

The Efficiency of Human Cognition Reflects Planned Information Processing

Mark K. Ho,¹ David Abel,² Jonathan D. Cohen,¹
Michael L. Littman,² Thomas L. Griffiths¹

¹ Department of Psychology, Princeton University, Princeton, NJ

² Department of Computer Science, Brown University, Providence, RI

{mho, jdc, tomg}@princeton.edu

{dmabel, mlittman}@cs.brown.edu

Abstract

Planning is useful. It lets people take actions that have desirable long-term consequences. But, planning is hard. It requires thinking about consequences, which consumes limited computational and cognitive resources. Thus, people should plan their actions, but they should also be smart about how they deploy resources used for planning their actions. Put another way, people should also “plan their plans”. Here, we formulate this aspect of planning as a meta-reasoning problem and formalize it in terms of a recursive Bellman objective that incorporates both task rewards and information-theoretic planning costs. Our account makes quantitative predictions about how people should plan and meta-plan as a function of the overall structure of a task, which we test in two experiments with human participants. We find that people’s reaction times reflect a planned use of information processing, consistent with our account. This formulation of planning to plan provides new insight into the function of hierarchical planning, state abstraction, and cognitive control in both humans and machines.

Introduction

Suppose you have to catch a flight in three hours, and you are in your bedroom packing. How would you plan your next move? You might reason, “I will finish packing my bag. Then, I will get a car to take me to the airport.” This plan seems intuitive and straightforward. However, many details have been left out: Which taxi or ride-sharing company will you use? How are you going to get from security to your gate? If you reach the gate and get hungry, can you pick up a snack in time to still make your flight? Clearly, you did not imagine every possible contingency from now until your flight departs. Instead, you sketched out a partial plan that considers the information most relevant to your current circumstance, and delayed thinking about other details to when they become more relevant. For instance, rather than thinking of a detailed path from your bedroom to the flight gate right now, you might *plan to later think about* the route to the gate *once you get through airport security*.

As this example suggests, human decision making not only involves planning one’s actions, but also planning one’s

plans. But why plan one’s plans? Why not just plan one’s actions? We propose that planning one’s plans emerges from two aspects of human decision making. First, plans can include different details at different times and change as time moves on. Second, representing detailed plans is itself costly (in terms of time and memory). Thus, people should consider what details they include in a plan and when to include those details. That is, they should plan their plans.

Here, we develop and empirically evaluate this idea in humans. We first discuss how planning one’s plans relates to existing ideas in psychology and machine learning. Then, we formalize the relation between planning one’s plans and the cost of planning by defining a general Bellman objective that includes both task rewards and information-theoretic planning costs. We then discuss several qualitative features of solutions to this objective. Finally, we report the results of two new human experiments that compare participant reaction times during problem solving to the predictions of our model and alternatives. These empirical findings support our normative account of planning one’s plans and raise new questions about the nature of boundedly rational decision making in both people and machines.

Background

Put simply, planning involves finding a good sequence of actions given a particular problem representation (Newell, Shaw, and Simon 1958). In computer science, many approaches have been developed to facilitate planning. These include classical planning methods such as depth-limited search (Korf 1990), heuristic search (Pearl 1984), and Monte-Carlo tree search (Abramson 1987) as well as the use of data structures like hierarchies (Sacerdoti 1974; Kaelbling and Lozano-Perez 2010), temporal abstractions (McDermott et al. 1998), and state abstractions (Givan, Dean, and Greig 2003). At the same time, psychologists have long recognized that people also rely on heuristics (Tversky and Kahneman 1974; Gigerenzer and Goldstein 1996) and use abstractions to organize their thoughts and behaviors (Lashley 1951; Rosenbaum, Inhoff, and Gordon 1984; Solway et al. 2014).

But why do people use heuristics and abstractions, and why do we build these structures into our algorithms to use? A simple reason is that naïve planning is prohibitively

costly (Bellman 1957; Littman, Dean, and Kaelbling 1995), so such aids focus limited computational resources on the most important, urgent, or relevant parts of a problem. This observation motivates the following question: What if usefully and efficiently representing plans were an explicit goal of making decisions over time?

To develop a model of adaptively planning with the costs of planning in mind (i.e., planning one’s plans), we draw on ideas from several lines of research. One is work on boundedly optimal intelligence (Russell and Subramanian 1995), anytime algorithms (Horvitz 1990; Dean and Boddy 1988), and human rational meta-reasoning (Griffiths, Lieder, and Goodman 2015), which articulates the value of managing computational costs during decision making. Another is work on the psychology of intertemporal choice (Berns, Laibson, and Loewenstein 2007) and representation (Trope and Liberman 2003), which studies how people’s construal of different aspects of the world changes as a function of distance, time, and context. Finally, we draw on formal tools from information-theoretic approaches to bounded rationality (Tishby and Polani 2011; Rubin, Shamir, and Tishby 2012; Ortega and Braun 2013), which provides a general framework for characterizing the relationship between environmental rewards and decision-making costs.

Planned Information Processing

We begin by formalizing the objective of planning one’s plans in terms of planned information processing. To model how an agent should plan at different points in time given a task and a cost of planning, we treat the decisions of what to plan and when to plan as a sequential decision-making problem (Puterman 1994). To capture a general, architecture-agnostic notion of planning costs, we use a formulation of partial planning and information-theoretic action costs (Tishby and Polani 2011; Rubin, Shamir, and Tishby 2012; Ortega and Braun 2013). Finally, at the end of this section and in the supplementary materials, we present a gradient-based algorithm for solving this objective.

Markov Decision Processes

Sequential decision-making can be described in terms of a Markov Decision Process (MDP) (Puterman 1994). A discrete MDP M is a tuple $\langle S, A, T, R, \gamma \rangle$, where S is the state space; A is the action space; $T : S \times A \rightarrow \Delta(S)$ is a transition function that defines a probability distribution over next states $s' \in S$ having taken action $a \in A$ in state $s \in S$; $R : S \times A \times S \rightarrow \mathbb{R}$ is a reward function that defines agent payoffs; and $\gamma \in [0, 1)$ is a discount rate.¹

An agent’s *policy* describes its behavior in an MDP. Formally, a policy $\pi : S \rightarrow \Delta(A)$ is a mapping from states $s \in S$ to probability distributions over actions $a \in A$. A policy combined with a transition function defines how an agent will move within the state space over time. Additionally, the *normalized discounted occupancy* of a policy π in MDP M is $\rho^\pi(s) \propto \sum_{t=0}^{\infty} \gamma \Pr\{s = s_t\}$.

We are interested in the *value function* for a particular policy, $V^\pi : S \rightarrow \mathbb{R}$, which is the expected discounted cu-

mulative reward that an agent receives by following π from state s onward. We are especially interested in the *optimal value function* defined by the unique fixed point of the Bellman equation (Bellman 1957), where for all $s \in S$:

$$V^*(s) = \max_{a \in A} \sum_{s' \in S} T_{s,s'}^a [R_{s,s'}^a + \gamma V^*(s')]. \quad (1)$$

The optimal value function describes the best that an agent can expect to do in terms of maximizing future discounted rewards from each state. It is also useful to define the optimal state-action value function $Q^*(s, a) = \sum_{s' \in S} T_{s,s'}^a [R_{s,s'}^a + \gamma V^*(s')]$, for all $s \in S, a \in A$. Finally, an optimal policy π^* is any policy π such that $V^\pi(s) = V^*(s)$ for all $s \in S$.

Partial Planning via Soft Planning

In an MDP, *planning* corresponds to finding a policy that yields high value by doing computations over a *model* of the task. For instance, an optimal policy π^* generates plans that perfectly maximize value. But, we may not always want a perfect plan. Rather, since perfection is costly, it is often useful to express imperfect *partial plans*.

We introduce two ideas to formalize partial plans. First, we distinguish between the ground MDP, M , and a “simulated” MDP, \tilde{M} . Here, we focus on the relationship between ground states S and simulated states \tilde{S} , while also assuming that $\tilde{M} = M$. This distinction allows us to express different quantities (e.g., action probabilities, expected values) over the same simulated state space but from the perspective of different ground states. For example, we denote the planned probability of taking action a at simulated state \tilde{s} from ground state s as $\tilde{\pi}(a \mid \tilde{s}; s)$. In the airport example, you are simulating what you would do at the airport, \tilde{s} , from your bedroom, s . Note the special case of $\tilde{\pi}(a \mid s; s)$ (i.e., $\tilde{s} = s$), which defines the actions an agent plans on taking at their current state.

Second, we introduce a *soft planning parameterization* of partial policies that controls the allocation of planning to different simulated states. Formally, an *inverse temperature assignment* from state $s, \beta(\cdot; s) : \tilde{S} \rightarrow \mathbb{R}_{\geq 0}$, assigns a positive real inverse temperature to each simulated state $\tilde{s} \in \tilde{S}$. Given this assignment, we define soft-Bellman equations over simulated states $\tilde{s} \in \tilde{S}$ from a ground state s :

$$\tilde{\pi}^\beta(a \mid \tilde{s}; s) \propto \exp \left\{ \tilde{Q}^\beta(\tilde{s}, a; s) \beta(\tilde{s}; s) \right\}, \quad (2)$$

$$\tilde{V}^\beta(\tilde{s}; s) = \sum_a \tilde{\pi}^\beta(a \mid \tilde{s}; s) \tilde{Q}^\beta(\tilde{s}, a; s), \quad (3)$$

$$\tilde{Q}^\beta(\tilde{s}, a; s) = \sum_{s'} T_{s,s'}^a \left[R_{s,s'}^a + \gamma \tilde{V}^\beta(\tilde{s}'; s) \right]. \quad (4)$$

Intuitively, the inverse temperature assignment captures how much attention is paid at each simulated state when constructing a partial plan. Larger inverse temperatures entail more attention at a particular state, and the interaction of temperatures induces a partial plan $\tilde{\pi}^\beta(\cdot \mid \cdot; s)$.²

¹ $\Delta(X)$ denotes the simplex over discrete elements $x \in X$.

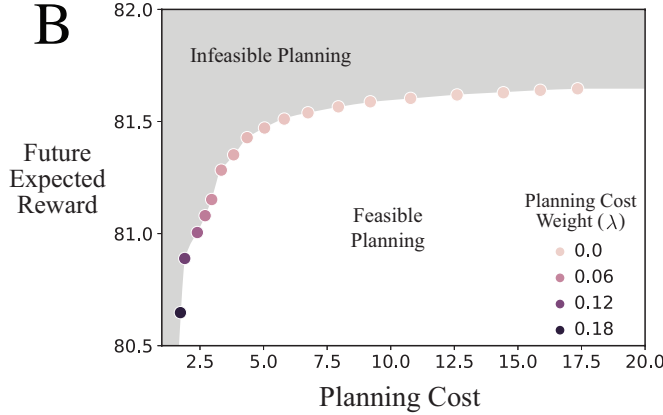
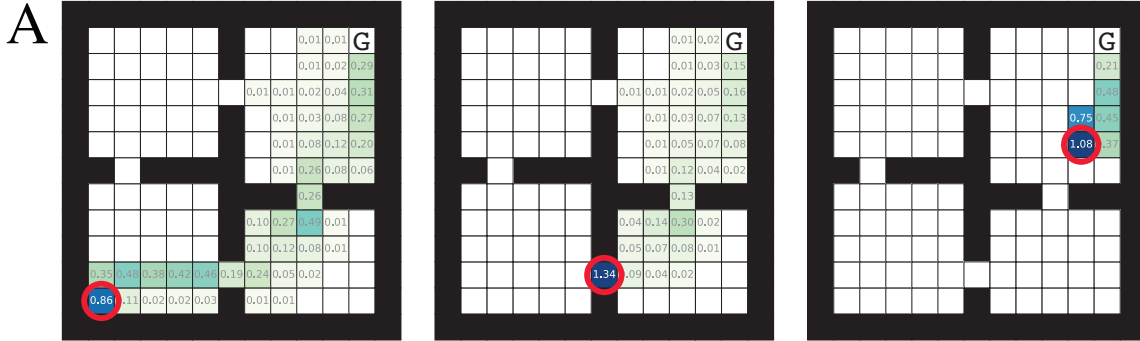


Figure 1: Results on Four Rooms Domain. (A) *Partial Planning Costs*. Shown are the planning costs associated with each simulated state \tilde{s} from three ground states s ($\tilde{s} = s$ are circled in red). Numbers correspond to the KL-divergence of the partial plan from the default plan at each simulated state. The total computational cost for a state is the sum of divergences over the simulated state space (see Equation 5). Higher values occur at states where the decision is relevant to the current decision, such as the entrances to doors. Values less than 0.005 are not shown. G denotes the goal. (B) *Reward-Planning Pareto frontier*. Planning to plan involves both task rewards and planning costs. Depending on the particular task, certain combinations of task rewards and planning costs may be infeasible. Optimizing Equation 6 for any particular λ finds a point where expected rewards cannot be increased without increasing planning costs and planning costs cannot be decreased without decreasing rewards. Plotted are the planning costs and future value (without planning costs) from the lower-left state in Four Rooms after running our algorithm for a range of λ values. The induced Pareto frontier is analogous to an information-theoretic rate-distortion tradeoff (Shannon 1948; Tishby and Polani 2011).

Information-Theoretic Planning Costs

How can we quantify the cost of a partial plan? Our proposal borrows the idea of information theoretic costs for actions (Tishby and Polani 2011; Rubin, Shamir, and Tishby 2012; Ortega and Braun 2013) and applies it to simulated (planned) actions. For instance, we can define the cost of a plan $\tilde{\pi}$ based on sum of Kullback-Leibler (KL) divergences, denoted D_{KL} , from a default policy $\bar{\pi}$ at each state:

$$C(\tilde{\pi}, \bar{\pi}) = \sum_{\tilde{s} \in \tilde{S}} D_{KL}[\tilde{\pi}(\cdot | \tilde{s}) || \bar{\pi}(\cdot | \tilde{s})], \quad (5)$$

where $D_{KL}[p||q] = \sum_x p(x) \log(\frac{p(x)}{q(x)})$ for distributions p and q with the same support (Cover and Thomas 1991). Here, we set $\bar{\pi}$ to be the uniform distribution at all states as it makes few assumptions and is justified by previous work (Ortega et al. 2015). However, our formulation

straightforwardly accommodates task-specific default policies.

Specifying planning costs as an information theoretic quantity has conceptual and practical advantages. First, planning costs can be characterized independently of an agent’s specific representation of a task. Second, it can be interpreted as the minimal cost in bits of using the old plan $\bar{\pi}$ to represent the action distributions in the new plan $\tilde{\pi}$ (Cover and Thomas 1991), which captures intuitions about planning as a costly process of information transformation and control. Finally, the cost term is differentiable with respect to the policy probabilities.

²Previous work has interpreted the soft-maximization in Eq. 2 in terms of the sampling process of a bounded rational decision-maker, where inverse temperature reflects noise (Train, 2003, p48; Ortega et al., 2015). Here, we mainly treat soft-planning as a way to parameterize partial policies and leave a more in-depth analysis

Planned Information Processing Bellman Objective

Given our formalization of partial planning and information theoretic planning costs, we can now define the objective of planning one’s plans. Put simply, we nest partial planning (Equations 2–4) inside a meta-planning problem that includes planning costs (Equation 5):

$$V_\lambda^*(s) = \max_{\beta(\cdot; s)} \left\{ \sum_{a \in A} \left[\tilde{\pi}^\beta(a|s; s) Q_\lambda^*(s, a) \right] - \lambda C(\tilde{\pi}^\beta, \bar{\pi}) \right\}, \quad (6)$$

where $Q_\lambda^*(s, a) = \sum_{s' \in S} T_{s, s'}^a \left[R_{s, s'}^a + \gamma V_\lambda^*(s') \right]$.

This equation extends the original Bellman equation (Equation 1) in two ways. First, ground actions are no longer directly chosen. Instead, a distribution over actions is induced at the current state s by influencing the computations of the planner with $\beta(\cdot; s)$.

Second, Equation 6 takes into account the immediate cost of planning via the term $\lambda C(\tilde{\pi}^\beta, \bar{\pi})$, where $\lambda \in \mathbb{R}$ is a planning cost weight. Note that this formulation also includes planning opportunities and costs *in the future* via the recursive nature of the Bellman equation. To understand the significance of this, consider again the airport example: While at home packing, you only vaguely consider how you will get from security to your gate. This is not because it does not matter; if you never make it to your gate and miss your flight, there is no point in packing. Rather, packing does not require you to think about the details of navigating the airport until later. In short, Equation 6 expresses how an agent should partially plan to act in the moment, taking into account that they will engage in planning in the future.

Optimal Planned Information Processing

To solve for the meta-planning objective in Equation 6, we implemented a gradient-based algorithm that solves for the optimal inverse temperature assignments (i.e., β^*) given an MDP M , planning cost weight λ , and default policy $\bar{\pi}$.

To understand optimal meta-planning and how it relates to previous work on hierarchical planning and learning, we ran our algorithm on the Four Rooms domain with deterministic transitions (Sutton, Precup, and Singh 1999). An absorbing goal state worth +100 points was placed in the upper-right corner. Step costs were included (−.1 points) and the discount rate was set to .99, with $\lambda = 0.01$. Planning iterations were chosen such that value iteration would converge ($H = 100$), while meta-planning parameters were solved using the Adam optimizer (Kingma and Ba 2014) for $N = 200$ iterations (see supplementary materials for details on the algorithm).

The optimization finds an inverse temperature assignment β^* that yields a partial plan from each ground state s —i.e., $\tilde{\pi}^{\beta^*}(\cdot | \cdot; s)$. The cost of each partial plan is determined by the entire simulated state space (Equation 5), but we can examine the contribution of the planned actions at each simulated state \tilde{s} to identify which ones are prioritized, shown

of the sampling interpretation of this model for future work.

as colors in Figure 1a. Two salient features emerge in the Four Rooms Domain. First, at any point in time, “doorways” leading to the goal are prioritized because the quality of the current decision depends on planning to go through the door and proceeding towards the goal. Second, actions at states that are closer along the path are represented in greater detail since they are more relevant to simulating the value of the current decision.

Additionally, Equation 6 implicitly defines a task-specific *Pareto frontier* where task rewards cannot be improved without worsening planning costs and vice versa. Since our algorithm seeks to find points that jointly maximize rewards and minimize planning costs, we can identify this curve for the Four Rooms task by running our algorithm with a range of λ values. Figure 1b shows the plot of this curve at the lower-left state, which divides the space of possible task rewards and planning costs into feasible and infeasible combinations of rewards and planning costs.

Experiment 1: Parametric Mazes

Our account makes quantitative predictions about how people should flexibly plan information processing on a task. To test these predictions, we examined how long it took people to make decisions when navigating through a set of parametrically generated 2D Gridworld mazes. For a given maze, our model predicts an optimal amount of partial planning at a state. For example, Figure 2a displays two mazes that our model predicts will have different immediate partial planning costs at the initial state. Since this cost reflects information processing, we can operationalize it in this task based on reaction times (RTs) as in previous work (e.g., Ortega and Stocker 2016). Specifically, we assume that the amount of time between when a participant is presented with a maze and when they take their first action reflects the cost of encoding a partial plan from the initial state.

Materials, Participants, and Procedure

The stimuli were a set of 50 different 9×9 Gridworld mazes in which the start state was in the lower right corner and the goal state was in the upper left corner. These were chosen by first randomly generating a batch of 2,000 mazes and then selecting a random subset that were predicted by the model to have a range of different optimal planning costs at the initial state.

We recruited 50 participants from Amazon Mechanical Turk and used the psiTurk framework (Gureckis et al. 2016). Each participant was paid a base pay of \$1.00. After reading the instructions and familiarizing themselves with the general mechanics of the task, participants started the main part of the task that included the 50 mazes. Each round, participants were first shown a blank 9×9 grid. When they pressed the spacebar, the maze for that round appeared immediately and they could move their agent (a blue circle) using the arrow keys. The initial-state RT measure was the time measured between the appearance of the maze on a round and their first action. When they reached the goal state, they received +100 points (50 points = 1¢; total bonus = \$1.00).

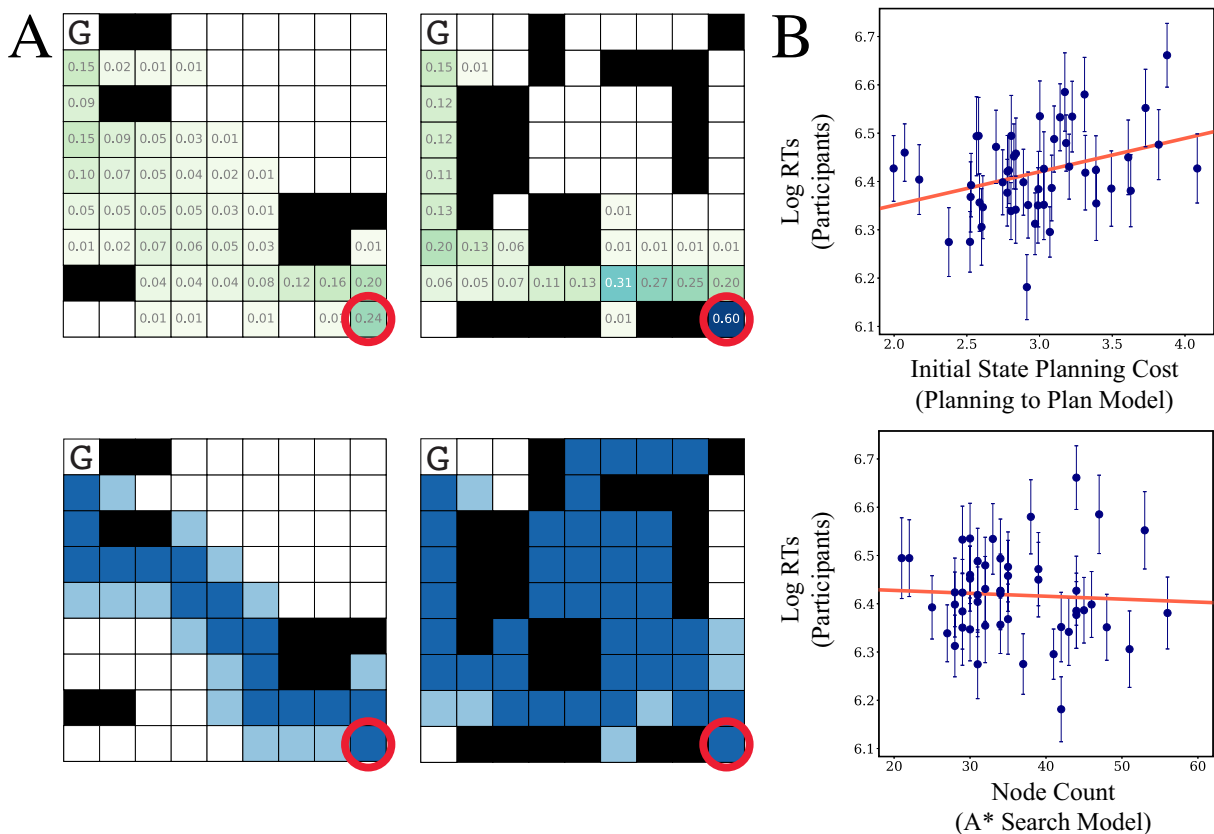


Figure 2: Experiment 1 examples and results. (A) Planned information processing (upper pair of examples) versus A* search (lower pair) on two of 50 mazes from Experiment 1. Partial plans at the initial lower-right state (red circles) are more specific (and costly) at future states relevant to the current decision, represented by darker green (top row). A* can plan with heuristics (e.g., Manhattan distance to goal) but does not adaptively allocate planning based on task structure (bottom row; dark and light blue are explored states and candidate states, respectively). If people are only planning, and not planning to plan, initial-state response times will reflect processes like A* and not partial planning. (B) Mean log-normalized RTs for the 50 grids used in the task as a function of partial planning costs (top) and A* Node Count. Red lines are regression lines for mean log RTs by item and model predictions (Top: $R^2 = 0.12$, $p < 0.05$; Bottom: $R^2 = 0.003$; $p = 0.71$). Error bars represent standard errors.

Planned Information Processing and Alternative Models

In this experiment, we are interested in the minimized information theoretic planning cost at the initial state predicted by planned information processing. Formally, this corresponds to the $C(\bar{\pi}^\beta, \bar{\pi})$ term in Equation 6 for an initial state s_0 . These values were calculated for each maze using our algorithm and the same parameters as in the simulation.

To assess whether people are not simply planning, but adaptively planning their information processing, we considered seven alternative planning-based metrics as predictors. First, we considered the length of the shortest path from the start state to the goal, calculated using value iteration (Optimal Plan Length). Second, we ran A* search (Hart, Nilsson, and Raphael 1968), a classical planning algorithm that finds a shortest path by maintaining a prioritized set of states to explore, starting with an initial state, and then iteratively exploring states and adding connected states to the exploration set until the goal state is reached. To facilitate bet-

ter exploration, we provided A* with a Manhattan distance heuristic to the goal. We considered the number of candidate states explored by the algorithm before termination (A* Node Count).

Third, we analyzed the action cost associated with the first step of the boundedly rational planning method proposed by Ortega et al., 2015 (Information theoretic bounded rationality). We set the information theoretic cost to be $\frac{1}{\alpha} = \lambda = .01$ to be commensurate with our planning to plan implementation. Fourth, we calculated the initial state entropy of a standard softmax over the optimal value function, with $\beta = 1$ (RL Softmax Entropy). Fifth, we calculated the initial state entropy of an optimal soft-Bellman policy, with $\beta = 1$ (Soft-Bellman Entropy). Sixth, we calculated the number of iterations of standard (planning) value iteration before convergence as a measure of planning computation (VI Iterations). Finally, as a heuristic measure of the complexity of a grid, we calculated the mean number of “turns” that occurred along a trajectory sampled from the optimal policy

Predictor	LL Ratio [$\chi^2(1)$]	p	β	SE
Partial Planning Cost	37.71	< .0001	0.15	0.02
A* Node Count	18.60	< .0001	0.00	0.00
Optimal Plan Length	14.06	< .001	-0.02	0.01
Information Theoretic Bounded Rationality	0.82	0.36	0.00	0.00
RL Softmax Entropy	1.50	0.22	0.32	0.26
Soft-Bellman Entropy	10.38	< .001	-0.67	0.21
VI Iterations	24.12	< .0001	0.01	0.00
Trajectory Turns	13.93	< .001	0.02	0.00

Table 1: Experiment 1 (Parametric Mazes) likelihood ratio tests and model estimates. Even when multiple planning metrics are included, the Partial Plan Cost derived from planning to plan is predictive of RTs.

(Trajectory Turns).

If people are only planning actions, then the planning-based metrics should be sufficient for predicting RTs. If people are planning information processing—that is, constructing a computationally inexpensive partial plan that provides a good action at their current state—then the partial planning cost in our model would separately predict RTs.

Results and Discussion

We analyzed our data by comparing the predictions of the models to participants’ initial-state RTs. Two participants had substantial missing data, and outliers were excluded, which left 2373 initial-state RT measurements. To assess the relative predictive power of the eight models, we first fit a fully specified mixed effects linear model to log-normalized RTs. This included by-participant intercepts and round number slopes as random effects, and Partial Planning Costs as well as the seven planning-metrics as fixed effects. We then performed log-likelihood ratio tests with lesions versions that did not include each of the eight fixed effects. As shown in Table 1, although several of the planning metrics are significant predictors, Partial Planning Cost not only predicts RTs, it is the predictor with the highest log-likelihood ratio test statistic. Thus, Experiment 1 suggests that people engage in planned information processing.

Experiment 2: Probing Partial Plans

Experiment 1 tested planning information processing using initial-state RTs to measure planning costs. In Experiment 2, we more directly examine whether people’s partial planning is captured by our model. To probe partial plans, we used a technique of teleporting participants’ avatars in Gridworld mazes and measuring their reactions. For instance, imagine that, in the middle of packing your bag for a flight you are unexpectedly teleported to the main terminal of the airport with your bag. It is likely you would quickly know what to do next (e.g., pick up your boarding pass) if that were part of your partial plan prior to being teleported. In contrast, if you had not thought that far ahead, then it would likely take you longer to determine your next action. Thus, *post-teleportation reaction times* can be used to measure the divergence between a pre-teleportation plan and a post-teleportation plan.

Materials, Participants, and Procedure

The experiment consisted of 64 rounds of Gridworld mazes. To generate the mazes, four base 12×12 mazes were generated such that the initial state was in the lower right corner and the goal state was in the upper left corner. These were then transformed using the eight symmetries of a square, yielding a total of 32 perceptually distinct mazes. Each of the 32 mazes appeared twice. Half of the rounds were Normal rounds while the other half were Teleportation rounds. On the Teleportation rounds, a random number n between 1 and the length of an optimal path for a maze was chosen, and on the n -th trial, the agent was hidden for 750ms and could not be controlled. It then reappeared in a randomly chosen location in the maze and could be controlled immediately. The amount of time between the reappearance of the circle and the participant’s response was the post-teleportation RT measure that is the focus of this study.

Sixty participants from MTurk were recruited for our experiment and given the same familiarization procedure as in Experiment 1. Participants were paid a base pay of \$1.00 and received a bonus of \$1.28 for completing the 64 trials.

Model-Predictor Variables

Our goal is to explain post-teleportation RTs as a function of pre- and post-teleportation states and task structure. Our account provides partial plans at the pre- and post-teleportation states. If these map onto people’s partial plans, then RTs will reflect a process of updating the pre-teleportation plan into the post-teleportation plan. To quantify this updating (i.e., re-planning) process, we calculated the state–action divergence between the pre-teleportation partial plan $\tilde{\pi}$ and post-teleportation partial plan $\tilde{\pi}'$, $D_{\text{KL}}[p^{\tilde{\pi}'}(a, s) || p^{\tilde{\pi}}(a, s)]$, where $p^{\pi}(a, s) = \pi(a | s)\rho^{\pi}(s)$. This “Partial-Plan Divergence” reflects the cost to encode the state–action distribution of the partial plan at the post-teleportation state starting from the one at the pre-teleportation state. It should thus reflect participants’ “new” planning at the post-teleportation state. The same parameters as in the model in Experiment 1 were used to calculate the partial plans.

We calculated several alternative planning measures. First, we calculated the length of the optimal path from the post-teleportation state to the goal. Second, we calculated the number of A* nodes from the post-teleportation state (A* Destination Nodes). Third, we calculated an A* Node Difference score, corresponding to the additional nodes that A*

Predictor	LL Ratio [$\chi^2(1)$]	p	β	SE
Partial-Plan Divergence	24.70	< .001	0.001	0.0003
A* Destination Nodes	0.25	.62	-0.002	0.005
A* Node Difference	2.80	.09	-0.003	0.002
Optimal Path Length	0.02	.89	0.001	0.01

Table 2: Experiment 2 (Probing Partial Plans) likelihood ratio tests and model estimates.

explores at the post-teleportation state, taking into account those already explored at the pre-teleportation state.

Results and Discussion

To assess the influence of the different predictors on post-teleportation log-normalized RTs, we used a similar mixed-effects linear model as in Experiment 1 (random effects were by-participant intercepts and round number, by-maze intercepts; fixed effects were model predictors). As summarized in Table 2, the Partial-Plan Divergence significantly predicted log-transformed RTs on post-teleportation trials. The planning-based predictors did not account for how quickly people reacted after being teleported.

Additionally, we conducted a separate analysis that included teleportation distance as a fixed effect in the full model. Note that unlike the planning and partial planning models, teleportation distance is not an explicit model of decision-making. In this new model, partial planning is significant but weakened ($\chi^2(1) = 4.8, p < .05$). Additional details are included in the Supplementary Materials.

Thus, overall, the results of Experiment 2 suggest that people’s RTs are explained by planned information processing via partial planning and not simply by planning actions.

General Discussion

This paper asks two questions. First, why plan one’s information processing? We argue that meta-planning lets agents adaptively capitalize on the benefits of planning while regulating planning costs. To make this precise, we formalize the general notion of *partial plans* that prioritize planning in different parts of a simulated model and define an *information-theoretic encoding cost* for partial plans, enabling us to define a novel recursive Bellman objective that includes both task rewards and planning costs. This model provides a point

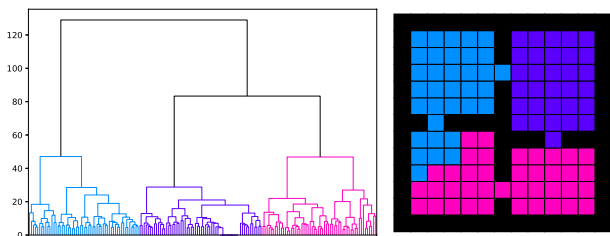


Figure 3: States clustered by partial planning similarity resemble option representations.

of departure for future normative accounts of human meta-planning.

Second, do people plan their information processing? We reported two human experiments that test our formal account of planned information processing. Experiment 1 demonstrates that adaptive partial planning explains people’s initial reaction times when navigating parametrically generated 2D mazes. Experiment 2 used unexpected teleportations while navigating mazes to probe partial planning representations. The optimal partial plans generated by our model explain human responses even when accounting for action planning.

People plan because planning is useful. But, planning is hard, so people make planning easier by being selective about what and when they plan. In other words, people should plan their planning. For the most part, current decision-making algorithms plan, albeit with the help of good heuristics and abstractions provided by computer scientists. But, ideally, algorithms would learn how to make planning easier for themselves by planning their planning. Understanding planned use of computational resources can also provide insight into the nature and function of abstractions when *learning*. For instance, we performed hierarchical clustering over states in Four Rooms based on the similarity of their optimal partial plans (Figure 3; details in supplementary materials), which results in clusters resembling options from research on hierarchical reinforcement learning (Sutton, Precup, and Singh 1999; Dietterich 2000; Parr and Russell 1998; Botvinick 2008). In short, this work is an important step towards understanding the scale and sophistication of human meta-planning and applying such insights to the design of machines.

Acknowledgements

The authors would like to thank Daniel Reichman, Bill Thompson, Fred Callaway, and Rachit Dubey for their advice and feedback on this work. This research was supported by NSF grant #1544924, AFOSR grant #FA9550-18-1-0077, and grant #61454 from the John Templeton Foundation.

References

- [Abramson 1987] Abramson, B. D. 1987. *The Expected-Outcome Model of Two-Player Games*. Ph.D. Dissertation, Columbia University, New York, NY, USA.
- [Bellman 1957] Bellman, R. 1957. *Dynamic programming*. Princeton University Press.
- [Berns, Laibson, and Loewenstein 2007] Berns, G. S.; Laibson, D.; and Loewenstein, G. 2007. Intertemporal choice—

- toward an integrative framework. *Trends in Cognitive Sciences* 11(11):482–488.
- [Botvinick 2008] Botvinick, M. M. 2008. Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences* 12(5):201–208.
- [Cover and Thomas 1991] Cover, T. M., and Thomas, J. A. 1991. *Elements of Information Theory*. Wiley Series in Telecommunications. New York, USA: John Wiley & Sons, Inc.
- [Dean and Boddy 1988] Dean, T. L., and Boddy, M. S. 1988. An analysis of time-dependent planning. In *AAAI*, volume 88, 49–54.
- [Dietterich 2000] Dietterich, T. G. 2000. Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of Artificial Intelligence Research* 13:227–303.
- [Gigerenzer and Goldstein 1996] Gigerenzer, G., and Goldstein, D. G. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103(4):650–669.
- [Givan, Dean, and Greig 2003] Givan, R.; Dean, T.; and Greig, M. 2003. Equivalence notions and model minimization in markov decision processes. *Artificial Intelligence* 147(1-2):163–223.
- [Griffiths, Lieder, and Goodman 2015] Griffiths, T. L.; Lieder, F.; and Goodman, N. D. 2015. Rational Use of Cognitive Resources: Levels of Analysis Between the Computational and the Algorithmic. *Topics in Cognitive Science* 7(2):217–229.
- [Gureckis et al. 2016] Gureckis, T. M.; Martin, J.; McDonnell, J.; Rich, A. S.; Markant, D.; Coenen, A.; Halpern, D.; Hamrick, J. B.; and Chan, P. 2016. psiturk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods* 48(3):829–842.
- [Hart, Nilsson, and Raphael 1968] Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics* 4(2):100–107.
- [Horvitz 1990] Horvitz, E. J. 1990. *Computation and action under bounded resources*. Ph.D. Dissertation, Stanford University, California.
- [Kaelbling and Lozano-Perez 2010] Kaelbling, L. P., and Lozano-Perez, T. 2010. Hierarchical Planning in the Now. *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- [Kingma and Ba 2014] Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [Korf 1990] Korf, R. E. 1990. Real-time heuristic search. *Artificial Intelligence* 42(2):189–211.
- [Lashley 1951] Lashley, K. S. 1951. The problem of serial order in behavior. In *Cerebral mechanisms in behavior; the Hixon Symposium*. Oxford, England: Wiley. 112–146.
- [Littman, Dean, and Kaelbling 1995] Littman, M. L.; Dean, T. L.; and Kaelbling, L. P. 1995. On the complexity of solving markov decision problems. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, 394–402. Morgan Kaufmann Publishers Inc.
- [McDermott et al. 1998] McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1998. Pddl—the planning domain definition language.
- [Newell, Shaw, and Simon 1958] Newell, A.; Shaw, J. C.; and Simon, H. A. 1958. Elements of a theory of human problem solving. *Psychological Review* 65(3):151–166.
- [Ortega and Braun 2013] Ortega, P. A., and Braun, D. A. 2013. Thermodynamics as a theory of decision-making with information-processing costs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 469(2153):20120683–20120683.
- [Ortega and Stocker 2016] Ortega, P. A., and Stocker, A. A. 2016. Human decision-making under limited time. In Lee, D. D.; Sugiyama, M.; Luxburg, U. V.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems* 29. Curran Associates, Inc. 100–108.
- [Ortega et al. 2015] Ortega, P. A.; Braun, D. A.; Dyer, J.; Kim, K.-E.; and Tishby, N. 2015. Information-Theoretic Bounded Rationality.
- [Parr and Russell 1998] Parr, R., and Russell, S. J. 1998. Reinforcement learning with hierarchies of machines. In *Advances in neural information processing systems*, 1043–1049.
- [Pearl 1984] Pearl, J. 1984. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley.
- [Puterman 1994] Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.
- [Rosenbaum, Inhoff, and Gordon 1984] Rosenbaum, D. A.; Inhoff, A. W.; and Gordon, A. M. 1984. Choosing between movement sequences: A hierarchical editor model. *Journal of Experimental Psychology: General* 113(3):372–393.
- [Rubin, Shamir, and Tishby 2012] Rubin, J.; Shamir, O.; and Tishby, N. 2012. Trading Value and Information in MDPs. Springer, Berlin, Heidelberg. 57–74.
- [Russell and Subramanian 1995] Russell, S. J., and Subramanian, D. 1995. Provably Bounded-Optimal Agents. *Journal of Artificial Intelligence Research* 2:575–609.
- [Sacerdoti 1974] Sacerdoti, E. D. 1974. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence* 5(2):115–135.
- [Shannon 1948] Shannon, C. E. 1948. A mathematical theory of communication. *Bell system technical journal* 27(3):379–423.
- [Solway et al. 2014] 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.
- [Sutton, Precup, and Singh 1999] Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence* 112(1-2):181–211.

- [Tishby and Polani 2011] Tishby, N., and Polani, D. 2011. Information Theory of Decisions and Actions. In *Perception-Action Cycle*. New York, NY: Springer New York. 601–636.
- [Train 2003] Train, K. 2003. *Discrete Choice Methods with Simulation*. Cambridge University Press.
- [Trope and Liberman 2003] Trope, Y., and Liberman, N. 2003. Temporal construal. *Psychological review* 110(3):403.
- [Tversky and Kahneman 1974] Tversky, A., and Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. *Science* 185(4157):1124–1131.

Supplemental Materials

Solving Planning to Plan

Algorithm 1: Gradient-based Planning to Plan

Input : MDP M , Planning cost weight λ , Default policy $\bar{\pi}$,
 Planning to Plan Iterations N , Partial Plan Horizon H

Output: Temperature allocation β , Partial Plans $\tilde{\pi}^{\beta^*}$

```

1  $\tilde{S} \leftarrow S$ ;
  /* Initialize temperatures */ */
2  $\beta_0(\tilde{s}; s) \leftarrow$  Randomly initialize for all  $s \in S, \tilde{s} \in \tilde{S}$ ;
  /* Planning to plan loop */ */
3 for  $n = 0$  to  $N$  do
  /* Partial planning loop */ */
4    $\tilde{q}_0^{\beta_n}(\tilde{s}, a; s) \leftarrow 0 \quad \forall a, s, \tilde{s}$ ;
5   for  $t = 0$  to  $H$  do
6      $\tilde{\pi}_t^{\beta_n}(a | \tilde{s}; s) \propto \exp\{\tilde{q}_t^{\beta_n}(\tilde{s}, a; s)\beta_n(\tilde{s}; s)\} \quad \forall a, s, \tilde{s}$ ;
7      $\tilde{v}_t^{\beta_n}(\tilde{s}; s) \leftarrow \sum_a \tilde{\pi}_t^{\beta_n}(a | \tilde{s}; s)\tilde{q}_t^{\beta_n}(\tilde{s}, a; s) \quad \forall s, \tilde{s}$ ;
8      $\tilde{q}_{t+1}^{\beta_n}(\tilde{s}, a; s) \leftarrow \sum_{\tilde{s}'} T_{\tilde{s}, \tilde{s}'}^a [R_{\tilde{s}, \tilde{s}'}^a + \gamma \tilde{v}_t^{\beta_n}(\tilde{s}; s)] \quad \forall a, s, \tilde{s}$ 
9   end
  /* Calculate information theoretic costs of partial plans */ */
10   $c_n(\tilde{s}; s) \leftarrow D_{\text{KL}}[\tilde{\pi}_H^{\beta_n}(\cdot | \tilde{s}; s) || \bar{\pi}(\cdot | \tilde{s})] \quad \forall s, \tilde{s}$ ;
11   $c_n(s) \leftarrow \sum_{\tilde{s}} c_n(\tilde{s}; s) \quad \forall s$ ;
  /* Calculate Planning to Plan Value (i.e. Policy Evaluation) */ */
12   $V_\lambda^{\beta_n}(s) \leftarrow \sum_a \left[ \tilde{\pi}_H^{\beta_n}(a | s; s) \sum_{s'} T_{s, s'}^a [R_{s, s'}^a + \gamma V_\lambda^{\beta_n}(s')] \right] - \lambda c_n(s) \quad \forall s$ ;
  /* Calculate Loss and Gradient, then Update Temperatures */ */
13   $L \leftarrow -\sum_s V_\lambda^{\beta_n}(s)$ ;
14   $\delta_n \leftarrow \frac{\partial L}{\partial \beta_n}$ ;
15   $\beta_{n+1} \leftarrow \text{Adam}(\beta_n, \delta_n)$ 
16 end
17  $\beta^* \leftarrow \beta_N$ ;
18  $\tilde{\pi}^{\beta^*} \leftarrow \tilde{\pi}_H^{\beta_N}$ ;
19 return  $\beta^*, \tilde{\pi}^{\beta^*}$ 

```

Solving the planning-to-plan objective allows us to understand its qualitative features as well as derive predictions of human decision-making. Algorithm 1 describes a gradient-based procedure for solving this problem, and conveys the main ideas of our account procedurally. The output of an inner loop of partial planning (lines 4 to 9) is evaluated based on the information-theoretic cost of plans (lines 10 and 11) and sequential decision-making rewards (line 12). This is nested within an outer Planning to Plan loop (lines 3 to 16) that optimizes the partial plans.

Deriving Option-like Representations

Our account concerns planning, but also sheds light on the nature of representations that facilitate good learning. We examined how option-like representations can emerge from the optimal partial planning process by clustering states based on their planning similarity.

To examine how the plans at different states in Four Rooms relate to one another, we calculated a *symmetric planning distance* for each pair of states. Specifically, for each pair of ground states s_A and s_B , we calculated

$$d(s_A, s_B) = \sum_{\tilde{s} \in \tilde{S}} D_{\text{KL}}[\tilde{\pi}^{\beta^*}(\cdot | \tilde{s}; s_A) || \tilde{\pi}^{\beta^*}(\cdot | \tilde{s}; s_B)] + D_{\text{KL}}[\tilde{\pi}^{\beta^*}(\cdot | \tilde{s}; s_B) || \tilde{\pi}^{\beta^*}(\cdot | \tilde{s}; s_A)]. \quad (7)$$

We then performed hierarchical clustering using Ward’s method with the distance matrix d . Figure 4 shows the results of hierarchical clustering. Figure 5 shows the largest three clusters. The first cluster is the room containing the goal, while the

other two clusters each contain one of the two intermediate rooms and half of the starting room. Although our goal here is to better understand *meta-planning*, we note that these clusters are highly reminiscent of options used in hierarchical reinforcement learning.

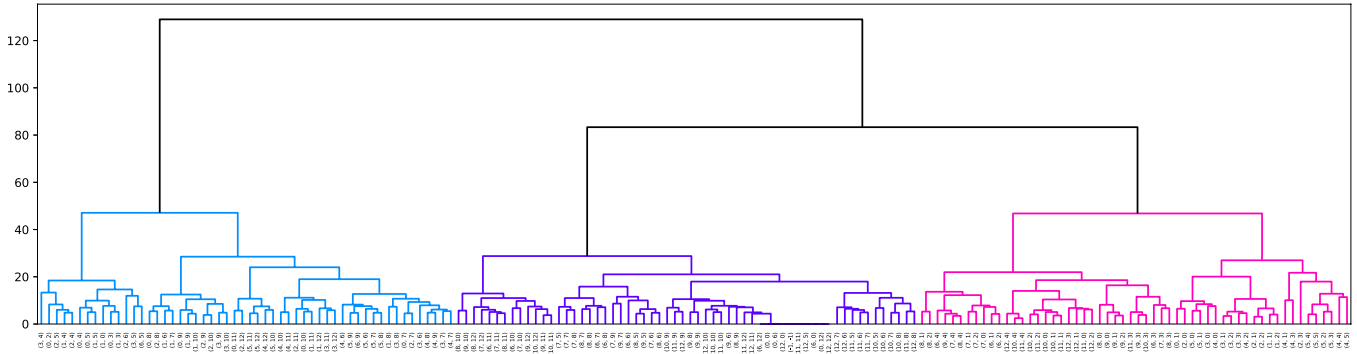


Figure 4: State dendrogram based on planning similarity in Four Rooms example

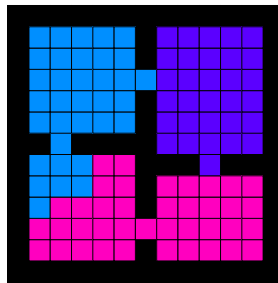


Figure 5: Top three clusters.

Experiment 2 Additional Analysis

Predictor	LL Ratio [$\chi^2(1)$]	p	β	SE
Partial-Plan Divergence	4.79	< .05	0.0007	0.0003
A* Distance to Goal	0.18	.67	-0.002	0.005
Teleportation Distance	15.83	< .001	0.01	0.003
A* Node Difference	5.38	< .05	-0.004	0.002
Optimal Path Length	0.23	.63	0.005	0.01

Table 3: Experiment 2 (Probing Partial Plans) likelihood ratio tests and model estimates where full model includes teleportation distance.

For Experiment 2, we conducted a secondary analysis in which the full mixed-effects linear model included the partial-plan divergence, A* Distance to Goal, A* Node difference, optimal path length, and teleportation distance. Unlike the other metrics, which are based on a planning model, the teleportation distance metric is derived by taking the Euclidean distance between the pre-teleportation and post-teleportation state. In this model, teleportation distance is strongly predictive, suggesting that future work is needed to disentangle the contributions of planning, meta-planning, and other processes in predicting human RTs.

Experimental Stimuli

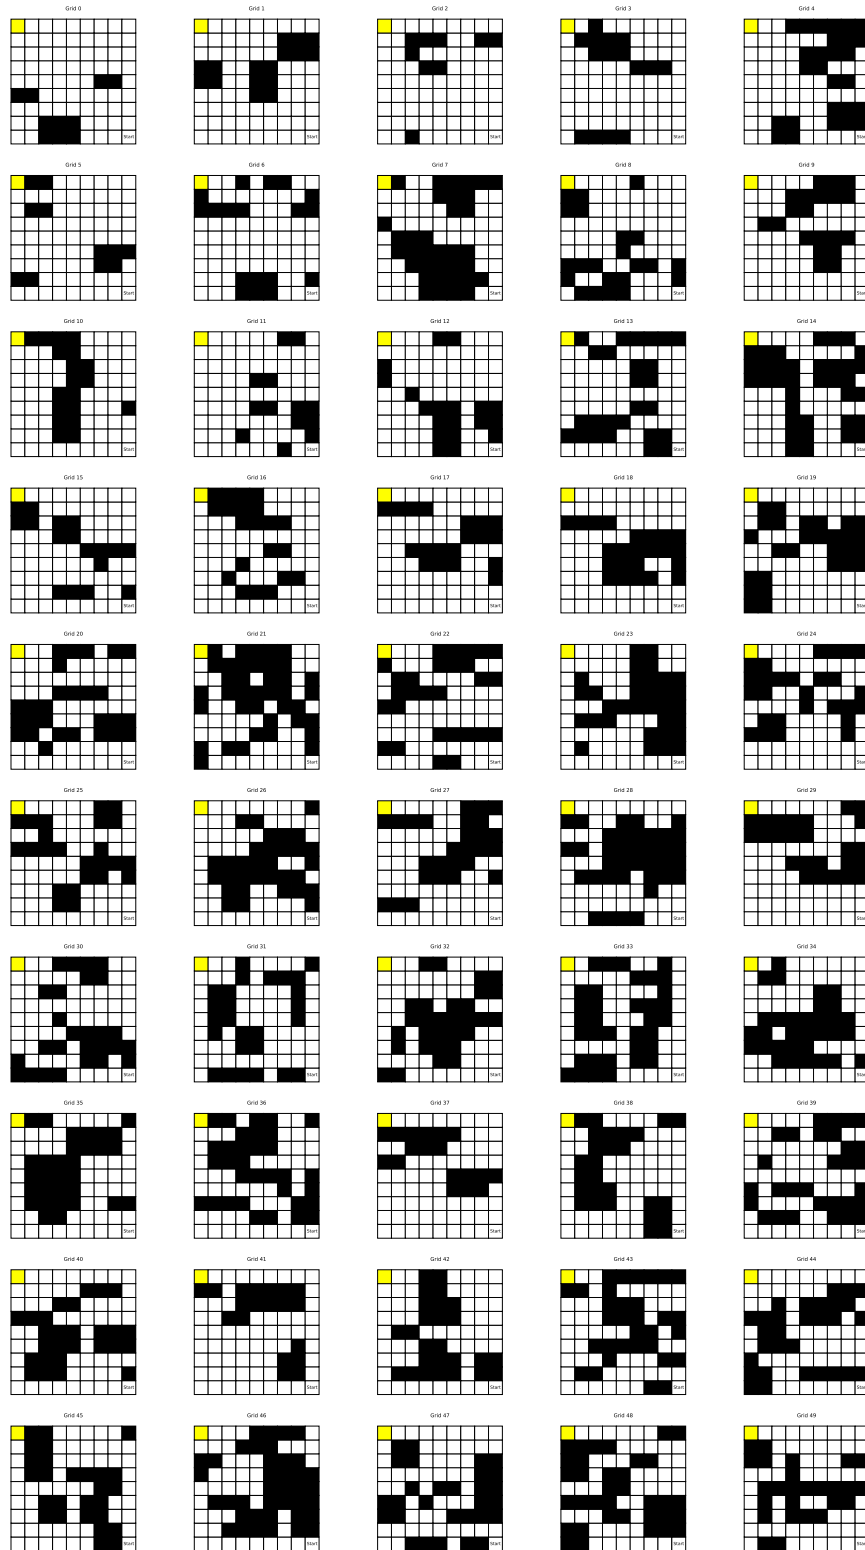


Figure 6: Mazes used in experiment 1.

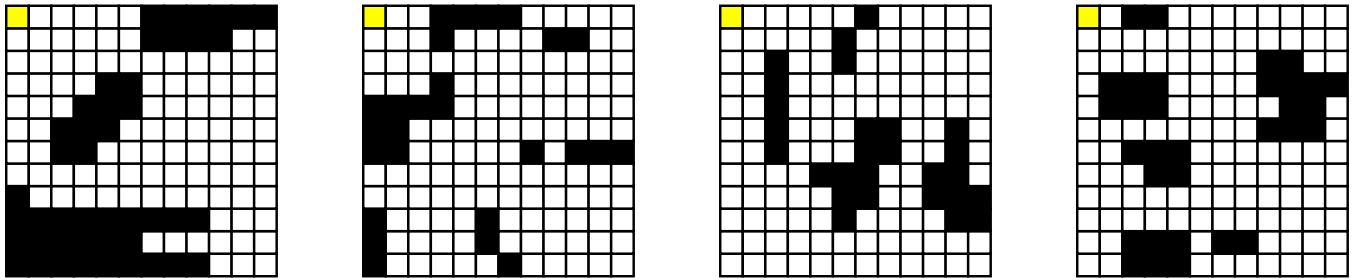


Figure 7: Mazes used in experiment 2. Start state was in the lower right. Note that the complete set of stimuli included eight transformations of these mazes, resulting in 32 unique mazes, each of which appeared once as a Teleportation round and once as a Normal round.