Gaze dynamics prior to navigation support hierarchical planning

Jeremy Gordon[1], John Chuang[2], and Giovanni Pezzulo[2,*]

[1]School of Information, University of California, Berkeley, California, USA
[2]Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy
[*]Corresponding author: giovanni.pezzulo@istc.cnr.it

May 31, 2025

Abstract

The task of planning future actions in the context of an uncertain world results in massive state spaces that preclude exhaustive search and other strategies explored in the domains of both human decision-making and computational agents. One plausible solution to this dimensionality explosion is to decompose the task into subgoals that match the information geometry of the task at hand. However, how individuals identify a productive hierarchy, and perceive and select subgoals suitable to planning, is not well understood. To investigate this topic, we designed a virtual-reality based behavioral experiment which collected eye movements during a pre-navigation planning phase. By capturing gaze dynamics correlated with the simulative processes used in planning, we were able to identify the spatiotemporal evolution of visual search under uncertainty. Our results highlight gaze dynamics indicative of a search process that exhibits hierarchical structure. These include a decreasing trend seen in gaze distance from origin and a broad to narrow shift (with reducing saccade distances and longer fixation durations) as plans are established. In line with prior work, critical tiles to which landscape connectivity is most sensitive were the strongest predictors of visual attention. We also find that deeper planning was correlated with success only on the most complex maps (e.g. those with a larger number of information-nodes, higher branching factor, and more forks, according to an info-graphical map analysis). This study highlights the role of embodied visual search during planning, and the skill-dependence of the specific subgoals and hierarchical decomposition used which unlocked successful performance.

Keywords: planning; spatial navigation; virtual reality; eye movements

Introduction

When planning future actions, e.g. finding the best route to a novel destination or the best sequencing for our daily tasks, we are often confronted with very large state spaces that we cannot search exhaustively. These exist among a class of problems we frequently encounter as humans which involves searching for a sequence of candidate actions likely to reliably result in a desired "goal" state, while also contending with limited (cognitive and motor) resources. Prospective search problems of this type are common in the realm of spatial navigation, but also apply to much more abstract domains of cognition (Newell and Simon, 1972; Hills et al., 2015; Pezzulo et al., 2018).

From a formal perspective, searching in large state spaces is different from searching in small ones that permit exhaustive analysis, and as a result, requires a different class of algorithms (Browne et al., 2012). Indeed, in discrete state spaces with known dynamics, it is possible to use simple decision trees that capture a complete set of possible action sequences, to identify paths resulting in a desired state, as well as an evaluation of their relative efficiency (e.g. by weighted length). However, in large state spaces (including tasks where states are continuous rather than discrete), the combinatorial nature of possible action sequences results in a dimensionality explosion that makes exhaustive search intractable. Even when loops are excluded (to avoid an infinite search space), the number of sequences to check exceeds the computational capabilities of both biological and artificial problem solvers. In machine learning, the problem of searching in large state spaces has lead to the development of various effective techniques to *sample* the most promising candidate actions (Browne et al., 2012). For example, Monte Carlo Tree Search (MCTS) algorithms, when coupled with deep neural networks, provide effective solutions to challenging challenging games such as chess and GO (Silver et al., 2016). In parallel, the identification of algorithms and heuristics organisms use to overcome these challenges has become a rich target of inquiry for cognitive science and psychology (Mattar and Lengyel, 2022; Miller and Venditto, 2021).

When faced with large search problems, living organisms (as well as artificial systems) must contend with the *breadth-depth dilemma*, in which the problem solver must decide how to direct limited resources (e.g. exploratory actions which consume time and energy). Typically, spending more time reasoning about (simulating) sequences within a particular region of a state space offers more precise estimates of its value (i.e. in pathfinding, the likelihood of particular positions being part of an efficient route to destination), at the cost of reduced precision in other regions. A number of studies have attempted to model and empirically measure human handling of this dilemma. In a study in which participants had to select between 121 options that provided different rewards, Wu et al. found that people sometimes adopt a sampling strategy similar to the ones used in machine learning and that they are able to use spatial structure (when present) to generalize effectively (Wu et al., 2018). In a problem solving study resembling the traveling salesman problem, Eluchans et al. found that people adapt their planning strategy to the minimal depth required by the problem (Eluchans et al., 2023). In a patch foraging-like experiment where the allocation of samples needed to be determined prior to sample feedback, Moreno et al. found that individuals prefer breadth when capacity is low, and depth when capacity is high (Moreno-Bote et al., 2020). Interestingly, even when sufficient capacity is available to sample from all states, optimal behavior involved ignoring

the majority of states in order to more deeply explore a subset. In a related study, the number of states sampled increased with capacity approximately according to a power law (with exponent $\approx \frac{3}{4}$) (Vidal et al., 2022).

Other studies have shown that when planning routes in novel mazes, people focus their attention and search resources on *salient* locations or landmarks, which provide more information gain. For example, during virtual navigation, people display more deliberative behavior (as indexed by a behavioral index called *vicarious trial-and-error* (VTE) and firstly studied by Tolman during rodent navigation (Tolman, 1948)) in locations of the maze where making optimal decisions required using more information (Santos-Pata and Verschure, 2018; Lancia et al., 2023). When people have access to the search space (e.g., a map of the maze) prior to navigation and can therefore form a prospective plan, they again focus their resources for encoding with high fidelity only the subset of map elements that are relevant for their (optimal) plan, therefore constructing a simplified mental representation of the problem (Ho et al., 2022). Similarly, another study suggests that people adaptively *compress* simulations of potential routes during prospective route planning (Arnold et al., 2016), using mechanisms that are potentially analogous to the (time-compressed) '(p)replay' of experience found in the rodent hippocampus that might support spatial planning (Buzsaki, 2019; Foster, 2017; Pfeiffer and Foster, 2013; Pezzulo et al., 2014, 2017; Ekman et al., 2023; De Cothi et al., 2022; Momennejad, 2020).

Another useful approach to search in large state spaces consists of decomposing the problem into smaller and more manageable sub-problems, thus using a hierarchical approach to planning as opposed to "flat" tree search. Various studies indicate that people use such hierarchical planning strategies when the environment permits hierarchical decomposition (Balaguer et al., 2016; Donnarumma et al., 2016; Schick et al., 2019; Fernandez Velasco et al., 2024; Correa et al., 2023). During navigation in a maze structured similarly to the "Hanoi tower" problem (Solway et al., 2014), participants faced with a choice between two paths of the same length preferred the path that crossed fewer community boundaries (as identified by graph theoretic measures), which is "shorter" as quantified by information-theoretic measures that reflect the hierarchical clustering of space (Donnarumma et al., 2016; McNamee et al., 2016). Another study reported that people spontaneously organizes space into clusters that support hierarchical planning and subgoaling (Tomov et al., 2020). A neuroimaging study suggested that when planning routes in a virtual subway network, people exploit the problem structure—and in particular the division into lines and stations—to form hierarchically organized plans (Balaguer et al., 2016). Interestingly, besides hierarchical problem decomposition, individuals often exhibit sensitivity to the structure of the environment in their planning and navigation strategies. For example, a large scale virtual navigation study showed that people are better at wayfinding in mazes having the same general structure as the place where they grew up (Coutrot et al., 2022).

Many of the studies reviewed above measure navigation behavior to infer planning strategies. However, planning is a largely covert process, implying that key dynamics may be missed by studying only wayfinding decisions. An alternative approach to study planning in large state spaces consists of recording eye movements while people visually scan the environment prior to beginning navigation. Given its flexibility and speed, the visual system affords a rapid virtual simulation of

candidate routes prior to actual movement. Therefore, looking at eye movements may provide privileged access to the covert stages of planning (Land et al., 1999; Hodgson et al., 2019). Along these lines, in a recent study of visual search prior to a navigation task, Zhu et al. showed that participants' eyes made fast 'sweeps' that traced the trajectories that participants later selected, and that gaze targets balanced fixations on goal locations and critical transitions (walls or openings) in the environment—findings which resemble results in neural replay (Zhu et al., 2022). Other studies of eye movements during movement preparation showed that visual attention was selectively allocated to goal-related information, such as the location of movement targets (Baldauf et al., 2006) and the planned trajectory around obstacles (Baldauf, 2018). Another study of visual search in humans has shown that eye movements themselves exhibit behavior consistent with multi-step planning (in this case, planning of a gaze sequence matched to the shape of a given search space), rather than greedy selection of informationally rich or otherwise salient locations (Hoppe and Rothkopf, 2019). Multi-step planning of eye movements has been reported also during tasks requiring participants to fixate sequence of stationary targets (Zingale and Kowler, 1987) and to find paths in a grid that collect multiple gems without passing through the same node twice (Eluchans et al., 2024). In a primate study, Lakshminarasimhan et al. showed that eye movements could be used to infer latent beliefs such as the location of a hidden goal during virtual spatial navigation (Lakshminarasimhan et al., 2020). Controlling fixations also showed detrimental effects on performance, indicating the involvement of free visual search in effective navigation.

Taken together, this literature indicates that when planning in large state spaces that cannot be searched exhaustively, people (and other mammals) use a range of effective strategies: they allocate limited resources by privileging information-rich states, they attend to the most salient aspects of the problem, and they split problems into more manageable sub-problems. These and other strategies can be interpreted in terms of a rational use of limited (bounded) resources (Simon, 1997; Bhui et al., 2021; Callaway et al., 2018,?; Lieder and Griffiths, 2020; Gershman et al., 2015; Callaway et al., 2022). However, despite these achievements, we still lack a comprehensive account of the planning strategies that people adopt when confronted with the exploration-exploitation dilemma inherent to visual exploration of large, novel and uncertain environments. Interestingly, while behavioral studies have found signatures of hierarchical planning during navigation, we still do not know how these hierarchical navigation plans might be established during visual exploration prior to movement.

In this work, we address three main questions. First, we ask whether gaze dynamics prior to movement show signatures of hierarchical planning and subgoaling – perhaps as a way to establish a hierarchical navigation plan. If planning resembles a flat search in a decision tree, as in MCTS, we should see a series of broader-to-finer-grained series of forward gazes, from the start to the potential goal locations. If instead planning is hierarchical, we should find that people first gaze and resolve uncertainty at the highest level (e.g., about goal and subgoal locations) and subsequently focus at the lower level, to select specific routes from one subgoal to the next – or backward, from the final goal to the intermediate subgoals to the start location. Second, in case of hierarchical planning, we ask how people select subgoal locations – and whether subgoal selection reflects critical transition points and information-theoretic principles of optimal plan decomposition, as suggested by

theoretical models (Donnarumma et al., 2016; McNamee et al., 2016; Solway et al., 2014; Lakshminarasimhan et al., 2024). Third, we ask which of these behavioral patterns and planning dynamics led to success planning under uncertainty. That is, we aim to answer the question: what sets top-performing planners apart?

To investigate these questions, we designed a novel Virtual Reality-based spatial navigation task where planning is explicitly required, and collected eye tracking data while participants searched the landscape prior to navigating through it. In contrast to standard wayfinding experiments in which only one goal location is present, the maps were designed with multiple candidate goal locations ("caches"), and participants did not know in advance which ones would be filled in with actual goals. This uncertainty is intended to increase planning demands, and align our task with the inherent uncertainty found within most real world planning contexts.

Methods

Participants

Participants were undergraduates, graduate students and staff at a university in the United States. Participants were invited by email to complete a screening questionnaire. We excluded participants who were deemed high risk for adverse effects of using VR. This includes those who have used VR in the past and experienced severe dizziness or nausea during usage, as well as those with epilepsy, a history of seizures, heart ailments, or those who respond that they are prone to motion sickness. Participants who are left handed or ambidextrous were also excluded in order to ensure standardized analysis of controller motion in the dominant hand. Participants were also required to be 18 years old or older, and speak English (native English was not a requirement).

Respondents qualifying for the study were then invited to participate in the main in-person study.

After giving informed consent, 40 participants began the study, but 6 had to stop early due to motion sickness. In total, 34 participants completed the study, and form the sample for this analysis. Mean participant age was 21.3 $\pm$ 3.0. 75% were female, 21% were male, and 4% declined to state. 62% identified as Asian, 21% as White, 3% as Black or African American, and 8% identified as multiple races or Other.

Experimental design

The study used a single-condition within subjects design. Maps were procedurally generated (see Supporting information for details), and balanced between high uncertainty (2 true goals among 6 possible goal locations), and low uncertainty (2 true goals among 3 possible goal locations). The same group of 33 maps was shown to all participants. The order of the first 30 (the procedurally generated set) was randomized uniformly. The final 3 hand-designed maps were shown at the end of the study, in the same order for all participants. See Fig 1 for an overview of study design.

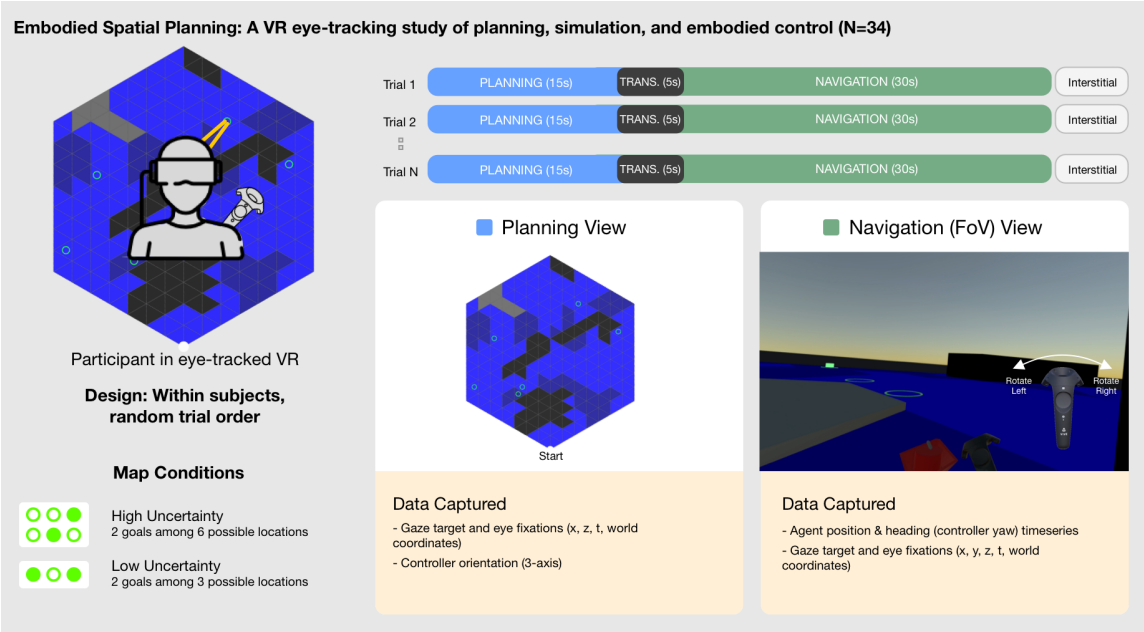

Figure 1: Schematic of experiment design. Participants moved through a series of trials, each split into a planning phase (15s), in which they viewed the map top down, and a navigation phase (30s), in which they moved through the environment with the goal of finding each of two goals. A 5-second transition phase marked the switch between the two main phases. Raw and target-aware eye-tracking data was captured during all three phases. Traces of agent location and orientation were captured during the navigation phase. Each trial presented a unique procedurally generated map in one of two conditions (high or low uncertainty, implemented by the number of goal cache locations).

Apparatus

An HTC Vive Pro Eye head mounted display (HMD), which was preconfigured with two Tobii eye-tracking cameras, was the primary apparatus supporting this study. The HMD was driven by an Alienware laptop running experiment scenes in Unity. A single HTC Vive controller was given to each participant to hold in the dominant hand. A "room-scale" setup with two lighthouses positioned on either side of the participant's seat was used.

Task & procedure

Prior to beginning the experiment, participants were briefed, signed an informed consent form, and were fitted with the VR headset. An eye-tracking calibration exercise was performed prior to starting the experiment. Then, a series of 3 practice trials began. In each, a simple map was shown to the participant, to ensure familiarity with the interaction and the task. We explained the interaction, meaning of colored tiles, and task objective (to find and collect both goals before time expires) as the participant completed the three practice trials, answering any questions that were raised.

Following the three practice trials, the main experiment begins. Each trial is constituted by a unique map, and progresses in three phases. In the *planning phase* (15 seconds), participants see a top-down view of the landscape in front of them. This map view shows the starting location of a boat—an agent which the participant controls—at the bottom of the map, as well as the colored features of the landscape. Possible goal locations are shown as green rings. Tall obstacles are shown as black tiles, low obstacles (which can be seen over but not passed through) are shown as gray tiles, and blue and dark blue tiles indicate water. During the planning phase, the participant is able to view the map, but cannot yet take any navigation action. 3 chimes indicate the approach of the end of the planning phase. Next, in the *transition phase* (5 seconds), the map is removed and participants wait for the beginning of the next phase. In the third and final phase, the *navigation phase* (30 seconds), participants' view shifts and they are now situated in the landscape just behind the boat they control. They can look in any direction to explore the landscape from the agent's perspective. By rotating the controller around the vertical axis (yaw), the participant controls the direction of the boat. By holding down the trigger, the participant is able to move the boat forward in its present orientation. In this way, participants are able to navigate the landscape with the aim of reaching and collecting both goals in the map. If the boat contacts either type of obstacle, or falls off the edge of the map, the trial ends immediately. If both goals are collected, the trial ends early, several seconds after the second goal was collected. If both goals are not collected in time, the trial ends when the 30-second navigation timer is exhausted (which is again identified by a series of 3 chimes).

Following the end of each trial, an interstitial screen appears telling the participant that they can pull the trigger to proceed to the next trial. The study continues this way, until all trials are completed. At the end of the study, a screen appears providing feedback on the participant's performance. This screen summarizes the number of goals collected, as well as the bonus compensation earned based on the participant's performance. Base compensation was $20. Bonus was added as

follows. 0-20% of all possible goals: \$0; 21-40%: \$1; 41-60%: \$2; 61-80%: \$3; 81-100%: \$5. After having read this information, we helped the participant remove the head-mounted display.

Finally, participants respond to a brief survey using a laptop. When finished, we released the participant to receive their total compensation, and exit the study room.

The experimental protocol was approved by the Committee for Protection of Human Subjects, under CPHS protocol number 2021-06-14394. Study methods and provisional hypotheses were registered and can be viewed on the study's OSF project page (https://osf.io/5tacn/). A video capture of the point of view of a sample participant can be downloaded from the "Participant POV" folder on the project page.

Analyses performed

Performance metric and other dependent variables

The simplest indicator of performance is the score obtained during the experiment (equal to the number of goals collected). Our analysis indicates that most participants devised planning strategies that were effective for the goal collection task. Mean score overall was 1.50 ± 0.63. Mean score per map varied, and in some maps, no participants successfully collected both goals. Hence, to account for this diversity, we defined a more fine-grained metric of performance that also takes into account map difficulty, and spatial progress towards the second goal when not collected (see Table 1 for details).

Based on this latter performance metric, we split participants into three equally sized percentile-based performance group bins (bottom, mid, and top-performers), and compare bottom to top-performers. For these analyses, mid-level performers were ignored. The top-performing group show higher performance metrics on both 3 and 6-cache maps (bottom-performers: 3-cache map mean performance: 1.16, 6-cache map mean performance: 1.13. Top-performers: 3-cache map mean performance: 1.57, 6-cache map mean performance: 1.52).

Furthermore, to study the strategies followed by participants, we defined additional dependent variables, which consider the sum of eye movements, mean fixation duration, and the percent of forward and reverse 'sweeps' of eye movements across possible maze paths (see Table 1).

DV	Definition
Performance	Computed based on goals collected (n) as proportion of optimal goal collection (n_{opt}), seconds remaining (τ), and path-progress from last goal (g_0) to remaining goal cache location (g_1). Formally: $performance = \frac{n}{n_{opt}} + \frac{\tau}{\tau_{opt}} + \frac{\lvert Path(g_0, agent_f)\rvert}{\lvert Path(g_0, g_1)\rvert}$. Optimal values are defined by highest performing individual participant.
Sum of eye movements	$\sum_i \lvert f_i - f_{i-1}\rvert$
Mean fixation duration	Mean duration of fixations, defined as gaze target remaining on the same object for over 50ms.

Table 1: Description of key dependent variables.

Salience of map features during planning

To identify which features of the map were considered to be more salient during the planning phase, we calculated an aggregate gaze heatmap for each map, across all participants, by summing tile-level gaze data during planning. Two methods of aggregation were explored: one based on fixation counts, and the other on total fixation duration. In both cases the resultant aggregate was standardized as an N-dimensional probability distribution (where N is the number of tiles in the base map). Notably, these aggregate heatmaps are sequence agnostic, and therefore are interpreted as a map of static visual salience during planning.

Then, we mapped the flattened gaze distributions to two set of features of the maps. The first set of features comprises the five tile types: water, muddy water, land, wall or goal. The second set of features corresponds to 9 feature maps that we identified, based on the spatial and informational geometry of each map. A description of each feature map tested is included in Table 2. Furthermore, Fig 2 provides example feature maps computed for map 2, as well as aggregated participant gaze on this map.

Feature Map	Definition	Mean EMD
Line-of-sight (LoS)	Tiles valued by number of goals visible (linear access)	22.5
SP-LoS	Shortest-path tiles weighted by goals visible	17.46
Goal Distance (2)	Tiles valued as per distance to second-nearest goal	18.4
Short Paths	Shortest paths to each goal location	19.27
Goals	Goal locations and adjacent tiles	25.8
Path Thru 2	Shortest path through origin and two goals	25.7
Path Thru 3	Shortest path through origin and three goals	19.3
Critical (GD-Origin)	Change in goal-distance at origin (see "Goal Distance (2)" above) when tile and adjacent have passability toggled	23.8
Critical (GD-Map)	Change in sum of map-wide goal-distance when tile and adjacent have passability toggled	16.5
Critical (Shortest-Path)	Number of changes to short paths map when tile and adjacent have passability toggled	17.6

Table 2: Description of feature maps used in prediction of planning gaze. The third column shows the mean cross-map Earth Mover Distance (EMD) between the feature map and the map-aggregate gaze heatmaps. See Fig 6 for visual comparison of feature-map EMDs.

At the map level, we compared the distributional distance between each feature-map and the aggregate gaze distribution using the 2D Euclidean earth mover distance (EMD), using the POT python optimal transport library (Flamary et al., 2021). The EMD provides a symmetrical measure of how far probability weights must be transported to match two distributions. This measure was chosen based on the intuition that gaze targets farther from predicted gaze density in 2D coordinates indicate a worse fit between the generated agent dynamics and empirical data. Feature maps with lower mean distance (across all maps) therefore are interpreted as supplying greater predictive power over preferred gaze locations.

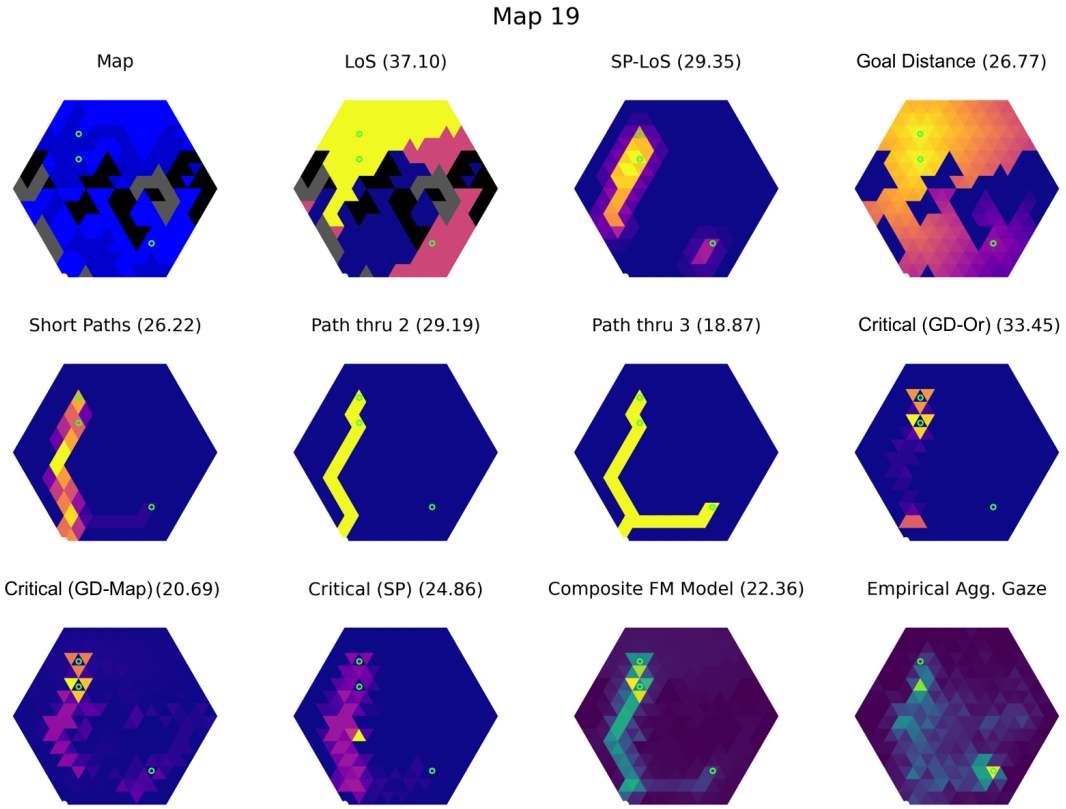

Figure 2: Sample feature and aggregated participant gaze for map 19. The first (top-left) figure shows map 19, in the format presented to participants. The subsequent 10 maps show different types of features, which comprise Line-of-Sight ("LoS"), Shortest-path Line-of-Sight ("SP-LoS"), distance to nearest goal ("Goal Distance"), Short Paths, Path through 2 goals, Path through 3 goals, Critical (via Goal-distance change at origin), Critical (via goal distance change map-wide), and Critical (via shortest path change). See Table 2 for details on each feature map. The mean of all feature maps is shown as the "Composite FM Model" in the second to last map. The last (bottom-right) figure shows the aggregated participants gaze for map 19. The numbers in parentheses above each map indicate the distributional distance between each feature-map and the aggregate gaze distribution using the 2D Euclidean earth mover distance (EMD).

Temporal gaze analysis and RQA

In addition to considering "static" gaze distribution, we extracted a number of temporal features from planning-time gaze data, including time-binned total saccadic distance (which approximates the total amount of eye movement during a trial), distance from origin, distance from nearest goal, multi-step sweeps, and a range of metrics derived from Recurrence Quantification Analysis (RQA).

RQA is an analytic method common in dynamical systems, and increasingly used to gain insights into temporal dynamics of eye tracking data (Anderson et al., 2013). We generated recurrence plots from the sequential tile fixation data during the planning phase, and extracted an array of common RQA metrics measuring attributes of the timeseries such as the likelihood of a fixation being

recurrent (recurrence rate, RR), the longest recurrent subsequence (L_{max}), and others. These were used as independent factors in the ANOVA analyses, to identify statistically significant differences correlated with both performance, as well as geometric map attributes.

Map analysis and classification

To address research questions about the influence of information and uncertainty-related attributes of particular maps on planning processes, we ran ANOVAs on a number of gaze-related attributes as dependent variables, against a list of map features which were hypothesized to be relevant to navigation planning. These included the length of the shortest path through 2, 3, and 4 goal cache locations (note that the last could only be computed for 6-cache maps), as well as a handful of metrics classified manually for each map which aimed to capture information access (e.g. where goal caches become visible in the landscape) and a simplified path-segment representation of each map. The coding process involved first producing a tree illustrating key segments and decision points, and then exhaustively considering goal cache observation outcomes (without knowledge of true goal locations).

We call the resultant representations "info-graphical," given their tree like graphical structure, and sensitivity to visual information availability in the landscape. For an example, consider map 7 shown on the right side of Fig 3. At the origin (as with all maps by construction), no goal cache locations are visible. Participants must remember to navigate to the left, and by the time they reach about the midpoint of their journey to the narrow corridor at the top of the map, they will have visual access to two of the goal caches (square info nodes). At planning time, they contend with two outcomes: either one present goal is observed, or both are. Both outcomes result in full goal location knowledge, but in the first case (bottom fork in info-graph), after collecting the first goal they'll need to continue around the low obstacle in the center of the map to collect the second. The target goal during this journey is fully observed, but some memory of the map structure is required to remember which way to navigate around the obstacle, and hence this latter segment is marked as memory-dependent (three hatch marks). Alternatively, if both closer goals are present at the info node, they can continue around and collect both with no significant memory dependence (un-hatched edge).

After producing info-graphical representations for all maps, these were used to calculate a number of key metrics, including: branching factor, number of forks, steps to certainty, and expected number of memory-dependent segments. See Fig 3 for sample info-graphical representations of selected maps, and Table 3 for a full list of attributes computed.

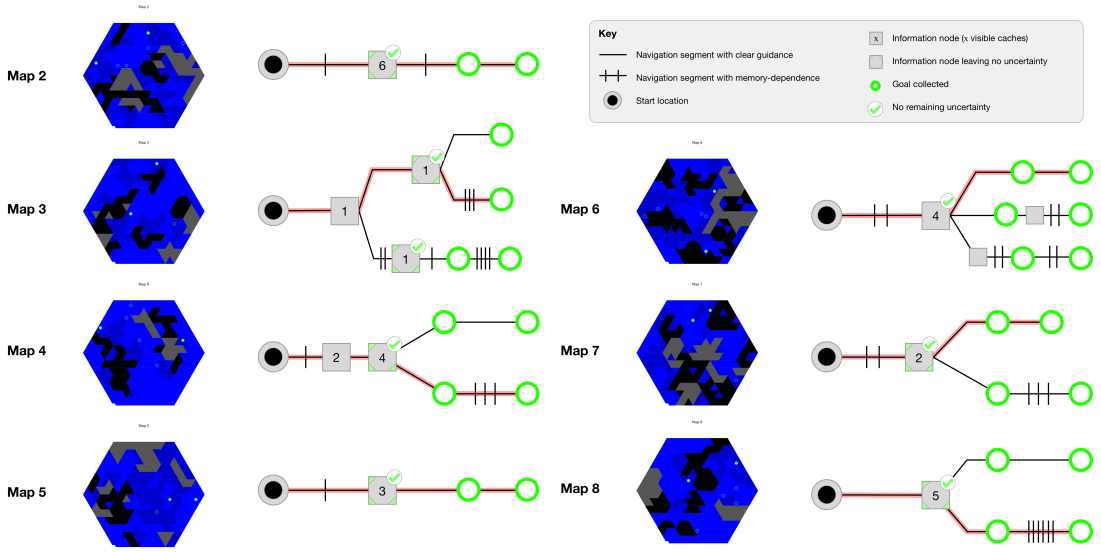

Figure 3: Map info-graphical classification for selected maps. See key for symbol semantics. Red line indicates expected path through graphical structure according to true goal locations.

Attribute	Definition
Branching factor	Average number of children of each node in the info-graphical path tree.
Number of forks	Count of forks in which a node splits into two or more children in the info-graphical path tree.
Number of information nodes	Number of informational nodes (represented as squares) in which goal-cache locations are hypothetically visible at navigation-time.
Steps to certainty	Expected number of nodes at which point all goal-cache presence statuses will have been observed.
Expected memory-dependent segments (planning)	Expected number of tree segments that require some memory of map structure to successfully navigate (indicated by hash marks through path segment).
Expected memory-dependent segments navigated	Expected number of tree segments to be navigated conditioned on true goal locations.

Table 3: Key map attributes analyzed.

Results

Trends in gaze dynamics through planning

The first question that we address in this study is whether gaze dynamics show signatures of hierarchical planning and subgoaling and whether there are differences between top and bottom performers. For this, we focus on temporal gaze dynamics prior to movement.

Figure 4 shows that consistent trends in gaze dynamics can be found across the 15-second planning phase. Total saccade distance decreases while fixation duration increases. Fixation targets converge to tiles closer to the path ultimately taken, with minimum distance seen at the end of planning (see also Supplementary Figure S1) for gaze dynamics in some example maps.

Additionally, a U-shaped trend is observed in distance from fixated tiles to the nearest goal location; tiles close to goals are more frequently fixated in the middle of the planning phase, while more distal tiles are viewed at the start and end of planning. These dynamics are largely consistent across top and bottom performers (top panels) and across 3 and 6-goal maps (bottom panels), but among top performers, gaze exhibits lower fixation durations and greater movement distances.

Finally, when gaze metrics were analyzed against map attributes, results show a consistent relationship between key map info-graphical attributes and eye movements. Specifically, maps with more forks, a higher branch-factor, and more memory-dependent segments are associated with longer eye movements and reduced fixation duration during planning. Furthermore, we found associations between map attributes and planning-time gaze recurrence metrics. The map branch factor, number of memory-dependent segments, and number of forks were associated with lower recurrence rate (RR). These results may indicate wider novel exploration (that therefore exhibit fewer repetitions during search) in more complex maps. See Table 4 for a summary of ANOVA results.

DV	Factor	DF	F	p	η_p^2	Dir.
Sum of eye move.	Branch factor	$F(2, 66)$	9.86	**< .001**	.011	+
Sum of eye move.	No. forks	$F(2, 66)$	13.3	**< .001**	.013	+
Sum of eye move.	No. memory-dep. segs.	$F(1, 33)$	4.9	**.033**	.004	+
Mean fix. dur.	Branch factor	$F(2, 66)$	4.99	**.010**	.006	-
Mean fix. dur.	No. forks	$F(2, 66)$	6.47	**.003**	.006	-
Mean fix. dur.	No. memory-dep. segs.	$F(1, 33)$	10.2	**.003**	.004	-
Planning RR	No. forks	$F(2, 66)$	8.2	**< .001**	.027	-
Planning RR	No. memory-dep. segs.	$F(1, 33)$	22.6	**< .001**	.032	-
Planning RR	Branch factor	$F(2, 66)$	4.1	**.020**	.015	-

Table 4: Summary of ANOVA results comparing gaze dynamics with map-level info-graphical metrics. "Dir" indicates whether the relationship is positive (+) or negative (-). Planning RR is gaze recurrence rate during the planning phase. Significant results (with $p \leq 0.05$) are marked in bold.

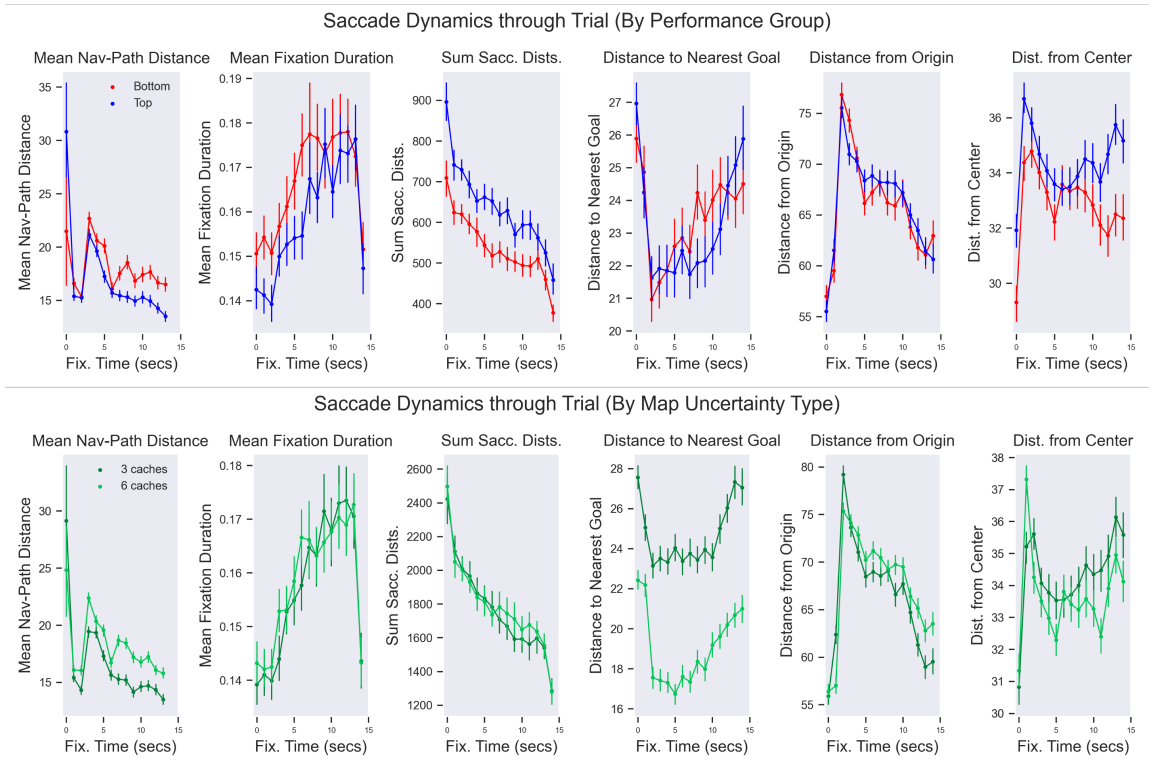

Figure 4: Trends in eye movement dynamics through planning phase. The six plots show the distance between gaze and navigation path, the saccade distance, saccade duration, distance from nearest goal, distance from origin, and distance from center. Top panel: comparison of top and bottom performers. Bottom panel: comparison of 3 and 6-goal maps.

Gaze distribution and the salience of map attributes

The second question that we address in this study is what features of the map are considered more salient during planning – either because they are candidate goal or subgoal locations or because they impair planning (e.g., obstacles that require detours) – and whether there are differences between top and bottom performers.

We considered the visual salience of map attributes through a static analysis of aggregate (cross-participant) gaze on each map. Figure 5 shows the fixation counts and duration in relation to the five different tile types (water, muddy water, land, wall or goal tiles). The value of zero in the plot indicates expected fixation duration and time if visual attention was perfectly uniform for a given map. Our analysis indicates that both goal locations and muddy water are disproportionately fixated at planning time. Goal locations are obviously salient for the task, whereas muddy water (and especially unavoidable muddy water along a chosen path) might require more careful consideration. Furthermore, we found that regular water tiles were under-fixated by both metrics, likely due to regions of each map easily identified as not relevant to the navigation task, and therefore pruned early on. We also found that fixations on muddy water and goal tiles was even more disproportionately favored on 6-cache as compared to 3-cache maps. Finally, we found decreasing attention on obstacles in the second half of the planning phase.

When comparing map attribute salience by performance group, we find a number of differences that may suggest attention-correlated sub-goaling strategies that were useful to performance. Bottom performers fixated longer on wall tiles as compared with top performers, especially during the second half of planning. We also find that bottom performers spent more time at the beginning of planning attending to "first fork" regions. First forks were defined by selecting a region around the first tile where any optimal path to goal diverged. These locations are often readily apparent as early points where a navigation decision will need to be made. Likewise, during the first 2 seconds of planning, top performers more quickly identified another key point of salience: the actual first goal they would collect.

We also analyzed another plausible sub-goal type critical to ambiguity reduction during navigation. The region of "first certainty" is defined as the closest location from which the exact location of all goals will be determinable. More specifically, first certainty is a line-of-sight based analysis that follows all paths to goals, and determines the earliest location where $n-1$ goal cache locations will be visibly accessible. From this point on the map, it is possible to infer the true location of each goal, by inspecting all (or all but one) goal caches. Top performers were found to attend to this region sooner than bottom performers on both 3 and 6-cache maps.

Figure 6 shows the fixation counts and duration in relation to the static feature-maps described in the Methods. We found that the feature maps were able to predict planning-time gaze in aggregate (across participants) significantly better than a uniform baseline. Feature-maps based on tile criticality, or the sensitivity of navigation paths and goal distances to the passability of particular locations, had lowest average EMD, implying they offer the highest predictive value (see Table 2 and Figure 6). See also Supplementary Fig S2 (and continuations in Fig S3, S4, and S5) for illustrative examples of how feature maps reflect gaze allocation.

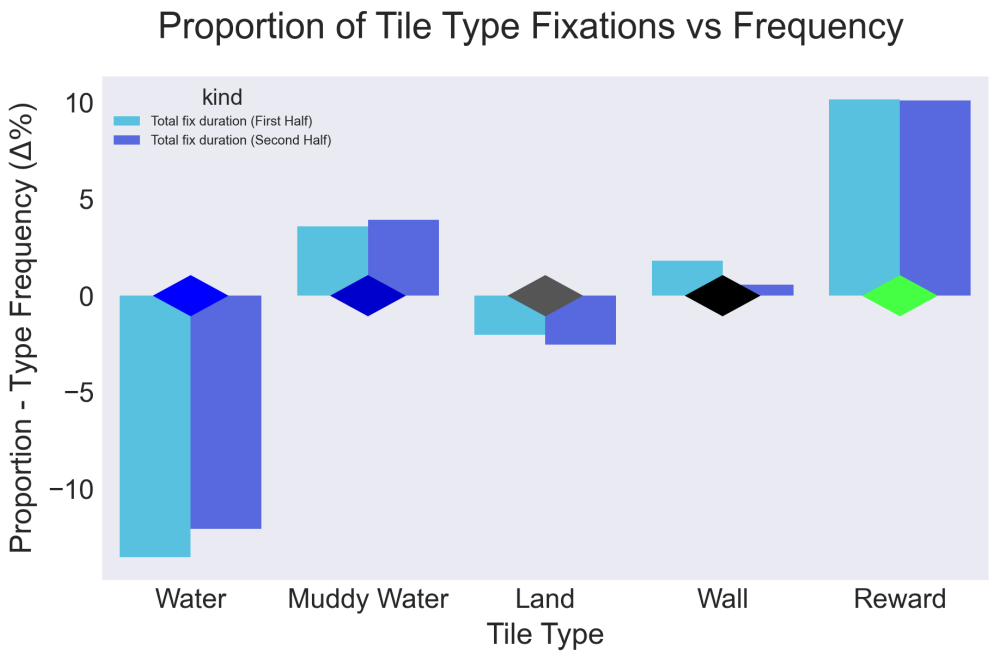

Figure 5: Tile type fixation proportionality. The value of zero in the plot indicates expected fixation duration and time if visual attention was perfectly uniform for a given map.

Relation between path planning dynamics and performance

The third question that we ask is whether and how the planning strategy of participants prior to movement influences their subsequent performance. We address this question by considering potential differences between path planning depth and initial movements during planning of top- and bottom-performers.

Relation between path planning depth and performance We first ask whether top- and bottom-performers differ in how much they plan ahead during the planning phase (i.e., "planning depth" (Russell and Norvig, 2016)).

Included in the feature-map analysis were maps generated by finding the shortest path through both the origin and either 1, 2, 3 or 4 goal locations (SP-1, SP-2, SP-3, and SP-4 respectively). Using these maps, we inferred each participant's "planning depth" – or the putative depth of the plans that participants developed during the planning phase – on each trial, by choosing the SP-depth map with the minimum distance (EMD) to trial-level planning gaze.

We therefore analysed the relationships between these four levels of planning depth (SP-1 to SP-4) and both performance and various map attributes. As shown in Figure 7, in three of the info-graphical attributes most directly correlated with map difficulty (number of forks, branch factor, and number of informational nodes), we see a relative agnosticism to depth class for simpler maps (e.g. those with less than 2 forks, or less than 3 info nodes). However, in the highest complexity group, a stronger performance correlation appears. Participants with planning depth most similar

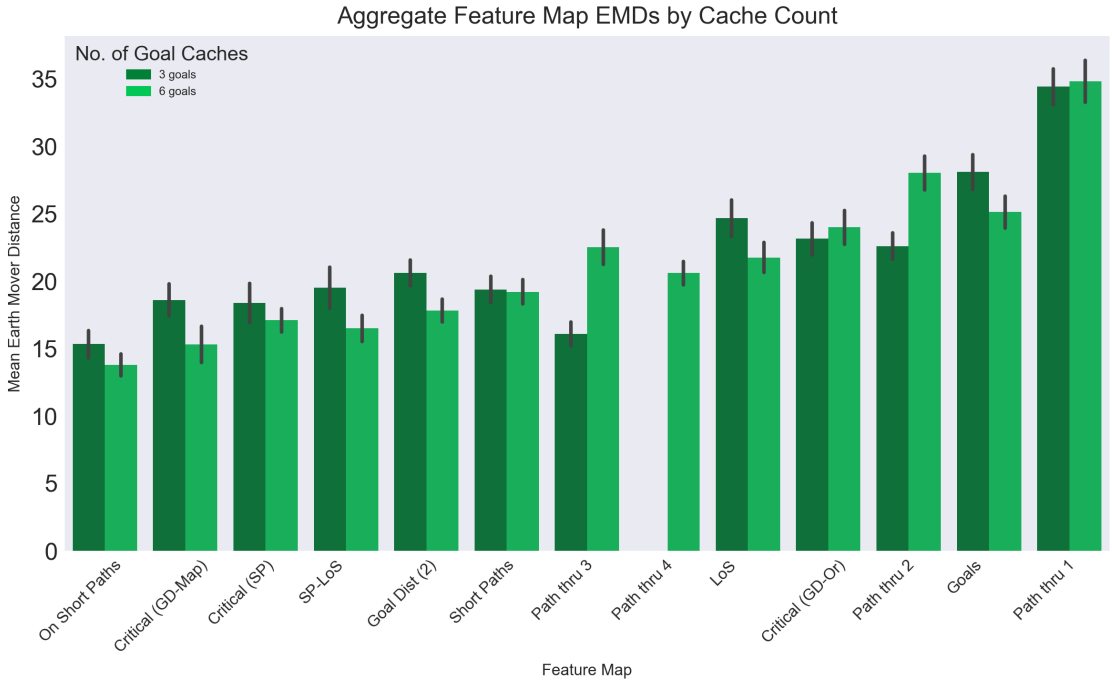

Figure 6: Feature-map predictors of aggregate gaze maps. Lower EMDs correspond to greater distributional similarity between feature map and aggregated participant gaze. Bars are plotted in ascending order from lower to higher distances.

to SP-3 performed best on maps with the largest number of forks, while SP-4 participants performed best on maps with the largest number of info nodes, and largest branching factor.

When bucketed by performance group, we find SP-3 and SP-4 depth classes were more common among top-performers. The SP-3 class was seen in 47.1% of bottom, as compared with 51.0% of top performers. The SP-4 class was seen in 26.7% of bottom, as compared with 30.1% of top performers.

Relation between arm movement during planning, initial heading, and performance Finally, we ask whether arm movements during planning are predictive of some of the very first decisions made—which direction to begin navigating – in bottom- and top-performers. For this analysis, we defined initial navigation direction based on the position of the agent in relation to the map center-line ("left" or "right"), 1 second after the beginning of navigation[1].

We find that, indeed, arm movements are directly correlated with initial heading, implying that the initial movement direction can be read out from the position of the yaw controller prior to the beginning of movement. This trend might reflect planning and preparatory processes leaking into the execution phase or a way planners offload part of their cognitive processing to physical activity (Risko and Gilbert, 2016; Clark et al., 2008; Kirsh and Maglio, 1994; Spivey, 2008). Notably, while

[1]The center-line relative position was chosen to capture the ultimate decision of participants to begin navigating to the left or right.

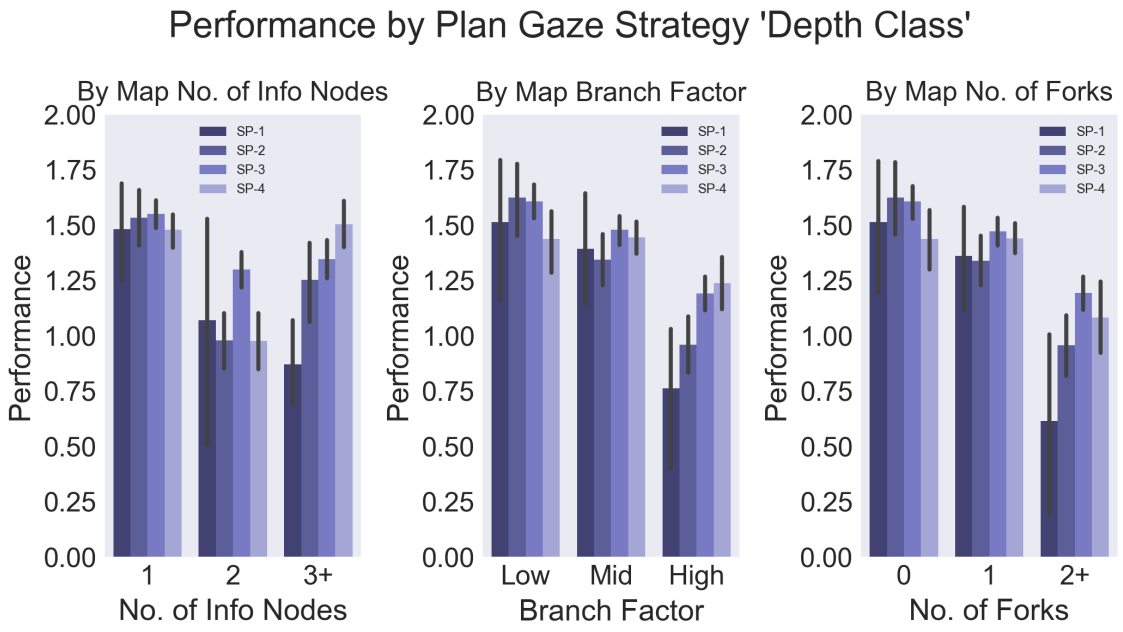

Figure 7: Performance and planning depth. Trial-level performance comparison showing the relationship between planning depth (as inferred from planning time gaze distance from shortest-paths through 1, 2, 3, and 4 goal caches), and info-graphical map attributes. A positive, nearly monotonic, relationship is seen in the most complex maps (as determined by number of forks, branch factor, and number of info nodes) between planning depth and performance.

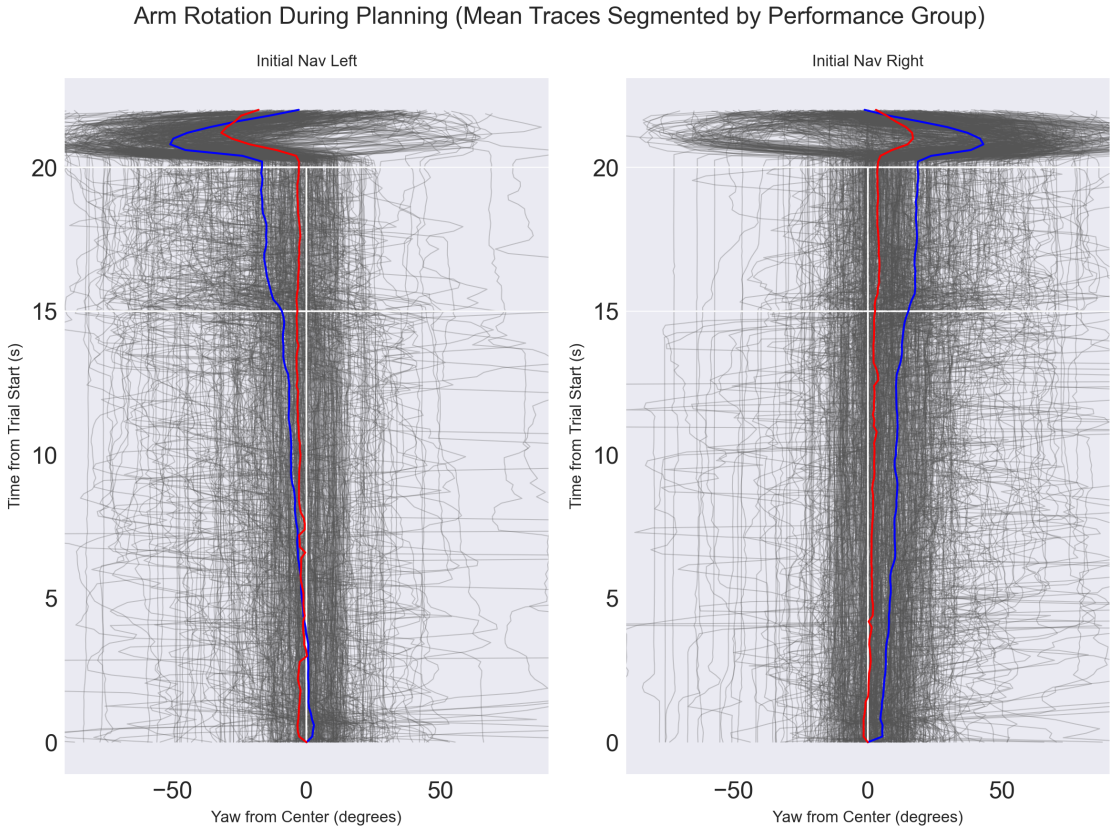

Figure 8: Arm movement traces as measured by controller yaw. Traces are segmented by initial navigation direction (left and right panels correspond to initial navigation to the left and right, respectively). Time from trial start progresses vertically from bottom. Navigation direction groups are defined as the sign of the offset from center-line 1000ms after navigation start. The blue and red lines show mean across trial-level traces among top and bottom-performers, respectively. Planning phase (bottom 15 seconds), transition (following 5 seconds), and first 5 seconds of navigation phase are shown, separated by white lines.

this trend is very clear in top-performers, it is much less pronounced in bottom-performers, in which the the average position of the yaw controller remains very close to the middle line. See controller yaw traces, segmented by initial navigation direction (left and right panels segment trials by initial navigation to the left or right, respectively), and performance group (blue and red lines show data from top and bottom-performers, respectively), in Figure 8. The higher-amplitude movements seen at the beginning of the navigation phase (top of plot) illustrate the larger movements required to control agent rotation as navigation began. As expected, this portion exhibits a heavy directional bias due to the definition of the "initial navigation direction" variable used for segmentation.

Discussion

In this human behavioral study, participants were asked to view an overview of a landscape for 15 seconds, prior to immersively navigating through it to find and collect up to two goals. Because the number of possible goal locations ("caches") exceeds the number of true goals, successful participants had to develop a navigation plan robust to this uncertainty. A primary motivation in this design is to better understand and characterize the dynamics of the internal representations developed during planning to guide future movement, and the handling of uncertainty during this process. However, these representations were not, themselves, directly measurable. Instead, we analyzed eye and arm movements collected during the planning phase of the task, which putatively reflect some dimensions of the cognitive processes of interest, and admit quantitative comparison with geometric, visual, and informational features of the landscapes studied.

This work pursued three aims: determining whether gaze dynamics prior to movement show signatures of hierarchical planning and subgoaling, identifying the most relevant features for subgoal selection, and assessing which strategic aspects of planning distinguish top- from bottom-performers.

Our spatial and temporal analyses, as well as correlations between gaze patterns prior to navigation and quantitative map-level attributes and task performance metrics, show signatures of a planning process that is both hierarchical and bi-directional. These include a decreasing trend seen in gaze distance from origin, and a broad to narrow shift (with reducing saccade distances and longer fixation durations) as plans are established. These signatures of hierarchical planning likely aid participants in efficiently transforming rich spatial information into a simpler, contingency-aware, plan of action. Furthermore, in line with prior work, we identified "critical tiles" to which landscape connectivity is most sensitive as the features that best predict visual attention. Finally, performance-group-level analyses show that the most effective participants (top-performers) navigated paths closer to regions of the map visually attended during planning, exhibited more frequent eye movements with shorter fixations, as well as earlier attention to ambiguity-reducing sub-goals, Top performers always showed greater focus on more peripheral tiles and deeper planning depth, and leveraged offloading of initial navigation direction through orientation of the arm. Below we discuss these results in more detail.

Gaze dynamics prior to navigation show signatures of hierarchical planning

In a temporal analysis of gaze dynamics, we find consistent trends indicative of hierarchical planning across the planning phase (see Fig 4). Total saccade distance decreases while (in line with this result) fixation duration increases. Fixation targets converge to tiles closer to the path ultimately taken, with minimum distance seen at the end of planning. A U-shaped trend is observed in distance from fixated tiles to the nearest goal location; tiles close to goals are more frequently fixated in the middle of the planning phase, while more distal tiles are viewed at the start and end of planning. Finally, gaze distance from origin shows a quick radiation outwards in the first several seconds, followed by a consistent negative trend for the remainder of the planning phase.

Together, these results illustrate that temporal gaze dynamics start broad and then spatially narrow as plans are established. Visual exploration first favors proximal locations, and is initially wide and fast, exhibiting saccades that move between easily-identified salient locations such as goal caches. Gaze then narrows (in terms of spatial spread), and deepens (seen as longer fixation durations) through the planning phase as a navigation plan is converged upon. These gaze dynamics could reflect a covert planning process that exhibits a hierarchical nature. Initial exploration of key landmarks may help to establish navigational affordances based on abstract map connectivity. Then, participants might establish a plan hierarchically, by enhancing an initially abstract path through high level navigation goals with lower level details of plausible routes between them. These results are consistent with previous findings on hierarchical organization of action execution (Balaguer et al., 2016; Donnarumma et al., 2016; Schick et al., 2019; Fernandez Velasco et al., 2024; Correa et al., 2023; Eluchans et al., 2024) and shed light on the temporal and visual dynamics enabling hierarchical plans to be established before movement.

Salient tiles and critical regions predict attention during planning

Findings from map-based spatial prediction of gaze preferences suggest that visual search strategies are highly sensitive to not only goal locations and other gross aspects of spatial geometry, but also to more subtle features and aspects of graphical connectivity and information geometry (e.g. goal presence information available at a particular tile).

Overall, we find that goal locations and muddy water were most disproportionately fixated. We also observe increasing attention on water, and decreasing attention on obstacles in the second half of the planning phase (see Fig 5). This finding may indicate that early saccades to obstacles support the development of a simplified graphical representation of map connectivity (and information access), while later gaze sequences support evaluation of candidate paths through the structure identified. These results are in agreement with other studies showing that people form simplified problem representations to plan (Ho et al., 2022) and that cognitive task representations can be distorted, removing task-irrelevant dimensions (D'Amato et al., 2024).

Gaze prediction analysis using static geometry-based feature-maps enabled a more nuanced analysis of the types of features that attracted visual attention. Feature-maps were effective predictors of planning-time gaze in aggregate (across participants). Comparison of predictive power for each feature-map analyzed is suggestive of the landscape features that most consistently attracted planning-time visual attention (see Table 2 and Fig 6). Feature-maps based on tile criticality, which identify regions for which a change to connectivity (i.e. switching a passable region for an obstacle) have particularly outsized effects on key navigation considerations such as the lengths of efficient paths (Critical Shortest-Path), or mean goal distance map-wide (Critical GD-Map), were especially useful predictors. This finding is consistent with results related to edge "toggling" in (Zhu et al., 2022) and shows that it generalizes to tasks requiring participants to develop robust navigation plan.

Top-performers show distinctive planning strategies along multiple dimensions

Our results show that several strategic aspects of planning distinguish performance groups. First, we find that among top performers, gaze exhibits lower fixation durations and greater movement distances. These attributes are correlated with a higher throughput of information processing at planning time and might reflect a more efficient allocation of limited resources (Callaway et al., 2018).

Second, we find that attention in some key map regions, especially early attention in the first few seconds of planning, differ by performance groups. A bias in attention towards the first fork is seen in bottom-performers, which may indicate that these more superficial or readily apparent map features attracted more attention for this group, while top-performers may have more quickly moved to subtler decision points such as deeper cache-location dependent forks, or more distant clusters of plausible goals. Additionally, we find that "first certainy" sub-goal locations were attended earlier by top performers, possibly reflecting a greater attention to epistemically critical sub-goals conveying full-map information (the location of all goals), but which require a richer (perhaps simulative) forecast of visual and informational dynamics at navigation time. Together, we interpret these findings as indicating that different performance groups' attention were attracted to different kinds of sub-goals, and at different times. First forks are readily apparent from local (near origin) map structure, whereas identifying which goal is most plausible to capture, or which locations will provide maximal ambiguity reduction, requires deeper consideration, and perhaps more robust simulation. This may explain why attention in these regions better differentiates skill and performance

Third, we find that deeper planning gaze predicts success on high-uncertainty maps. Planning depth results showed that the SP-3 map was the most predictive of aggregate gaze across participants (on both 3- and 6-cache maps), indicating a possible balance of (cognitively costly) planning depth and the risk of planned caches being empty. Performance correlations with gaze depth-class (as determined by similarity with "SP-X" feature maps) indicate that some map attributes rendered deeper plans particularly helpful. In particular, planning gaze similar to SP-3 was associated with highest performance on maps with at least one fork. Similarly, maps with fewer information-nodes (locations at which goal existence information would be obtained) saw greatest performance when gaze resembled SP-3. In contrast, SP-4 performed best in the most epistemically complex maps with the largest number (3 or more) of informational nodes. See Fig 7 for comparisons, and the Methods for details on the determination of map classification attributes. Another perspective on this result can also be seen in the performance-binned gaze preferences in Fig S1, which show a bias for more peripheral, and later path segments among the top performing group. We interpret these findings as evidence that deeper plans predicted performance, but only on more challenging maps where multiple branches, or multiple stages of information access, required the development of a more robust plan (one with visits to more goal caches). On the other hand, a greedier (and riskier) strategy was satisfactory on maps with simpler information geometry and less path uncertainty.

Finally, we find that the behavior of top-performers before navigation better predicts their initial navigation direction compared to bottom-performers. The fact that it is possible to decode

initial navigation direction before navigation starts reflects the fact that there is not a clear cut separation between the planning and execution phases, but rather some aspects of the plan leak into the overt execution (Gordon et al., 2021; Cisek and Pastor-Bernier, 2014; Yoo et al., 2021; Spivey, 2008). However, this trend is much more prominent in top-performers, which may suggest that they *exploited* the continuous nature of planning and execution in the task at head. These findings may be signatures of offloading part of the decision memory to physical state, potentially alleviating cognitive burden (Risko and Gilbert, 2016; Clark et al., 2008; Kirsh and Maglio, 1994; Priorelli et al., 2024).

Limitations & future work

In this study, uncertainty was primarily manipulated by requiring participants plan with incomplete information about true goal locations. While this paradigm corresponds with certain naturalistic settings where rewarding states often fall within a set of discrete locations, other types of uncertainty are frequent in real-world situations requiring planning and navigation. As one example, path connectivity may differ from maps or other representations available when a route is being selected, e.g. due to unforeseen obstacles, the presence and influence of other individuals (e.g. crowds that slow one's progress), and other modifications to the environment (e.g. a transit line discovered to be down for maintenance). These environmental uncertainties likely require even more robust contingency planning, and perhaps prompt the development of different internal representations better suited to more continuous and dynamic changes. Such a design would likely limit the efficacy of heuristic pathfinding strategies such as "find a short path through 4 goal caches," likely used by some participants in this study, since path length would itself be unpredictable during planning. Eye movements and other biometric data would likely shed light on both correspondence and divergence in the planning processes used by individuals encountering these other types of uncertainty.

Similarly, while the plan-then-execute paradigm employed here reflects common daily scenarios where planning may progress prior to direct engagement in the environment, as discussed in (Gordon et al., 2021), more embodied planning and action scenarios require the constant interplay of perception and action as the state of an individual with respect to their environment continuously evolves. The meta-cognitive and resource-rational aspects of how to coordinate movement with sensory sampling, and when and how to leverage deeper simulations for more distal planning processes, even during ongoing activity, is not well understood; these questions offer many rich avenues for further study.

Conclusion

This work presented a virtual-reality based behavioral experiment which collected eye movements during a pre-navigation phase, with the aim of capturing aspects of covert planning processes related to pathfinding through a landscape with uncertain goal locations. Taken together, our results demonstrate the way gaze, and other biometric signals, can uncover hierarchical attributes of planning dynamics through time. Furthermore, we show that the nature of the hierarchy built

(e.g. the set of sub-goals chosen as its nodes) is skill-dependent – highlighting the capability of top-performers to allocate their cognitive resources to task-relevant features and to form deeper plans. These results showcase the importance of devising effective planning and active exploration strategies to solve challenging navigation tasks that defy exhaustive search.

Data availability

The data and materials for all experiments are available at `https://osf.io/5tacn/?view_only=9a2db28ca7744dc8a4de13d42271d99f`.

Acknowledgments

This research received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 952215 (TAILOR); the European Research Council under the Grant Agreement No. 820213 (ThinkAhead), the Italian National Recovery and Resilience Plan (NRRP), M4C2, funded by the European Union – NextGenerationEU (Project IR0000011, CUP B51E22000150006, "EBRAINS-Italy"; Project PE0000013, "FAIR"; Project PE0000006, "MNESYS"), and the Ministry of University and Research, PRIN PNRR P20224FESY and PRIN 20229Z7M8N. This work was also supported by the UC Berkeley Center for Long-Term Cybersecurity, and the National Science Foundation's Graduate Research Fellowship Program Grant No. 2019236659. The GEFORCE Quadro RTX6000 and Titan GPU cards used for this research were donated by the NVIDIA Corporation. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

Anderson, N. C., Bischof, W. F., Laidlaw, K. E., Risko, E. F., and Kingstone, A. (2013). Recurrence quantification analysis of eye movements. *Behavior research methods*, 45:842–856.

Arnold, A. E., Iaria, G., and Ekstrom, A. D. (2016). Mental simulation of routes during navigation involves adaptive temporal compression. *Cognition*, 157:14–23.

Balaguer, J., Spiers, H., Hassabis, D., and Summerfield, C. (2016). Neural mechanisms of hierarchical planning in a virtual subway network. *Neuron*, 90(4):893–903.

Baldauf, D. (2018). Visual selection of the future reach path in obstacle avoidance. *Journal of cognitive neuroscience*, 30(12):1846–1857.

Baldauf, D., Wolf, M., and Deubel, H. (2006). Deployment of visual attention before sequences of goal-directed hand movements. *Vision research*, 46(26):4355–4374.

Bhui, R., Lai, L., and Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences*, 41:15–21.

Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., and Colton, S. (2012). A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43.

Buzsaki, G. (2019). *The brain from inside out*. Oxford University Press, USA.

Callaway, F., Lieder, F., Das, P., Gul, S., Krueger, P. M., and Griffiths, T. (2018). A resource-rational analysis of human planning. In *CogSci*.

Callaway, F., van Opheusden, B., Gul, S., Das, P., Krueger, P. M., Griffiths, T. L., and Lieder, F. (2022). Rational use of cognitive resources in human planning. *Nature Human Behaviour*, 6(8):1112–1125.

Cisek, P. and Pastor-Bernier, A. (2014). On the challenges and mechanisms of embodied decisions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1655):20130479.

Clark, A. et al. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. OUP USA.

Correa, C. G., Ho, M. K., Callaway, F., Daw, N. D., and Griffiths, T. L. (2023). Humans decompose tasks by trading off utility and computational cost. *PLoS computational biology*, 19(6):e1011087.

Coutrot, A., Manley, E., Goodroe, S., Gahnstrom, C., Filomena, G., Yesiltepe, D., Dalton, R., Wiener, J. M., Hölscher, C., Hornberger, M., et al. (2022). Entropy of city street networks linked to future spatial navigation ability. *Nature*, 604(7904):104–110.

D'Amato, L., Lancia, G. L., and Pezzulo, G. (2024). The geometry of efficient codes: how rate-distortion trade-offs distort the latent representations of generative models. *arXiv preprint arXiv:2406.07269*.

De Cothi, W., Nyberg, N., Griesbauer, E.-M., Ghanamé, C., Zisch, F., Lefort, J. M., Fletcher, L., Newton, C., Renaudineau, S., Bendor, D., et al. (2022). Predictive maps in rats and humans for spatial navigation. *Current Biology*, 32(17):3676–3689.

Donnarumma, F., Maisto, D., and Pezzulo, G. (2016). Problem solving as probabilistic inference with subgoaling: explaining human successes and pitfalls in the tower of hanoi. *PLoS computational biology*, 12(4):e1004864.

Ekman, M., Kusch, S., and de Lange, F. P. (2023). Successor-like representation guides the prediction of future events in human visual cortex and hippocampus. *eLife*, 12:e78904.

Eluchans, M., Lancia, G. L., Maselli, A., D'Alessandro, M., Gordon, J., and Pezzulo, G. (2023). Adaptive planning depth in human problem solving. *bioRxiv*, pages 2023–05.

Eluchans, M., Maselli, A., Lancia, G. L., and Pezzulo, G. (2024). Eye and hand coarticulation during problem solving reveals hierarchically organized planning. *bioRxiv*, pages 2024–11.

Fernandez Velasco, P., Griesbauer, E.-M., Brunec, I., Morley, J., Manley, E., McNamee, D., and Spiers, H. (2024). Expert navigators deploy rational hierarchical priorization over predictive maps for large-scale real-world planning. *bioRxiv*, pages 2024–03.

Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N. T., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D. J., Tavenard, R., Tong, A., and Vayer, T. (2021). Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8.

Foster, D. J. (2017). Replay comes of age. *Annual review of neuroscience*, 40(1):581–602.

Gershman, S. J., Horvitz, E. J., and Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278.

Gordon, J., Maselli, A., Lancia, G. L., Thiery, T., Cisek, P., and Pezzulo, G. (2021). The road towards understanding embodied decisions. *Neuroscience & Biobehavioral Reviews*, 131:722–736.

Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., Group, C. S. R., et al. (2015). Exploration versus exploitation in space, mind, and society. *Trends in cognitive sciences*, 19(1):46–54.

Ho, M. K., Abel, D., Correa, C. G., Littman, M. L., Cohen, J. D., and Griffiths, T. L. (2022). People construct simplified mental representations to plan. *Nature*, 606(7912):129–136.

Hodgson, T. L., Hermens, F., Pennington, K., Pickering, J. S., Ezard, G., Clarke, R., Sharma, J., and Owen, A. M. (2019). Eye movements in the "morris maze" spatial working memory task reveal deficits in strategic planning. *Journal of Cognitive Neuroscience*, 31(4):497–509.

Hoppe, D. and Rothkopf, C. A. (2019). Multi-step planning of eye movements in visual search. *Scientific reports*, 9(1):1–12.

Kirsh, D. and Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive science*, 18(4):513–549.

Lakshminarasimhan, K., Zhu, S., and Angelaki, D. (2024). Planning by active sensing. In *Gaze Meets Machine Learning Workshop*, pages 125–141. PMLR.

Lakshminarasimhan, K. J., Avila, E., Neyhart, E., DeAngelis, G. C., Pitkow, X., and Angelaki, D. E. (2020). Tracking the mind's eye: Primate gaze behavior during virtual visuomotor navigation reflects belief dynamics. *Neuron*, 106(4):662–674.

Lancia, G. L., Eluchans, M., D'Alessandro, M., Spiers, H. J., and Pezzulo, G. (2023). Humans account for cognitive costs when finding shortcuts: An information-theoretic analysis of navigation. *PLOS Computational Biology*, 19(1):e1010829.

Land, M., Mennie, N., and Rusted, J. (1999). The roles of vision and eye movements in the control of activities of daily living. *Perception*, 28(11):1311–1328.

Lieder, F. and Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and brain sciences*, 43:e1.

Mattar, M. G. and Lengyel, M. (2022). Planning in the brain. *Neuron*, 110(6):914–934.

McNamee, D., Wolpert, D. M., and Lengyel, M. (2016). Efficient state-space modularization for planning: theory, behavioral and neural signatures. *Advances in neural information processing systems*, 29.

Miller, K. J. and Venditto, S. J. C. (2021). Multi-step planning in the brain. *Current Opinion in Behavioral Sciences*, 38:29–39.

Momennejad, I. (2020). Learning structures: predictive representations, replay, and generalization. *Current Opinion in Behavioral Sciences*, 32:155–166.

Moreno-Bote, R., Ramírez-Ruiz, J., Drugowitsch, J., and Hayden, B. Y. (2020). Heuristics and optimal solutions to the breadth–depth dilemma. *Proceedings of the National Academy of Sciences*, 117(33):19799–19808.

Newell, A. and Simon, H. A. (1972). *Human problem solving*. Prentice-Hall, Englewood Cliffs, N.J.

Pezzulo, G., Kemere, C., and Van Der Meer, M. A. (2017). Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition. *Annals of the New York Academy of Sciences*, 1396(1):144–165.

Pezzulo, G., Rigoli, F., and Friston, K. J. (2018). Hierarchical active inference: a theory of motivated control. *Trends in cognitive sciences*, 22(4):294–306.

Pezzulo, G., Van der Meer, M. A., Lansink, C. S., and Pennartz, C. M. (2014). Internally generated sequences in learning and executing goal-directed behavior. *Trends in cognitive sciences*, 18(12):647–657.

Pfeiffer, B. E. and Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447):74–79.

Priorelli, M., Stoianov, I. P., and Pezzulo, G. (2024). Embodied decisions as active inference. *bioRxiv*, pages 2024–05.

Risko, E. F. and Gilbert, S. J. (2016). Cognitive offloading. *Trends in cognitive sciences*, 20(9):676–688.

Russell, S. J. and Norvig, P. (2016). *Artificial intelligence: a modern approach*. Pearson.

Santos-Pata, D. and Verschure, P. F. M. J. (2018). Human Vicarious Trial and Error Is Predictive of Spatial Navigation Performance. *Frontiers in Behavioral Neuroscience*, 12.

Schick, W., Halfmann, M., Hardiess, G., Hamm, F., and Mallot, H. A. (2019). Language cues in the formation of hierarchical representations of space. *Spatial Cognition & Computation*, 19(3):252–281.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489.

Simon, H. A. (1997). *Models of bounded rationality: Empirically grounded economic reason*, volume 3. MIT press.

Solway, A., Diuk, C., Córdova, N., Yee, D., Barto, A. G., Niv, Y., and Botvinick, M. M. (2014). Optimal behavioral hierarchy. *PLoS computational biology*, 10(8):e1003779.

Spivey, M. (2008). *The continuity of mind*. Oxford University Press.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological review*, 55(4):189.

Tomov, M. S., Yagati, S., Kumar, A., Yang, W., and Gershman, S. J. (2020). Discovery of hierarchical representations for efficient planning. *PLoS computational biology*, 16(4):e1007594.

Vidal, A., Soto-Faraco, S., and Moreno-Bote, R. (2022). Balance between breadth and depth in human many-alternative decisions. *Elife*, 11:e76985.

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., and Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature human behaviour*, 2(12):915–924.

Yoo, S. B. M., Hayden, B. Y., and Pearson, J. M. (2021). Continuous decisions. *Philosophical Transactions of the Royal Society B*, 376(1819):20190664.

Zhu, S., Lakshminarasimhan, K. J., Arfaei, N., and Angelaki, D. E. (2022). Eye movements reveal spatiotemporal dynamics of visually-informed planning in navigation. *Elife*, 11:e73097.

Zingale, C. M. and Kowler, E. (1987). Planning sequences of saccades. *Vision research*, 27(8):1327–1341.

Supporting information

Map generation

Maps for the task were generated in a three step procedure: generate map pool, filter based on criteria, and finally, select an experimental sample. Below, we detail the procedure for each.

Procedurally generate map pool

Map geometry was defined similarly to (Zhu et al., 2022) but with boundaries defined by tile types, rather than walls drawn along tile edges. Specifically, a hexagonal grid composed of equilateral triangular tiles, and a side length of 7, formed the base map. Each tile could be set to one of four types: water (blue), muddy water (dark blue), land (gray), or obstacle (black). Land tiles were low altitude permitting vision, but not navigational access, while obstacle tiles were tall, and blocked visual access. To generate unique maps with variation in structure and connectivity, the algorithm below was followed (See Algorithm 1).

Algorithm 1: Procedural map generation

```
Data: tileTypes[N] = WATER
;                                        /* Set all tiles initially to water */
Result: tileTypes[N], goalCaches[6]
n ← Uniform(10, 40) ;                    /* Choose n strokes to add */
i ← 0;
while i ≤ n do
    k ← Uniform(10, 40) ;                /* Choose k stroke length */
    j ← 0;
    if i ≤ n/2 then
        type ← MWATER;
    else
        type ← choice(LAND, OBSTACLE);
    tileCursor = randomTile() while j ≤ k do
        tileTypes[tileCursor] = type;
        tileCursor = successors(tileCursor) ;   /* Successors is a list of adjacent
         tiles with prior visited tile removed */
        j ← j + 1;
    i ← i + 1;
goalCaches[] = {};
r ← 0;
while r ≤ 6 do
    goalCaches[r] = randomTile(WATER);
```

Filter pool

To filter the pool of map candidates (defined by an assignment for each tile type, and an array of 6 goal cache locations), we confirmed that each one satisfies the following criteria:

1. No goal locations visible (line of sight) from origin
2. All goal locations reachable from origin
3. No obstacles adjacent to origin

Select experimental sample

Finally, with the filtered pool of eligible map candidates, we selected 30 maps to compose the main experiment. We then randomly assigned half of the maps to each uncertainty group (low vs high). For both groups, we randomly chose 2 goals among the 6 goal caches to be present in the map. For the low uncertainty group, we then randomly removed 3 of the unused goal caches. The end result was a list of maps, where each contained 2 present goals, among either 3 or 6 possible goal cache locations.

Finally, we hand designed two practice maps with trivial geometric structures, to be used at the start of the experiment during the tutorial, as well as three maps testing other information geometries of interest. Map order was randomized for all participants, while ensuring that these latter three maps were always shown at the end of the experiment, to avoid any influence on the primary, procedurally generated maps.

Four sample maps from the main experiment can be seen in Fig S1 (lower-right figures show actual map as represented to participants during planning, with navigation trajectories overlayed).

Supplementary figures

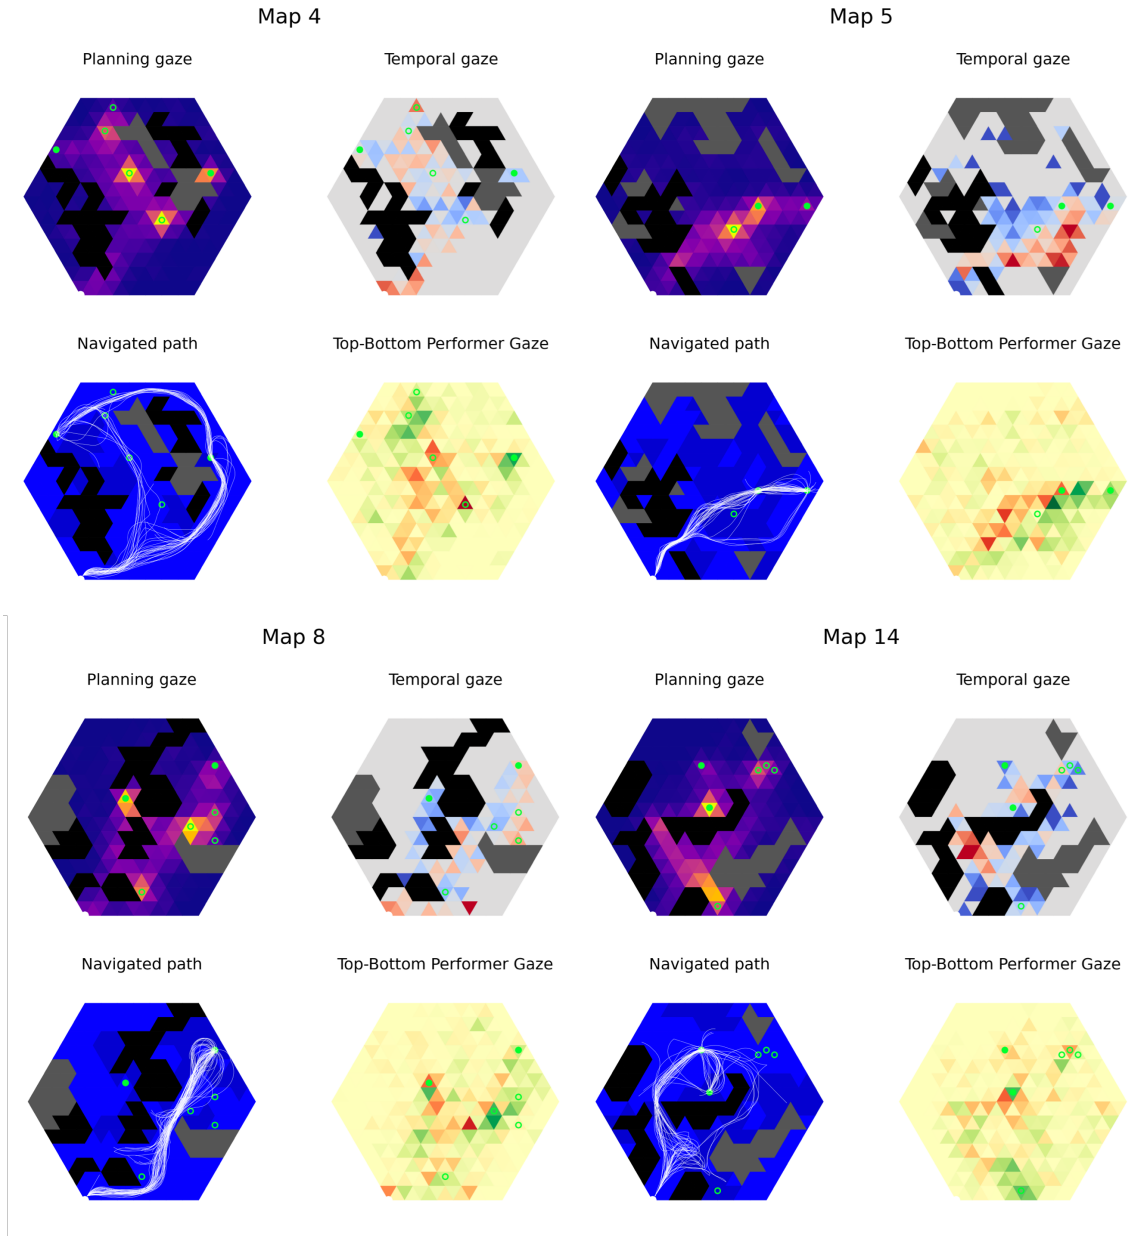

Figure S1: Sample gaze and navigation analyses for four maps. For each map, the following analyses are visualized: aggregated planning gaze (upper left), temporal gaze shift (upper right; blue tiles viewed earlier during planning, red later), path navigation trajectories (lower left), and performance-group differences in gaze density (lower right; green tiles show greater density among the top-performing group, and red tiles are those with greater density among the bottom-performing group.)

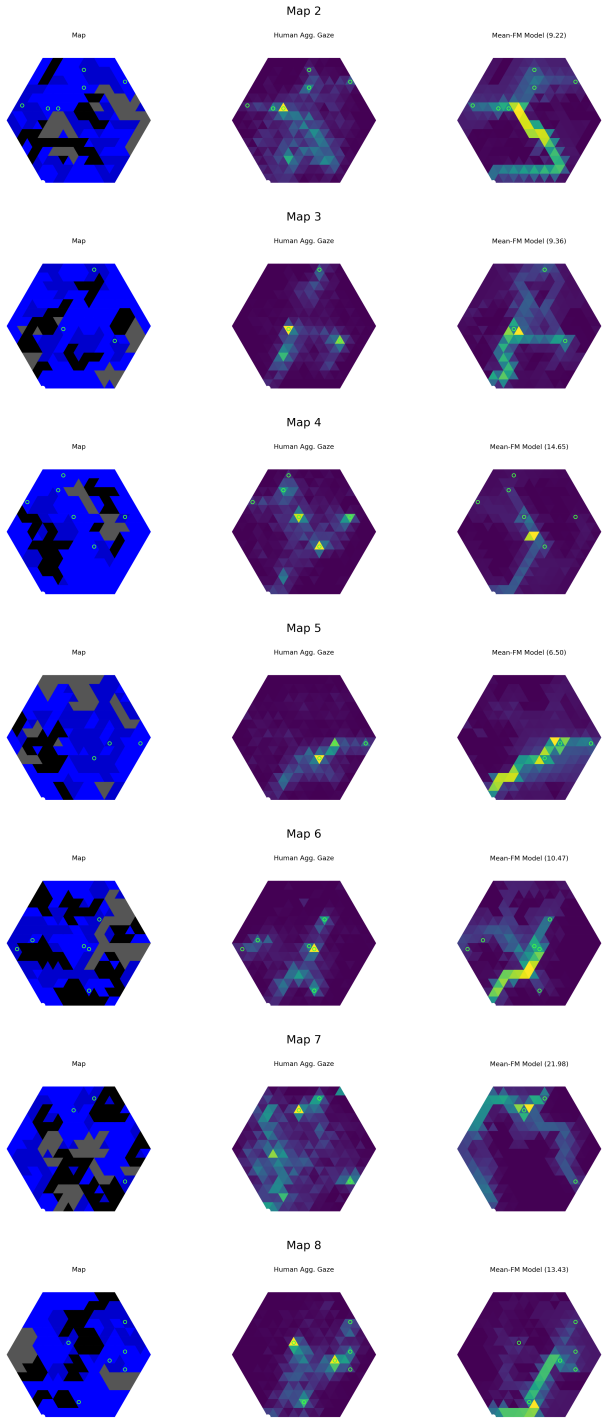

Figure S2: Map-level gaze summary of: a) map geometry, b) aggregate human gaze, c) uniform linear combination of feature maps. Number in title of FM model and agent gaze reports EMD from human aggregate gaze heatmap. Maps 2-8.

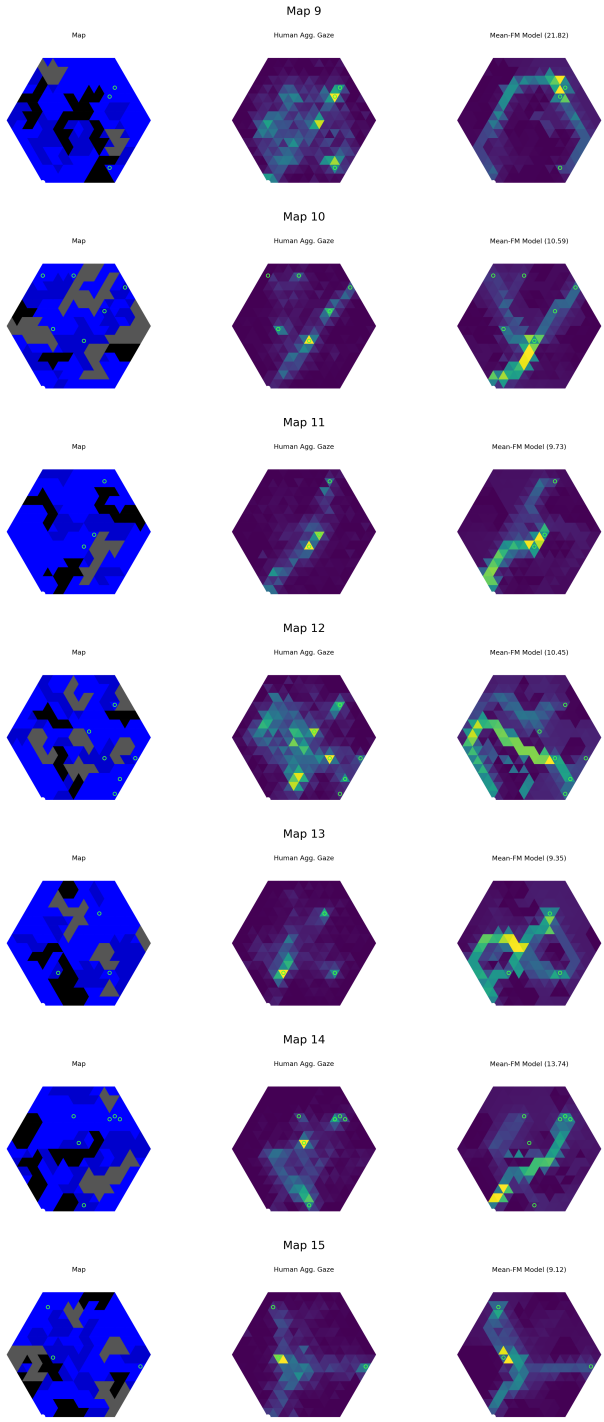

Figure S3: Continuation of S2 Fig. Maps 9-15.

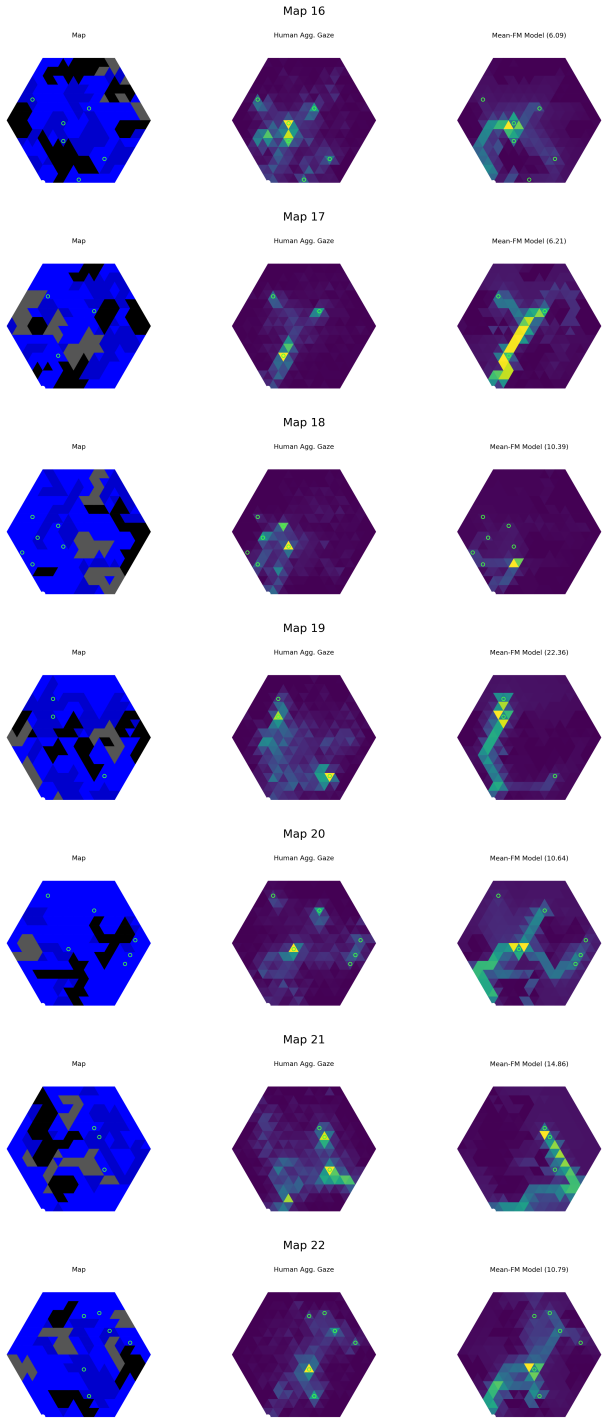

Figure S4: Continuation of S2 Fig. Maps 16-22.

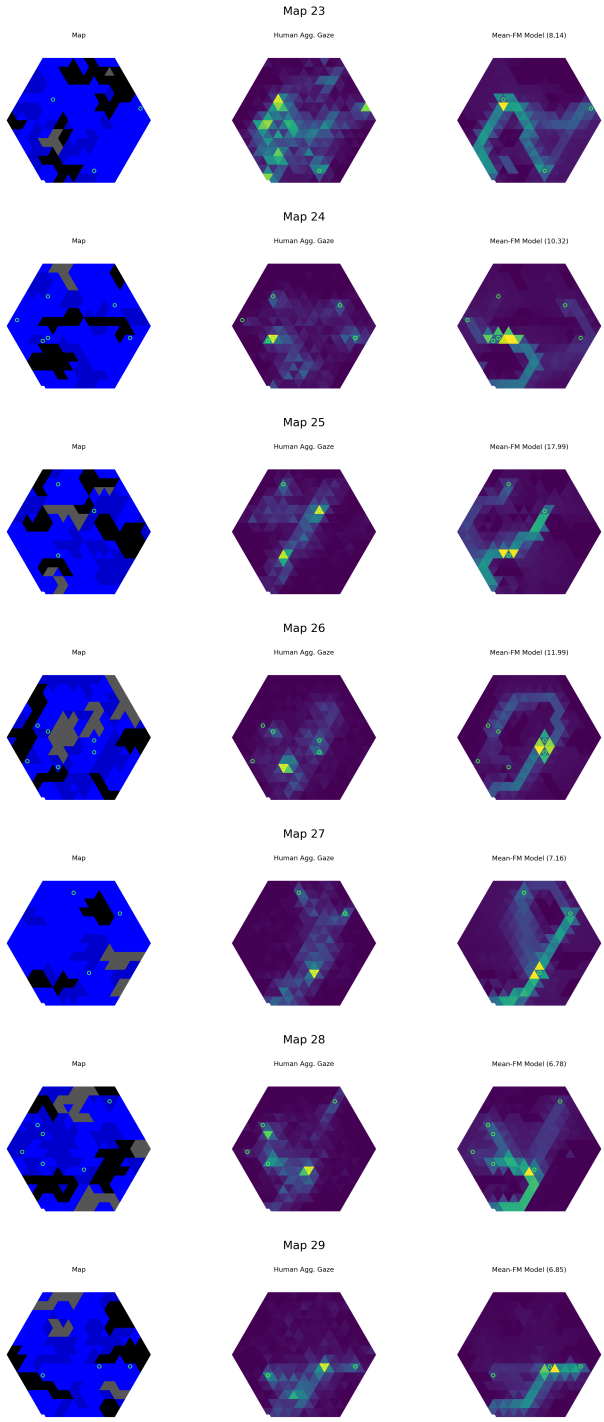

Figure S5: Continuation of S2 Fig. Maps 23-29.

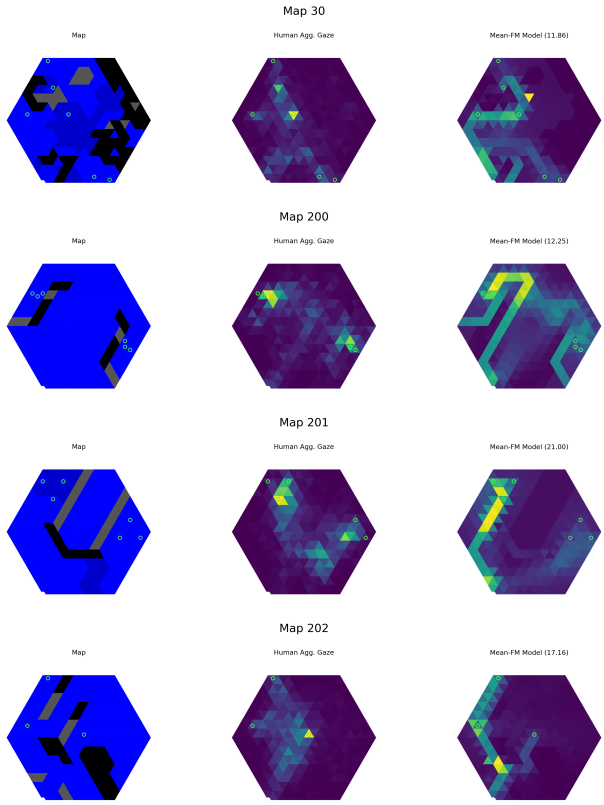

Figure S6: Continuation of S2 Fig. Map 30 and hand designed maps 200-202.