning algorithm in the future. Figure 12 demonstrates the difference in the formulation of DHCG due to navigation pattern. The variation in cluster formation is caused by features F (i.e., frequency of visit between adjacent nodes) and R (i.e., the last visited time of nodes). Path distance and visibility features remain the same in both scenarios.

8.1 Distortion due to Category Adjustment

Figure 11 (left) presents the difference in paths using FTC and FTC+CA for one navigation task. An important measure that decides the level of distortion in a cluster is the memory (i.e., landmarks salience) of the visited nodes for that cluster. In the East-To-West exploration direction, the nodes marked as zone A were visited earlier than the nodes marked as Zone B, resulting in higher salience for zone B nodes due to recency effect [58]. This is similar to how humans remember the landmarks strongly if visited last. Due to higher memory values, nodes on the right side of the grid may be more likely to become clustered together and thus distort the center of the abstract nodes towards zone B nodes at level two of the DHCG. As a consequence, the agent selected to make a late turn in the presence of CA-based distortion. The same behavior occurs again at the later stage of the path-planning (marked as zone C).

		Left to right exploration			Right to left exploration		
Algorithm	**Navigation Tasks**	**Distance**	**Path Edit Distance**	**Turns**	**Distance**	**Path Edit Distance**	**Turns**
A*	1	87.6	N.A.	0	87.62	N.A.	0
	5	121.6	N.A.	4	121.6	N.A.	4
	10	161.5	N.A.	4	161.5	N.A.	4
FTC	1	119	11	5	118.3	9	8
	5	129.2	8	4	152.3	13	4
	10	167.3	13	6	170.6	8	4
FTC+CA	1	120.3	11	5	120.9	9	8
	5	147.7	9	4	154.2	13	4
	10	175.5	14	5	172.9	8	4
FTC+SE	1	123.1	9	4	87.7	0	0
	5	131.6	10	4	135.7	11	4
	10	172.5	14	7	174.4	15	4
FTC+CA+SE	1	123.8	12	5	124	9	8
	5	133.1	8	6	157.3	13	4
	10	192.7	14	5	176.8	8	4

TABLE 7: Comparison of various metrics for three navigation tasks for various combinations of the proposed cognitive path-planning algorithm on Manhattan grid. Two opposite exploration directions were used to build the agent's cognitive map.

8.2 Distortion due to Sequence Order Effect

Figure 11 (right) highlights the difference in paths using FTC and FTC+SE for one navigation task. The landmarks (i.e., nodes) visited earlier have a more substantial impact on the cluster center due to primacy effect [58], causing it to be distorted towards the initially visited nodes as per equation 1. In the West-To-East exploration direction, nodes marked as zone E were visited earlier than the nodes marked as zone D, which were visited earlier in East-To-West exploration. Thus, we see a series of distortions towards the left and right side of the Manhattan grid over the entire path for the wayfinding trajectories generated by the agent who explored the environment in East-To-West and West-To-East directions, respectively.

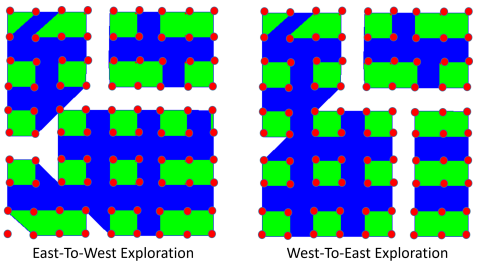

Fig. 12: Comparison between clusters formed using East-To-West and West-To-East exploration direction on a Manhattan grid. Level 1 landmarks are shown as point in red, and Level 2 and 3 clusters are visualized in green and blue, respectively.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

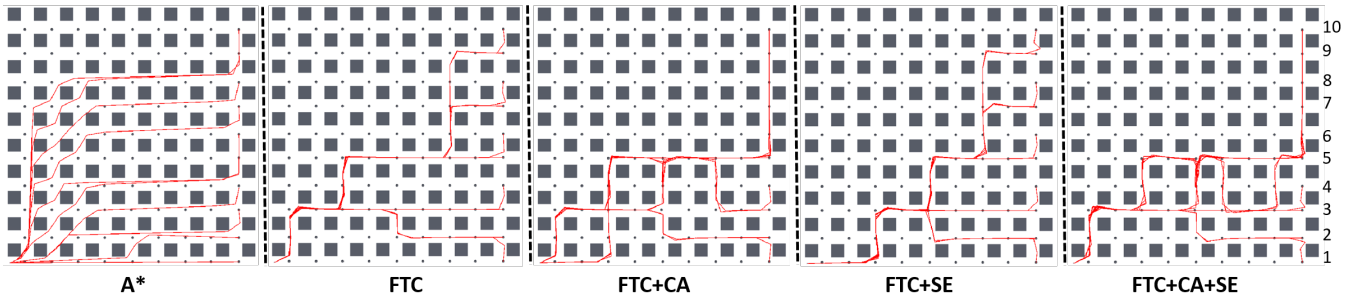

Fig. 10: Visualization of agent's trajectories using various combinations of FTC with distortions. A* based path-planning is used for comparison. Agent's start location is at bottom left corner of the grid. Trajectories from ten different navigation tasks were recorded. The destination locations are marked using 1 to 10.

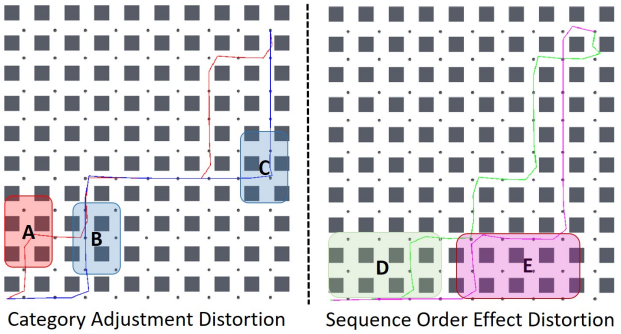

Fig. 11: Visualization of distortion in path-planning caused by category adjustment and sequence effect. (left) Agent's path computed using FTC (red) and FTC + CA (blue) for navigation task 10. (right) Agent's path computed with sequence order effect (FTC+SE) after learning the environment in East-To-West (green) and West-To-East (magenta) exploration directions.

9 Conclusion & Future Works

Unlike prior static and comprehensive representations of environment-based path-planning models, this paper proposes a computational model of cognitive path-planning to simulate human-like wayfinding behavior, which incorporates spatial memory distortions. The proposed model can produce a unique/subjective mental representation of the environment, encode the spatiotemporal sequence effects of landmark visits, and generate paths based on memory decay and approximation. A desktop-based VR experiment is performed with 41 participants to parameterize and validate the proposed model. There are three significant contributions of this work. (1) We propose a Dynamic Hierarchical Cognitive Graph representation (DHCG) to encode the environment structure based on the agent's exploration pattern that models a subjective mental representation of an environment. (2) We extend the most commonly employed wayfinding heuristic FTC to accommodate CA and SE distortions while systematically varying memory decay to formulate a novel cognitive path-planning model. (3) We ground the parameterization of the proposed data structure on behavioral data collected from a VR-based study. The results obtained in our simulations showcase the effectiveness of the proposed model for closely replicating human-like route decisions and overall trajectories based on identical exploration patterns. The proposed cognitive path-planning model is sufficiently fast for real-time path-planning in a large and complex environment. We believe that the proposed model's results are the first steps towards building a realistic human-like wayfinding model that captures spatial memory distortion in route selection during path-planning.

Limitations. Each agent must maintain its own independent spatial memory, which both increases the computational complexity and requires frequent updates (i.e., each navigational decision point/intersection) for acquiring and decaying information.

Future Works. We would specifically like to understand and compute the memory decay rate on spatial memory with elapsed time. Most theories in spatial navigation treat path integration and cognitive map as two distinct and detached mechanisms for human navigation. We hypothesize that path integration may play a role in the formulation of cognitive maps. In the future, we will like to investigate this underlying mechanism, if any.

10 Acknowledgements

The research was conducted at the Future Cities Laboratory at the Singapore-ETH Centre: FI 370074016. The research was supported in part by NSF awards: IIS-1703883, IIS-1955404, and IIS-1955365. This work was also partially supported through funding by the Leonhard Obermeyer Center at the Technical University of Munich.

Ethical Approval: All procedures performed in studies involving human participants were in accordance with the ethical standards of the ethical committee of ETH Zurich (EK 2019-N-79) and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

References

[1] T. Quandt and S. Kröger, *Multiplayer: The social aspects of digital gaming*. Routledge, 2013.

[2] M. Kapadia, N. Pelechano, J. Allbeck, and N. Badler, "Virtual crowds: Steps toward behavioral realism," *Synthesis lectures on visual computing: computer graphics, animation, computational photography, and imaging*, vol. 7, no. 4, pp. 1–270, 2015.

[3] I. Karamouzas, N. Sohre, R. Narain, and S. J. Guy, "Implicit crowds: Optimization integrator for robust crowd simulation," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–13, 2017.

[4] P. M. Kielar, D. H. Biedermann, A. Kneidl, and A. Borrmann, "A unified pedestrian routing model for graph-based wayfinding built on cognitive principles," *Transportmetrica A: transport science*, vol. 14, no. 5-6, pp. 406–432, 2018.

[5] S. Lee, M. Park, K. Lee, and J. Lee, "Scalable muscle-actuated human simulation and control," *ACM Transactions On Graphics (TOG)*, vol. 38, no. 4, pp. 1–13, 2019.

[6] A. Botea, M. Müller, and J. Schaeffer, "Near optimal hierarchical path-finding," *Journal of game development*, vol. 1, no. 1, pp. 7–28, 2004.

[7] L. van Elswijk, I. Sprinkhuizen-Kuyper, and F. Wiedijk, "Hierarchical path-finding theta," 2013.

[8] S. C. Hirtle and J. Jonides, "Evidence of hierarchies in cognitive maps," *Memory & cognition*, vol. 13, no. 3, pp. 208–217, 1985.

[9] S. Nayak, V. Mishra, and A. Mukerjee, "Towards a cognitive model for human wayfinding behavior in regionalized environments," in *2011 AAAI Fall Symposium Series*, 2011.

[10] H. Voicu, "Hierarchical cognitive maps," *Neural Networks*, vol. 16, no. 5-6, pp. 569–576, 2003.

[11] E. C. Tolman, "Cognitive maps in rats and men." *Psychological review*, vol. 55, no. 4, p. 189, 1948.

[12] M. Peer, I. K. Brunec, N. S. Newcombe, and R. A. Epstein, "Structuring knowledge with cognitive maps and cognitive graphs," *Trends in Cognitive Sciences*, 2020.

[13] E. R. Chrastil and W. H. Warren, "From cognitive maps to cognitive graphs," *PloS one*, vol. 9, no. 11, 2014.

[14] T. P. McNamara, "Knowledge representation," in *Thinking and problem solving*. Elsevier, 1994, pp. 81–117.

[15] E. K. Sadalla, L. J. Staplin, and W. J. Burroughs, "Retrieval processes in distance cognition," *Memory & Cognition*, vol. 7, no. 4, pp. 291–296, 1979.

[16] R. Briggs, "Urban cognitive distance," *Image and environment: Cognitive mapping and spatial behavior*, pp. 361–388, 1973.

[17] P. W. Thorndyke, "Distance estimation from cognitive maps," *Cognitive psychology*, vol. 13, no. 4, pp. 526–550, 1981.

[18] D. B. Haun, G. L. Allen, and D. H. Wedell, "Bias in spatial memory: A categorical endorsement," *Acta Psychologica*, vol. 118, no. 1-2, pp. 149–170, 2005.

[19] B. Tversky, "Distortions in cognitive maps," *Geoforum*, vol. 23, no. 2, pp. 131–138, 1992.

[20] N. H. Feldman, T. L. Griffiths, and J. L. Morgan, "The influence of categories on perception: explaining the perceptual magnet effect as optimal statistical inference." *Psychological review*, vol. 116, no. 4, p. 752, 2009.

[21] L. E. Crawford and S. Duffy, "Sequence effects in estimating spatial location," *Psychonomic Bulletin & Review*, vol. 17, no. 5, pp. 725–730, 2010.

[22] J. Huttenlocher, L. V. Hedges, and S. Duncan, "Categories and particulars: Prototype effects in estimating spatial location." *Psychological review*, vol. 98, no. 3, p. 352, 1991.

[23] J. M. Wiener and H. A. Mallot, "'fine-to-coarse'route planning and navigation in regionalized environments," *Spatial cognition and computation*, vol. 3, no. 4, pp. 331–358, 2003.

[24] B. Kuipers, "Modeling spatial knowledge," *Cognitive science*, vol. 2, no. 2, pp. 129–153, 1978.

[25] N. A. Schmajuk and A. Thieme, "Purposive behavior and cognitive mapping: a neural network model," *Biological cybernetics*, vol. 67, no. 2, pp. 165–174, 1992.

[26] M. E. Jefferies and W.-K. Yeap, *Robotics and cognitive approaches to spatial mapping*. Springer, 2008.

[27] P. Beeson, J. Modayil, and B. Kuipers, "Factoring the mapping problem: Mobile robot map-building in the hybrid spatial semantic hierarchy," *The International Journal of Robotics Research*, vol. 29, no. 4, pp. 428–459, 2010.

[28] B. Kuipers, R. Browning, B. Gribble, M. Hewett, and E. Remolina, "The spatial semantic hierarchy," *Artificial intelligence*, 2000.

[29] C. K. Wong, J. Schmidt, and W. K. Yeap, "Using a mobile robot for cognitive mapping." International Joint Conferences on Artificial Intelligence, 2007.

[30] S. S. Sohn, S. DeStefani, and M. Kapadia, "Dynamic cognitive maps for agent landmark navigation in unseen environments," in *Proceedings of the 11th Annual International Conference on Motion, Interaction, and Games*. ACM, 2018, p. 14.

[31] A. Stevens and P. Coupe, "Distortions in judged spatial relations," *Cognitive psychology*, vol. 10, no. 4, pp. 422–437, 1978.

[32] T. P. McNamara, "Mental representations of spatial relations," *Cognitive psychology*, vol. 18, no. 1, pp. 87–121, 1986.

[33] T. Kohonen, P. Lehtiö, J. Rovamo, J. Hyvärinen, K. Bry, and L. Vainio, "A principle of neural associative memory," *Neuroscience*, vol. 2, no. 6, pp. 1065–1076, 1977.

[34] S. M. Graham, A. Joshi, and Z. Pizlo, "The traveling salesman problem: A hierarchical model," *Memory & cognition*, vol. 28, no. 7, pp. 1191–1204, 2000.

[35] S. Timpf, *Hierarchical structures in map series*. Department of Geoinformation, Technical Univ., 1998.

[36] T. P. McNamara, "Theories of priming: I. associative distance and lag." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 18, no. 6, p. 1173, 1992.

[37] R. G. Golledge *et al.*, *Wayfinding behavior: Cognitive mapping and other spatial processes*. JHU press, 1999.

[38] D. R. Montello, "Cognitive map-design research in the twentieth century: Theoretical and empirical approaches," *Cartography and Geographic Information Science*, vol. 29, no. 3, pp. 283–304, 2002.

[39] R. A. Epstein, E. Z. Patai, J. B. Julian, and H. J. Spiers, "The cognitive map in humans: spatial navigation and beyond," *Nature neuroscience*, vol. 20, no. 11, p. 1504, 2017.

[40] J.-Y. Lee and W. Yu, "A coarse-to-fine approach for fast path finding for mobile robots," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2009, pp. 5414–5419.

[41] C. Zhong, S. Liu, Q. Lu, B. Zhang, and S. X. Yang, "An efficient fine-to-coarse wayfinding strategy for robot navigation in regionalized environments," *IEEE transactions on cybernetics*, vol. 46, no. 12, pp. 3157–3170, 2016.

[42] L. Tai and M. Liu, "Towards cognitive exploration through deep reinforcement learning for mobile robots," *arXiv preprint arXiv:1610.01733*, 2016.

[43] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, and J. Malik, "Cognitive mapping and planning for visual navigation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2616–2625.

[44] H. Li, Q. Zhang, and D. Zhao, "Deep reinforcement learning-based automatic exploration for navigation in unknown environment," *IEEE transactions on neural networks and learning systems*, vol. 31, no. 6, pp. 2064–2076, 2019.

[45] T. Madl, S. Franklin, K. Chen, R. Trappl, and D. Montaldi, "Exploring the structure of spatial representations," *PloS one*, vol. 11, no. 6, p. e0157343, 2016.

[46] R. K. Dubey, S. S. Sohn, T. Thrash, C. Hoelscher, and M. Kapadia, "Identifying indoor navigation landmarks using a hierarchical multi-criteria decision framework," in *Motion, Interaction and Games, MIG 2019, Newcastle upon Tyne, UK, October 28-30, 2019*, 2019, pp. 15:1–15:11. [Online]. Available: https://doi.org/10.1145/3359566.3360066

[47] A. Solway, B. B. Murdock, and M. J. Kahana, "Positional and temporal clustering in serial order memory," *Memory & cognition*, vol. 40, no. 2, pp. 177–190, 2012.

[48] P. Barrouillet, A. De Paepe, and N. Langerock, "Time causes forgetting from working memory," *Psychonomic Bulletin & Review*, vol. 19, no. 1, pp. 87–92, 2012.

[49] S. Lewandowsky, K. Oberauer, and G. D. Brown, "No temporal decay in verbal short-term memory," *Trends in cognitive sciences*, vol. 13, no. 3, pp. 120–126, 2009.

[50] E. M. Altmann and C. D. Schunn, "Decay versus interference: A new look at an old interaction," *Psychological Science*, vol. 23, no. 11, pp. 1435–1437, 2012.

[51] W. Zhang and S. J. Luck, "Sudden death and gradual decay in visual working memory," *Psychological science*, vol. 20, no. 4, pp. 423–428, 2009.

[52] S. Koenig and M. Likhachev, "D^* lite," *Aaai/iaai*, vol. 15, 2002.

[53] B. B. Murdock, "Todam2: a model for the storage and retrieval of item, associative, and serial-order information." *Psychological review*, vol. 100, no. 2, p. 183, 1993.

[54] J. S. Reitman and H. H. Rueter, "Organization revealed by recall orders and confirmed by pauses," *Cognitive Psychology*, vol. 12, no. 4, pp. 554–581, 1980.

[55] I. Jeantet, Z. Miklós, and D. Gross-Amblard, "Overlapping hierarchical clustering (ohc)," in *Advances in Intelligent Data Analysis XVIII*, M. R. Berthold, A. Feelders, and G. Krempl, Eds. Cham: Springer International Publishing, 2020, pp. 261–273.

[56] C.-T. Kuo and I. Davidson, "On the equivalence of tries and dendrograms-efficient hierarchical clustering of traffic data," *arXiv preprint arXiv:1810.05357*, 2018.

[57] J.-H. Kim, J. Lee, and S.-J. Kim, "Navigating non-playable characters based on user trajectories with accumulation map and path similarity," *Symmetry*, vol. 12, no. 10, p. 1592, 2020.

[58] C. W. Mayo and W. H. Crockett, "Cognitive complexity and primacy-recency effects in impression formation." *The Journal of Abnormal and Social Psychology*, vol. 68, no. 3, p. 335, 1964.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TVCG.2022.3163794, IEEE Transactions on Visualization and Computer Graphics

Rohit K. Dubey received the Ph.D. degree in computer Science from ETH Zurich in October 2020. He is currently a postdoctoral researcher in the chair of Computational Modeling and Simulation at the Technical University of Munich. He is currently developing a cognitively grounded computational framework of microscopic pedestrian simulation that models the uncertainty and fusion of potentially multiple conflicting information sources for wayfinding decision-making. His research interests are spatial cognition, computer graphics, and reinforcement learning.

Samuel S. Sohn is a Ph.D. candidate in the Department of Computer Science at Rutgers University. His research has involved developing cognitively grounded models of human navigation in built environments, particularly focused on memory and perception. His current focus is on leveraging machine learning methods to accelerate crowd simulations well beyond their current limits in terms of computational speed.

Tyler Thrash was a postdoctoral researcher in cognitive science at ETH Zurich. His work focuses on spatial cognition and navigation from a largely ecological/Gibsonian perspective. With this approach, he attempts to explain higher-level cognition (e.g., biases in spatial memory) in terms of lower-level, perceptual processes (e.g., visual exposure to environmental structure). Dr. Thrash also develops simulations for interactive virtual environments and mathematical models for understanding human behavior.

Christoph Hölscher is Full Professor of Cognitive Science in the D-GESS at ETH Zürich since 2013, with an emphasis on Applied Cognitive Science. Since 2016 Christoph is a Principal Investigator at the Singapore ETH Center (SEC) Future Cities Laboratory, heading a research group on 'Cognition, Perception and Behaviour in Urban Environments'. Christoph is the Program Director of Future Resilient Systems FRS at the SEC since 2019, leading the current FRS 2 phase (2020-2025). He holds a PhD in Psychology from University of Freiburg, served as honorary senior research fellow at UCL, Bartlett School of Architecture, and as a visiting Professor at Northumbria University Newcastle. Christoph has several years of industry experience in Human-Computer Interaction and usability consulting. The core mission of his research groups in Zurich and Singapore is to unravel the complex interaction of humans and their physical, technical and social environment with an emphasis on cognitive processes and task-oriented behavior.

Andre Borrmann received the PhD degree from Technical University of Munich, Germany, in 2007. He is now a full professor with the Department of Civil and Environmental Engineering at Technical University of Munich, and speaker of the TUM Center of Digital Methods for the Built Environment. His research interests lie on the application of computational methods in all areas of construction, ranging from automating the design process of buildings, over monitoring and steering construction works using sensors and robots, to the optimized operation of facilities by simulating pedestrian dynamics.

Mubbasir Kapadia is the Director of the Intelligent Visual Interfaces Lab and an Associate Professor in the Computer Science Department at Rutgers University. Previously, he was an Associate Research Scientist at Disney Research Zurich and the Assistant Director of the Human Modeling Simulation Lab at University of Pennsylvania. Kapadia's research lies at the intersection of artificial intelligence, visual computing, and human-computer interaction, with a mission to develop intelligent visual interfaces to empower content creation for human-aware architectural design, digital storytelling, and serious games. He has published more than 100 journal and conference papers at premier venues in Computer Graphics, Computer Vision, and Artificial Intelligence. Kapadia's research is funded by DARPA and NSF, and through generous support from industrial partners including Disney Research, Autodesk Research, Adobe Research, and Unity Labs. He received his PhD in Computer Science at University of California, Los Angeles.